AI; Autonomation; Anticipation

Technology trajectory & strategic choices for an accelerating world

Peter Coffee – VP for Strategic Research, Salesforce
pcoffee@salesforce.com

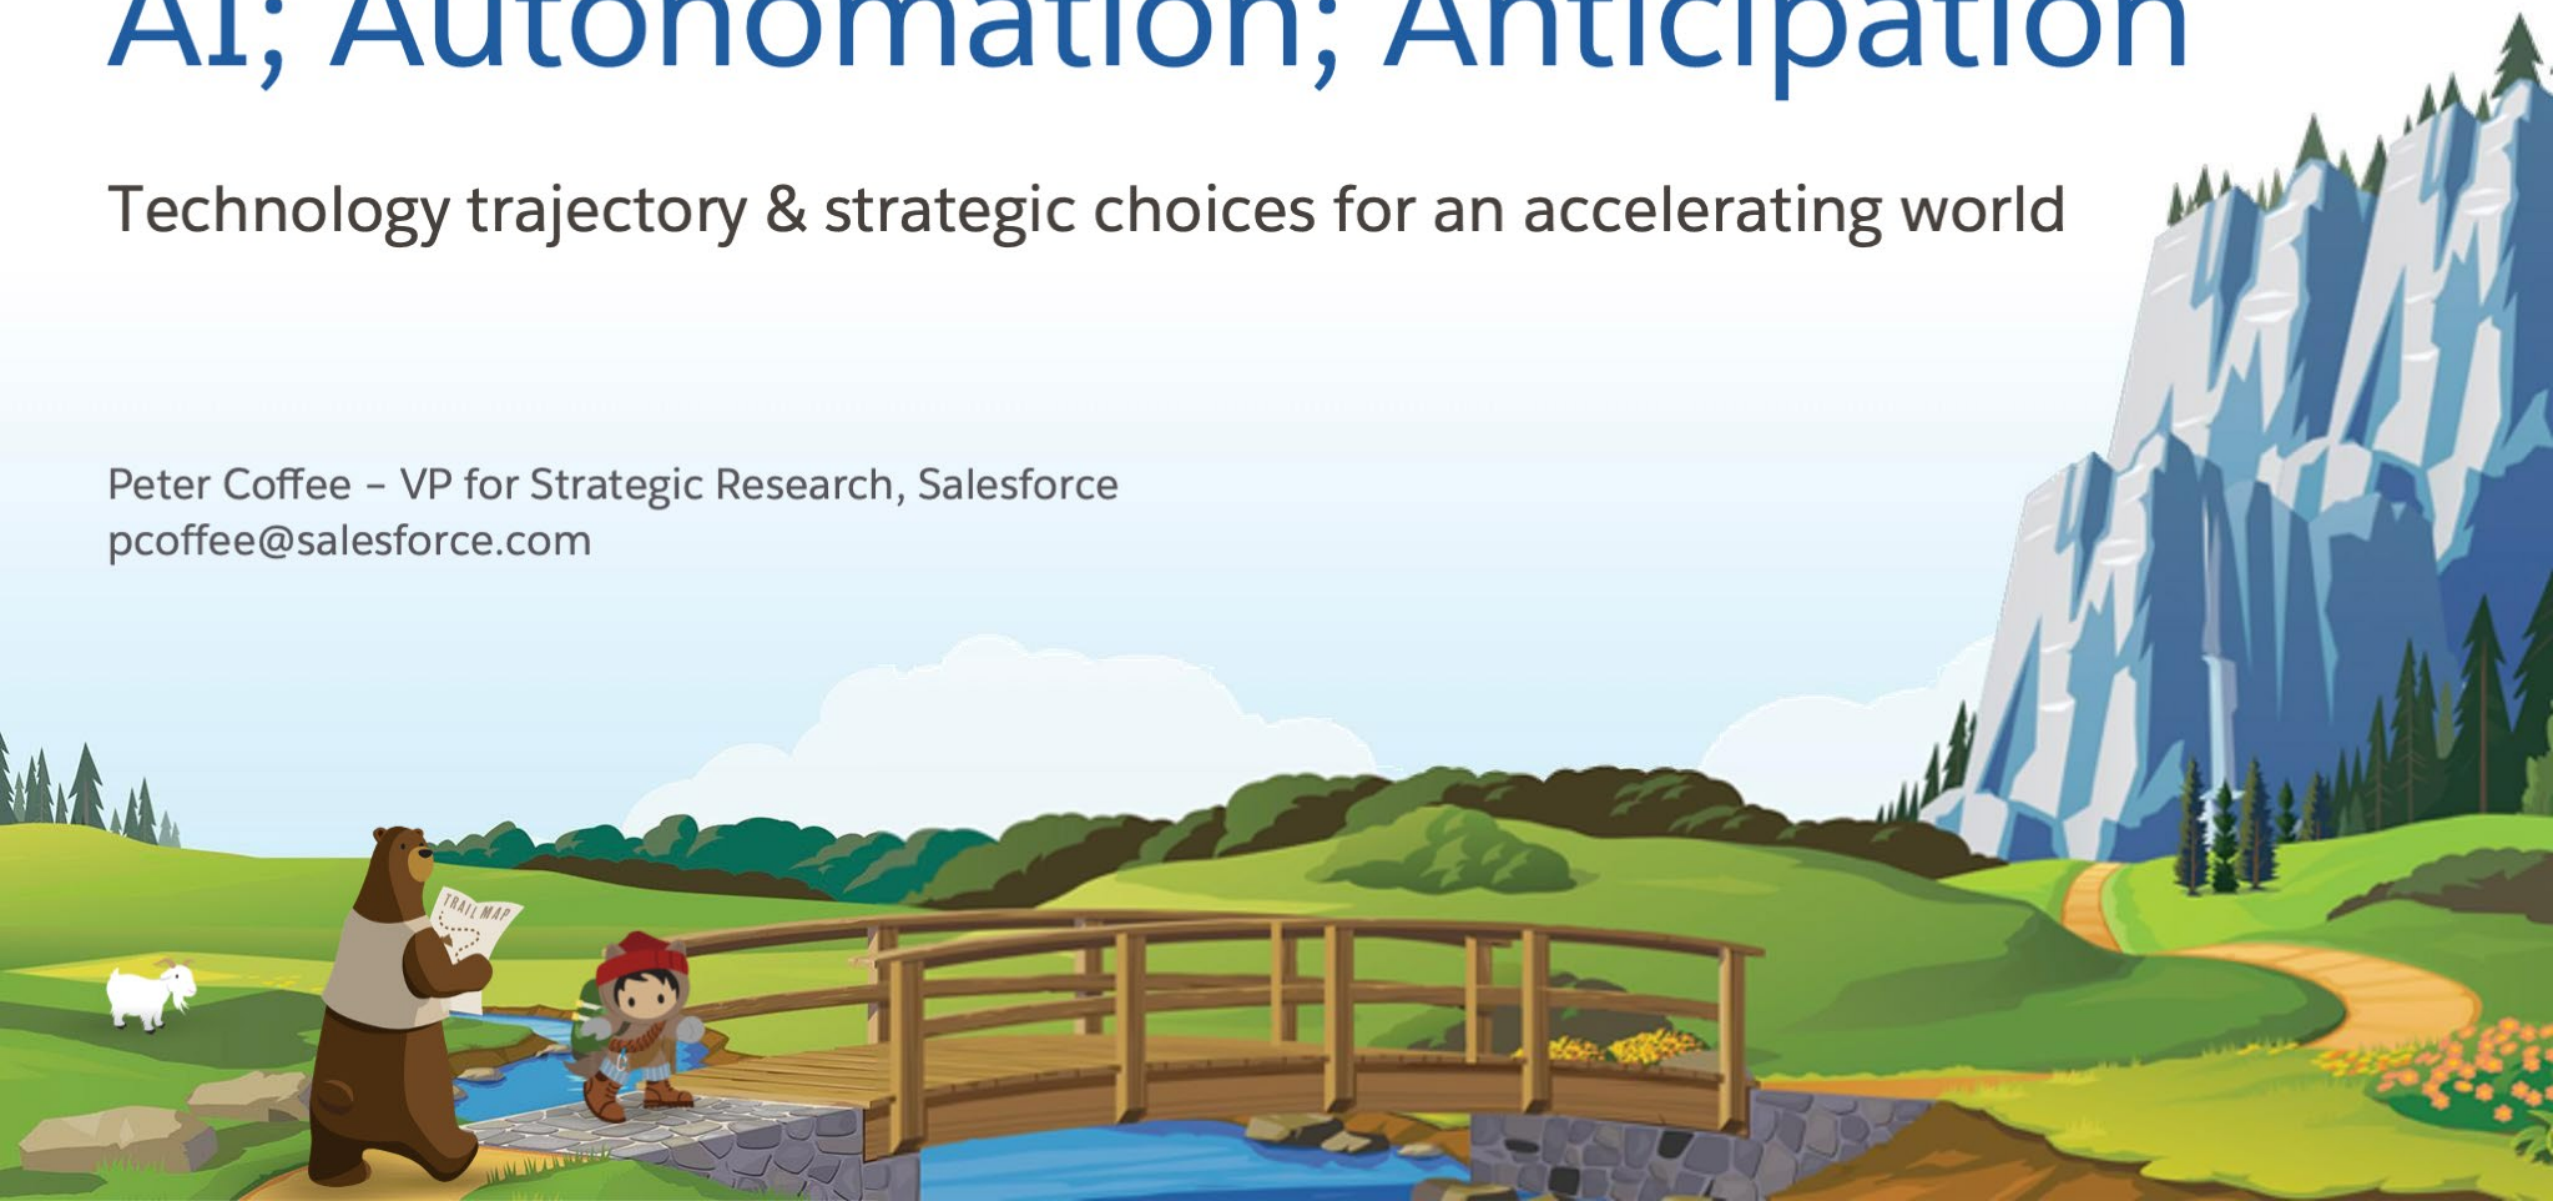

Why is AI *able* today to transition from "research" to "ROI"?

The challenge of “AI” was somewhat…underestimated

A PROPOSAL FOR THE DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

J. McCarthy, Dartmouth College
M. L. Minsky, Harvard University
N. Rochester, I.B.M. Corporation
C.E. Shannon, Bell Telephone Laboratories

August 31, 1955

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

- Automatic Computers: “The major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have.”
- Neuron Nets: “How can a set of (hypothetical) neurons be arranged so as to form concepts.”
- Self-Improvement: “Probably a truly intelligent machine will carry out activities which may best be described as self-improvement. Some schemes for doing this have been proposed and are worth further study.”

The challenge of “AI” was somewhat…underestimated

A PROPOSAL FOR THE DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

J. McCarthy, Dartmouth College
M. L. Minsky, Harvard University
N. Rochester, I.B.M. Corporation
C.E. Shannon, Bell Telephone Laboratories

August 31, 1955

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

- Automatic Computers: “The major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have.”

The challenge of “AI” was somewhat…underestimated

A PROPOSAL FOR THE DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

J. McCarthy, Dartmouth College
M. L. Minsky, Harvard University
N. Rochester, I.B.M. Corporation
C.E. Shannon, Bell Telephone Laboratories

August 31, 1955

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

- Automatic Computers: “The major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have.”

- Neuron Nets: “How can a set of (hypothetical) neurons be arranged so as to form concepts.”

The challenge of “AI” was somewhat…underestimated

A PROPOSAL FOR THE DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

J. McCarthy, Dartmouth College
M. L. Minsky, Harvard University
N. Rochester, I.B.M. Corporation
C.E. Shannon, Bell Telephone Laboratories

August 31, 1955

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

- Automatic Computers: “The major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have.” ❌
- Neuron Nets: “How can a set of (hypothetical) neurons be arranged so as to form concepts.” 🤔
- Self-Improvement: “Probably a truly intelligent machine will carry out activities which may best be described as self-improvement. Some schemes for doing this have been proposed and are worth further study.” ✅

This was a 1980s notion of adequate “AI hardware”

4404
CPU: 68010@10 MHz + NS32081 FPU
RAM: 1 – 4 MB
Resolution: 640 x 480
CRT: 13″
HDD: 45 MB
FDD: 320 kB
Interfaces: RS232C, LAN, SCSI, centronics
OS: Unix-like OS
Initial date: 1984
Initial price: $14950

4404
CPU: 68010@10 MHz + NS32081 FPU
RAM: 1 – 4 MB
Resolution: 640 x 480
CRT: 13″
HDD: 45 MB
FDD: 320 kB
Interfaces: RS232C, LAN, SCSI, centronics
OS: Unix-like OS
Initial date: 1984
Initial price: $14950

A crucial new abundance: computing power

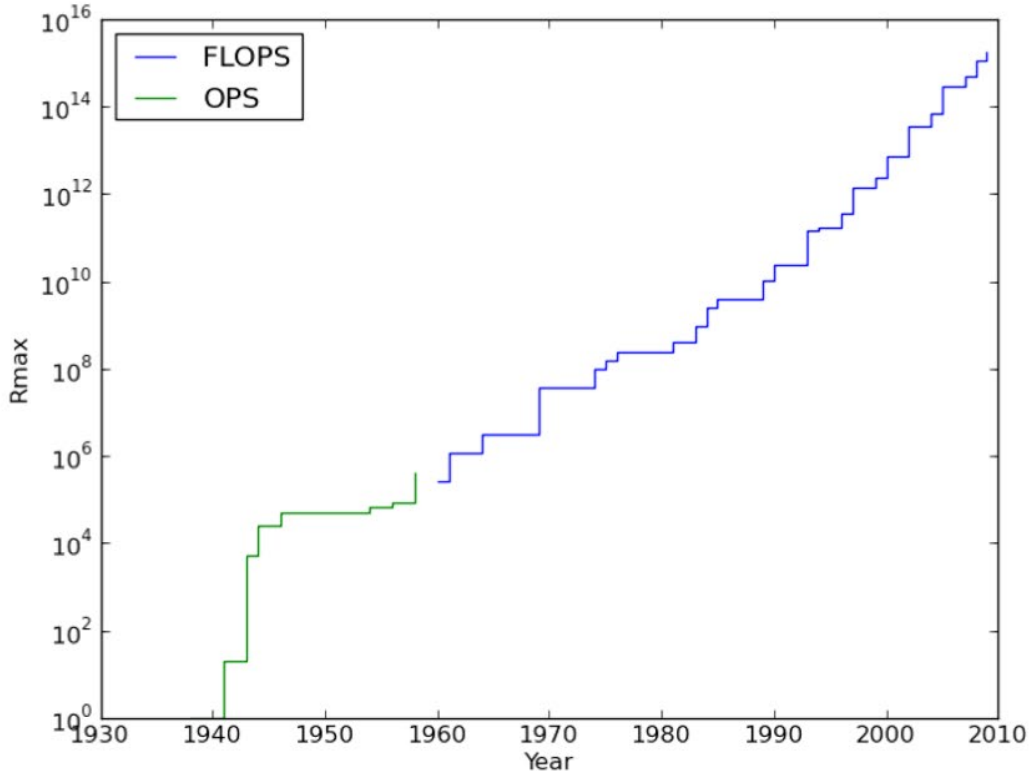

Note: these are *logarithmic* vertical scales: a straight line represents exponential growth. Showing aggregate compute power on a linear scale makes it look as if we started building them in 2005...

