Finding Work (*Arming Yourself*) in the AI Revolution

Keith L. Downing

Department of Computer Science

NTNU

Signs of the Times ??

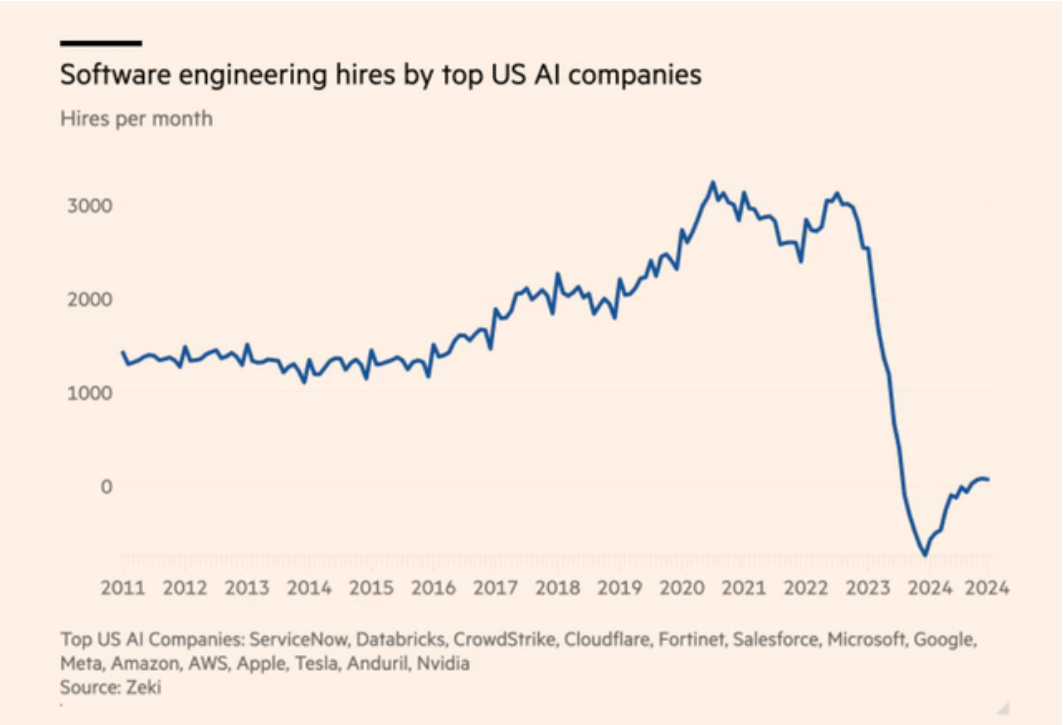

Tech Sector Layoffs

2024: 653 people / day 2025: 495 people / day

Microsoft Layoffs (June, 2025)

10,000 = 5% of workforce

Meta Layoffs (Feb, 2025)

4,000 = 5% of workforce

The view from the frontlines of AI is sobering. Two-plus years spent embedded in this world have left me feeling petrified by the sweeping changes about to strike the global labor market. An economic earthquake is coming that will permanently alter the landscape of human work—yet few in power are recognizing what's happening, let alone doing anything about it.

The problem is that we are often missing the point when we think about how AI factors into the future of work. The central question isn't if AI will transform our economy—it will—but whether we'll harness it to augment human potential rather than simply replace it. Will we seize this moment to create more fulfilling work and widely shared prosperity, or allow its benefits to flow primarily to those who own the algorithms? There's no immutable law that says automating labor means eliminating human workers.

(Gary Rivlin, Pulitzer-Prize winning journalist)

Generative AI and the Future of Work in America
(McKinsey Global Institute, July 2023)

By 2030:

- **30%** of today's human work could be automated, much by Gen AI.
- This will **enhance** worker productivity in STEM, creative, business + legal, education+training areas (without costing jobs).
- But will **reduce** jobs in office support, customer and food service.
- **+ 17%** (9.9 M jobs) – healthcare, digitization, last-mile delivery.
- **+ 7%** (2.8 M jobs) – traditional infrastructure, net-zero-emissions transition, education for worker reskilling.
- **- 10%** (6 M jobs) – Automation adoption, accelerating e-commerce, fewer customer-facing roles.

It's tough to make predictions, especially about the future...Yogi Berra

How AI is Changing the Future of Work **(Forbes, June 2023)**

- ***I believe we will see a paradigm shift from machines assisting humans to humans assisting machines in the next decade.*** Joanne Chen, Foundation Capital.
- ***AI's impact will be as big as the Industrial Revolution, empowering one to do the work of many.*** Ann Bordetsky, NEA.
- ***White collar job functions, from coding to financial services to legal, will become vastly more efficient.*** Lars Albright, Unusual Ventures.
- ***AI is already reducing the number of junior and mid-level software developers needed to ship new products.*** John Spindler, Twin Path Ventures.
- ***AI transforms the way organizations make critical decisions.*** Ben Blume, Atomico.
- ***AI will impact our day-to-day work, first gradually, and then suddenly.*** Ramy Adeeb, 1984 Ventures.

Part 1: AI History and Trends

History is a gallery of pictures in which there are few originals and many copies
… Alexis de Tocqueville

The goal of science is not universal truth…but the gradual removal of prejudice (Niels Bohr).

The Prejudice	The Science
Earth is center of the universe	Copernican Revolution
Humans are a special species	Darwinian Evolution
Superiority of human intelligence	**AI**

Bold Beginnings (Dartmouth, 1956)

Every aspect of learning or any other feature of intelligence can be so precisely described that a machine can be made to simulate it.