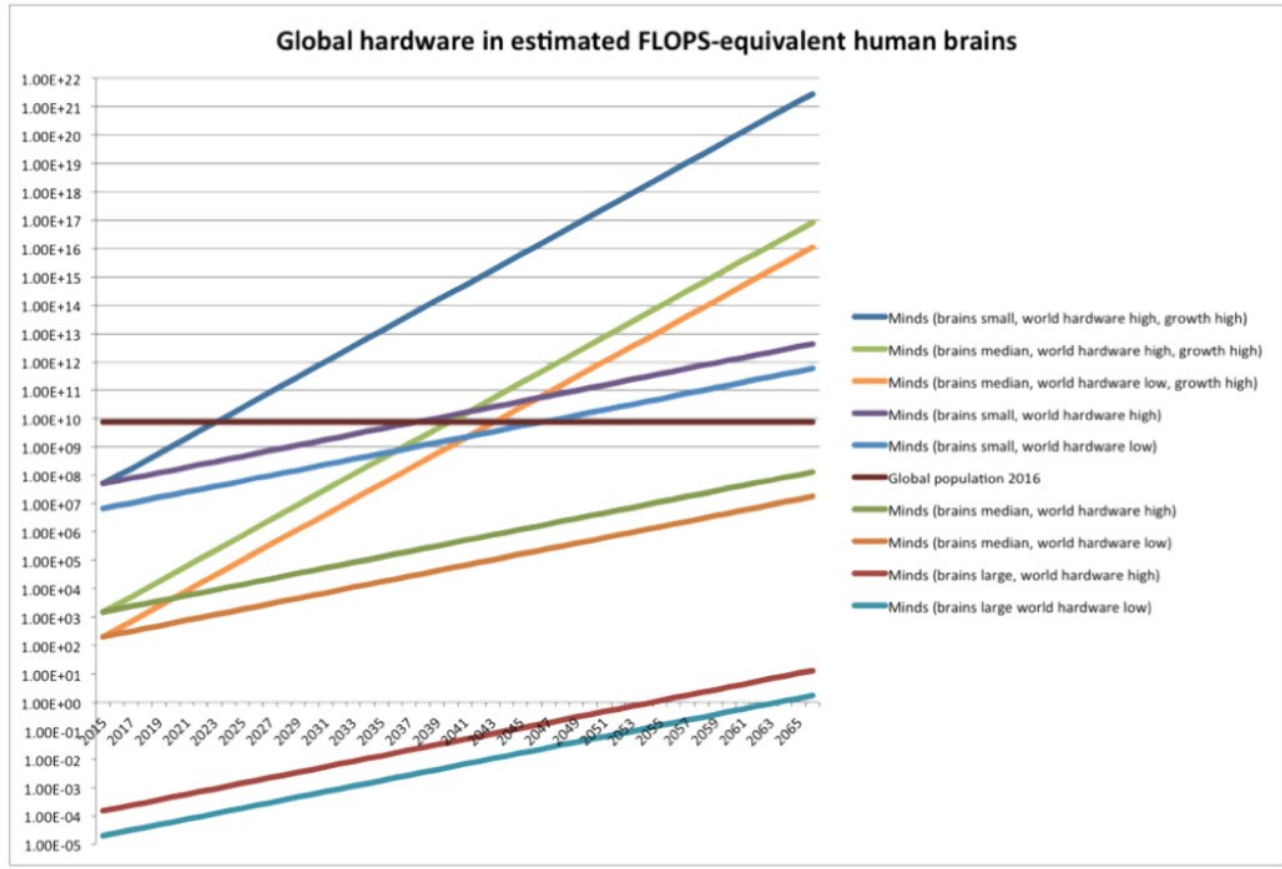

Figure: Projected number of human brains equivalent to global hardware under various assumptions. For brains, 'small' = 3 x 10^ 13, 'median' = 10^18, 'large' = 10^25. For 'world hardware', 'high' =2 x 10^20, 'low' = 1.5 x 10^21. 'Growth' is growth in computing hardware, the unlabeled default used in most projections is 25% per annum (our estimate above), 'high' = 86% per annum (the apparent growth rate in ASIC hardware in around 2007).

Eleven orders of magnitude improvement

www.economist.com/technology-quarterly/2019/09/12/drastic-falls-in-cost-are-powering-another-computer-revolution

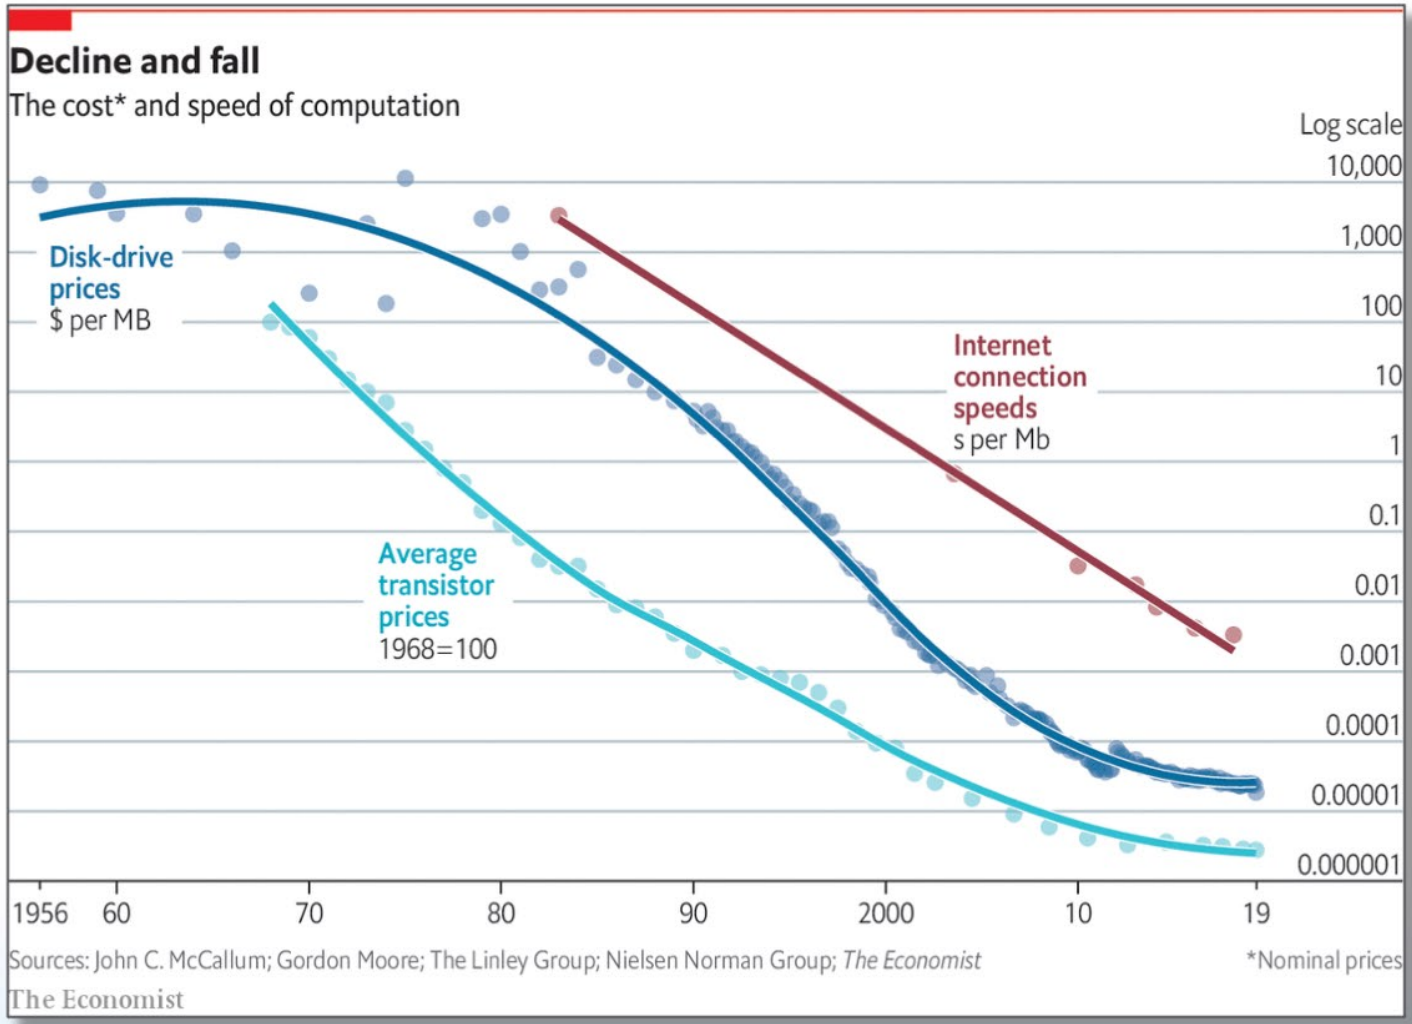

"The price of computation today is roughly one hundred-millionth what it was in the 1970s, when the first microprocessors became commercially available...**between 1950 and 2010 the amount of number-crunching possible with a kilowatt-hour of energy grew roughly a hundred-billion-fold.**"

– The Economist, September 2019

Hardware capacity enables non-trivial “neural nets”

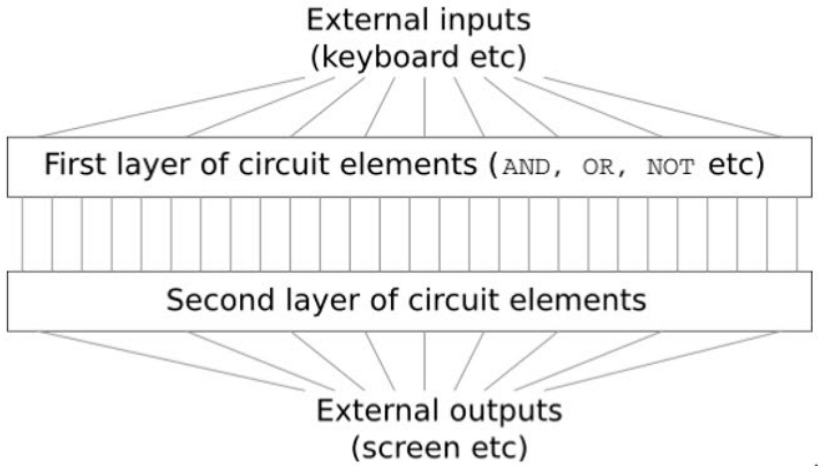

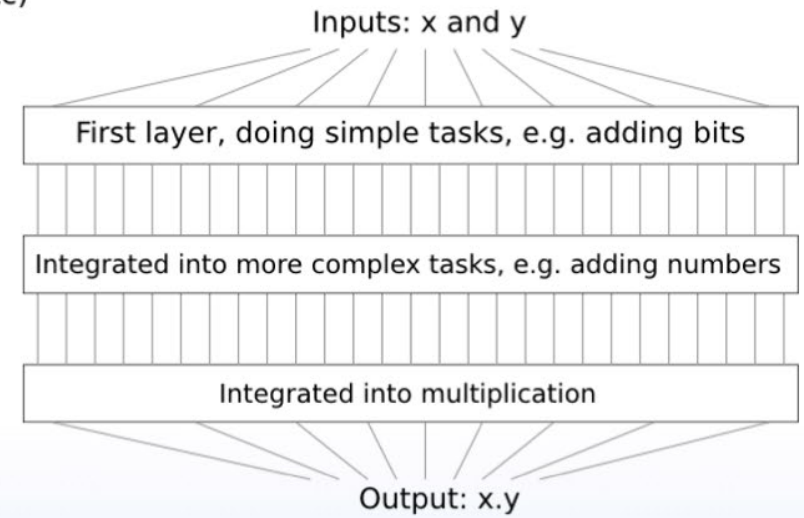

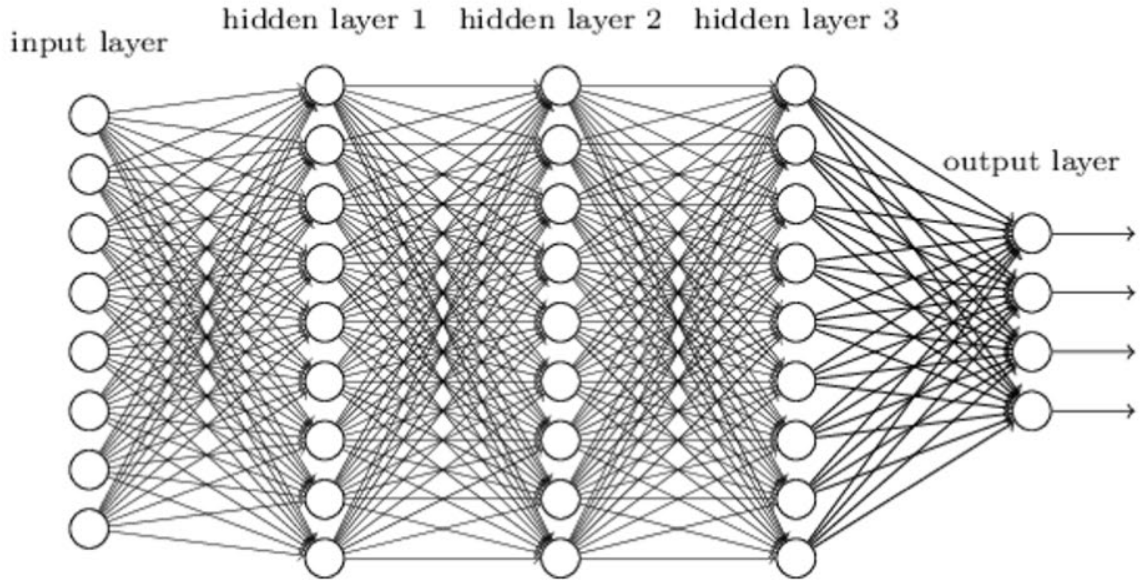

“Multiple layers of abstraction seem likely to give deep networks a compelling advantage in learning to solve complex pattern recognition problems… Deep networks are intrinsically more powerful than shallow networks.”

neuralnetworksanddeeplearning.com/chap5.html

Self-improvement is essential

ELSEVIER

Data & Knowledge Engineering 25 (1998) 161–197

DATA & KNOWLEDGE ENGINEERING

Knowledge Engineering: Principles and methods

[b] *Artificial Intelliger*

[c] *Dept. of Social S*

Abstract

This paper gives an ov
the paradigm shift from a
in Knowledge Engineeri
evolved in recent years w
is supplemented by disc
knowledge-based systen
Knowledge Engineering to Software Engineering, Information Integration and Knowledge Management.