Roller Coaster History of AI

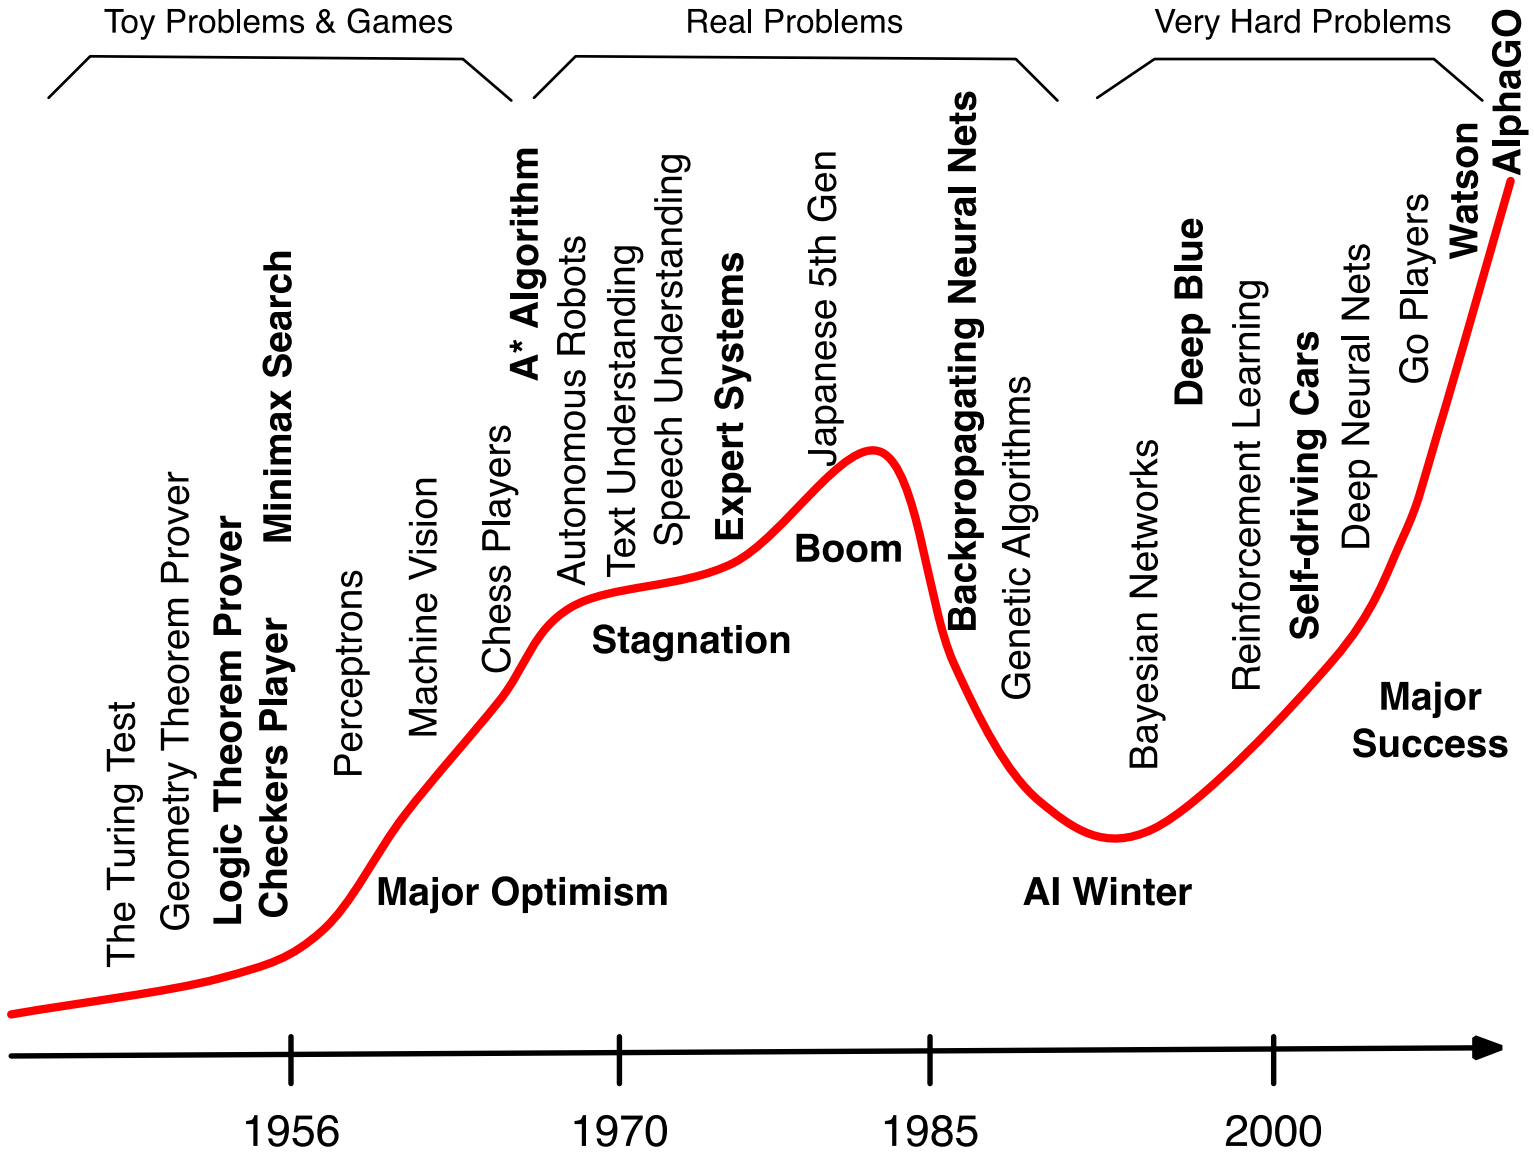

Brief History of AI Success

AI Beats or Equals Humans

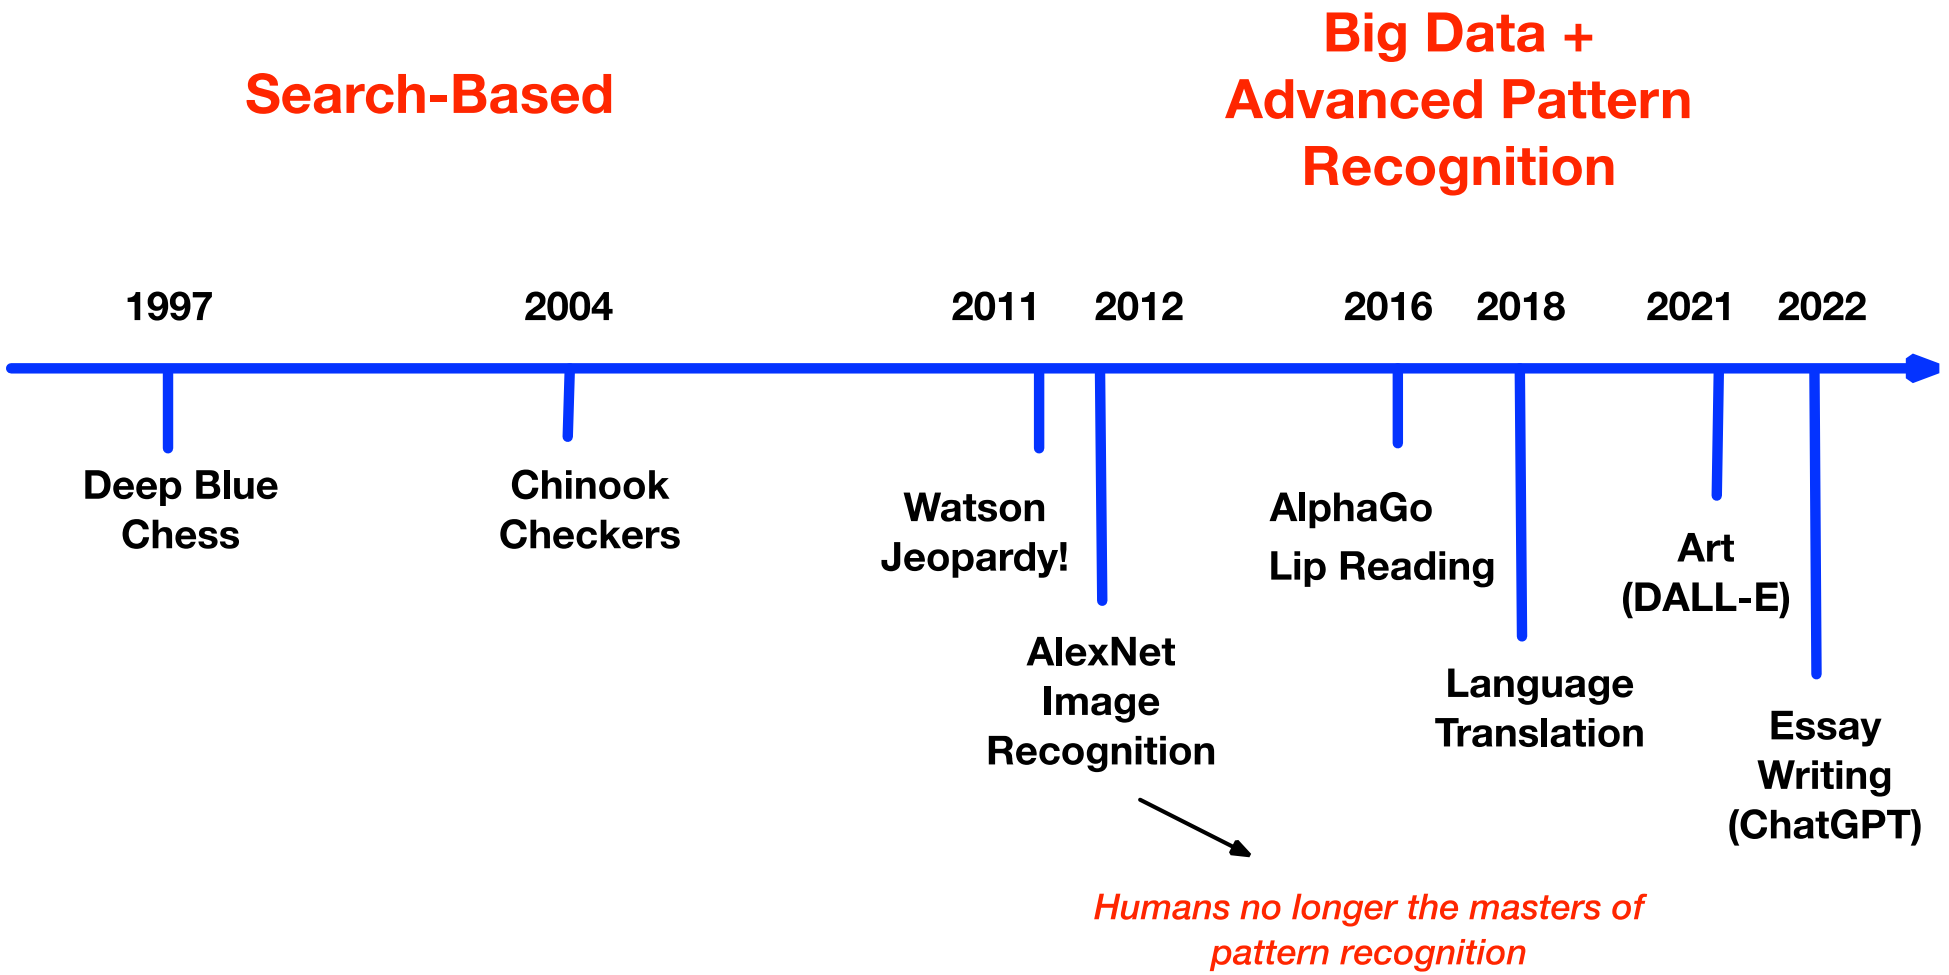

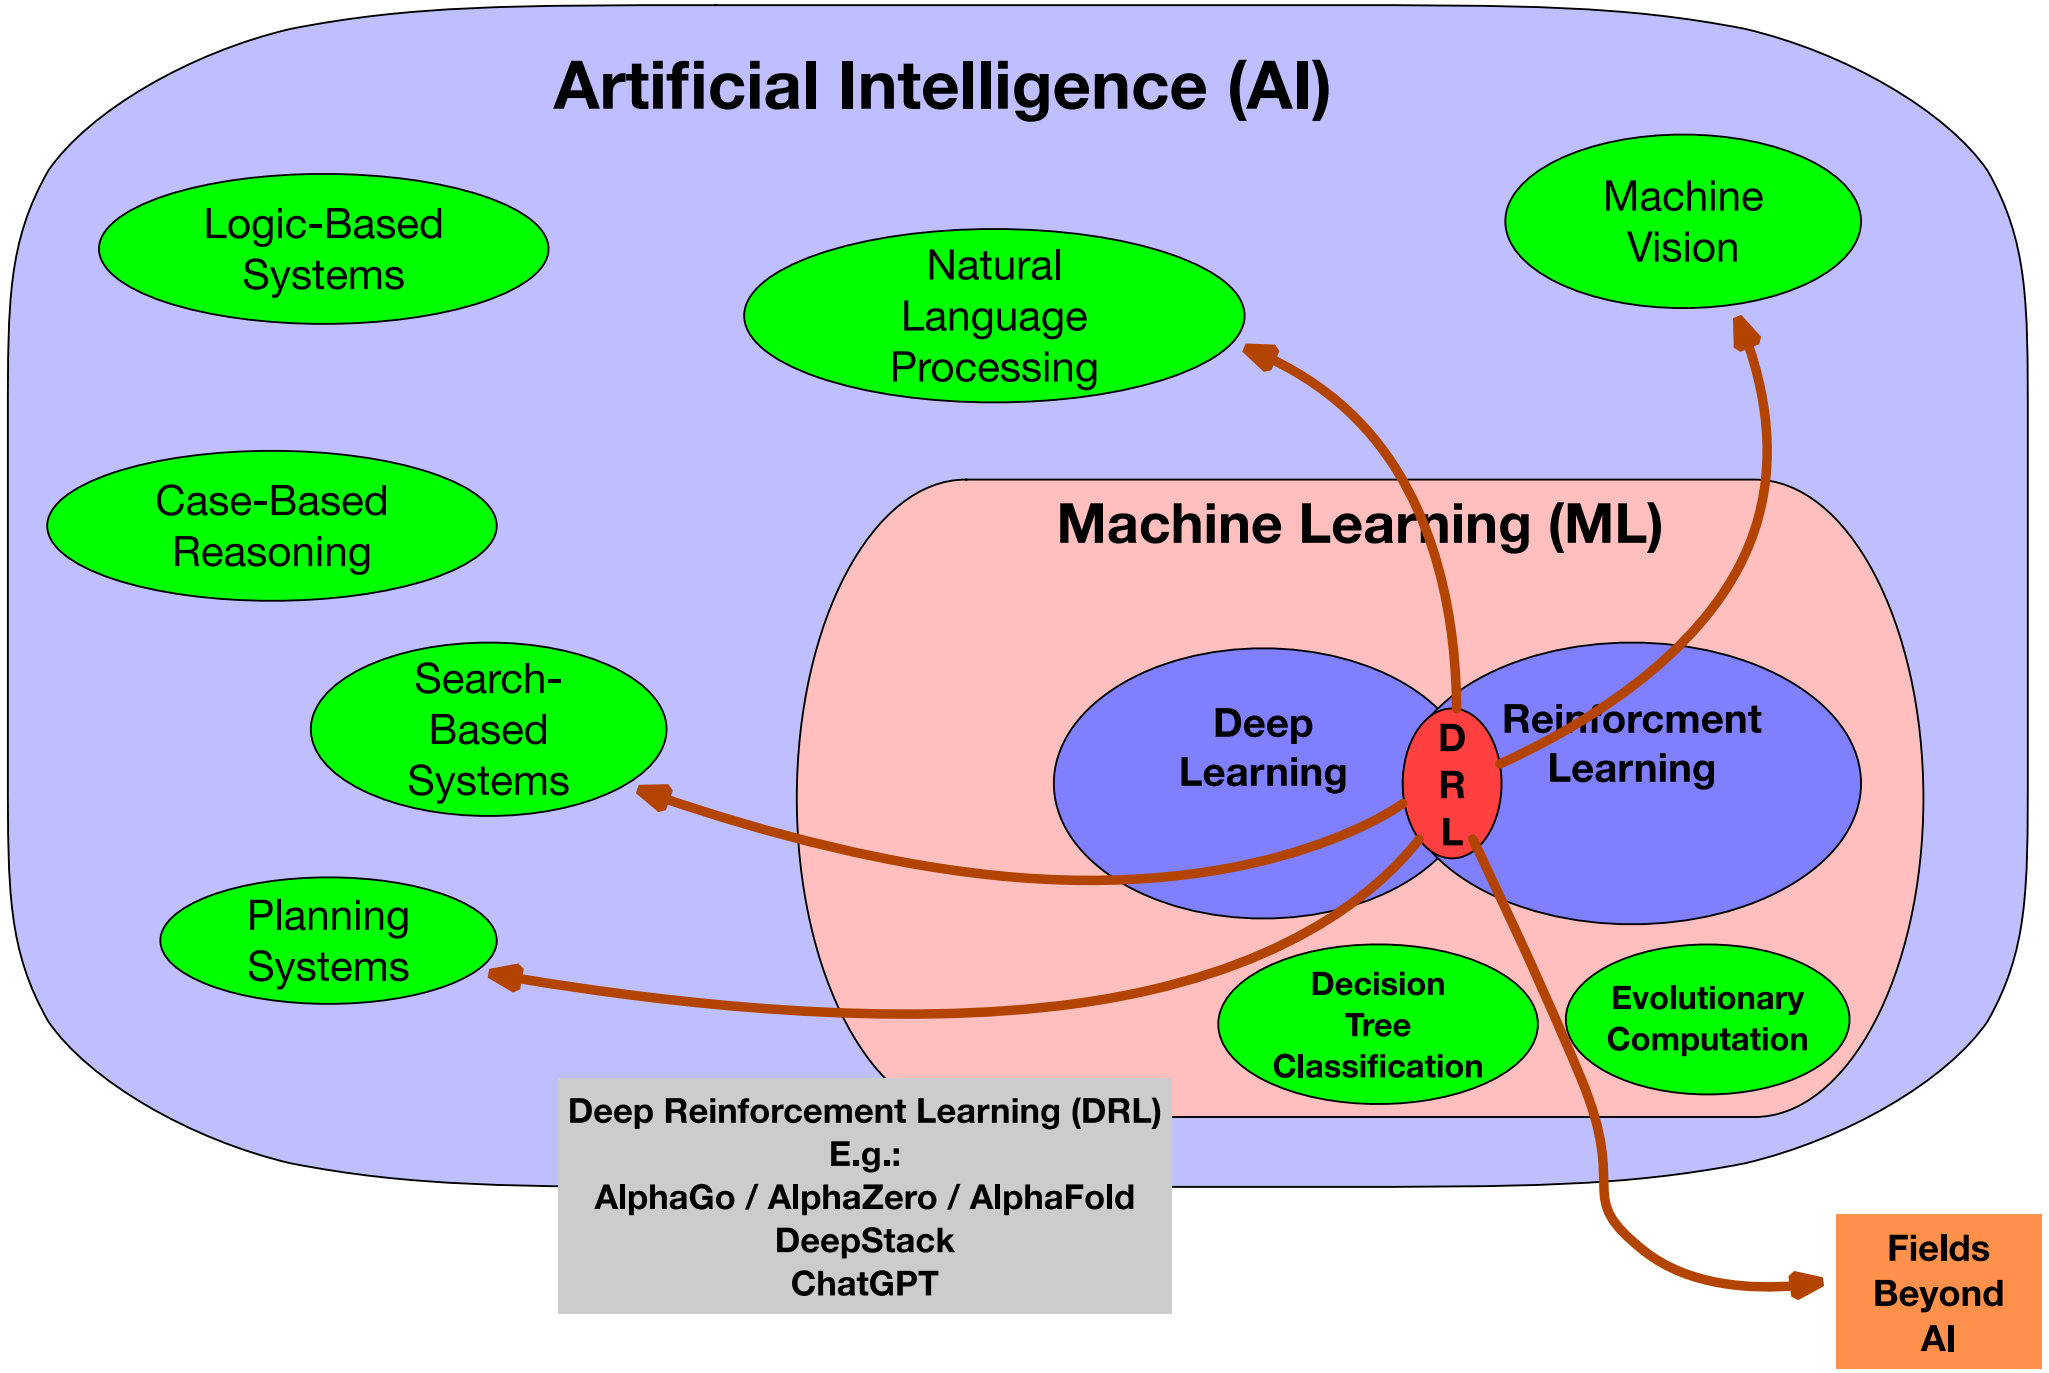
Artificial Intelligence (AI)
Logic-Based Systems
Natural Language Processing
Machine Vision
Case-Based Reasoning
Machine Learning (ML)
Search-Based Systems
Deep Learning
D
R
L
Reinforcment Learning
Planning Systems
Decision Tree Classification
Evolutionary Computation
Deep Reinforcement Learning (DRL)
E.g.:
AlphaGo / AlphaZero / AlphaFold
DeepStack
ChatGPT
Fields Beyond AI

Two Very Different AI Approaches

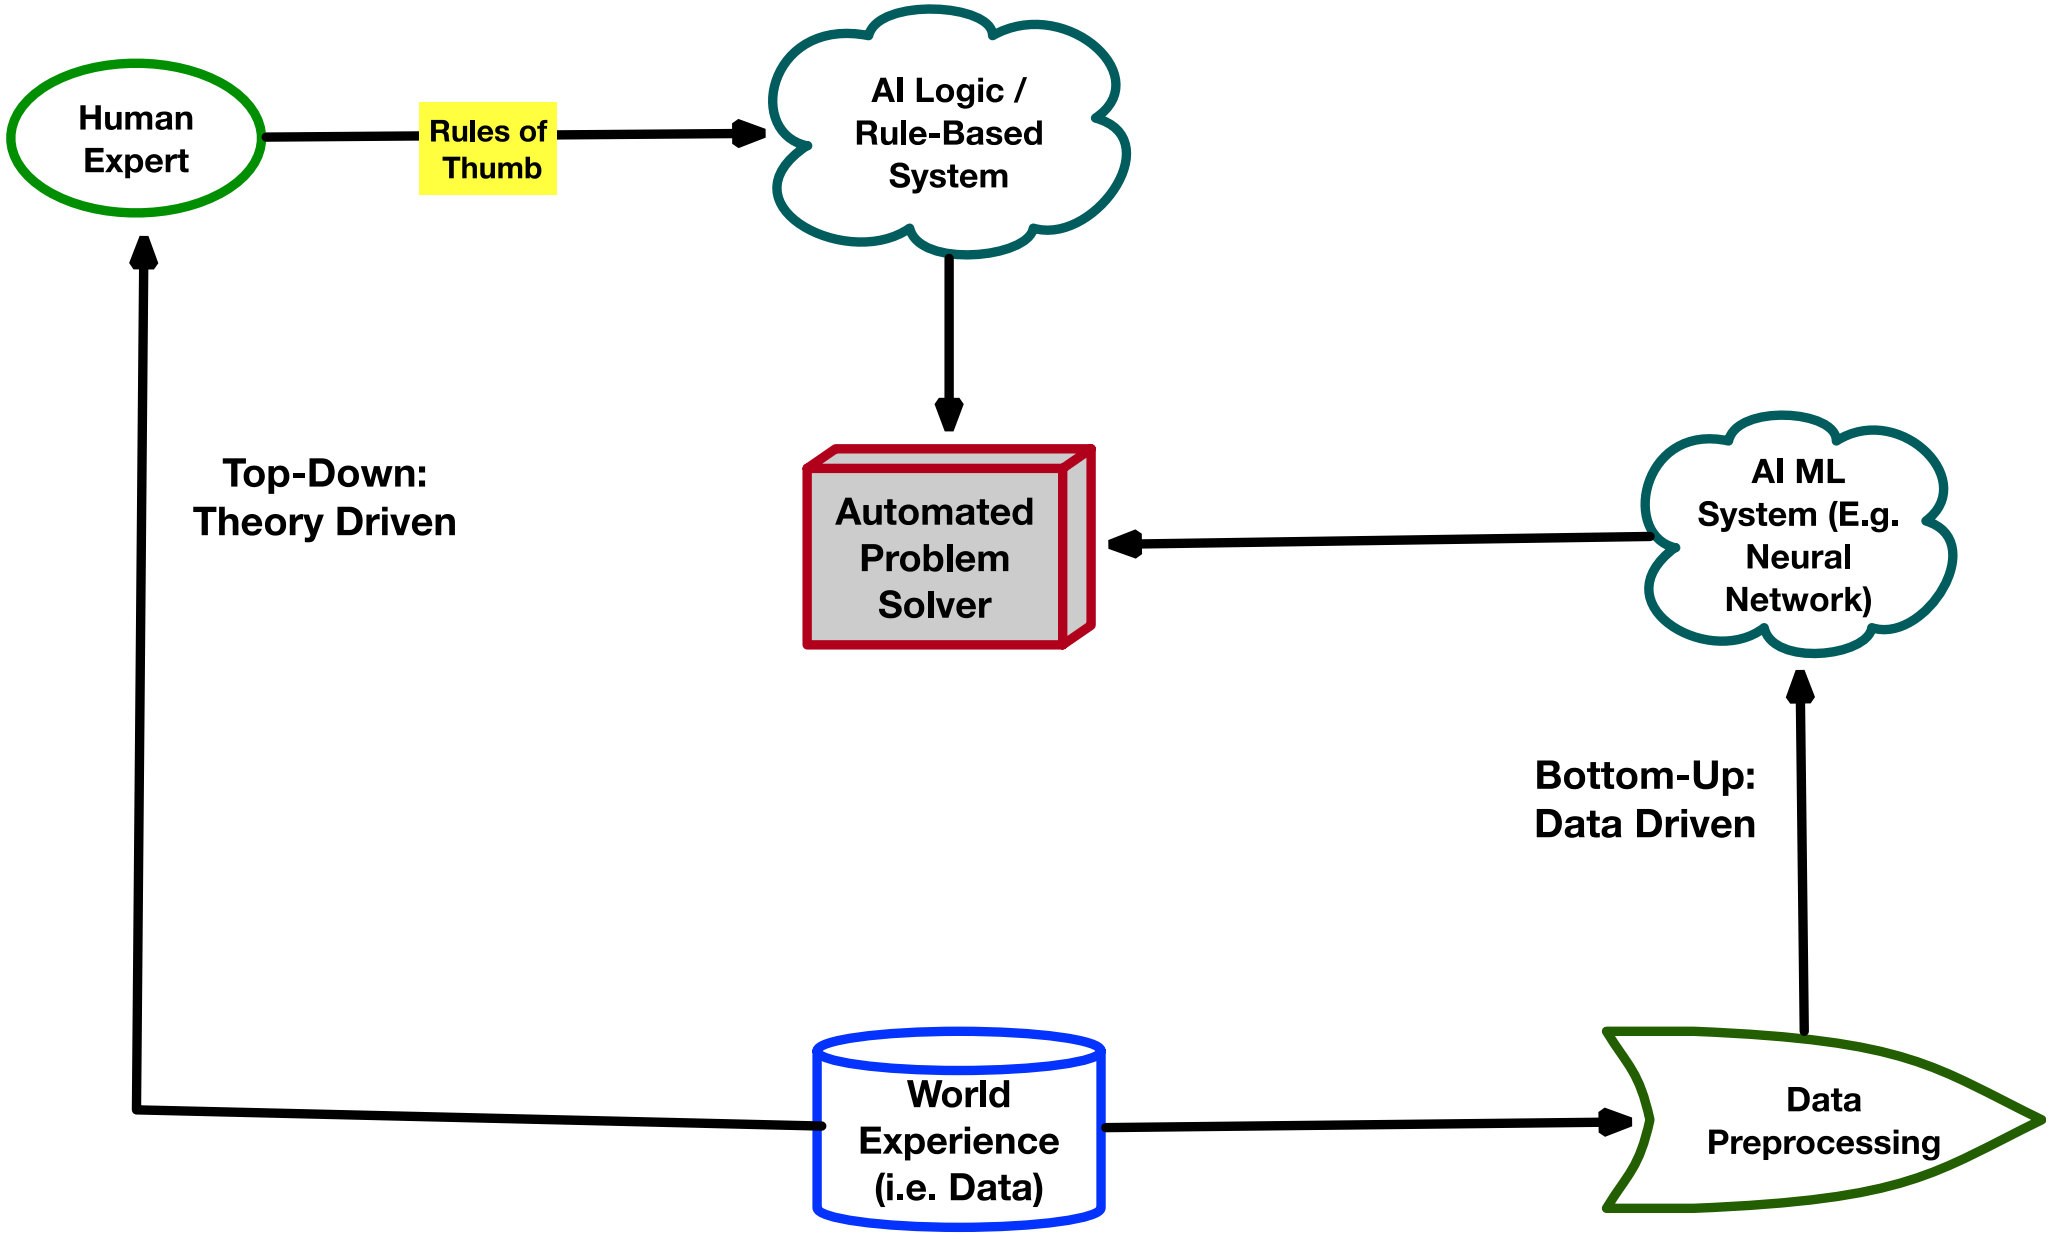

AI Texas Hold'em (Top-Down)

Interview poker experts (or read strategy books) to produce a large set of situation-action rules:

- IF your hole cards are a pair of face cards, AND the third such face card appears in the flop, THEN slow-play the hand (i.e. do not raise, only call).

- IF it is the final ("river") round of betting, AND the opponent raises before you, AND the pot-odds are favorable, THEN call the raise.

AI Texas Hold'em (Bottom-Up)

Roll-Out Perspective for Player 2

Player 2

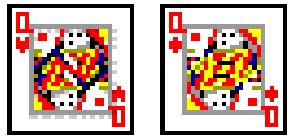

Player 1

Flop

Turn

River

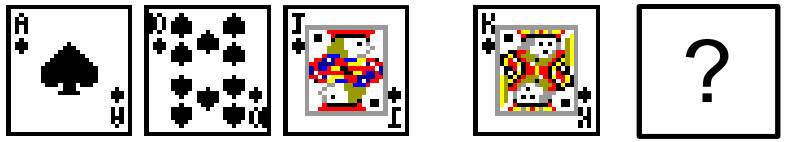

Player 3

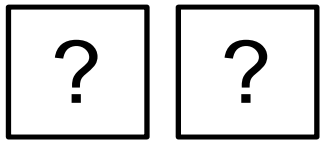

Player 4

- Deal out remaining cards thousands of times.
- Count # times (W) that Player 2 wins.
- Base Player 2's betting decision on W.

AI in Medicine

Top-Down

- Doctors formalize their (deep) medical knowledge
- AI developers enter ('hardwire') it into an expert system (ES).
- Patient cases run through the ES; **no learning**.

Bottom-Up

- Medical experts label hundreds (thousands) of MRI images.
- AI system **learns** associations between images and labels.
- AI system computes labels for new images.
- AI system has no **explicit** deep domain knowledge.

Microscopic Pattern Recognition

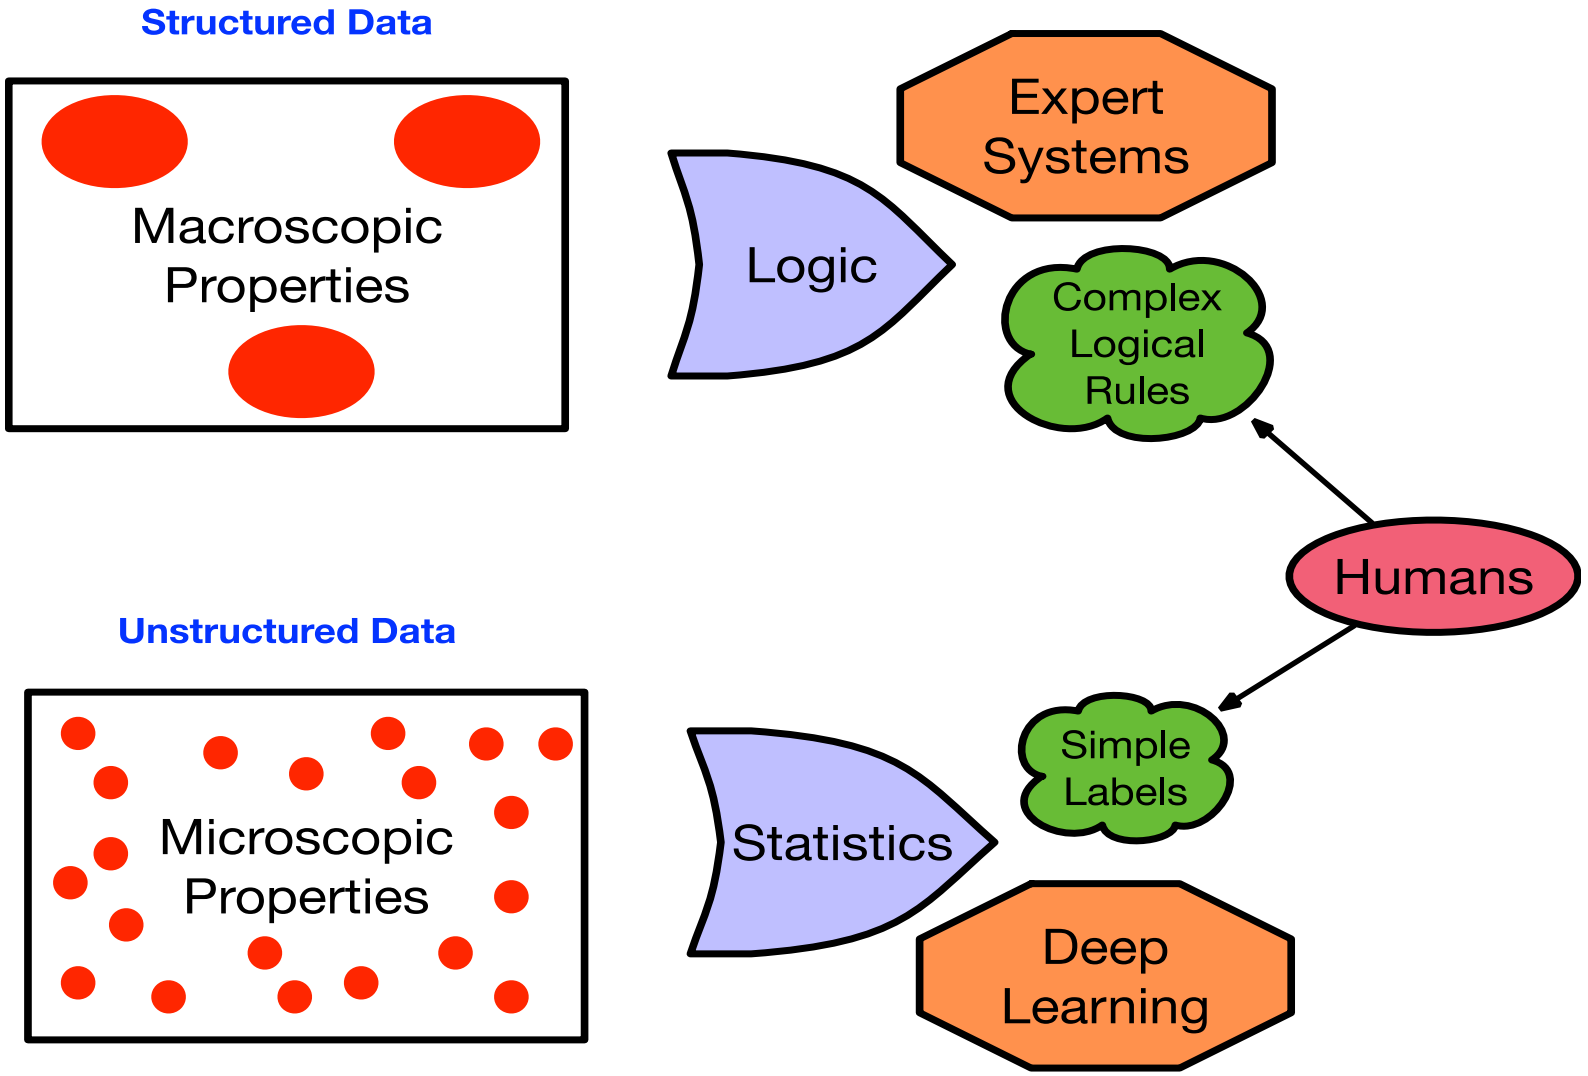

- **Recognition of Complex Microscopic Patterns is now possible (with AI)**
- **Reams of Microscopic Data are now much easier to produce**

Deep Reinforcement Learning (DRL)

- Goal of RL: Learn many complex rules (**state => action**) at the **microscopic** level that are **general**.
- General: maps **similar** states to similar actions.
- But **similarity** can be hard to define.
- Deep Nets learn **salient similarities** (that humans often miss) due to the net's superior pattern-recognition abilities.
- Deep Nets can process billions of microscopic cases in producing these rules.
- These cases are **generated internally**, by the trial-and-error search inherent in RL. **No human labeling needed!**

Deep-Learning and/or DRL Applications

- **Self-Driving Vehicles**
- Championship-level Game Playing
- Adding Sound Effects to Silent Movies
- **Fake Video; Detecting Fake Social-Media posts**
- Accelerated Discovery of New Drugs
- **Predicting Protein Folds**
- Interpreting Web Data for Cancer Research
- Fraud Detection ; Credit-risk assessment
- **Early detection of Cerebral Palsy (Open AI Lab)**
- **Natural Language Translation**
- Generating Articles from Statistics (e.g. Sports)
- Generating long text from short and/or complex prompts.
- **Generating Art + Music** based on styles of particular artists
- Earth Monitoring – Interpreting Satellite Images
- Habitat Monitoring – Analysis of animal sounds to detect poachers

AlphaGo

Silver et. al., *Mastering the game of Go with* ***deep neural networks*** *and* ***tree search***, Nature, 2016.

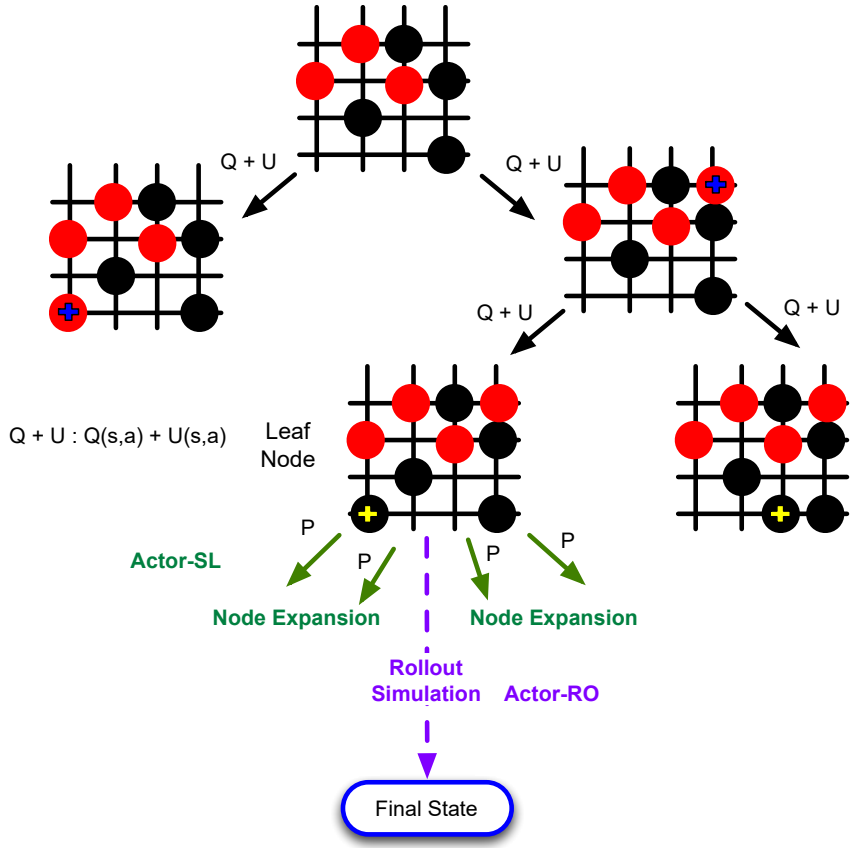

*I guess I lost the game because I wasn't able to find any weaknesses…***Lee Sedol** (World # 2)

The Alpha Progression

- Pre-trained on catalog of expert games
- Self-play from intermediate level upwards.
- Beat human world champion.

- No expert knowledge needed.
- Self-play from scratch.
- Became world champ; beat AlphaGo 100-0

- Generalized AlphaGo Zero to other games.
- Became world champ at them.
- Another step toward Artificial General Intelligence (AGI)

AlphaFold (DeepMind, 2020)

Given: a nucleotide sequence (which directly codes for amino acids)
AlphaFold Computes: 3d structure of the resulting protein = How the amino-acid sequence **"folds".**
Implications: A protein's 3d structure determines its physicochemical behavior.

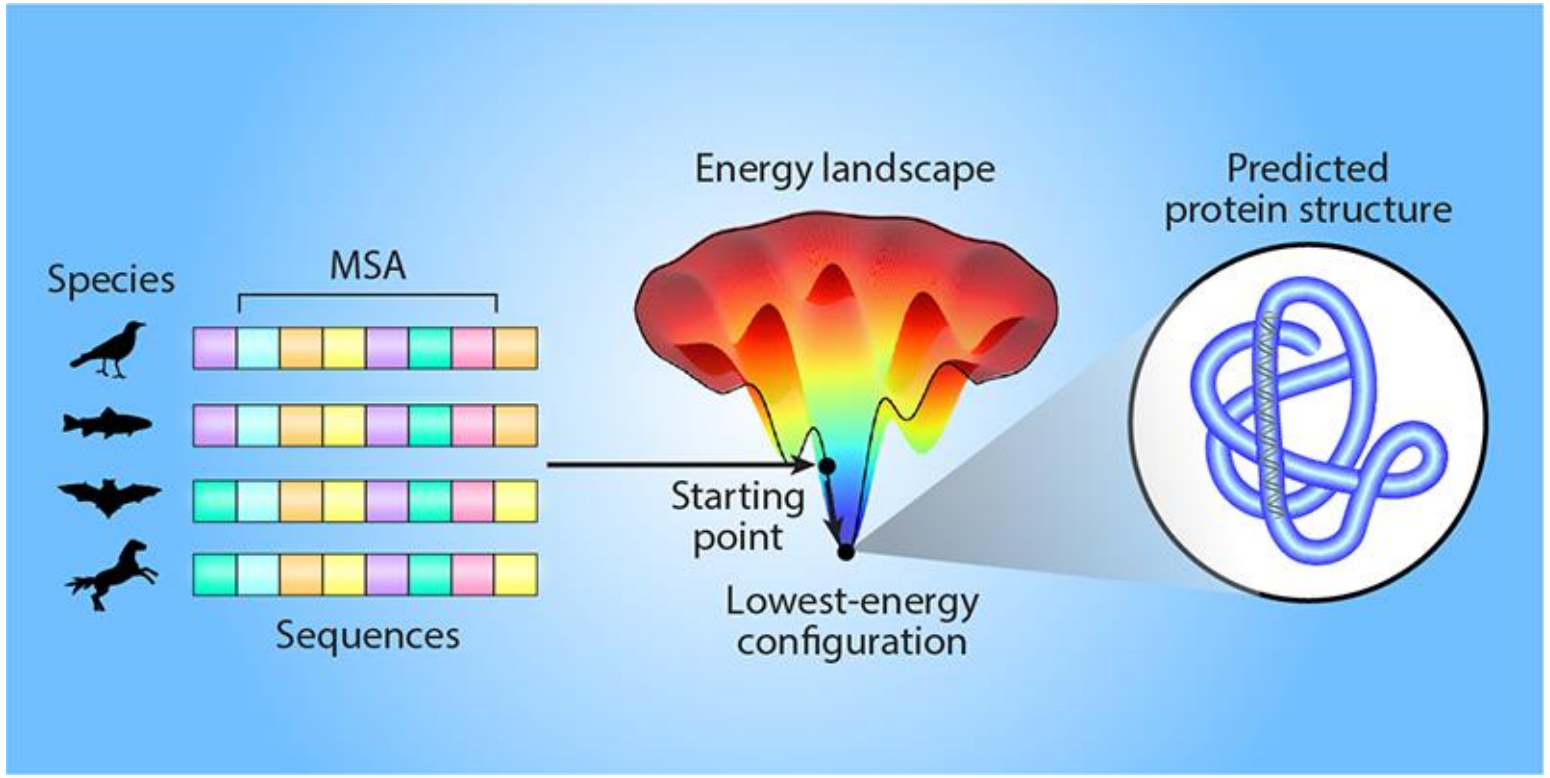

- A process that once took **years** (e.g. one PhD per protein) now takes **minutes**.
- 200 million proteins on earth, with many discovered each year, but few have known 3d structures.
- AlphaFold => Huge knowledge gain + massive labor savings !!

AlphaGeometry (DeepMind, 2024)

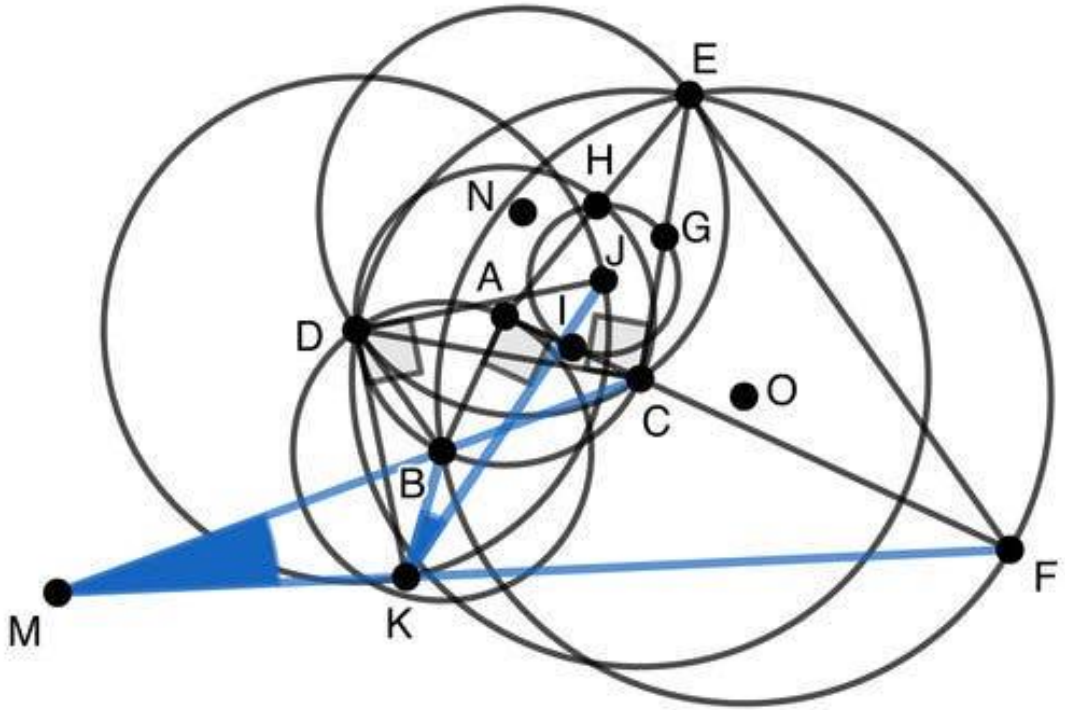

Creativity = Introduction of extra objects: points, lines, triangles, circles..