"The transfer of this technology into commercial use in order to build large Knowledge-Based Systems failed in many cases... The means to develop small academic prototypes did not scale up to the design and maintenance of large long-living commercial systems."

Abstract

XCON is a rule-based expert system that configures computer systems. Over 7 years, XCON has grown to 6,200 rules, of which approximately 50% change every year. While the performance of XCON is satisfactory, it is increasingly becoming more difficult to change. With the goal of facilitating maintenance, DEC has developed a new rule-based language, RIME, in which the successor to XCON, XCON-in-RIME, is being written. This paper evaluates the potential for enhanced maintainability of XCON-in-RIME over XCON.

I. Introduction: Motivation and Goals

The following properties of XCON, an expert system, make it a particularly interesting system to examine:

- *XCON performs a complex design task:* XCON configures computer systems for DEC; XCON is used in a production mode, day in, day out -- it has been used since January 1980.
- *XCON is a very large rule-based system:* currently there are approximately 6,200 rules in XCON, which draw on a database of approximately 20,000 parts.
- *XCON undergoes constant change:* 50% of the rules in XCON are changed each year.

Self-improvement is essential

ELSEVIER

Data & Knowledge Engineering 25 (1998) 161–197

DATA & KNOWLEDGE ENGINEERING

Knowledge Engineering: Principles and methods

[b] Artificial Intelliger

[c] Dept. of Social S

Abstract

This paper gives an ov
the paradigm shift from a
in Knowledge Engineeri
evolved in recent years w
is supplemented by disc
knowledge-based systen
Knowledge Engineering

"While the expert may consciously articulate some parts of his or her knowledge, he or she will not be aware of a significant part of this knowledge since it is hidden in his or her skills.

"This knowledge acquisition process is no longer seen as a transfer of knowledge into an appropriate computer representation, but as a model construction process."

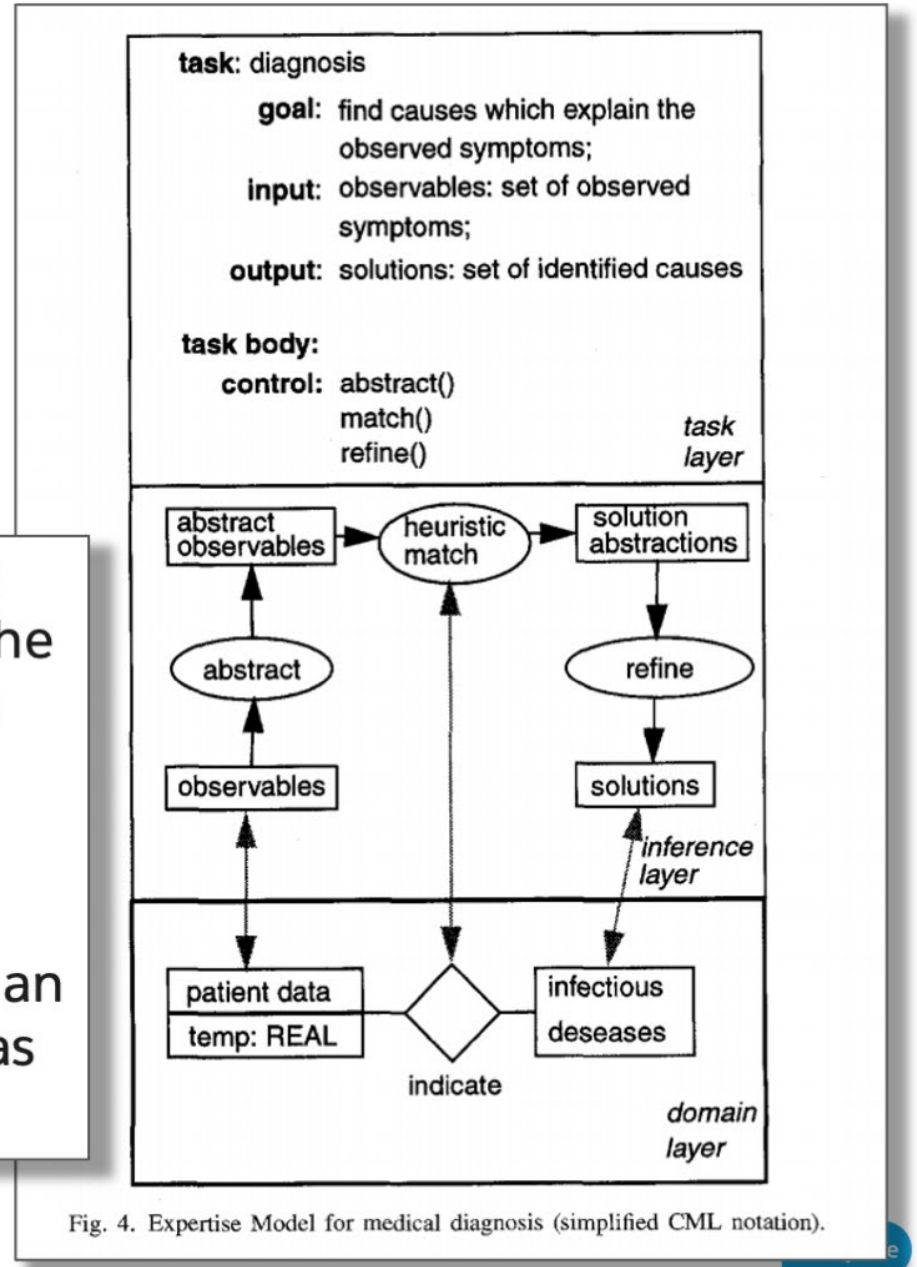

Fig. 4. Expertise Model for medical diagnosis (simplified CML notation).

Data flow enables “thinking” to become “knowing”

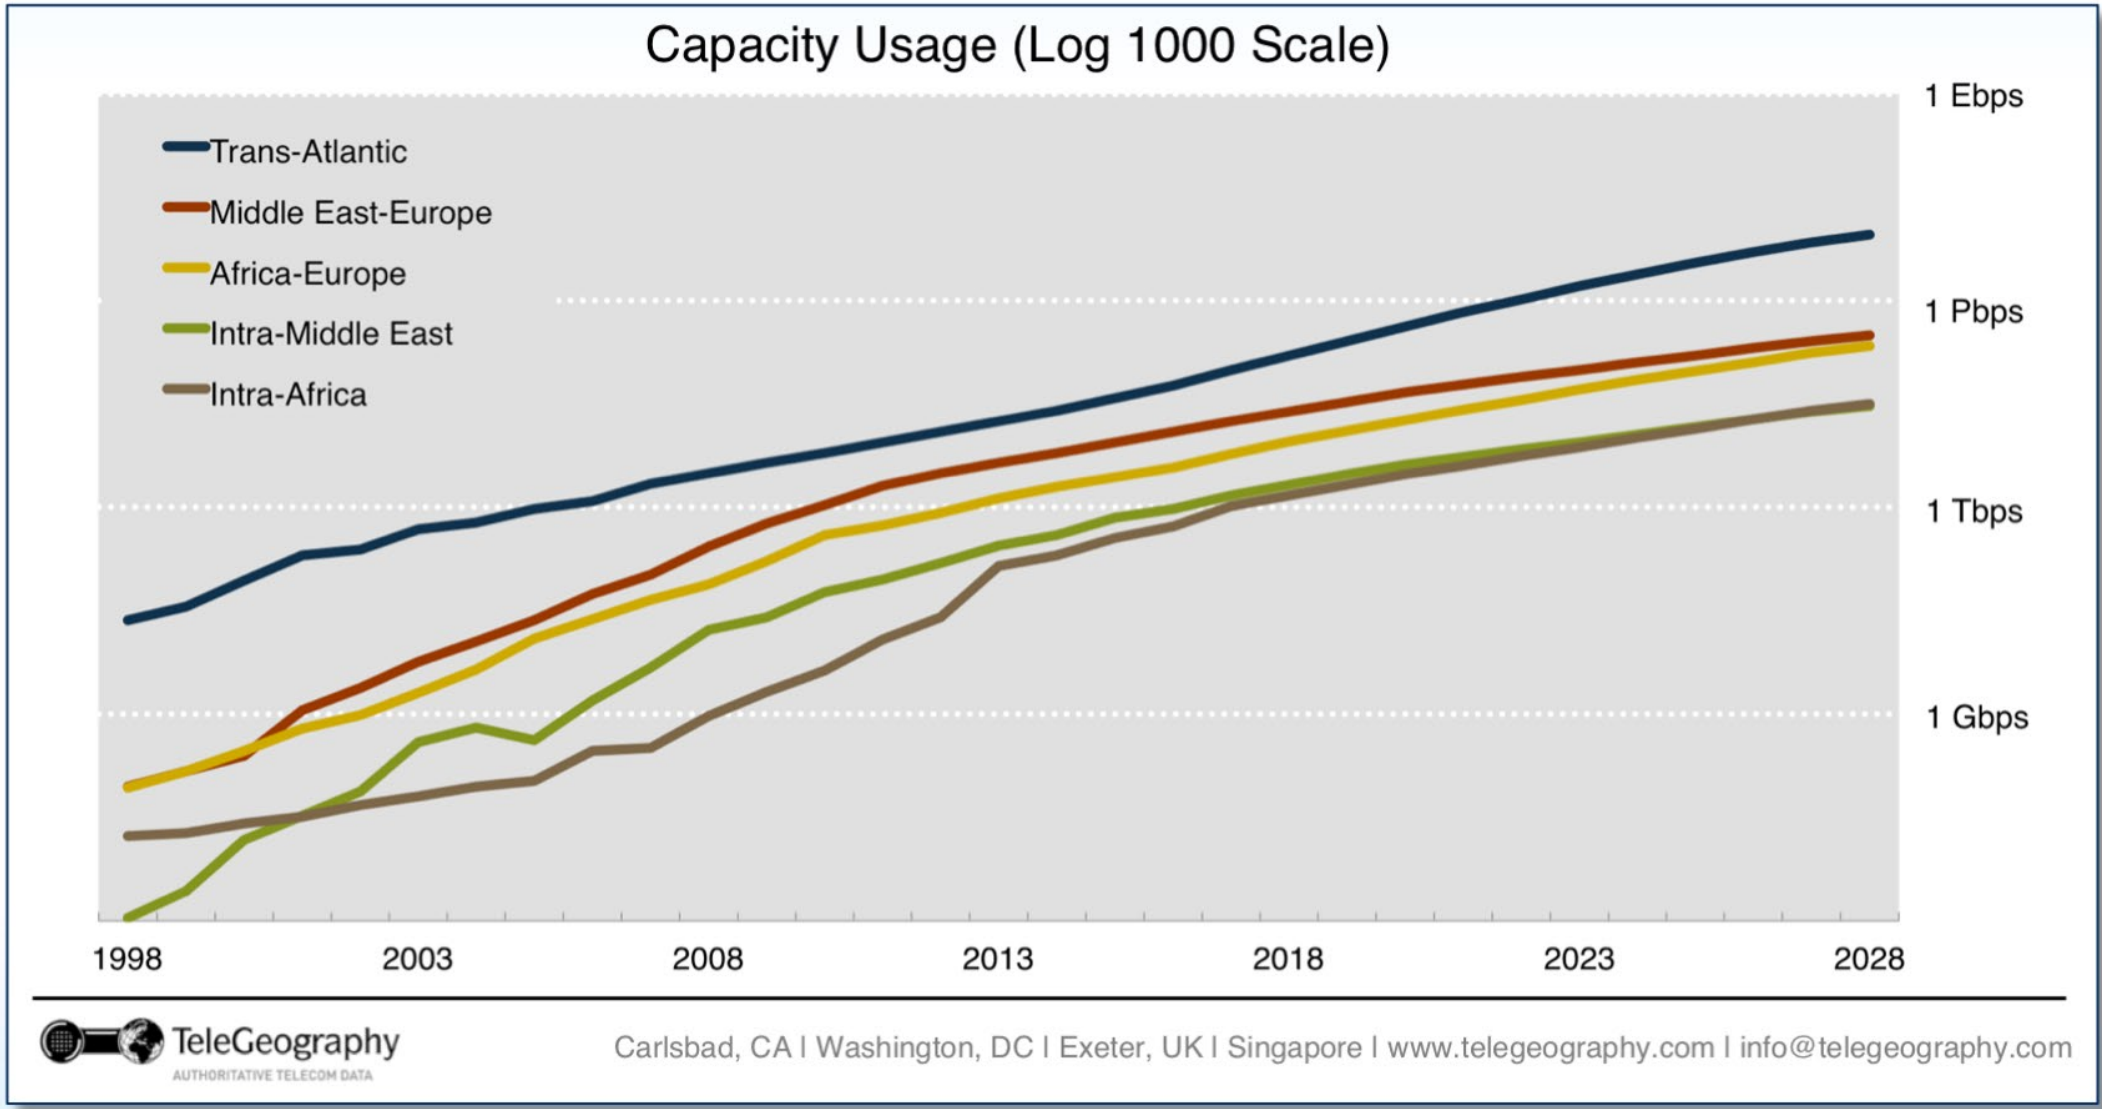

Why is AI *needed* today – enough to justify its disruptions?

It’s not about money

“Fed staff economists in Washington, who had been predicting a mild recession late this year, no longer expect one.”

– *The New York Times*, 27 July

“The U.S. economy may be headed toward what the Federal Reserve’s been hoping for - a soft landing. That means inflation has been brought under control without triggering a recession.”

– National Public Radio, 30 July

It’s not about health

“The number of Covid deaths has now dropped low enough that they are difficult to notice in the overall death data”

– *The New York Times*, 17 July

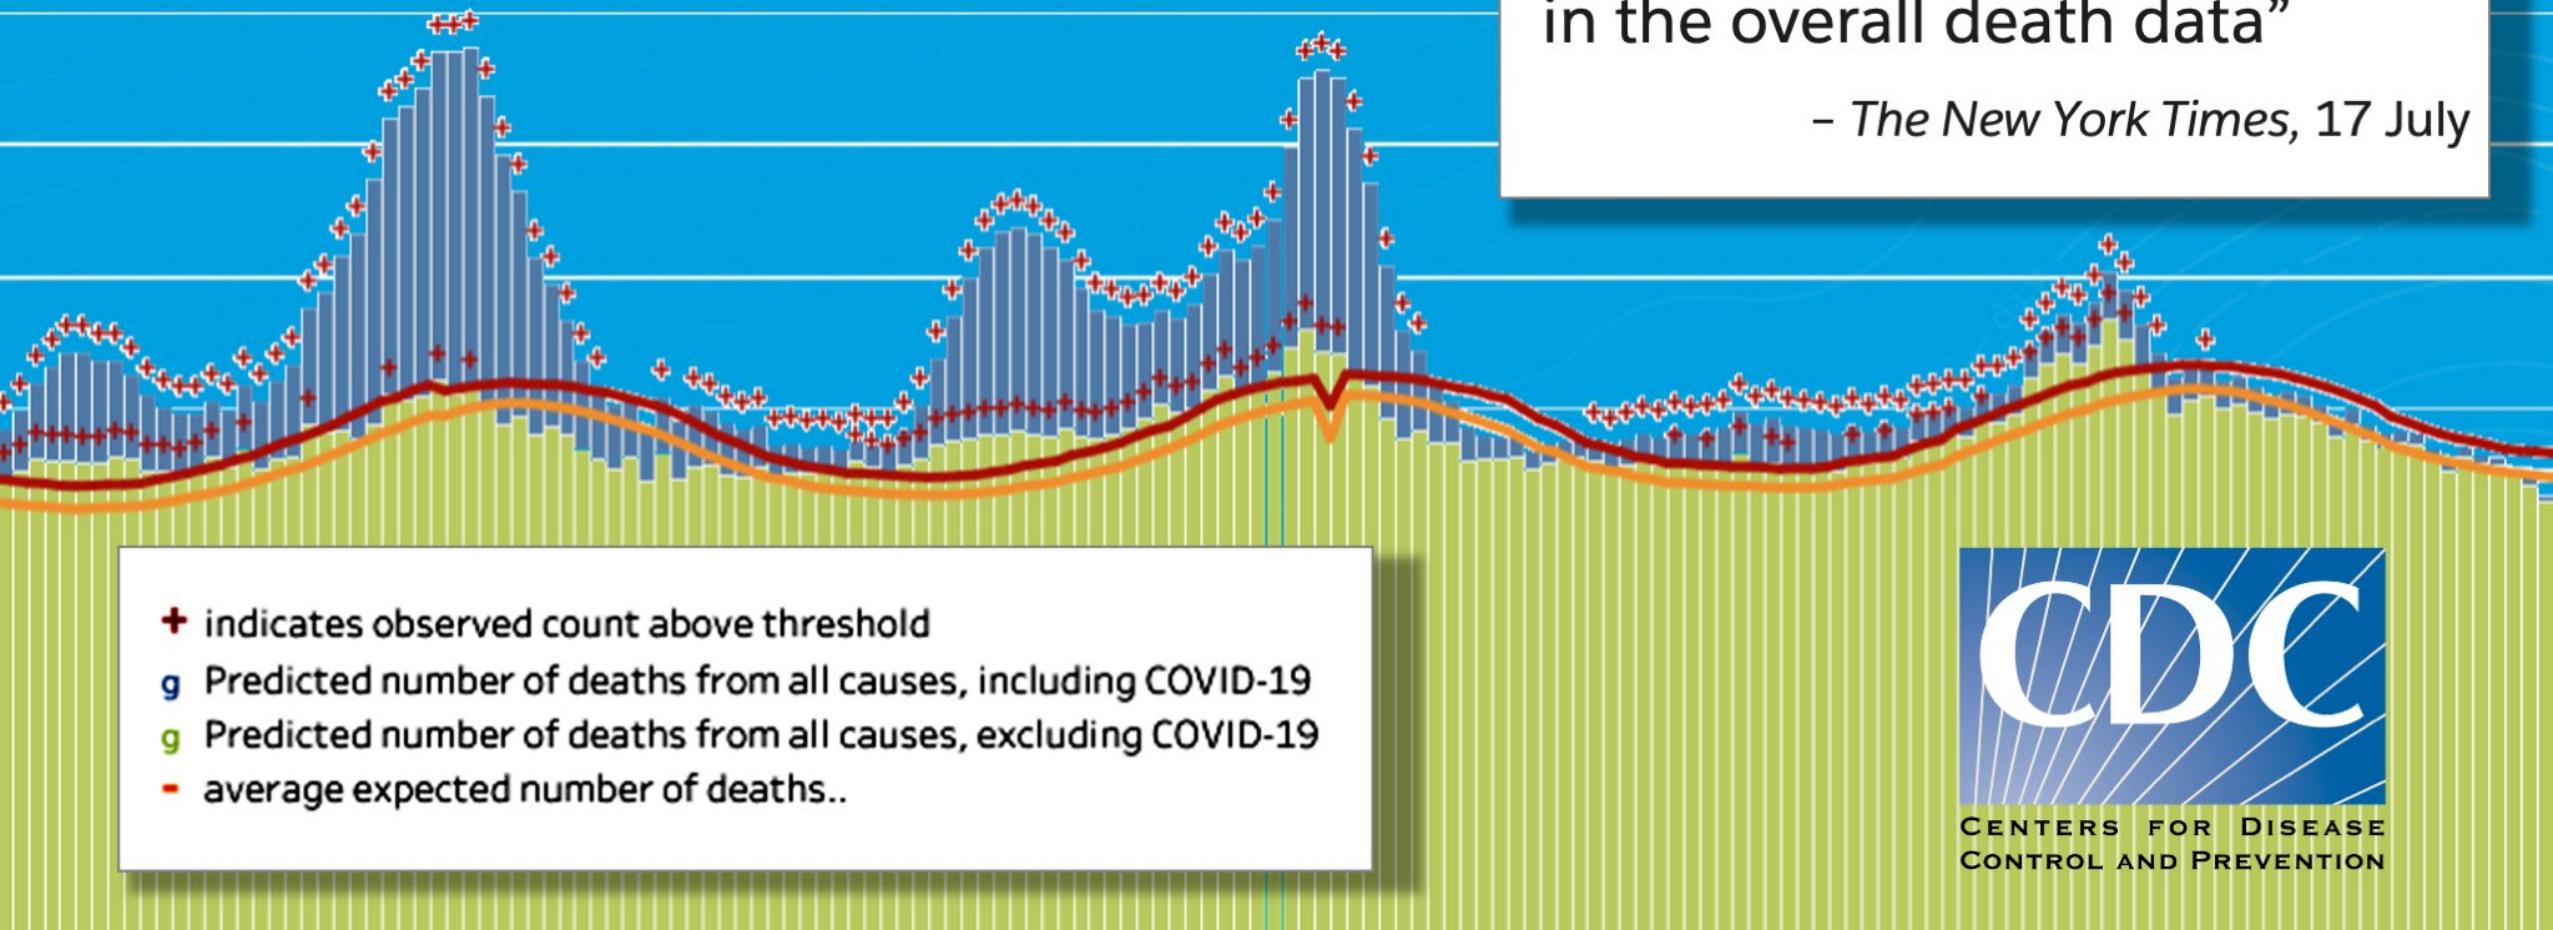

+ indicates observed count above threshold
g Predicted number of deaths from all causes, including COVID-19
g Predicted number of deaths from all causes, excluding COVID-19
- average expected number of deaths..

CDC
CENTERS FOR DISEASE CONTROL AND PREVENTION

The word of the **decade** is “resilience”