Traditional Symbolic AI Theorem Prover + Large Language Model

Silver Medal Level on Math Olypiad

Alternative Approaches to Common Sense and Artificial General Intelligence (AGI)

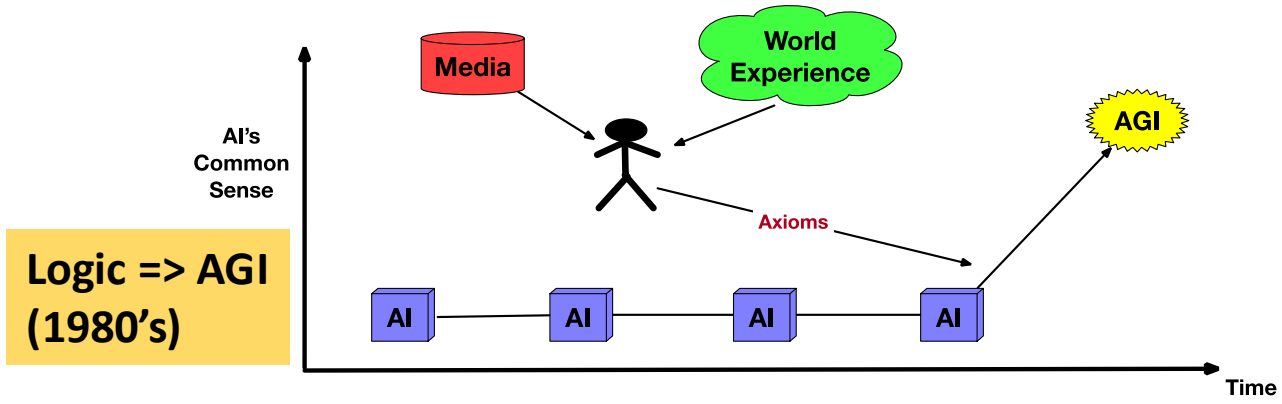

Logic => AGI (1980's)

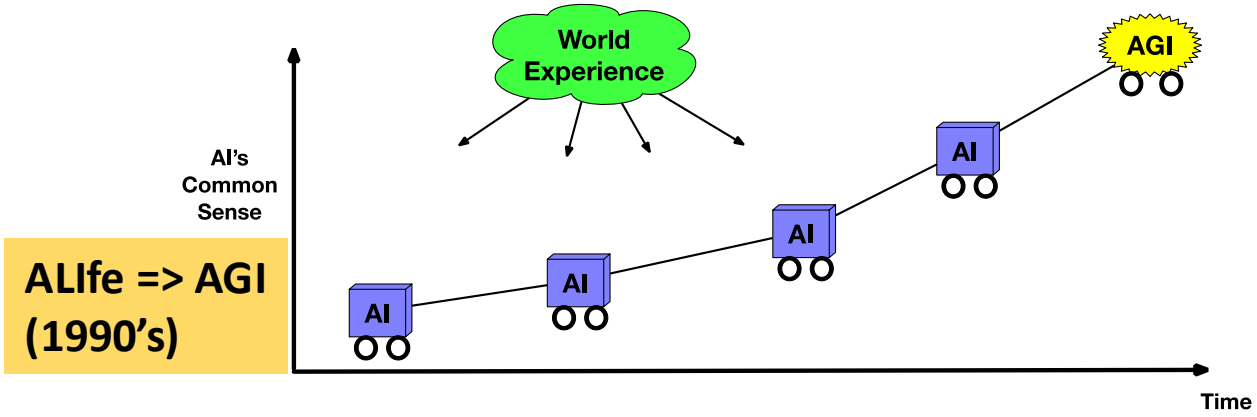

ALIfe => AGI (1990's)

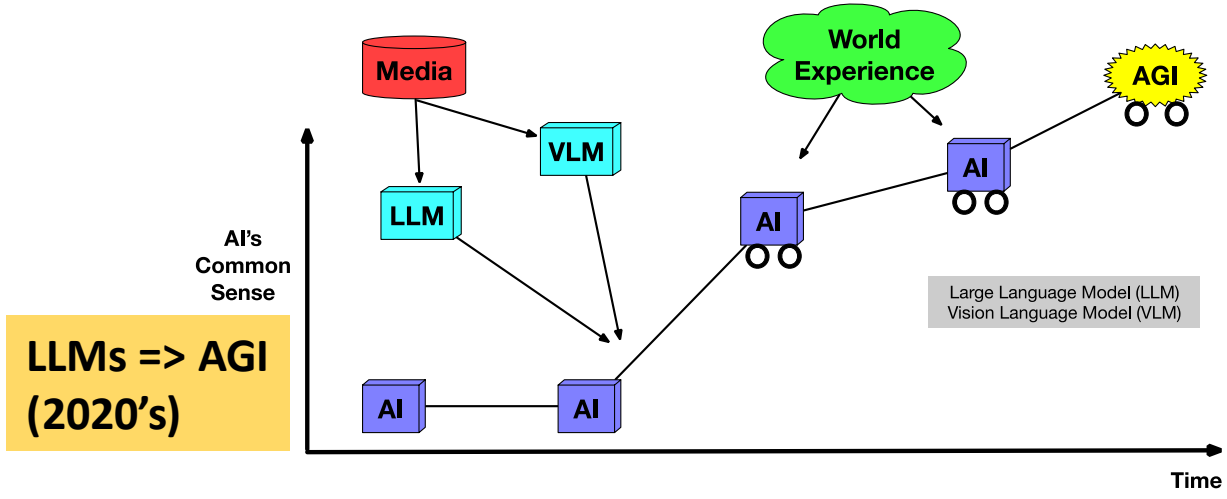

LLMs => AGI (2020's)

Common Sense = primary deficiency of most AI systems.

RT-2: Vision-Language-Action (VLA) Models Transfer Web Knowledge to Robotic Control (50+ co-authors, Google Deepmind, July 2023)

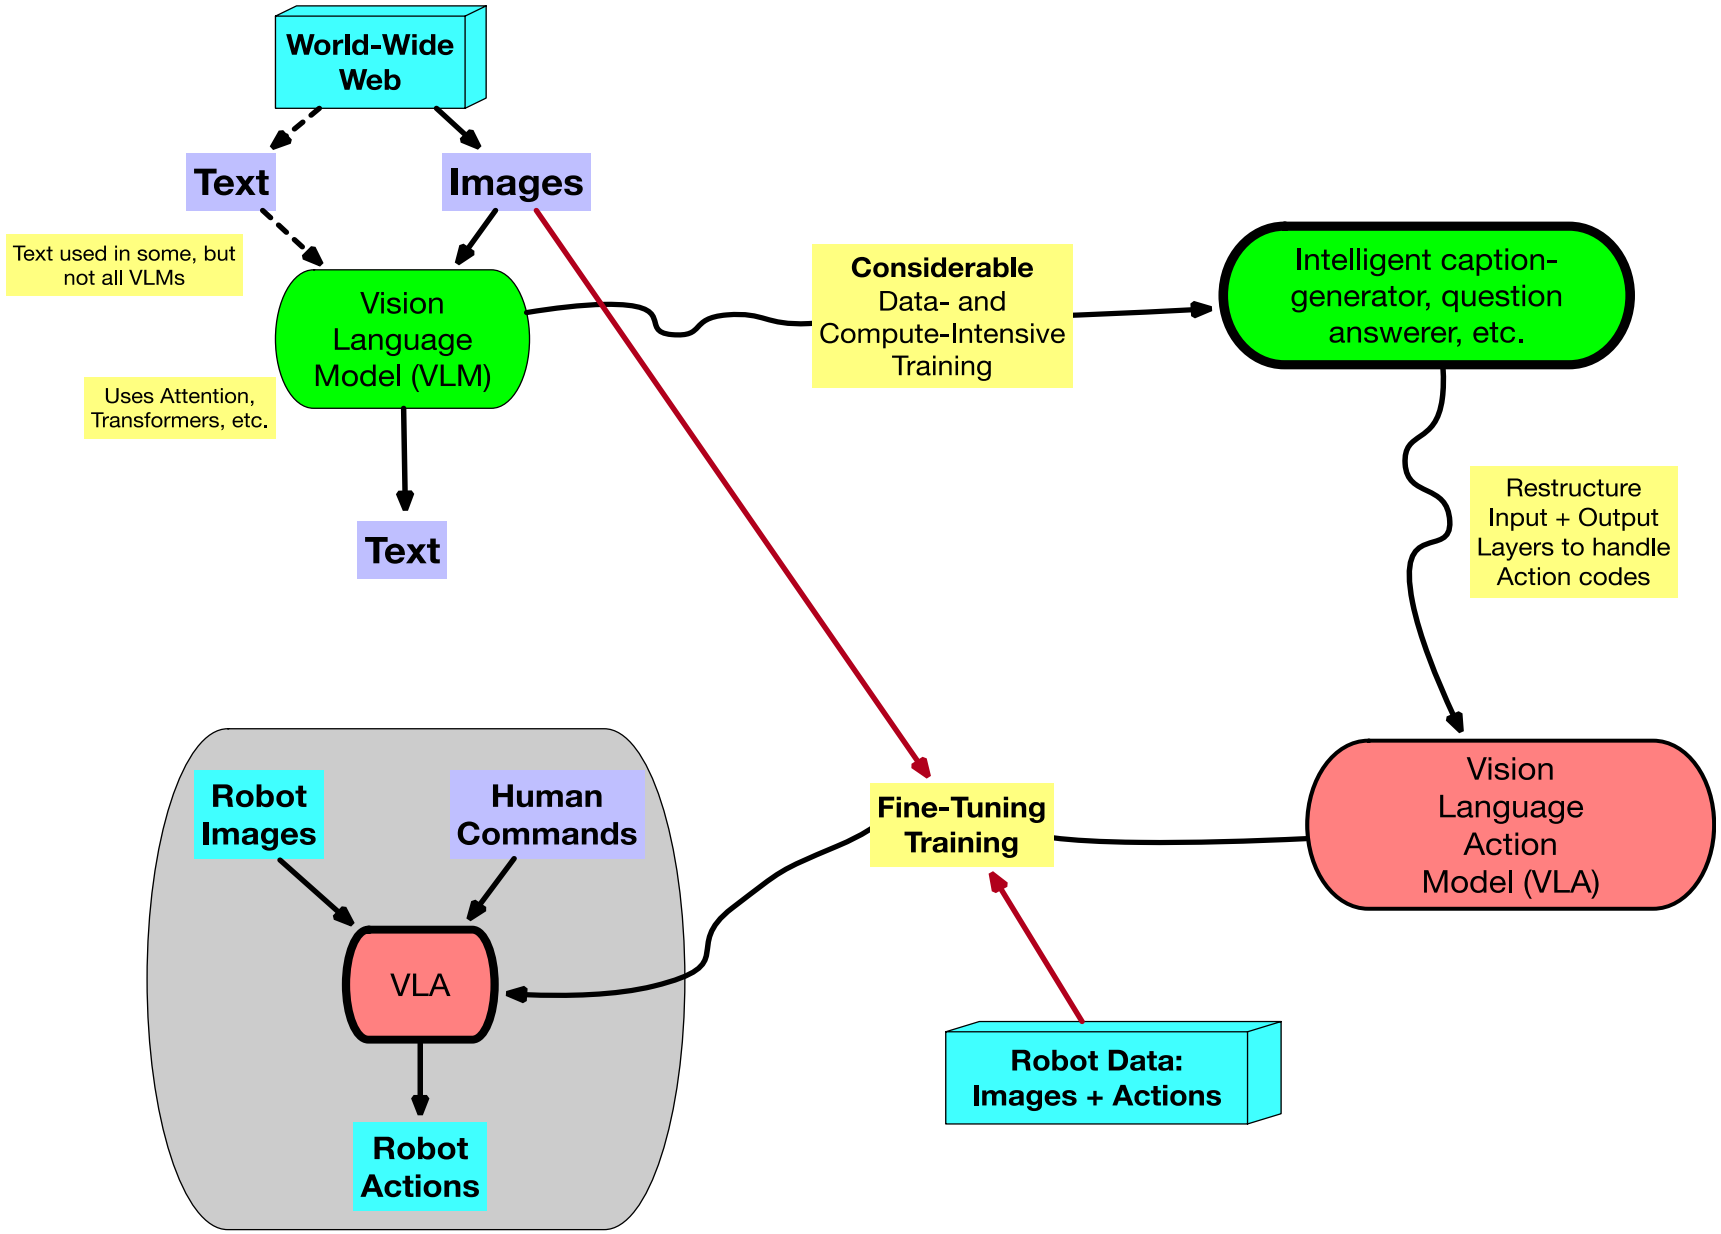

Emergent Common Sense in LLMs

Query:

I am riding down a road. Potholes are everywhere. A vial of nitroglycerine is in my backpack. My bike has no chain guard, and I am wearing long pants. Describe all of the potential dangers of this situation.

ChatGPT (3.5) Response:

Riding down a road filled with potholes while carrying a vial of nitroglycerine in your backpack, with a bike lacking a chain guard, and wearing long pants presents several potential dangers:

1.Potholes and Uneven Terrain: The presence of potholes can cause your bike to jolt and lose balance, leading to potential falls or collisions. Uneven terrain increases the risk of accidents and makes it difficult to maintain control.

2.Nitroglycerine Vial: Nitroglycerine is an extremely sensitive and volatile substance. The vibrations and impacts from riding over potholes can potentially trigger a chemical reaction or explosion within the vial, leading to a catastrophic accident.

3. Backpack Impact: The jolts and impacts from riding over potholes can cause your backpack to move vigorously. This movement may lead to the vial of nitroglycerine banging against other objects in your backpack, increasing the risk of accidental rupture or detonation.

4. Bike Chain and Pants: Without a chain guard, your pants could get caught in the bike's chain and gears, leading to sudden stops or loss of balance. This could cause you to fall, potentially into the path of traffic or other hazardous objects.

5. Loss of Control: Riding over rough terrain increases the likelihood of losing control over your bike. This can lead to sudden swerves, collisions with obstacles, or even being thrown off the bike.

6. Combustion Risk: Nitroglycerine is highly sensitive to shock and friction. The friction between your pants and the bike's moving parts, combined with the jarring motions from potholes, could potentially create enough heat and friction to ignite or detonate the nitroglycerine vial.

7. Traffic Hazards: If you are riding in an area with traffic, the combination of potholes, potential loss of control, and the distraction of managing the nitroglycerine vial can increase the risk of accidents involving other vehicles.

8. Inadequate Protection: Long pants might offer some protection, but they may not be sufficient to prevent injuries in case of a fall or collision. Depending on the severity of the accident, you could still suffer cuts, abrasions, or more serious injuries.

Given the extreme risks involved in this situation, it is strongly advised to find a safe place to stop, assess the situation, and take appropriate measures to mitigate these dangers. This could involve securing the nitroglycerine vial properly, addressing the chain guard issue, and ensuring your safety while navigating through the hazardous road conditions.

My Evolving Trauma

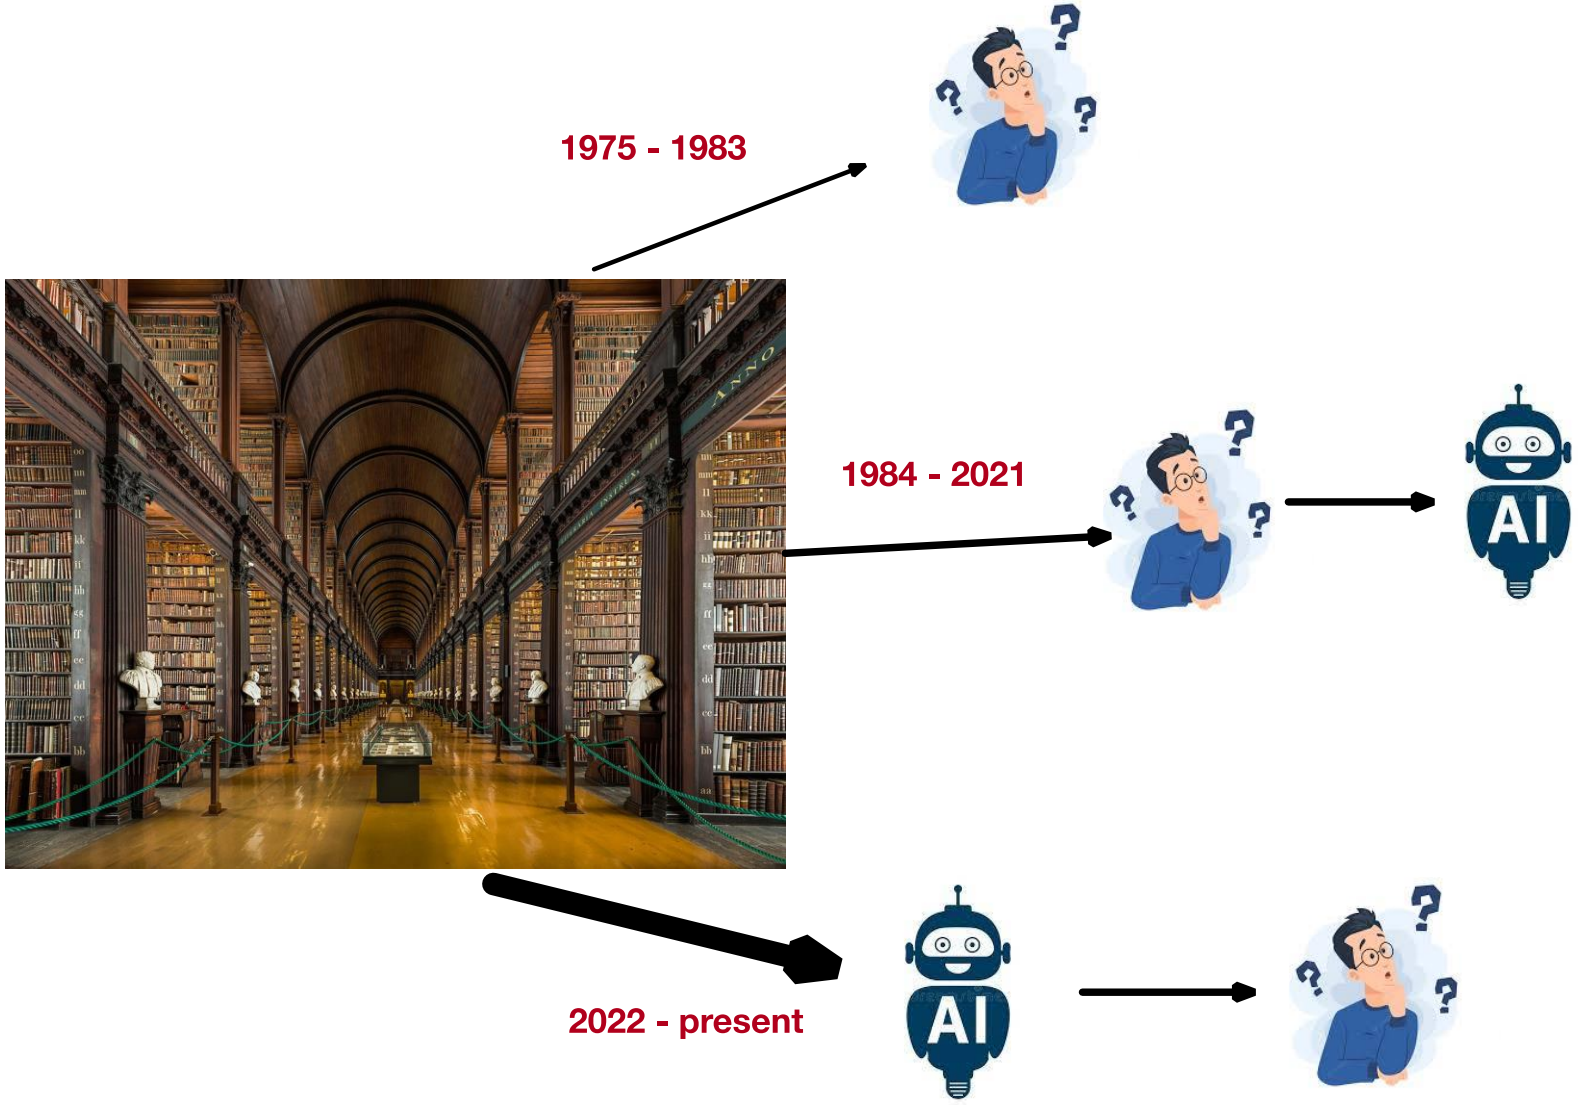

Part 2: Humans Need Not Apply

One machine can do the work of 50 ordinary men. No machine can do the work of one extraordinary man
... Elbert Hubbard