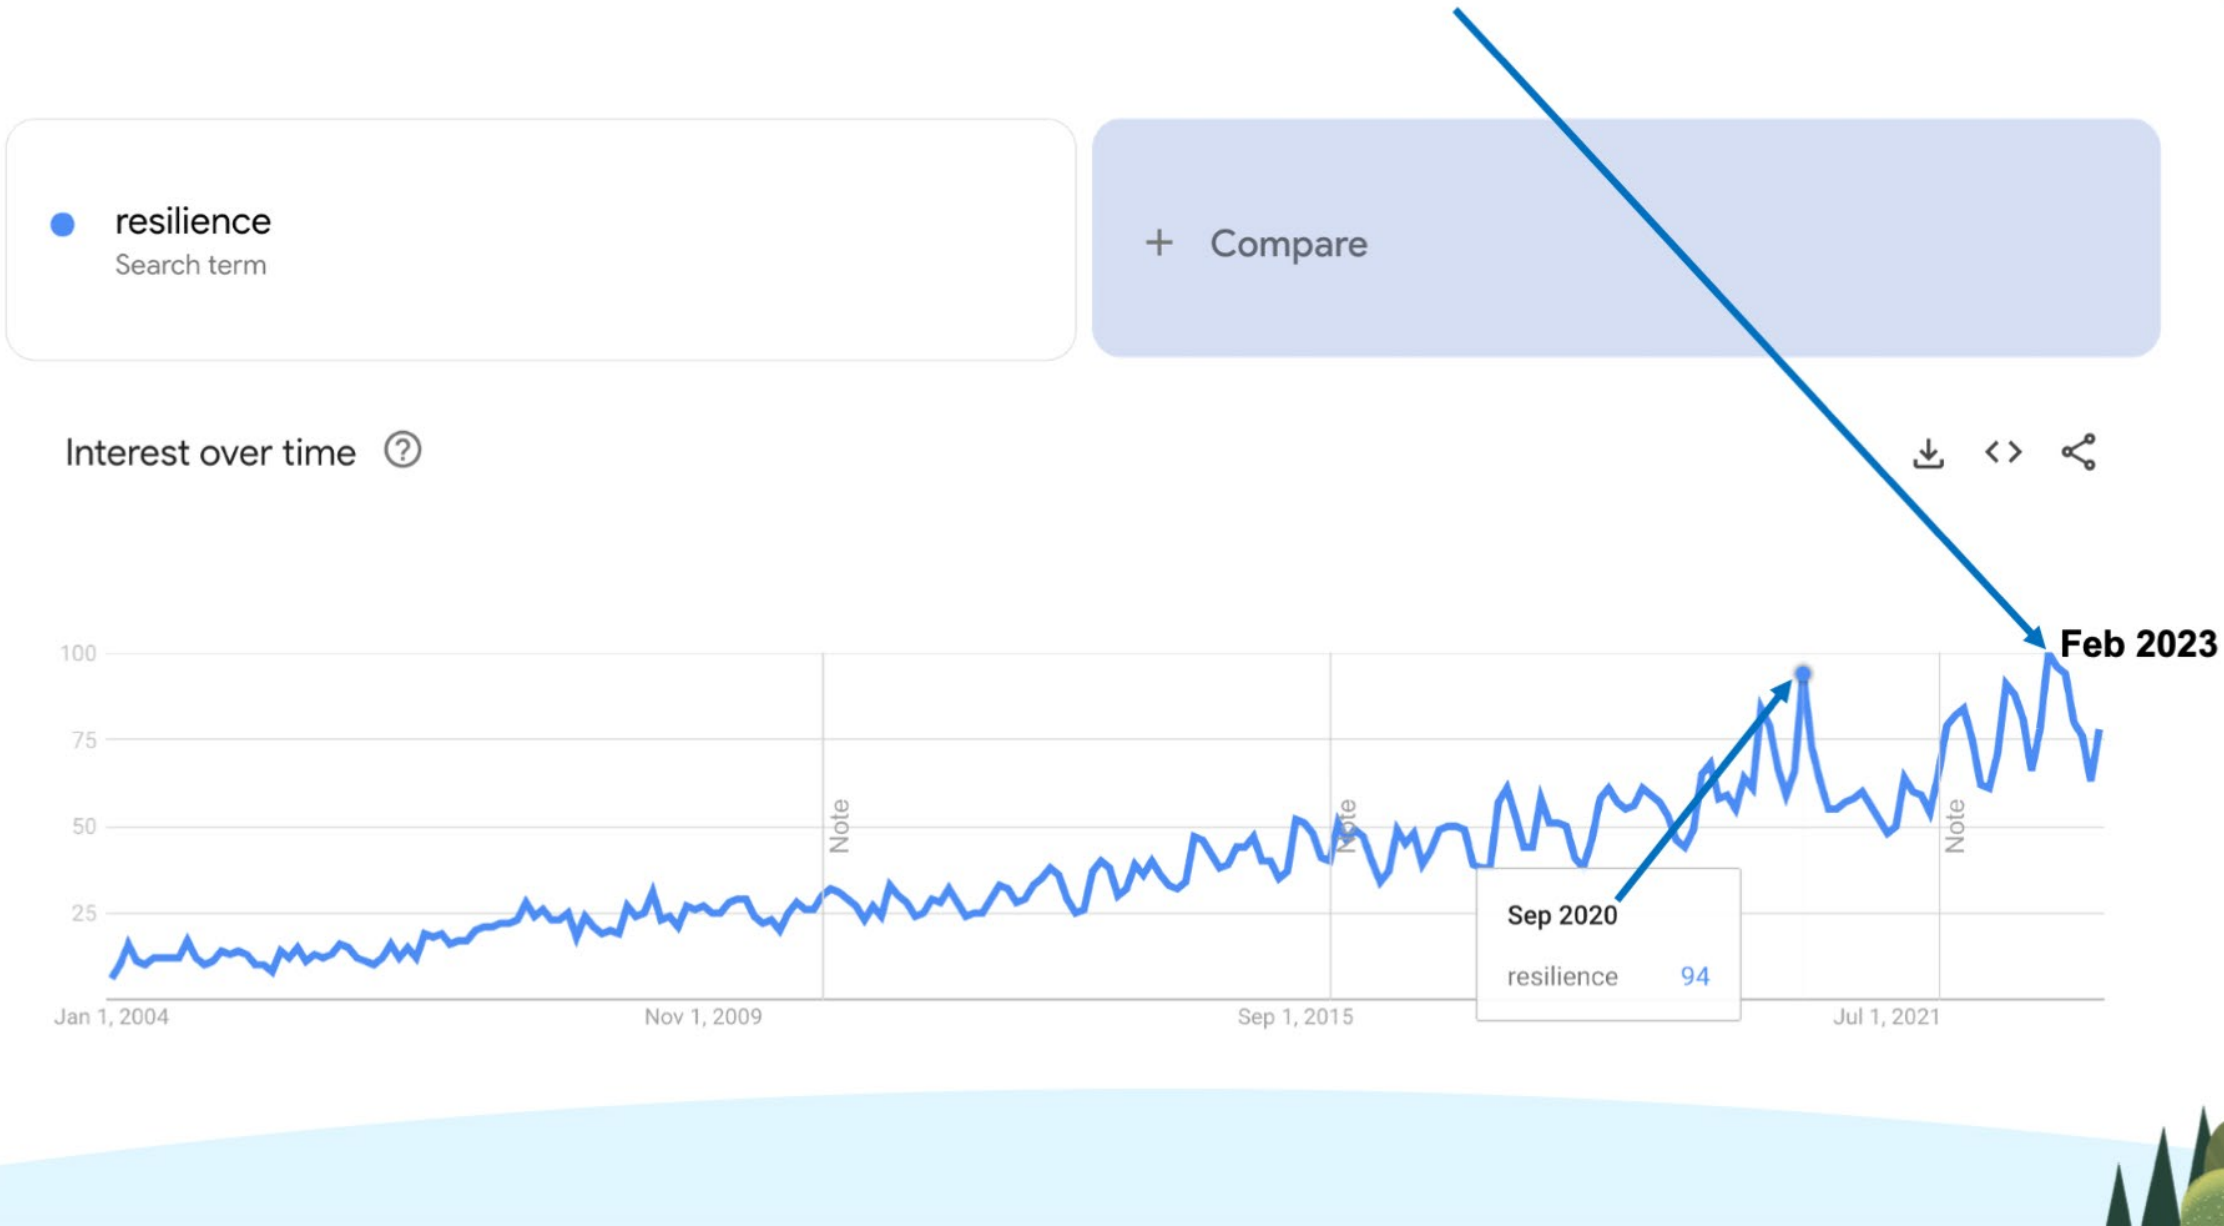

“Post-pandemic” or not, there is no “back to normal”

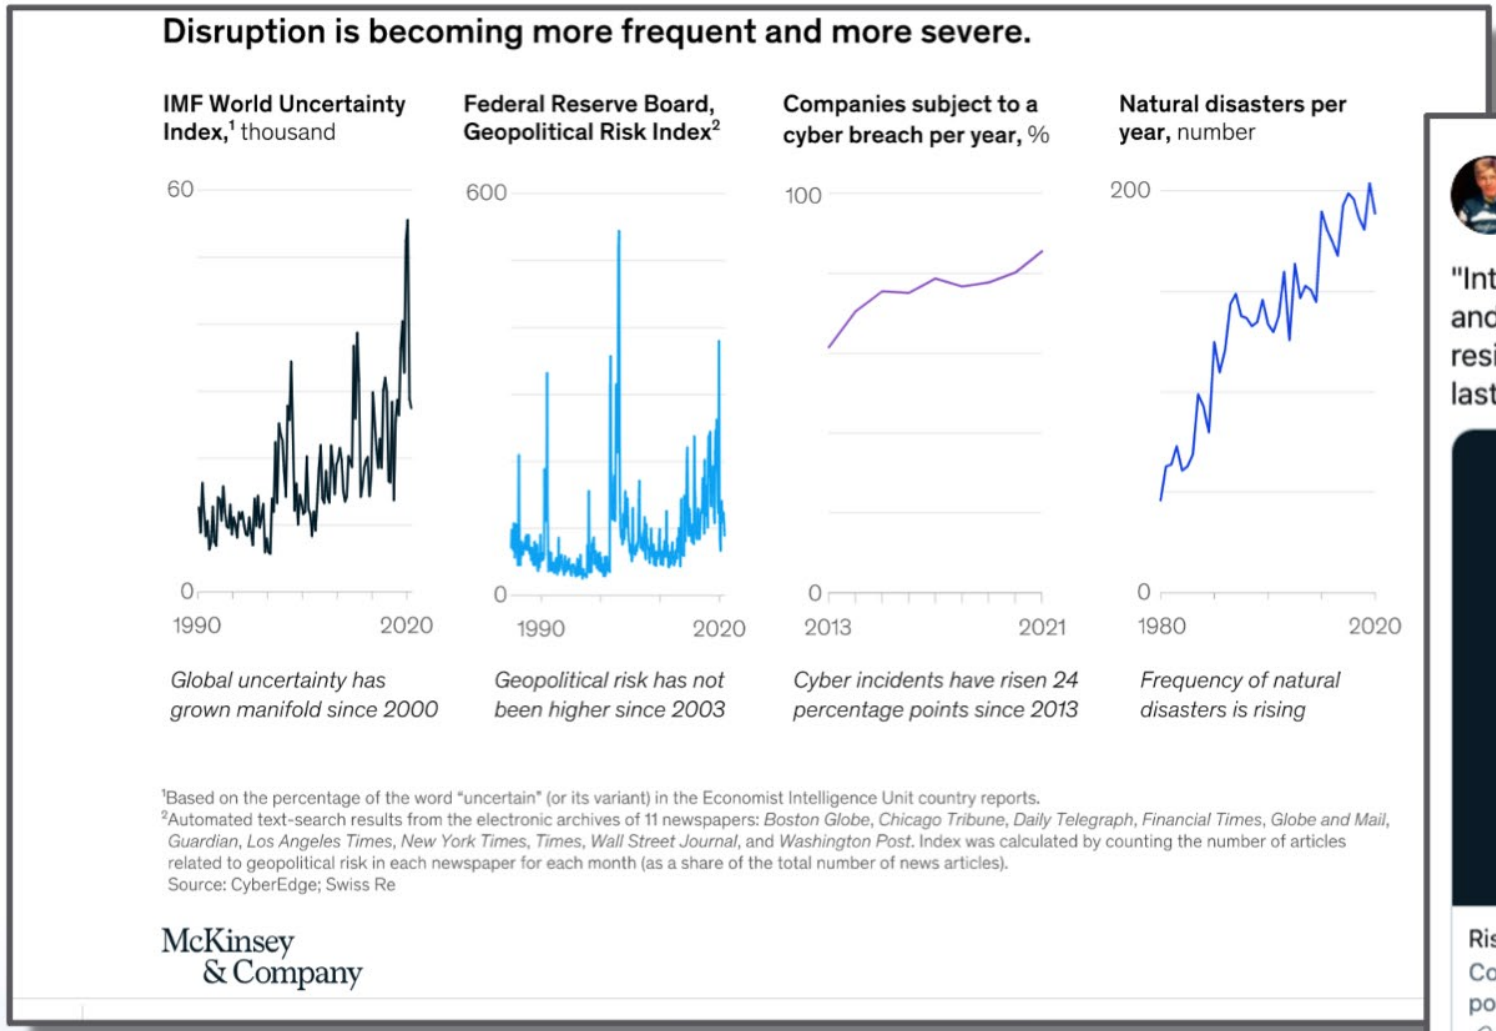

Resilience at scale is the new expectation from IT

Autonomic Computing helps to address complexity by **using technology to manage technology.** The term autonomic is derived from human biology. The autonomic nervous system monitors your heartbeat, checks your blood sugar level and keeps your body temperature close to 98.6°F without any conscious effort on your part. In much the same way, **self-managing autonomic capabilities anticipate IT system requirements and resolve problems with minimal human intervention.** As a result, IT professionals can focus on tasks with higher value to the business.

– IBM, "An architectural blueprint for autonomic computing," 3rd Ed. June 2005

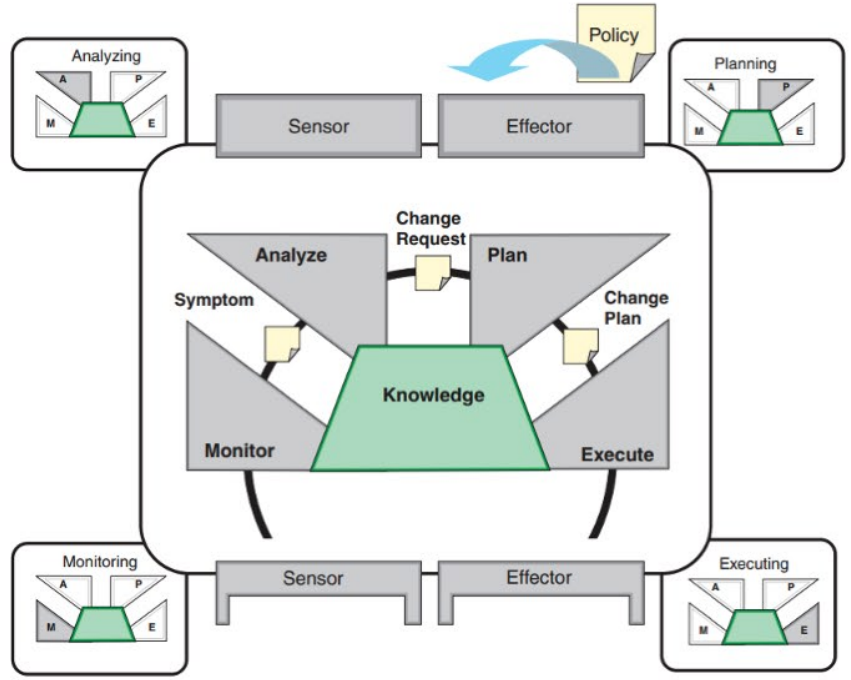

Figure 4. Functional details of the autonomic manager

Notions that “an AI” “thinks” are overblown

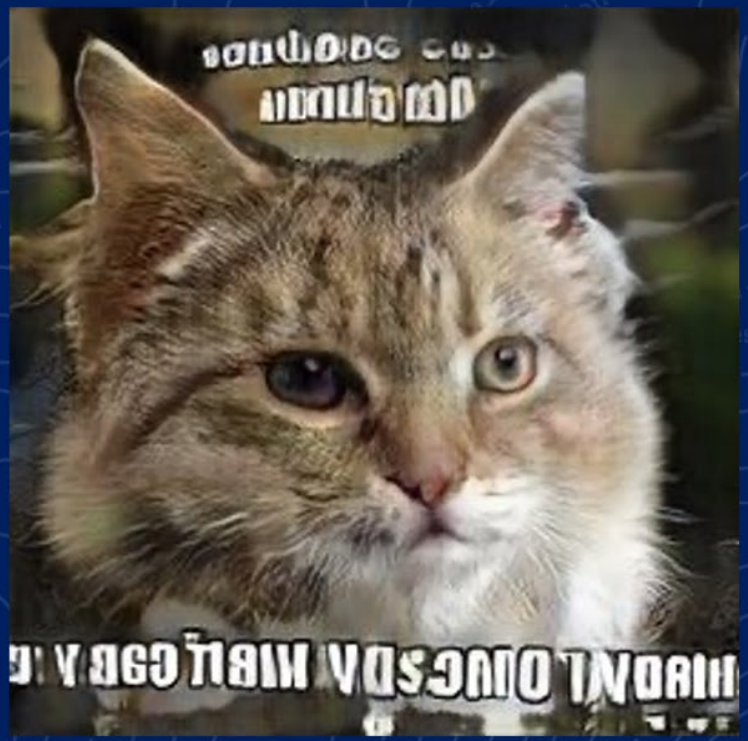

Texty squiggles are *not* part of the concept of “cat”

AI is *not* yet ready to be a “practitioner” administering drugs

IN THE HOUSE OF REPRESENTATIVES

SEPTEMBER 30, 2021

Mr. SCHWEIKERT introduced the following bill; which was referred to the Committee on Energy and Commerce

A BILL

To amend the Federal Food, Drug, and Cosmetic Act to clarify that artificial intelligence and machine learning technologies can qualify as a practitioner eligible to prescribe drugs if authorized by the State involved and approved, cleared, or authorized by the Food and Drug Administration, and for other purposes.

Be it enacted by the Senate and House of Representatives of the United States of America in Congress assembled,

SECTION 1. SHORT TITLE.

This Act may be cited as the “Healthy Technology Act of 2021”.

SEC. 2. PRESCRIPTION OF DRUGS BY ARTIFICIAL INTELLIGENCE OR MACHINE LEARNING TECHNOLOGIES.

Section 503(b) of Federal Food, Drug, and Cosmetic Act (21 U.S.C. 353(b)) is amended by adding at the end the following:

“(6) In this subsection, the term ‘practitioner licensed by law to administer such drug’ includes artificial intelligence and machine learning technology that are—

“(A) authorized pursuant to a statute of the State involved to prescribe the drug involved; and

“(B) approved, cleared, or authorized under section 510(k), 513, 515, or 564.”.

the term ‘practitioner licensed by law to administer such drug’ includes artificial intelligence —

Machine Learning 101

Overview [edit]

Tom M. Mitchell provided a widely quoted, more formal definition: "A computer program is said to learn from experience *E* with respect to some class of tasks *T* and performance measure *P* if its performance at tasks in *T*, as measured by *P*, improves with experience *E*."[12] This definition is notable for its defining machine learning in fundamentally operational rather than cognitive terms, thus following Alan Turing's proposal in his paper "Computing Machinery and Intelligence", that the question "Can machines think?" be replaced with the question "Can machines do what we (as thinking entities) can do?".[13] In the proposal he explores the various characteristics that could be possessed by a *thinking machine* and the various implications in constructing one.

Do we care if it's "thinking," as long as it usefully does what we have to "think" to do?

Let’s talk about how “AI” can help – and not “artificially”

Automate (and ***autonomate***) - delegate low-value, high-attention tasks to machines
Predict - use machine learning and pattern recognition to **anticipate rather than react**
Plan - lay out a sequence of tasks in ways that maximize results at minimum cost
Recognize - **minimize false positives & negatives** in diagnosis and troubleshooting
Optimize - explore every corner of feasible regions **without preconception bias**
Protect - detect and address the **errors and oversights** that put data and processes at risk
Rationalize - streamline and simplify processes
Iterate - continually **review and refine** operations based on new data and analysis
Actuate - focus sensory input through customer-experience lenses: **differentiate & delight**
Trust - elevate from black-box mystery to **informative explanation** of results
Enlighten - don't replace people; rather, **augment them and empower them**

APPROPRIATE Intelligence: because *verbs* have *effects*

Automate (and ***autonomate***) - delegate low-value, high-attention tasks to machines
Predict - use machine learning and pattern recognition to **anticipate rather than react**
Plan - lay out a sequence of tasks in ways that maximize results at minimum cost
Recognize - **minimize false positives & negatives** in diagnosis and troubleshooting
Optimize - explore every corner of feasible regions **without preconception bias**
Protect - detect and address the **errors and oversights** that put data and processes at risk
Rationalize - streamline and simplify processes
Iterate - continually **review and refine** operations based on new data and analysis
Actuate - focus sensory input through customer-experience lenses: **differentiate & delight**
Trust - elevate from black-box mystery to **informative explanation** of results
Enlighten - don't replace people; rather, **augment them and empower them**

This changes the prices of everything

Prices serve as signals of scarcity and abundance

"The number of **different ways to use, combine, and recombine resources** is unimaginably colossal. Almost all of these ways are useless. It would be a mistake, for example, to combine Arnold Schwarzenegger with medical equipment and have him perform brain surgery.

"**Only a tiny fraction of all the possible ways to allocate resources is useful. How can we discover these ways?** Market prices are vital because they condense, in as objective a form as possible, information on the value of alternative uses."

– econlib.org/library/Enc/InformationandPrices.html

When (relative) prices change, jobs change

Proceedings of the

EASTERN JOINT COMPUTER CONFERENCE

November 7-9, 1955 Boston, Mass.

Automatic Programming

One significant activity called automatic programming may reduce total general purpose system costs. The basic idea in this activity is that the computer should do part of the programming job thus cutting time and personnel costs.

However, the idea applied to business problems has not proven too successful as yet. There are too many variables and diffe
cation and another
so-called sub-progr
generalized automa
business applicatio
expense of a good d
time is required for
possible variations
another and to deci
be satisfied.

Machine Costs

On the larger computers, the cost of machine time devoted to automatic programming such as compiling and generating routines may amount to an appreciable portion of the cost of personnel programming for the specific application. This is true because machine time may be worth several hundred dollars per hour while a programmer's time is worth perhaps three to six dollars per hour. Thus the savings in programming costs will not be as high as one might expect and they will be partially canceled out by higher processing costs.

This is true because machine time may be worth several hundred dollars per hour while a programmer's time is worth perhaps three to six dollars per hour.

When college-level “tech” skills become free web sites

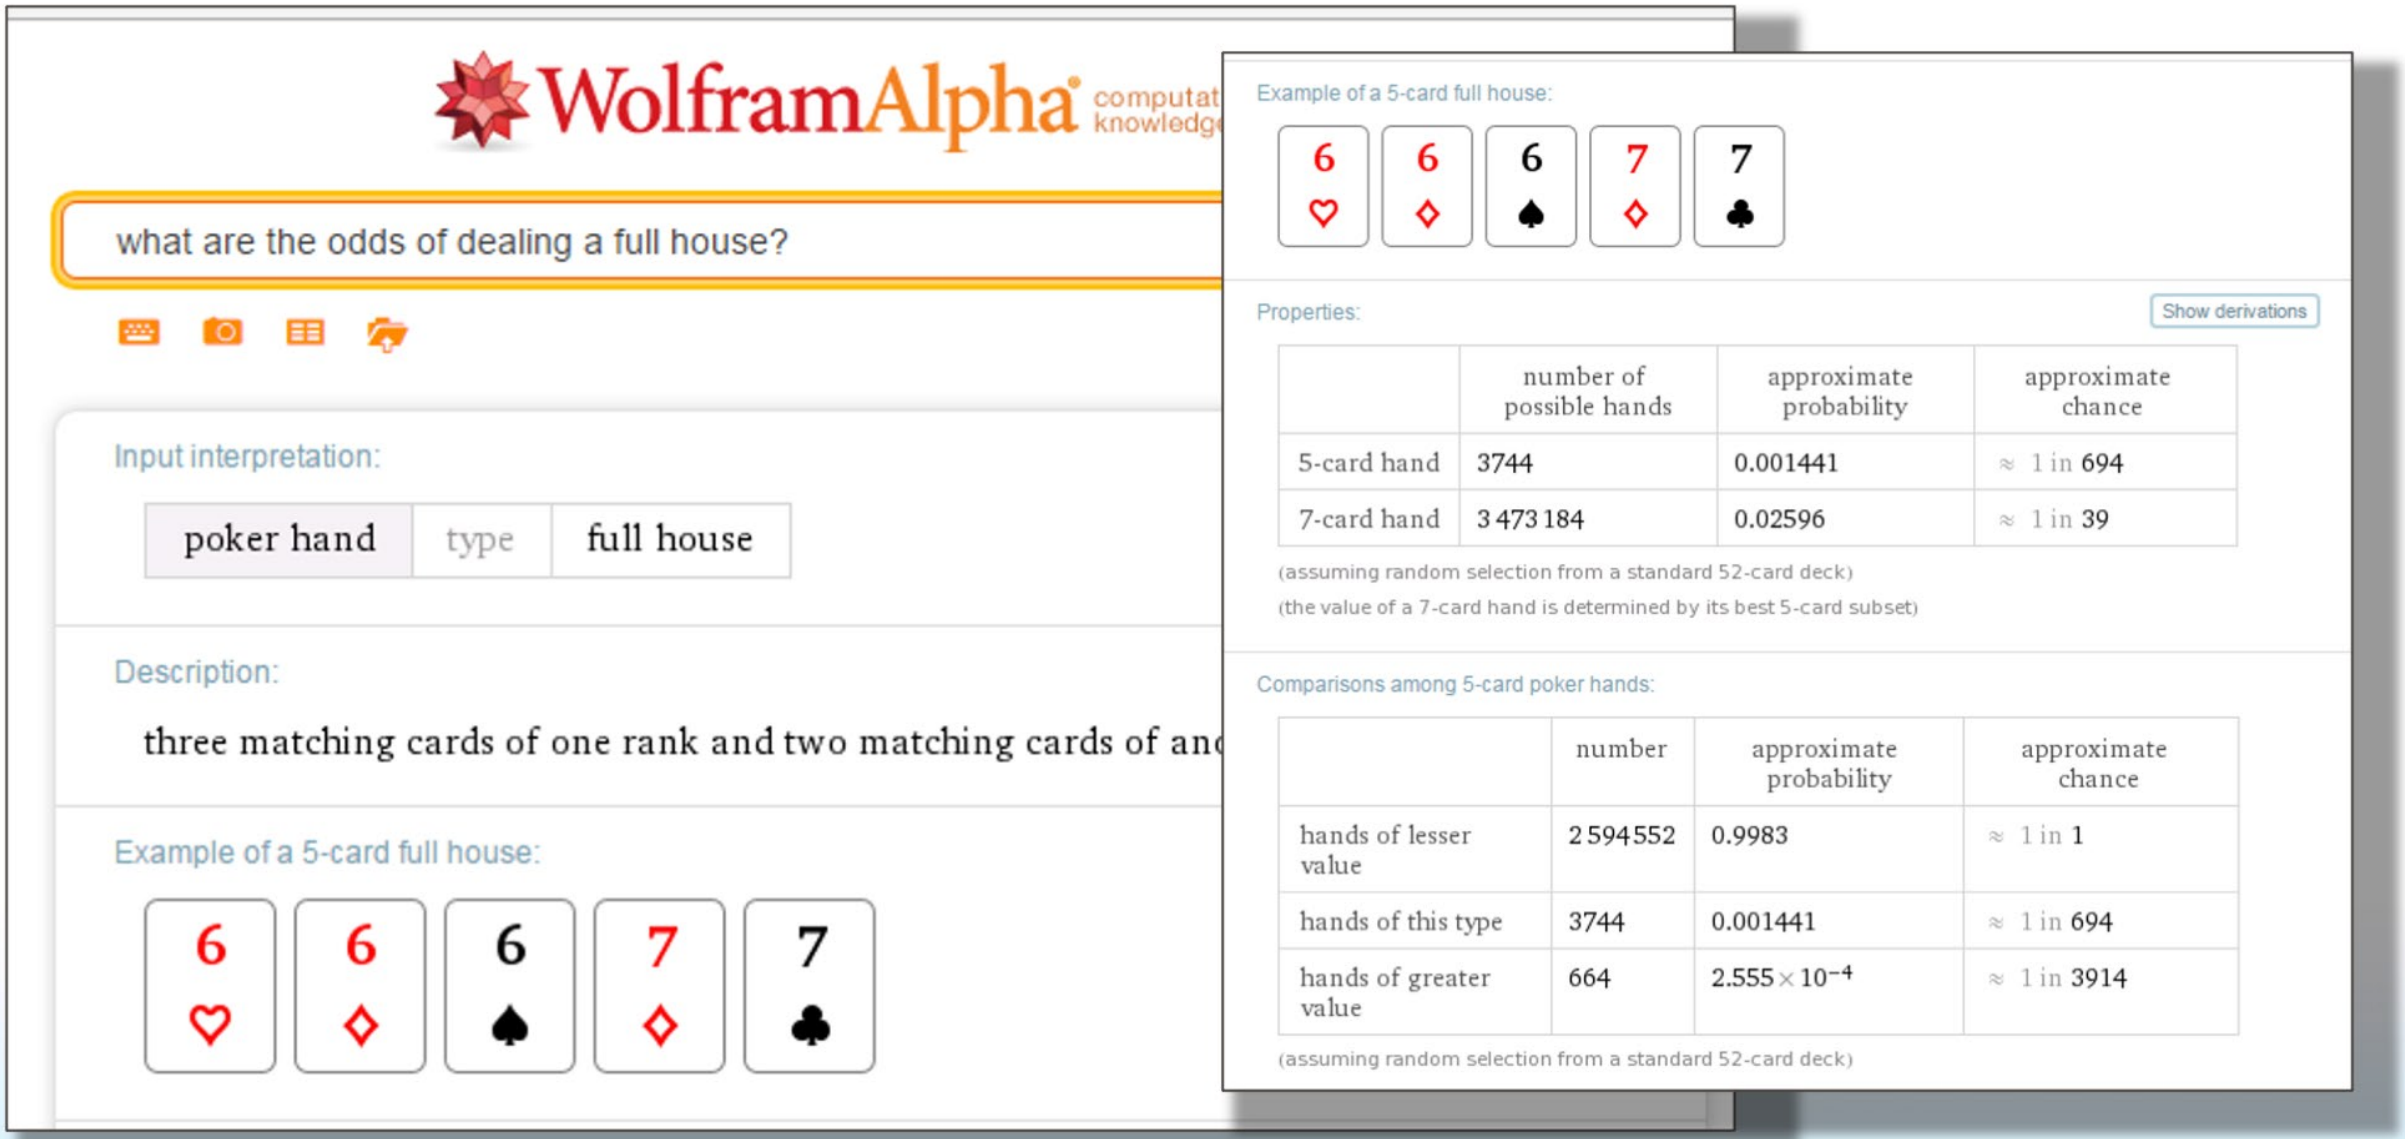

When formerly “human” skills become free web sites

Stable Diffusion Online

```
The Terminator
cyborg doing brain
surgery
```

When free web sites become interactive collaborators

Stable Diffusion Online

```
The Terminator
cyborg performing
brain surgery
```

When free web sites start to be “creative”

Stable Diffusion Online

```
The Terminator
cyborg performing
brain surgery in an
operating room
```

When iterating on the question becomes the key skill

Stable Diffusion Online

The Terminator cyborg holding a scalpel and wearing a surgical mask in an operating room

We'll pay more for insightful questions than we do for computable answers

Cue: what's wrong with the Beninia project?
It won't work

Cue: would it be fair to say you don't believe in what you've been told about Beninia?
I don't believe it

Postulate that the data given you about Beninia are true. Cue: what would be necessary to reconcile them with everything else you know?
That a force of unknown nature is acting on the population and causing them to behave differently from known patterns of human reaction under comparable circumstances elsewhere

"Shal," Chad said softly, "such a force exists, and is at present being investigated by experts to determine its nature. I tell you three times!"

"All right," he said wearily. "Try him now." The answer came back:
The estimated return for the Beninia project will be-

The man looked at Rex. "I think he's done it, sir," he exclaimed.