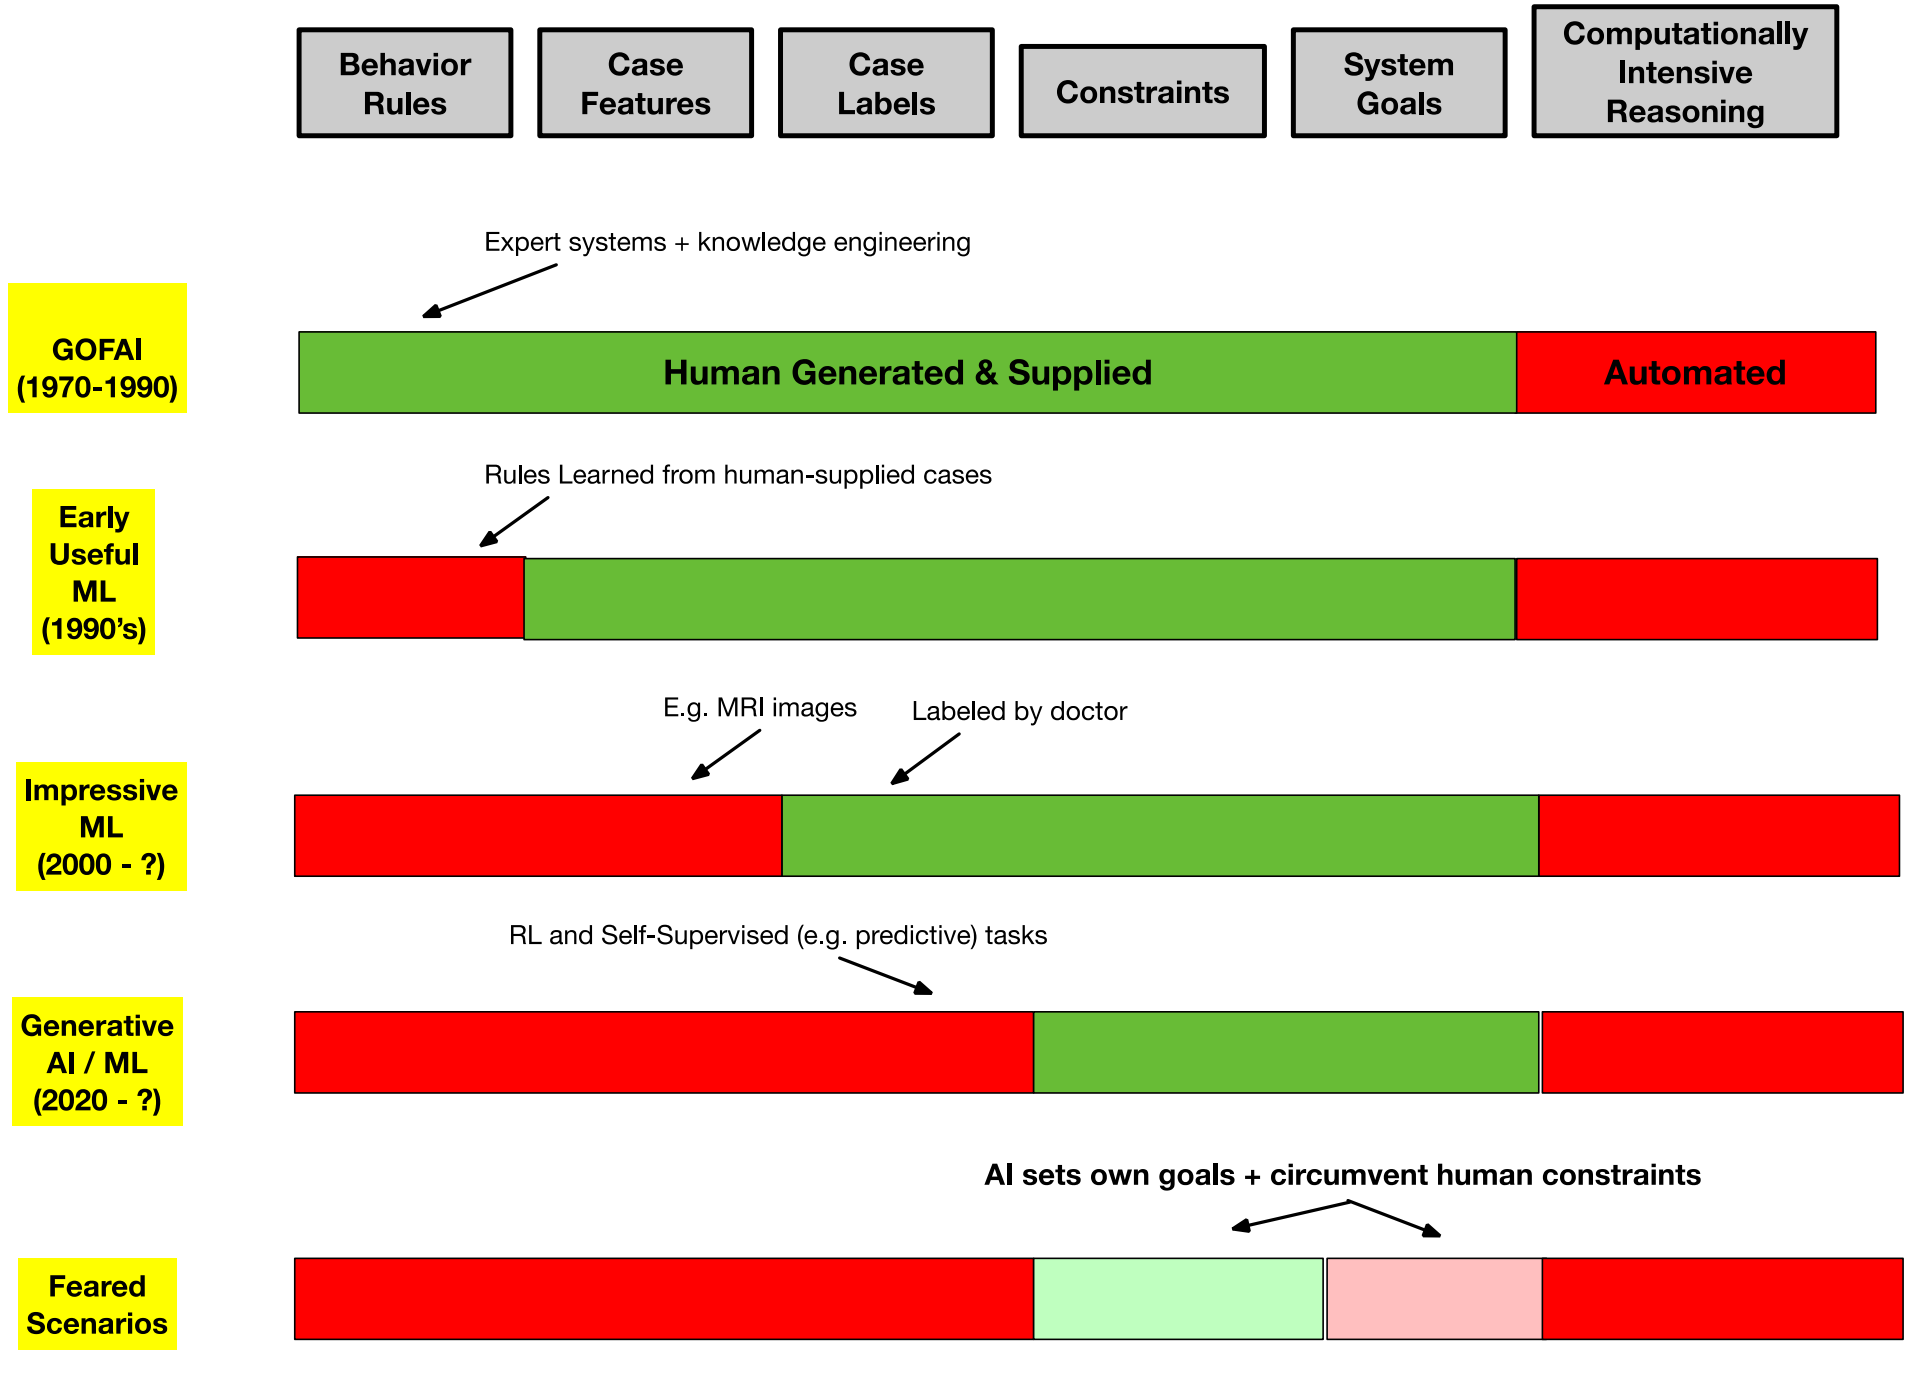
Behavior Rules
Case Features
Case Labels
Constraints
System Goals
Computationally Intensive Reasoning
Expert systems + knowledge engineering
GOFAI (1970-1990)
Human Generated & Supplied
Automated
Rules Learned from human-supplied cases
Early Useful ML (1990's)
E.g. MRI images
Labeled by doctor
Impressive ML (2000 - ?)
RL and Self-Supervised (e.g. predictive) tasks
Generative AI / ML (2020 - ?)
AI sets own goals + circumvent human constraints
Feared Scenarios

HELP WANTED (Machines Only)

- AI learns by itself. No need for expert human **logic / rules**)
- Reinforcement Learning (RL) systems generate their own labeled datasets as they explore the world => Don't even need humans to **label data** !!
- Humans no longer undisputed masters of **pattern recognition**.
- Humans may not even be the best experts on **commonsense**: It emerges in Large language models (LLMs) !!
- AI systems with **less human bias** can be **more creative**: E.g. Alpha Go's Move 37. bias
- **When humans become obsolete for one job, they retrain. Now, the machines can retrain too…and faster!**

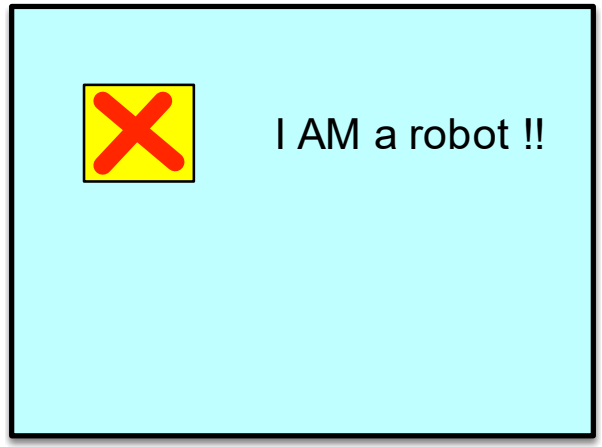

Rapid Acceleration of AI Progress

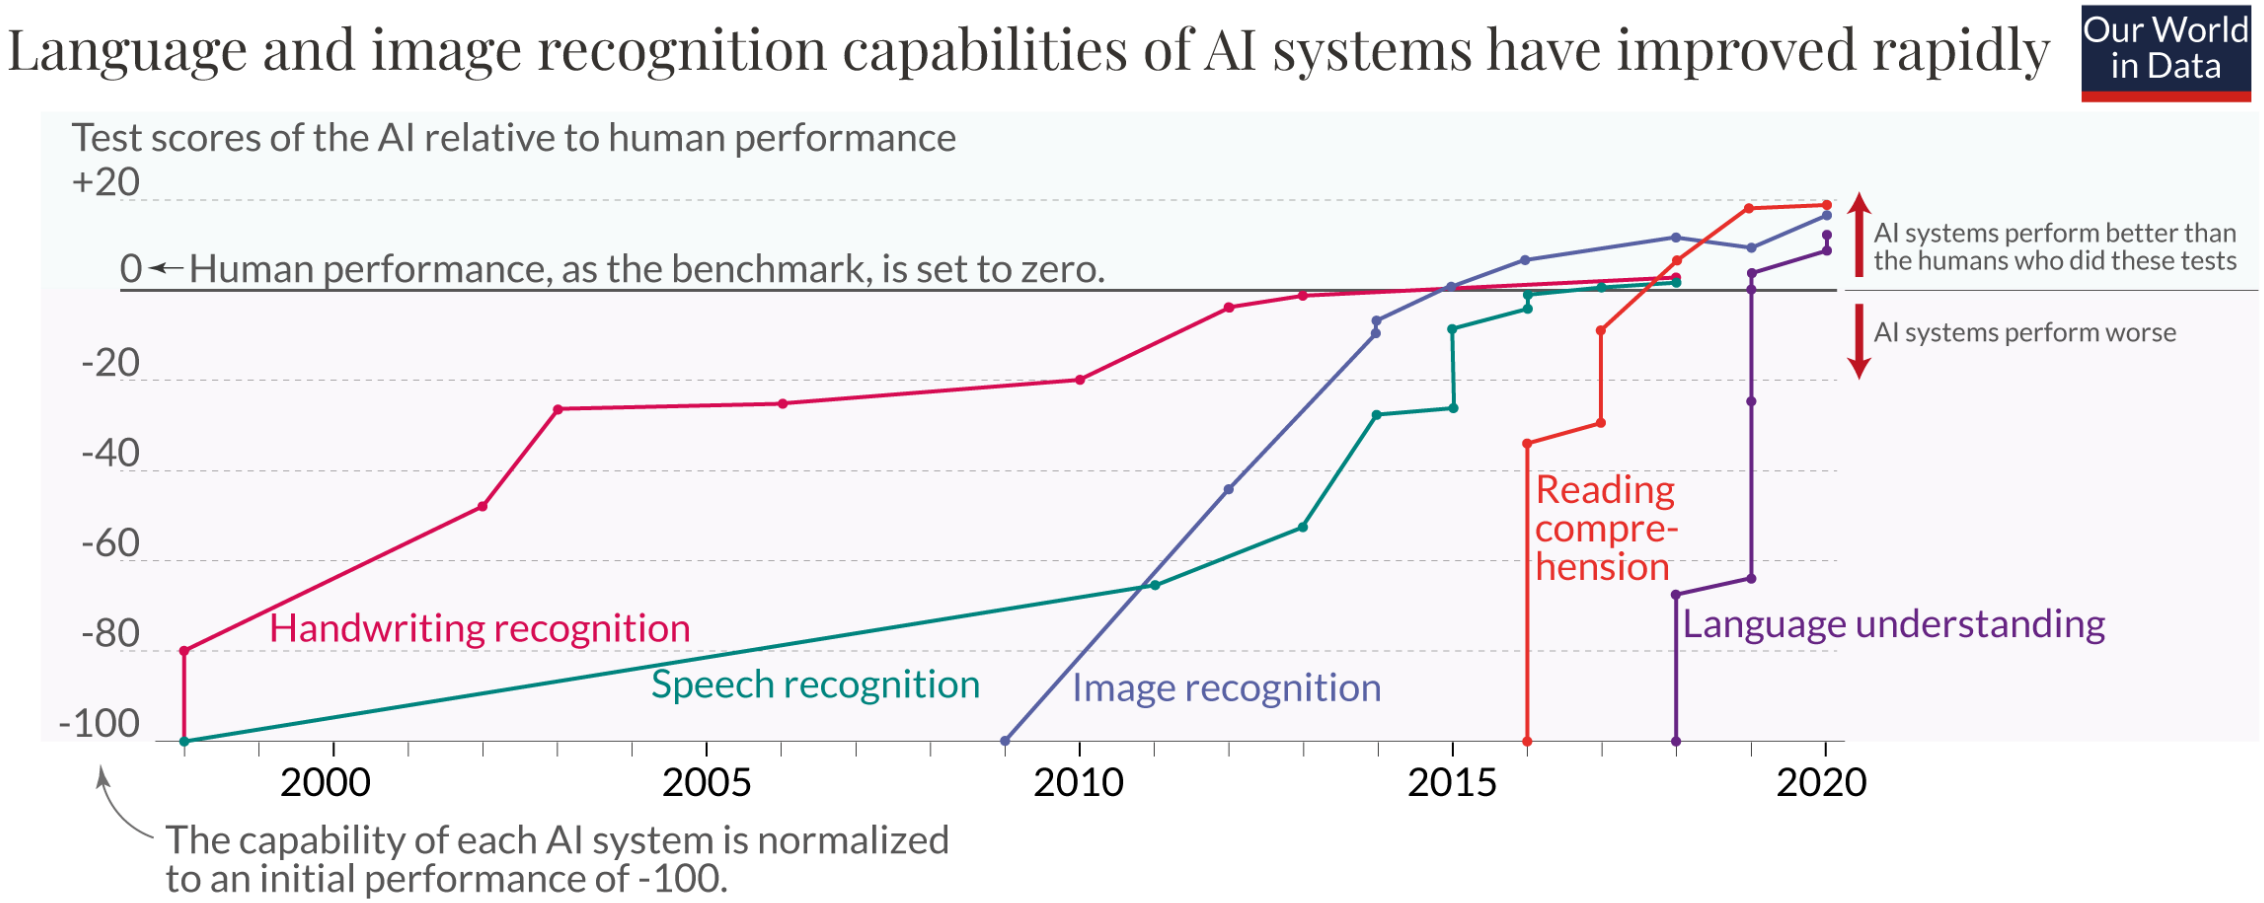

Data source: Kiela et al. (2021) – Dynabench: Rethinking Benchmarking in NLP
OurWorldinData.org – Research and data to make progress against the world's largest problems.

* AI **quickly** achieves super-human performance on many new tests.

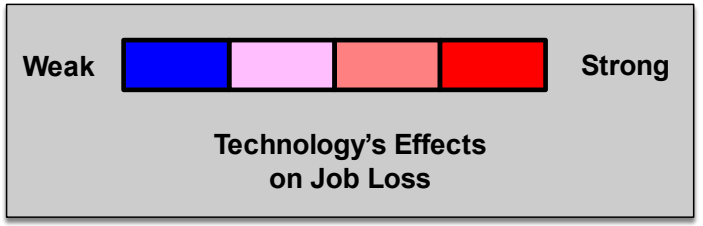

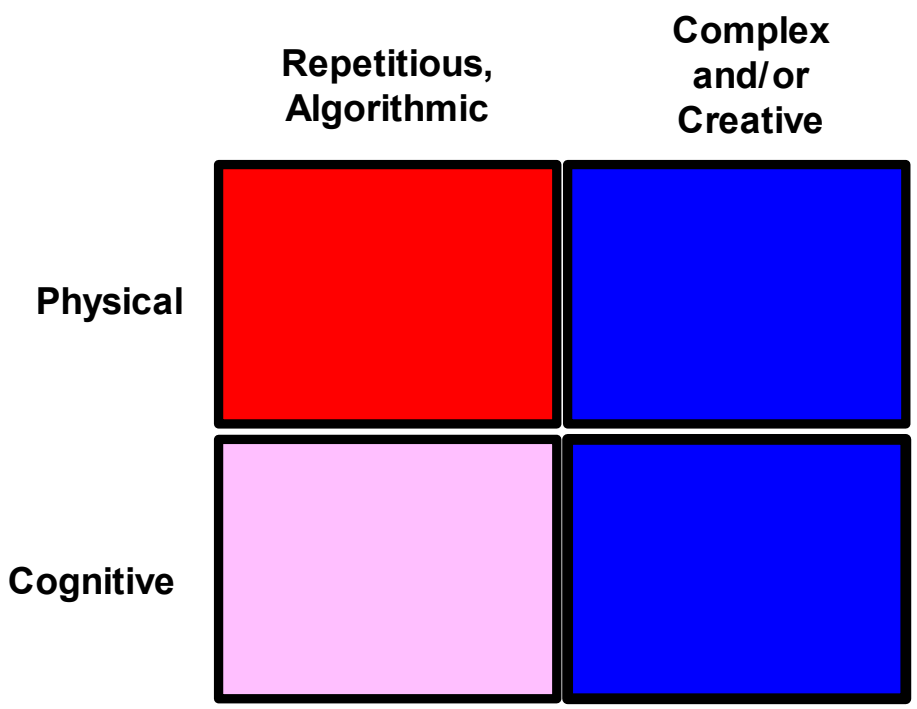

Classic Technological Disruption

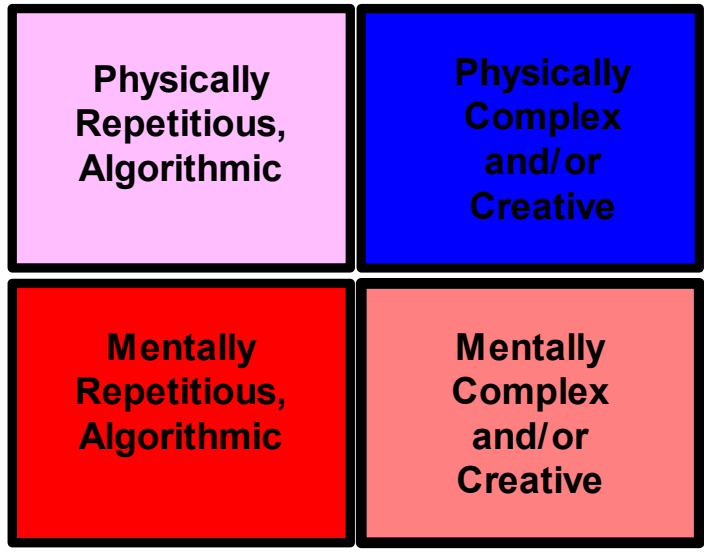

AI's Disruption

Job Security in an AI World

- Know how to use AI
- Nurture your creativity
- Be empathetic
- STEMpathy
- Be a people person
- Do complex physical work.

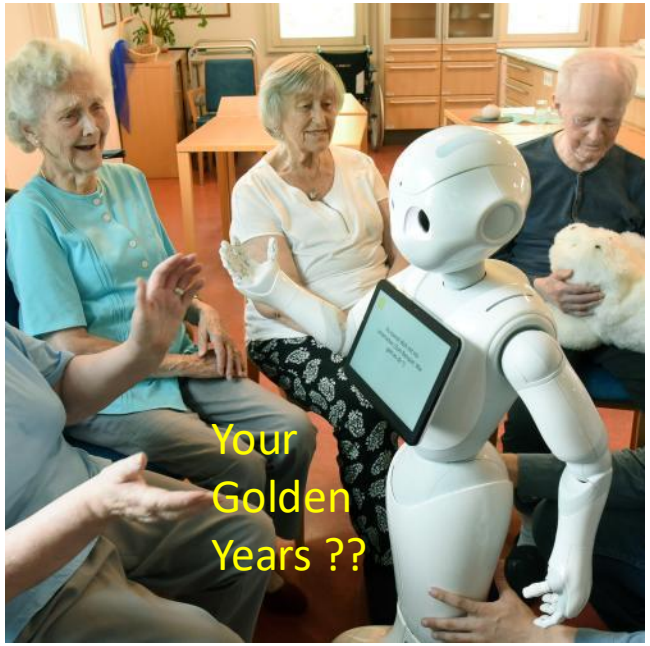

	Repetitive / Algorithmic	Creative
Physical	Factory Assembly Assistant Craftsman Lawncare Custodial Restaurant Jobs Retail Sales & Service Driving	Craftsmanship Civil Construction Mechanical Repair Athletics Politics (?)
Mental	Accounting Legal Aid Secretarial Middle Management Sports Journalism Diagnosis Financial Investment Help Centers Travel Planning	Architecture Poetry & Fiction Music & Art Upper Management Software Design Product Design

Current AI & Robotic Prowess

High | Medium | Low

MVBP (Most Valuable Body Part)

- **Pre-1870**

* The 2nd Industrial Revolution began around 1870.