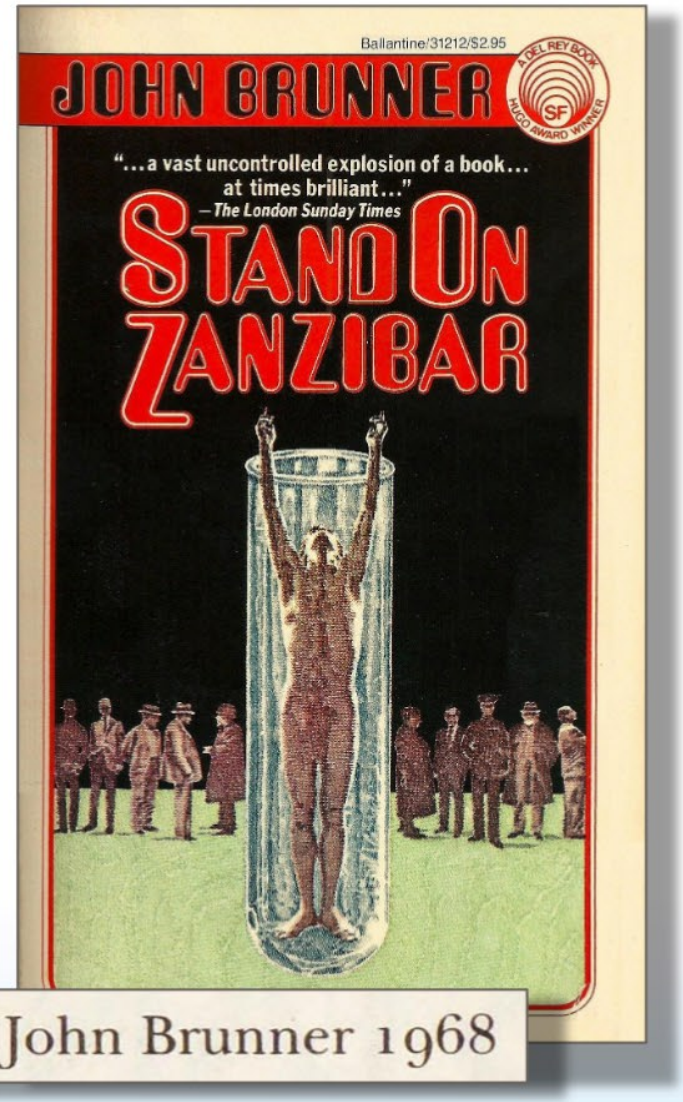

We will pay for world-building – not for “programming”

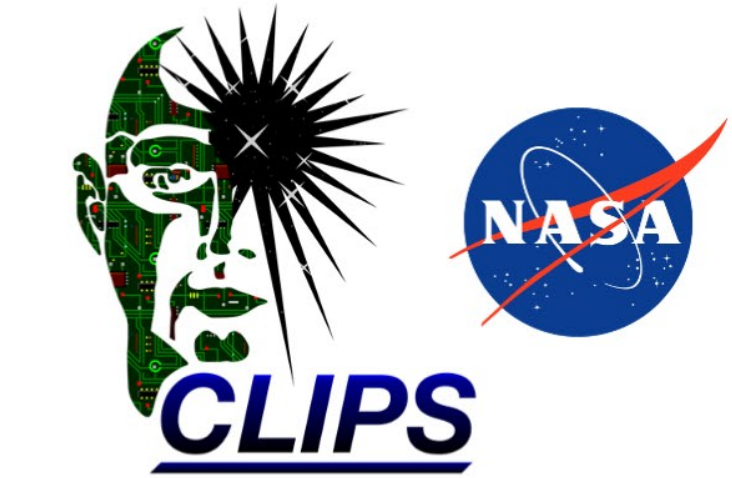

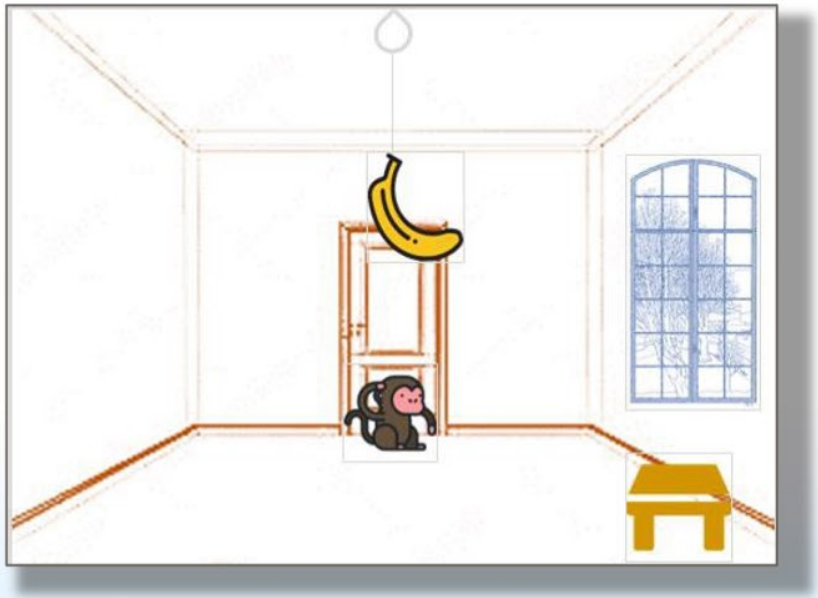

```
...**********************
;;;
;;;* INITIAL STATE RULE *
...**********************
;;;

(defrule startup ""
  =>
  (assert (monkey (location t5-7) (on-top-of green-couch) (holding blank)))
  (assert (thing (name green-couch) (location t5-7) (weight heavy)))
  (assert (thing (name red-couch) (location t2-2) (weight heavy)))
  (assert (thing (name big-pillow) (location t2-2) (on-top-of red-couch)))
  (assert (thing (name red-chest) (location t2-2) (on-top-of big-pillow)))
  (assert (chest (name red-chest) (contents ladder) (unlocked-by red-key)))
  (assert (thing (name blue-chest) (location t7-7) (on-top-of ceiling)))
  (assert (chest (name blue-chest) (contents bananas) (unlocked-by blue-key)))
  (assert (thing (name blue-couch) (location t8-8) (weight heavy)))
  (assert (thing (name green-chest) (location t8-8) (on-top-of ceiling)))
  (assert (chest (name green-chest) (contents blue-key) (unlocked-by red-key)))
  (assert (thing (name red-key) (location t1-3)))
(assert (goal-is-to (action eat) (arguments bananas))))
```

Implications for Practitioners

We must share understanding of exponential change

"This has to level off"

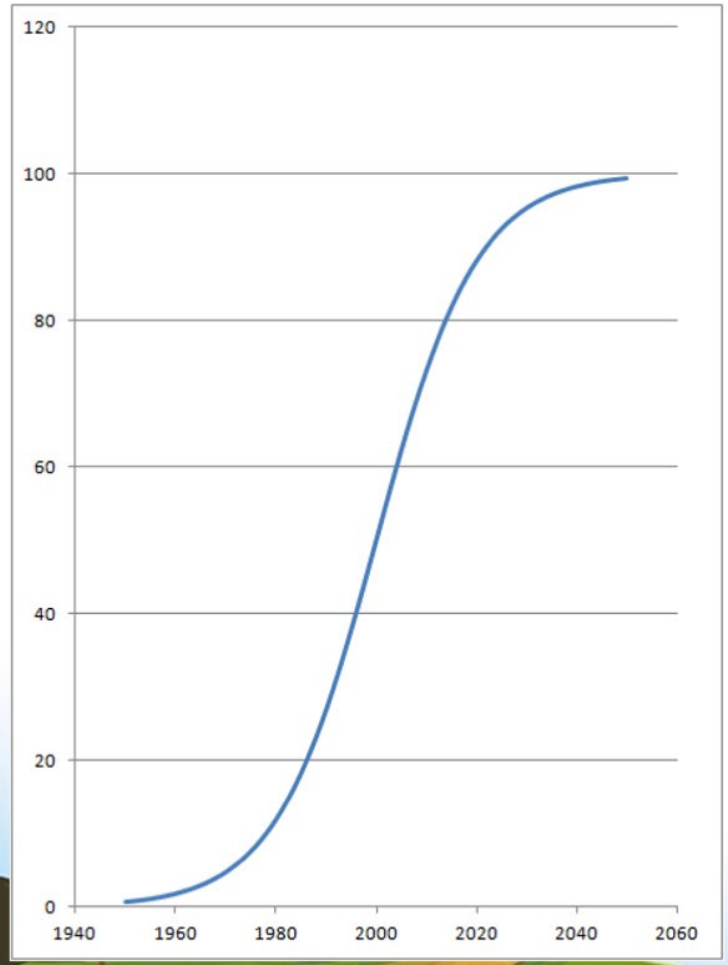

"That's too conservative"

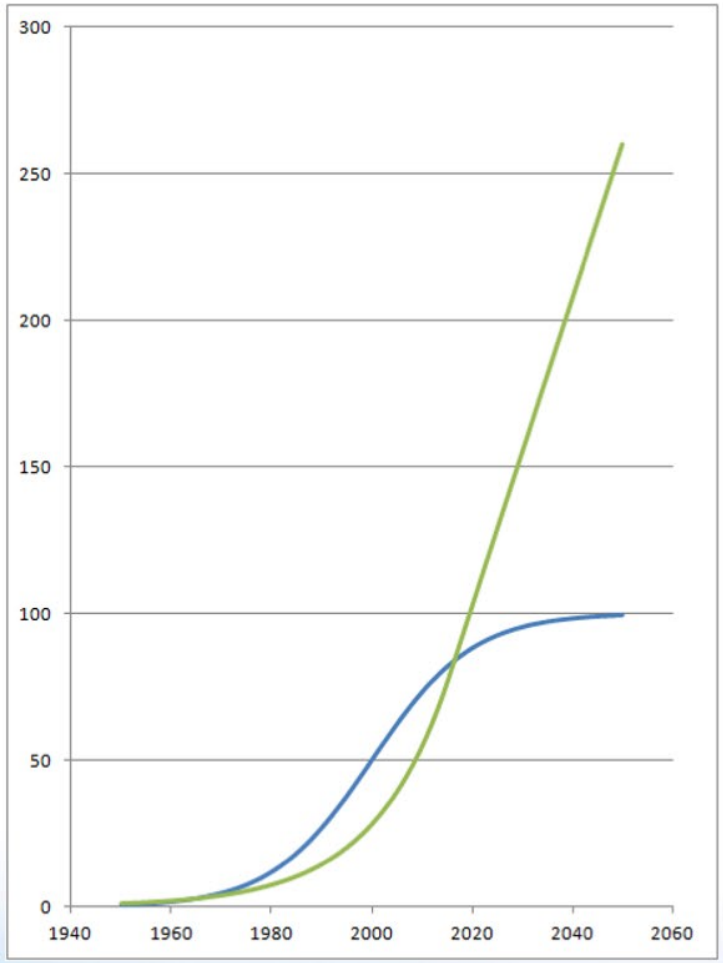

"*Much* too conservative"

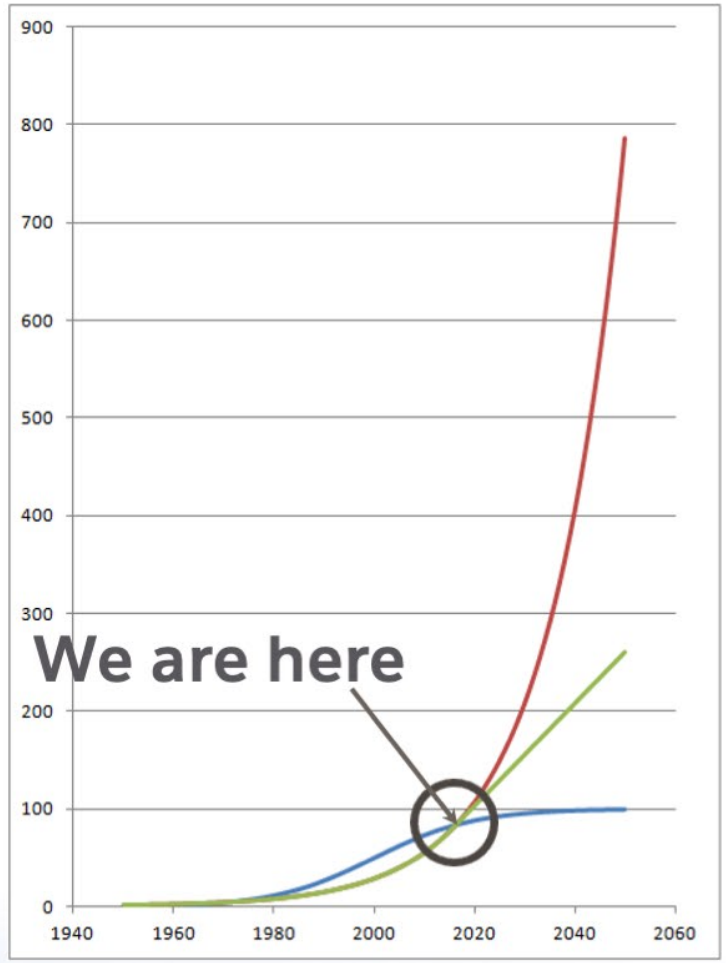

"Here's the science fact: the pace of change right now is the absolute slowest it will be for the rest of your life. Fasten your seatbelts. It's going to be a fascinating ride."

– Pat Gelsinger, VMware, September 2017

The pace of change is not just temporarily fast

"There's every chance IBM has just unveiled the blueprint for the future of AI development with an analog AI chip that's said to be up to 14 times more energy efficient than current industry-leading components.

"IBM claims its 14nm analog AI chip mimics the way in which a human brain would operate, with the microchip performing computations directly within memory."

Nvidia, beware! IBM has a new analog AI chip that could give the H100 a run for its money

News By Keumars Afifi-Sabet published 1 day ago

Try as we might to anticipate courageously, the things that "won't happen" or that "machines can't do"...turn out to be things they *can* do, sooner than expected

We will pay for what algorithms don't replace

ADAPTED FROM: Institute for the Future, Future Work Skills 2020, 2011

Transdisciplinarity: "We can no longer rely on just bringing together groups of specialists to solve our most complex problems"

Computational Thinking: "To understand the meaning, the trends and patterns of what the data is telling us becomes paramount."