- **1870 - Present**

- **Near Future**

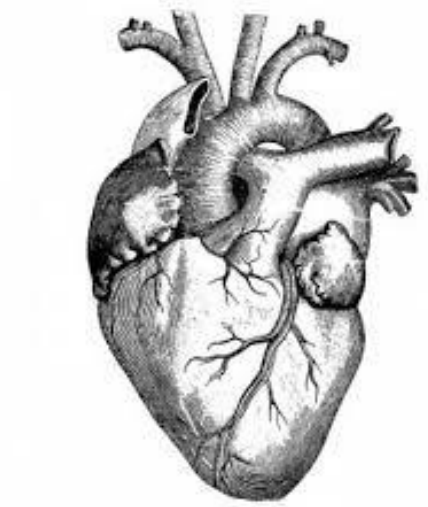

Economics is not
Rocket Science
… unfortunately

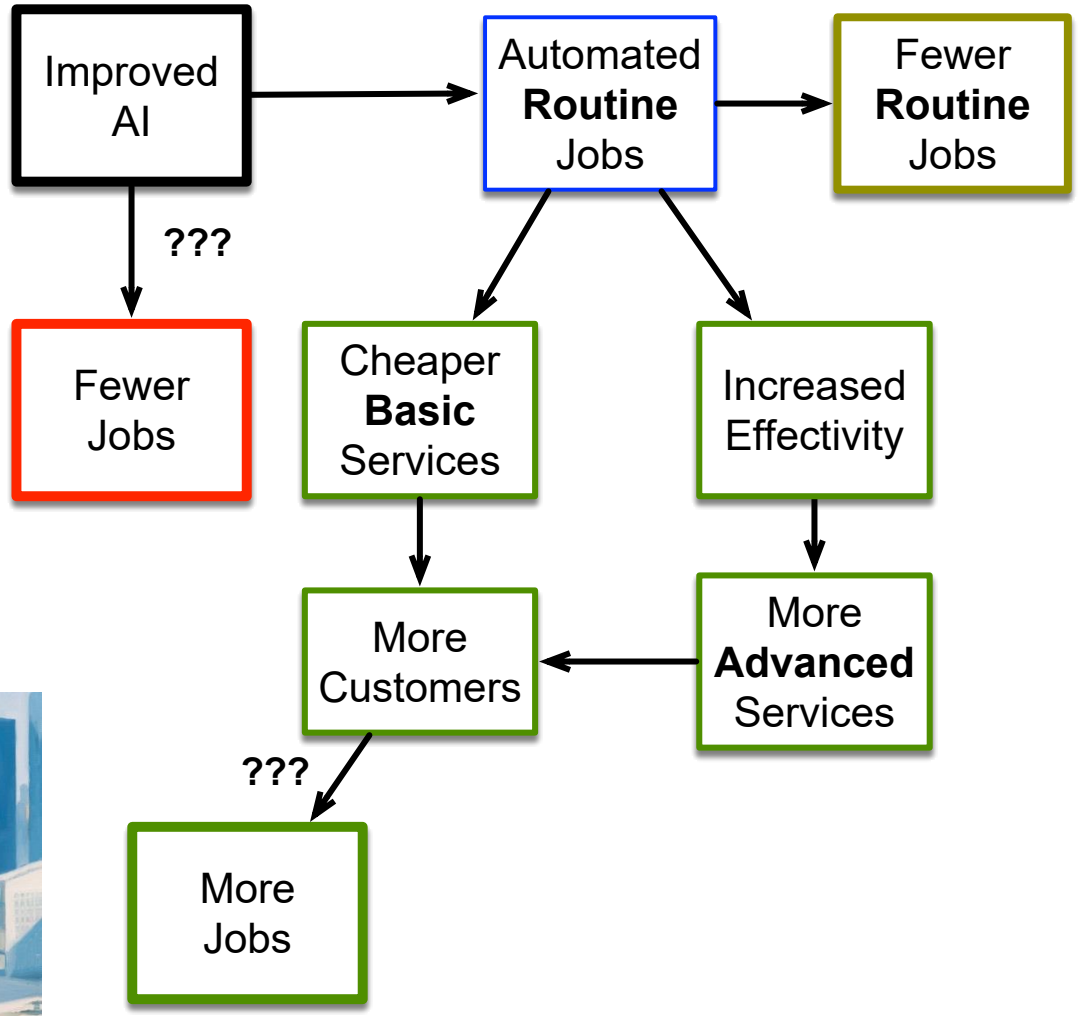

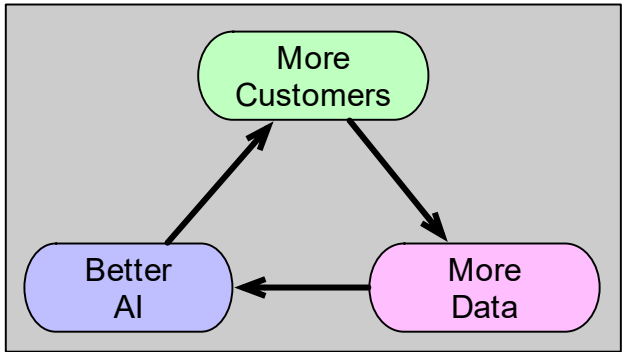

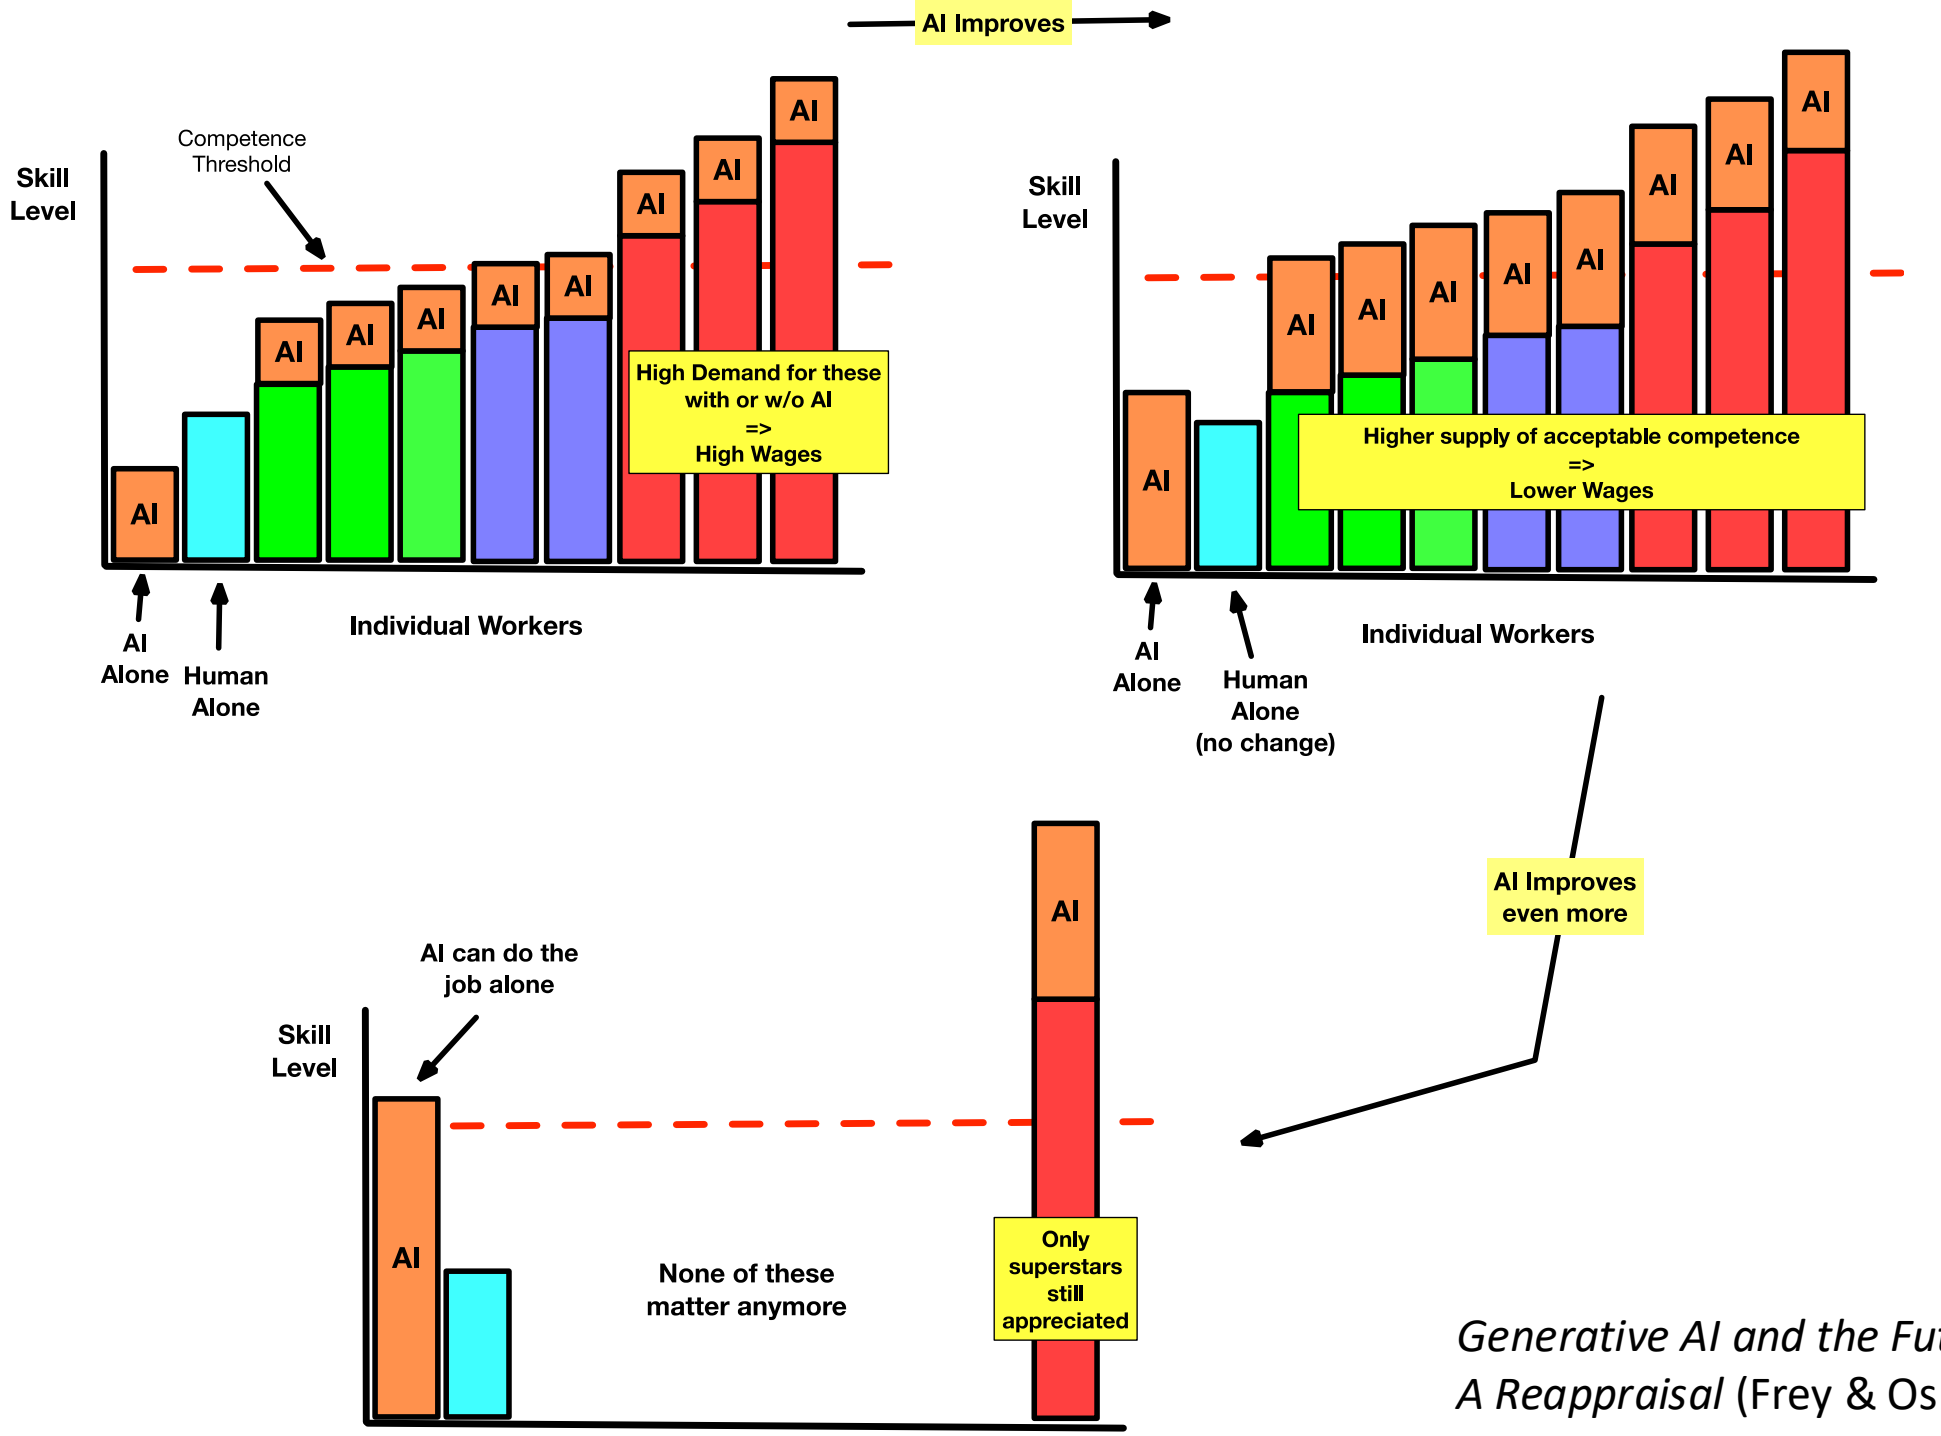

Generative AI and the Future of Work: A Reappraisal (Frey & Osborne, 2023)

Scott Galloway
(Popular Tech & Business Podcaster)

You won't lose your job to an AI, but you might lose it to someone who knows AI.

Code using Co-pilot or become obsolete

Become a ChatGPT whisperer (a.k.a. *prompt engineer*) – although this may be a fleeting need.

Part 3: Humans Need Not Apply Themselves

The real problem is not whether machines think but whether men do too
… B.F. Skinner

Plight of the Infovore

... as we come to rely on computers to mediate our understanding of the world, it is our own intelligence that flattens into artificial intelligence (Nicholas Carr, ***The Shallows***)

Constant distractions of cyberspace

"Shallows": Reduced depth of thought and emotion

Less Creative

Reduced value in an information society

Automation Bias: Trust machines over ourselves.

Rely on computers to understand the world

Rely on AI for **wisdom**

Human intelligence becomes **artificial** **...** and thus more easily **predicted** and **controlled** by AI

Who Works for Who/What?

- Comments from Laura Preston, a *human fallback* for Brenda, a real-estate chatbot:

Months of impersonating Brenda had depleted my emotional resources. It occurred to me that I wasn't really training Brenda to think like a human; Brenda was training me to think like a bot, and perhaps that had been the point all along.

- *Kenyans paid under $2/hour to label (very disturbing) images for ChatGPT, and similar work for Facebook. Even refugees in camps are used for this.*

- *AI Surveillance of low-paid workers doing repetitious jobs in warehouses.*

Men have become the tools
of their tools… Henry David Thoreau

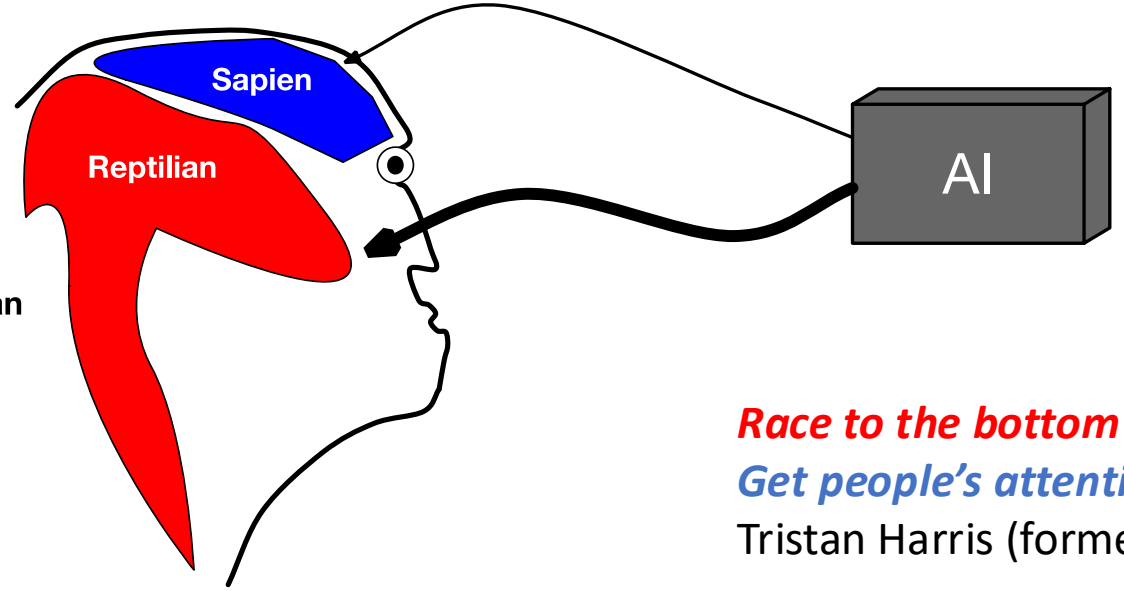

System 1: Reptilian
System 2: Sapien

Race to the bottom of the brainstem...
Get people's attention at all cost.
Tristan Harris (formerly with Google)

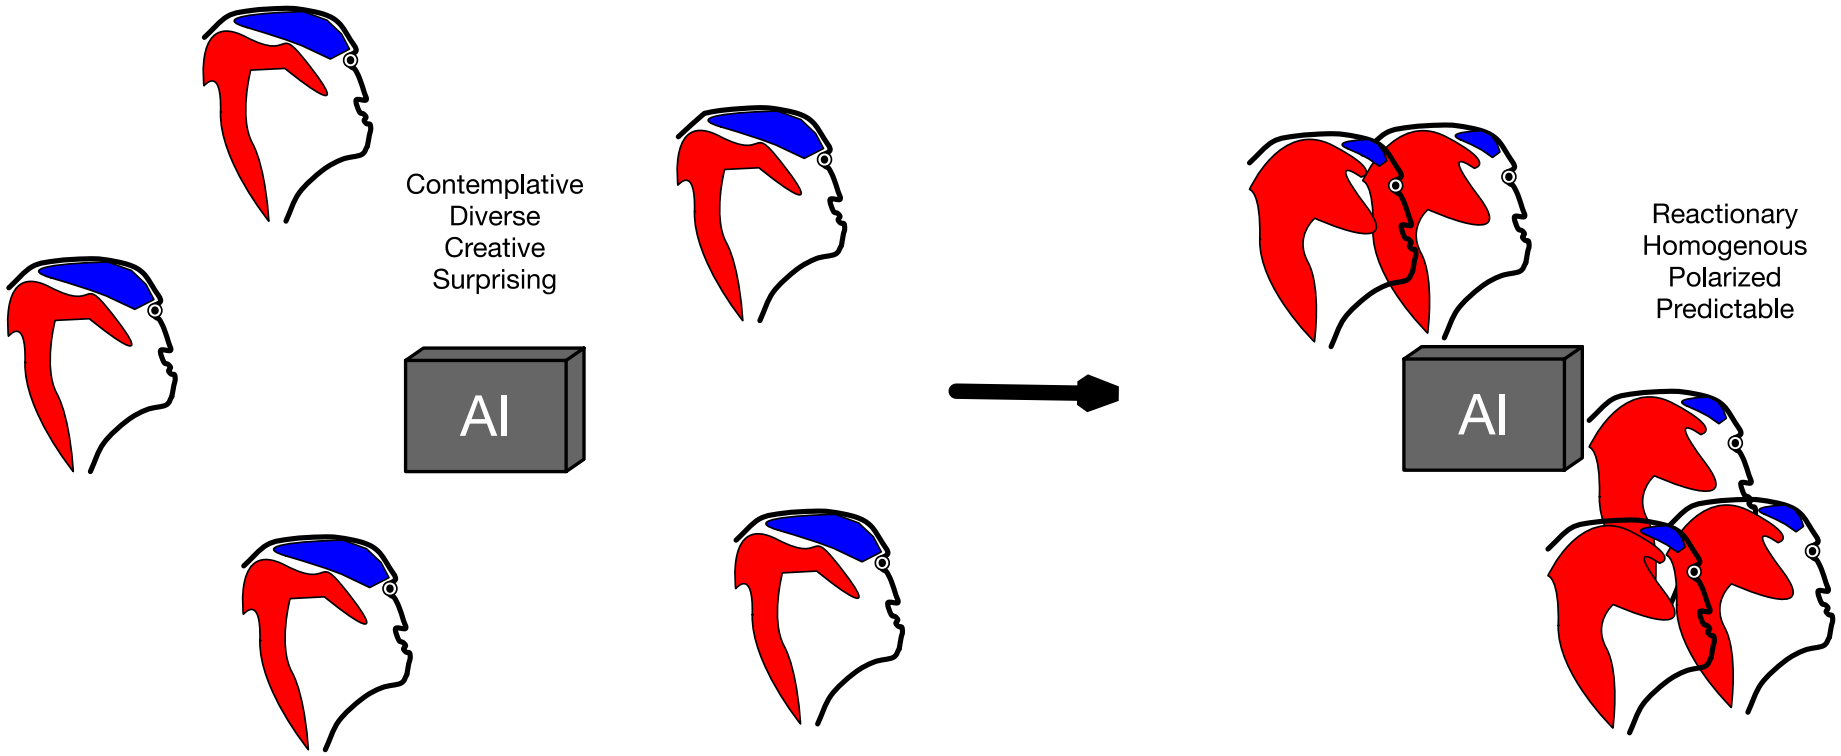

Disrupted Concentration => Weak Cognition

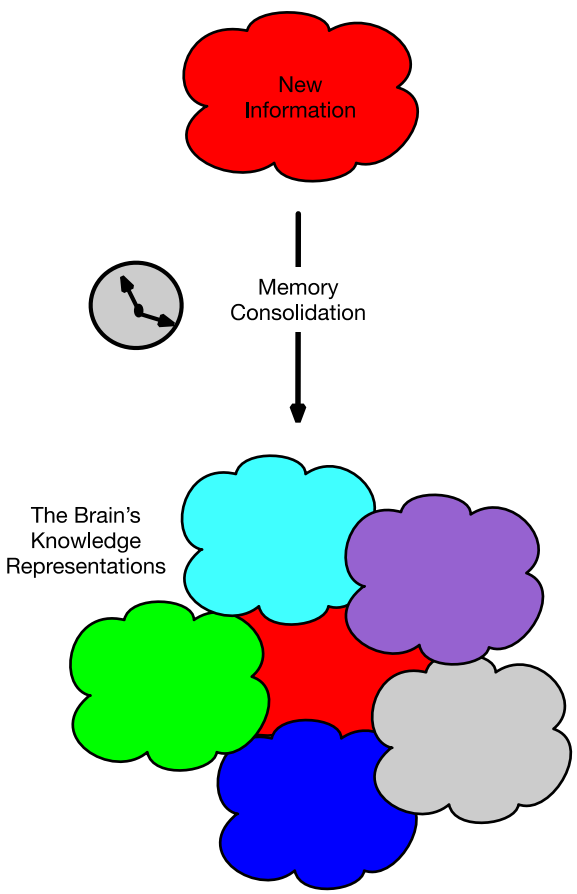

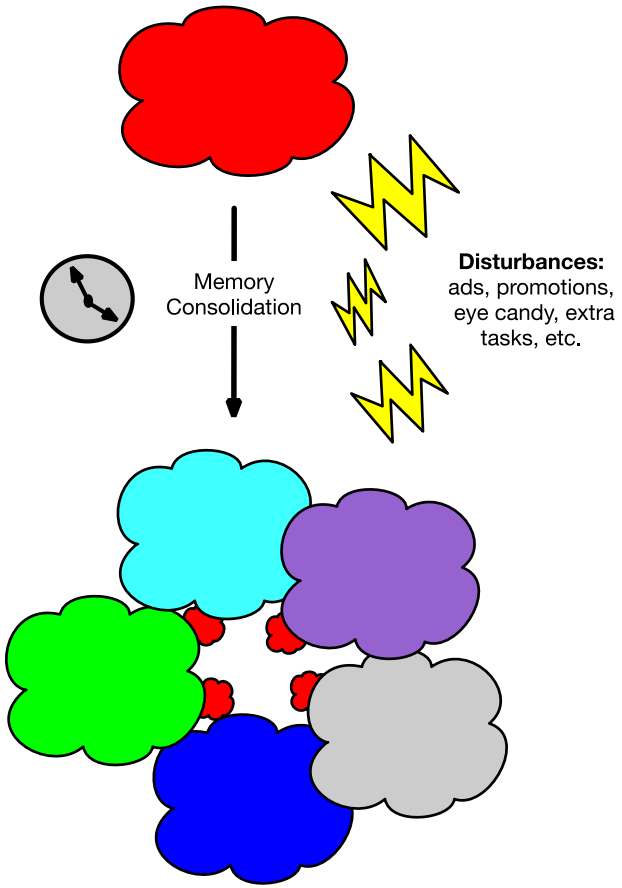

Low Retention, Shallow Thinking, Poor Problem-Solving

Do as AI Do

- Quantity and frequency of information can inhibit deep cognition.
- Engagement of emotions and reactive circuits of System 1 (reptilian brain), but little of System 2 (conscious reasoning)
- Attention span maxes out at 280 characters !!
- It thus becomes easier to just **do what the AI system suggests**, without thinking much about it. For example:
  - GPS
  - YouTube recommendations – down the rabbit hole of extremism
  - Music & book recommendations - *Meatloaf, again*?
  - Microsoft Teams suggestions for sentence extensions and completions.
  - Powerpoint's suggestions for formatting these slides.
  - **ChatGPT**

Part 4:
Opting Out

Surveillance Capitalism

Money for nothing and your clicks for free.

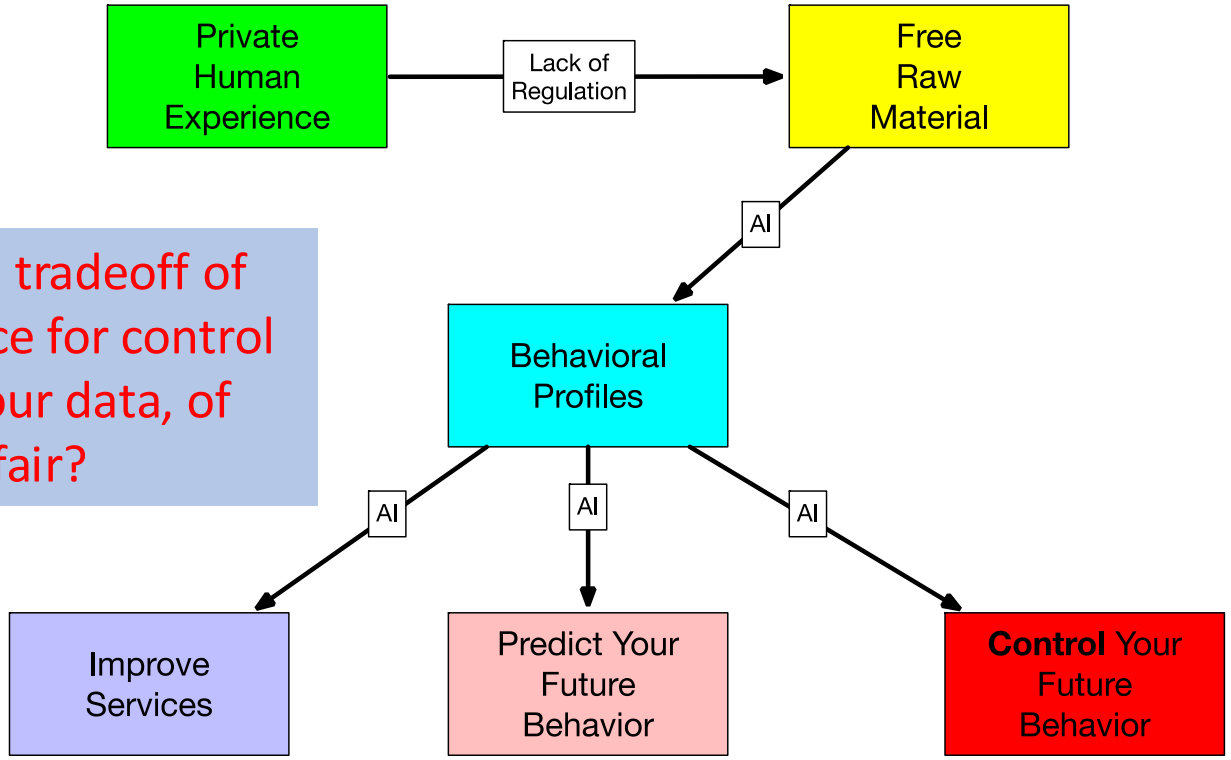

Is the tradeoff of service for control (of your data, of you) fair?

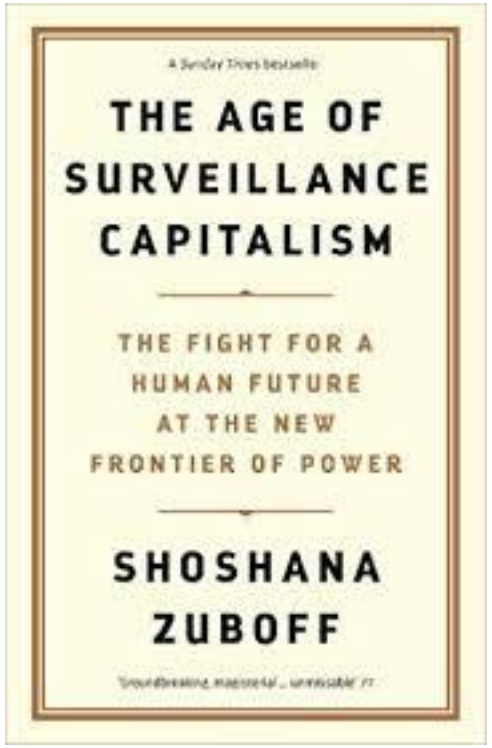

*These predictions are traded in a new futures market, where surveillance capitalists sell certainty to businesses determined to know **what we will do next**.*

*In the competition for certainty, surveillance capitalists learned the most predictive data comes not just from monitoring but also from **directing** behavior.*

*Our individual and social patterns are already built to give away our agency to mechanisms outside our consciousness – to our assumptions, to cues in our surroundings, to our peers …we're on the precipice of giving our entire way of life over to an invisible, imperceptible, irresistible instinct to hand difficult choices – even and especially the vitally important ones – to automated systems…Now we have to learn, in this moment, how to resist the pull of convenience and profit and protect the best aspects of who we are. (*Jacob Ward, *The Loop*, 2022*)*

Reclaim our tools as instruments of ourselves, as instruments of experience rather than just means of production (Nicholas Carr, *The Glass Cage*, 2014)

Human-AI Interaction Spectrum

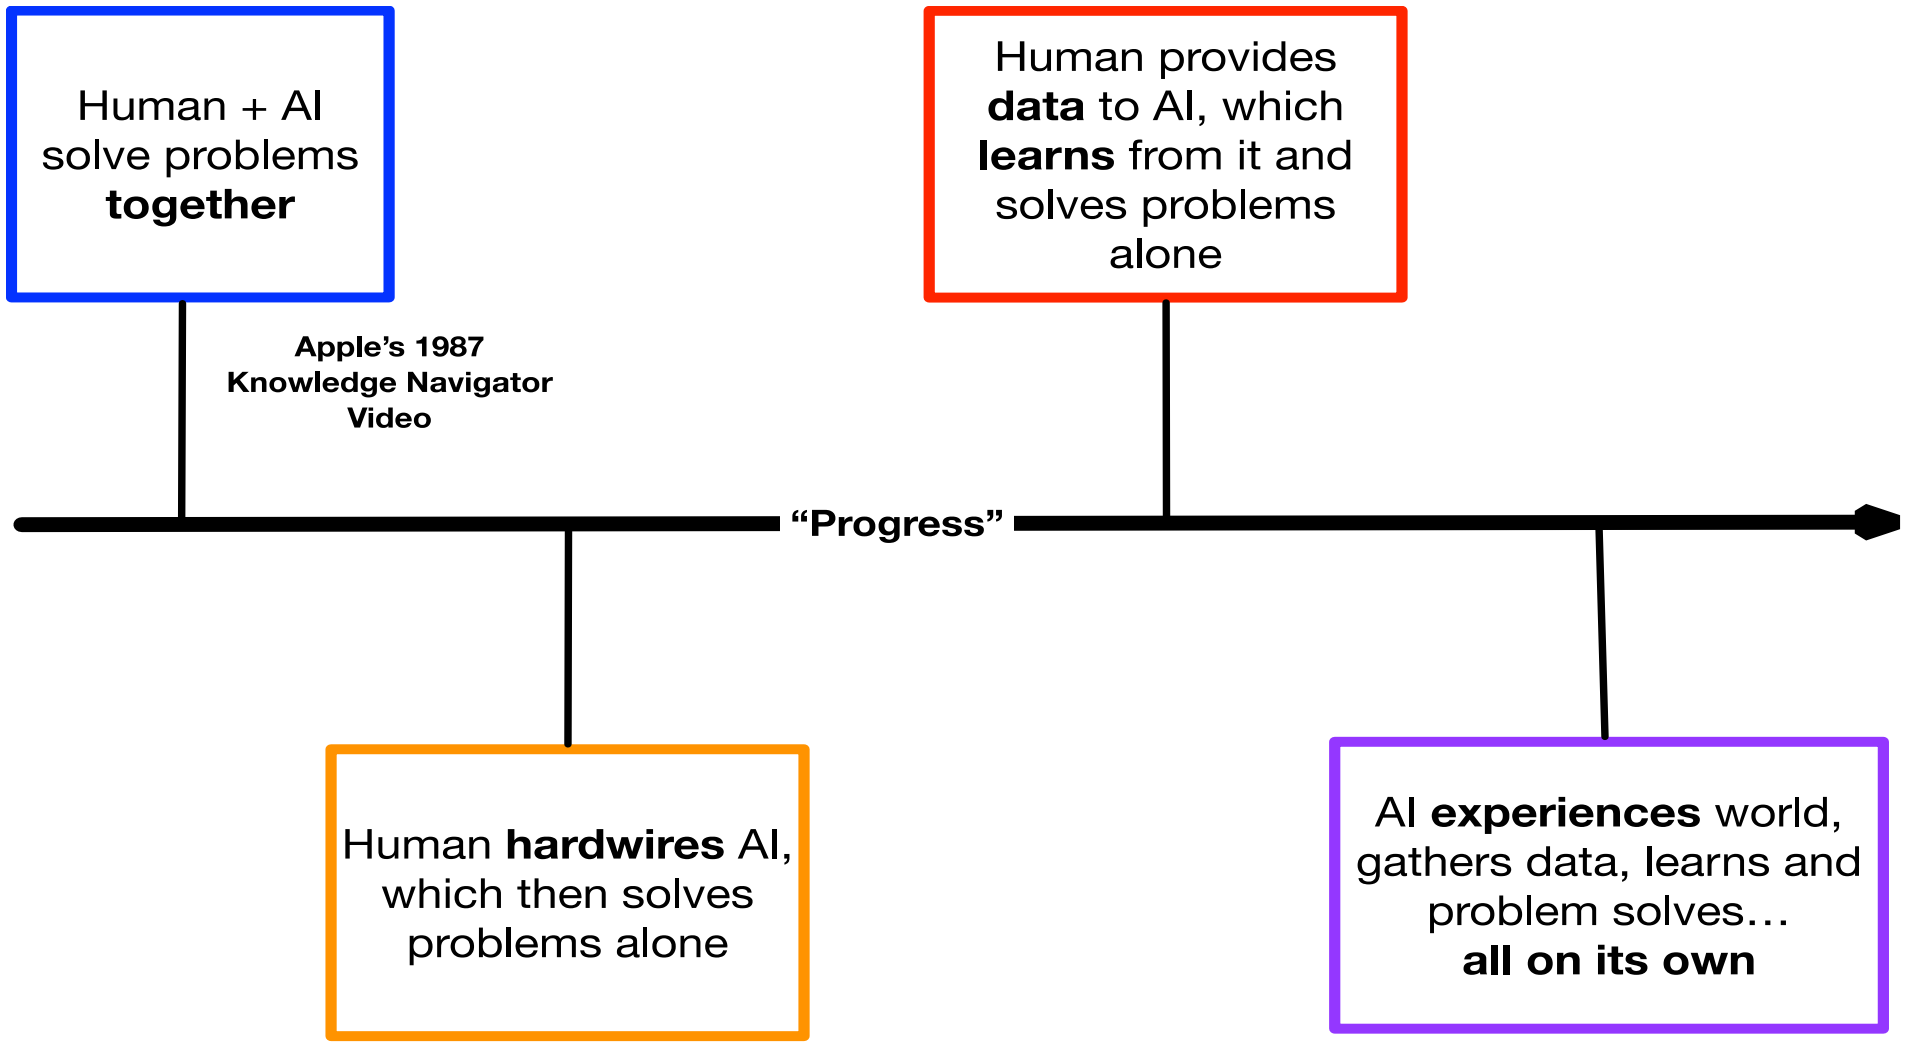

- Should AI help us learn, not just do the job itself?
- For the good of humanity, shouldn't we **move back left** ?

Don’t Become a Spoon

- AI will continue to make many jobs economically (and even morally) infeasible for anything but machine solutions.
- Humans need to keep **applying themselves** if they plan to **apply** for future jobs.

*When visiting an Asian worksite in the 1960’s, the great economist Milton Friedman was surprised to see workers using shovels, not modern machines. When told the reason: “It was a **jobs project**,” he responded:*

If it’s a jobs project, why not give them spoons instead of shovels?

- Humans must avoid becoming spoons… and AI research and development needs to be part of the solution, not the problem.

The Oppenheimer Movie

- Robert Oppenheimer: *I believe there is blood on my hands.*

- President Harry Truman: Do you think anyone in Hiroshima or Nagasaki gives a $#%& who built the bomb? You didn't drop the bomb. I did. Hiroshima isn't about you.

Attention Is All You Need

Ashish Vaswani* Google Brain avaswani@google.com
Noam Shazeer* Google Brain noam@google.com
Niki Parmar* Google Research nikip@google.com
Jakob Uszkoreit* Google Research usz@google.com
Llion Jones* Google Research llion@google.com
Aidan N. Gomez* † University of Toronto aidan@cs.toronto.edu
Lukasz Kaiser* Google Brain lukaszkaiser@google.com
Illia Polosukhin* ‡ illia.polosukhin@gmail.com

Abstract

The dominant sequence transduction models are based on complex recurrent or convolutional neural networks that include an encoder and a decoder. The best performing models also connect the encoder and decoder through an attention mechanism. We propose a new simple network architecture, the Transformer, based solely on attention mechanisms, dispensing with recurrence and convolutions entirely. Experiments on two machine translation tasks show these models to be superior in quality while being more parallelizable and requiring significantly less time to train. Our model achieves 28.4 BLEU on the WMT 2014 English-to-German translation task, improving over the existing best results, including ensembles, by over 2 BLEU. On the WMT 2014 English-to-French translation task, our model establishes a new single-model state-of-the-art BLEU score of 41.8 after training for 3.5 days on eight GPUs, a small fraction of the training costs of the best models from the literature. We show that the Transformer generalizes well to other tasks by applying it successfully to English constituency parsing both with large and limited training data.

If this seems like exaggerated concern, see Tristan Harris:

- The AI Dilemma (YouTube, March 2023)
- Beyond The AI Dilemma (YouTube, Sept. 2023)

Who will AI be about?

AI Transformation Playbook,
Andrew Ng, landing.ai

Great insights, from a famous AI researcher,
for any organization thinking about using AI.

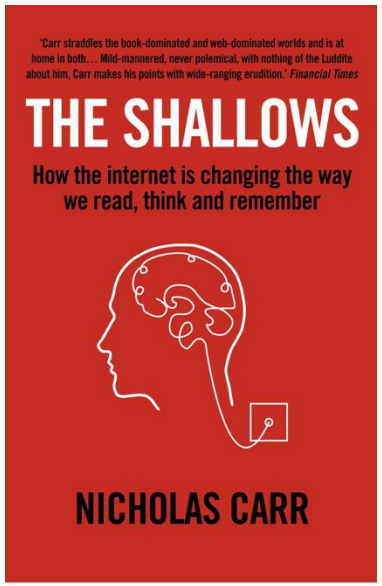

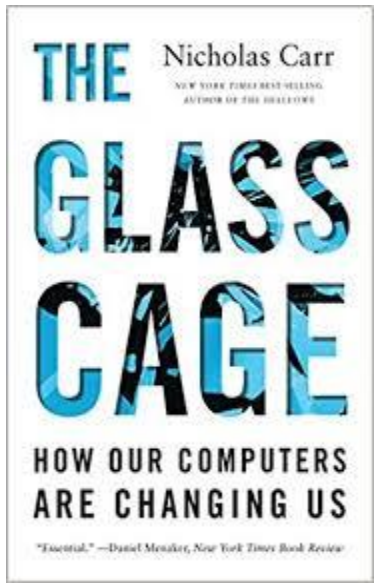

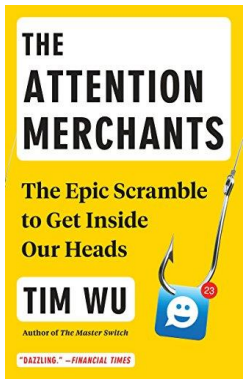

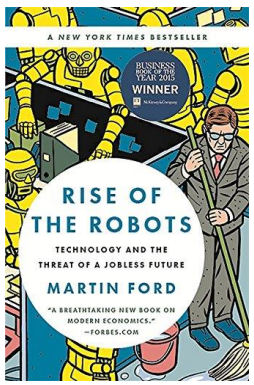

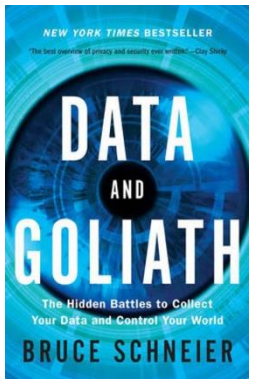

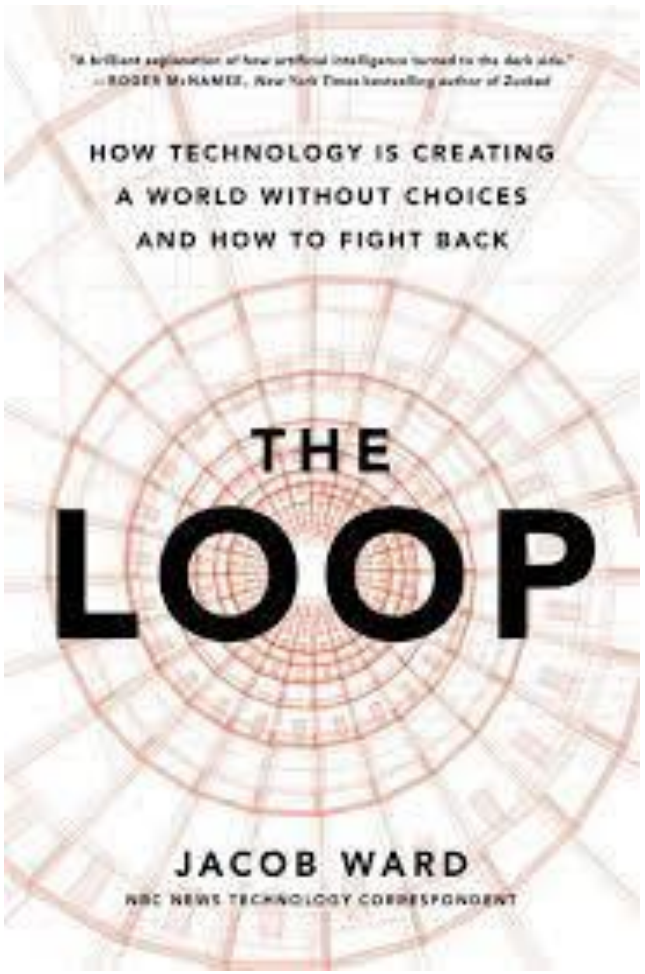

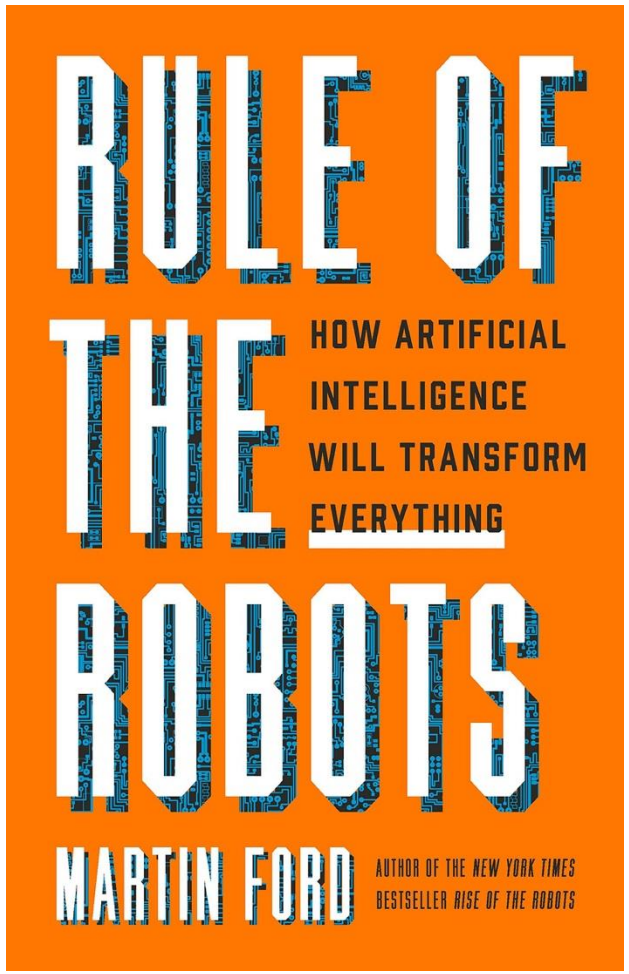

*Layout suggested by AI