– Institute for the Future

Sensemaking: "When you ask creative people how they did something, they feel a little guilty because they didn't really do it, they just saw something"

– Steve Jobs

Social Intelligence: "Influence and relationship-building will now come from asking the right questions, not necessarily having all the answers"

– John Hagel

- If you want a job tomorrow, you can have one
- It will be harder than the one you have now

"According to the World Economic Forum's Future of Jobs Report, 50% of all employees will need to reskill by 2025 as the adoption of technology increases."

hrforecast.com/a-guide-to-future-oriented-skills-skills-in-demand-to-watch-in-the-next-five-years/

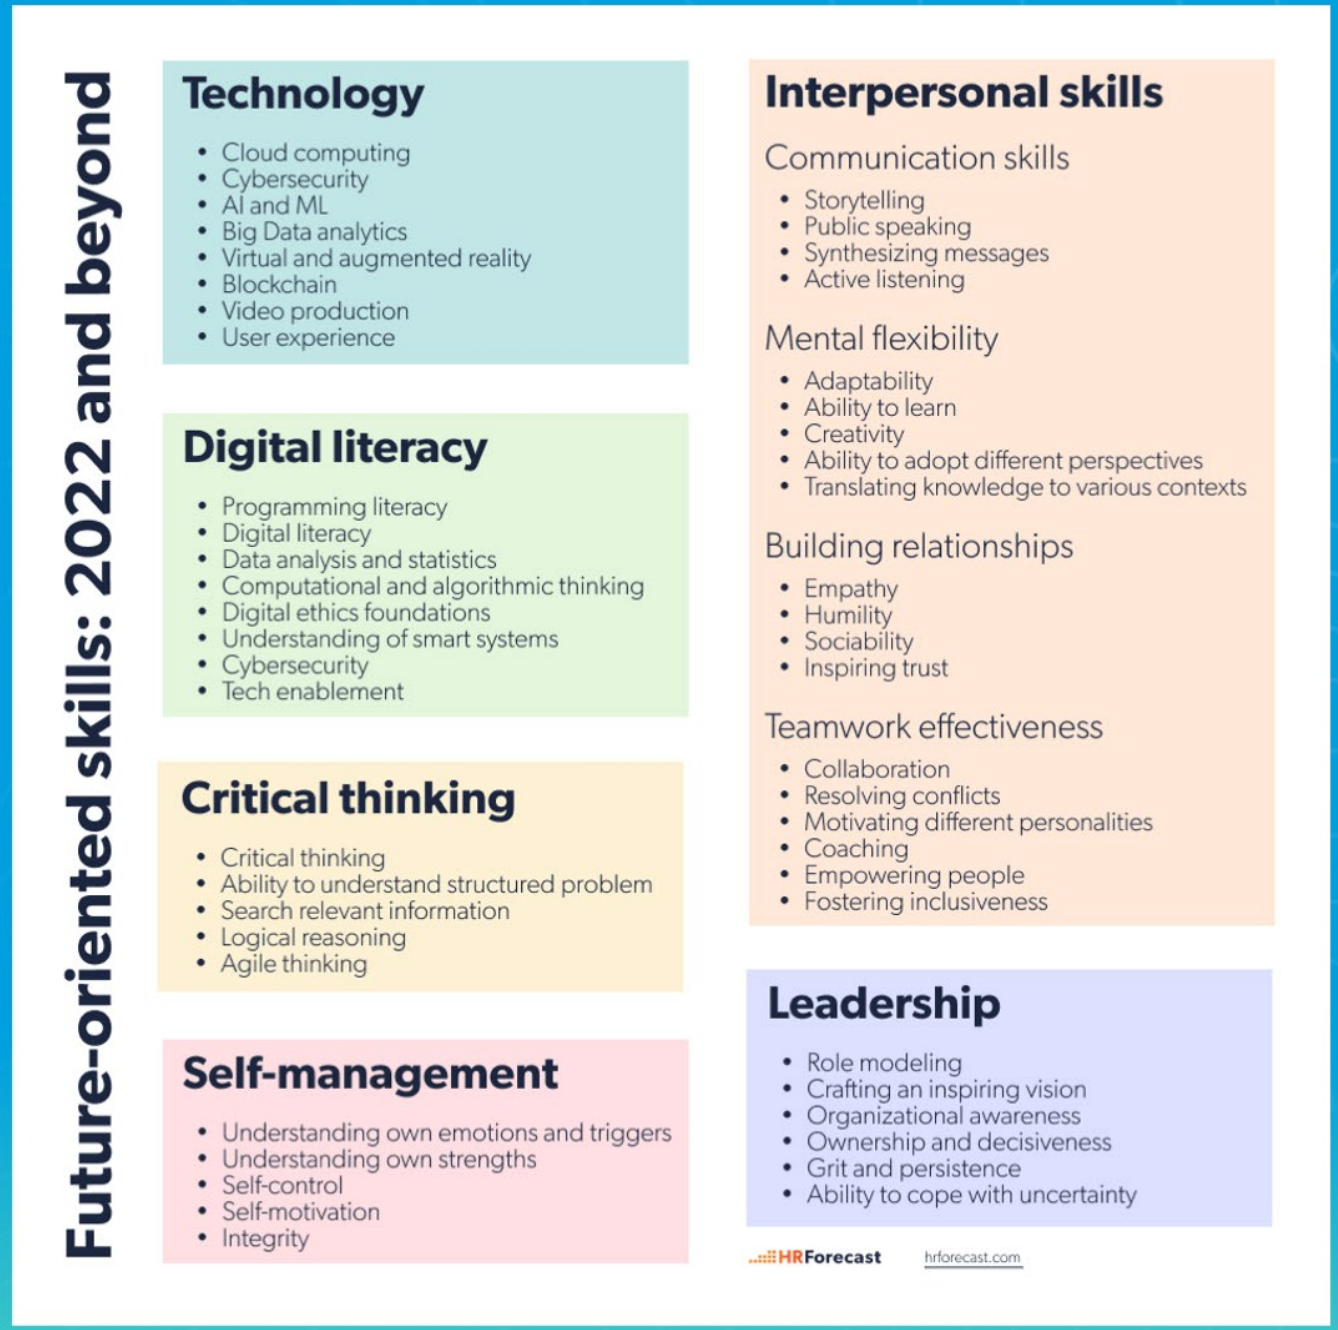

Implications for Organizations

Answers only matter if we are asking the right questions

- What has technology change made **possible**?
- What has behavior change made **necessary**?

(or)

What behavior change must become **accepted**?

- What must leadership do to **make it happen**?

Post-pandemic connections will be a challenge

Customers Now Expect Connected Experiences

More than a third (36%) of organizations provide a completely connected user experience across all channels, an increase from 30% in 2022. However integrating user experiences remain a challenge for many organizations.

Teams Across the Business Demand Automation

Automation adoption is on the rise. Most IT organizations centrally manage (67%) and track (59%) automation work, but as more non-IT roles request automation organizations are turning to no-code tools and approaches to meet demand.

2023 Connectivity Benchmark Report

MuleSoft

It's *not* just about *customer* experience

American Workers' Interest in Upskilling, by Occupation

How interested are you in participating in training/education to upgrade your skills or learn new skills that could help you advance your career?

% Very/Extremely interested

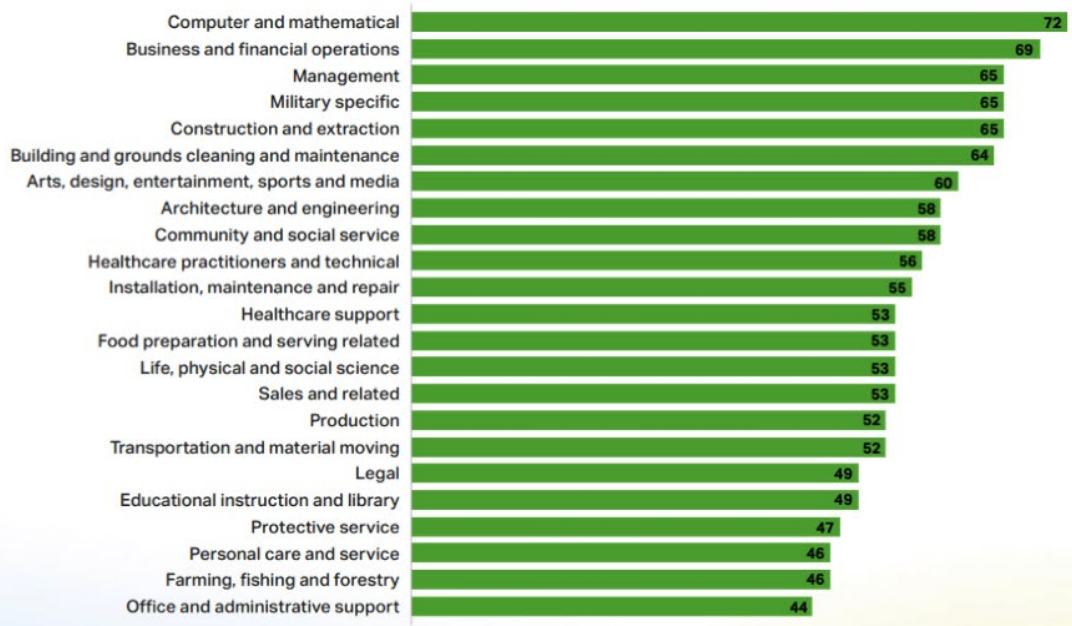

Workers' Interest in Upskilling Opportunities

How interested would you be in each of the following opportunities to upgrade your skills or learn new skills if they were provided at no cost to you by an employer?

% Extremely/Very interested

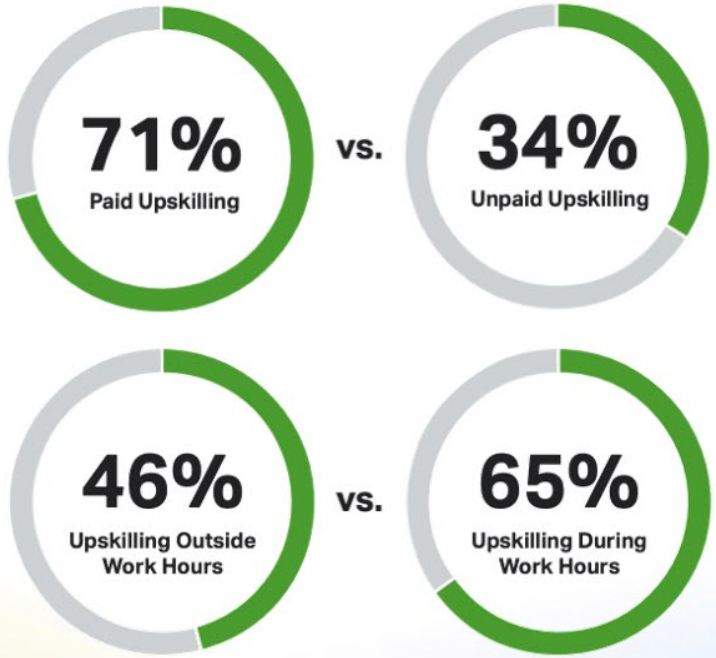

Recognize the U-shaped curve

Figure 3.3 **Changes in demand for jobs per ISCO skill level**

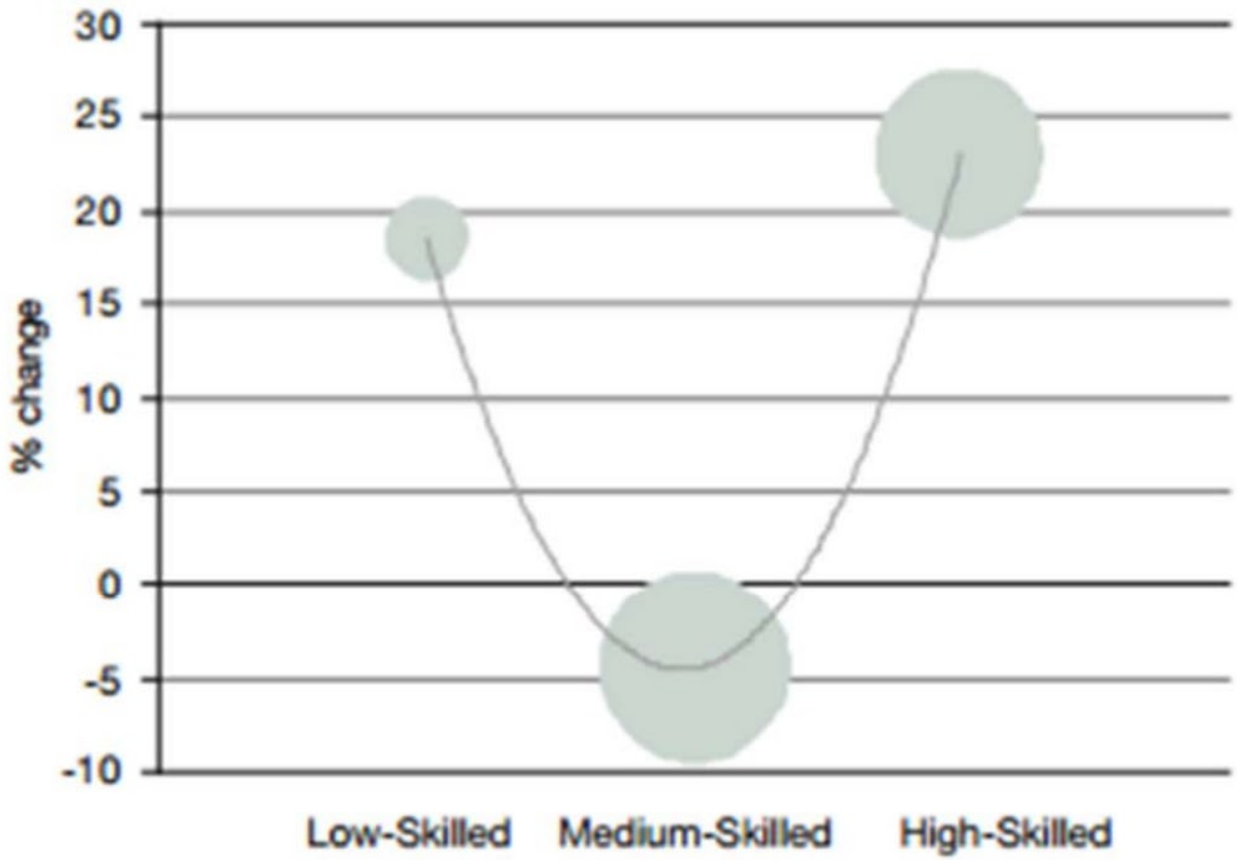

Address the new inequalities

Women have more to worry about than men from a coming wave of automation and artificial intelligence that could replace almost a third of hours worked across the US economy.

Women are 1.5 times more likely to need to move into a new occupation than men during that period. They're over-represented in the industries with lower-wage jobs that will be most impacted by automation, including office support and customer service. **Blacks and Hispanics will also be adversely affected** as demand for food and production workers shrinks.

We see generative AI enhancing the way STEM, creative, and business and legal professionals work rather than eliminating a significant number of jobs outright. Automation's biggest effects are likely to hit other job categories. Office support, customer service, and food service employment could continue to decline.

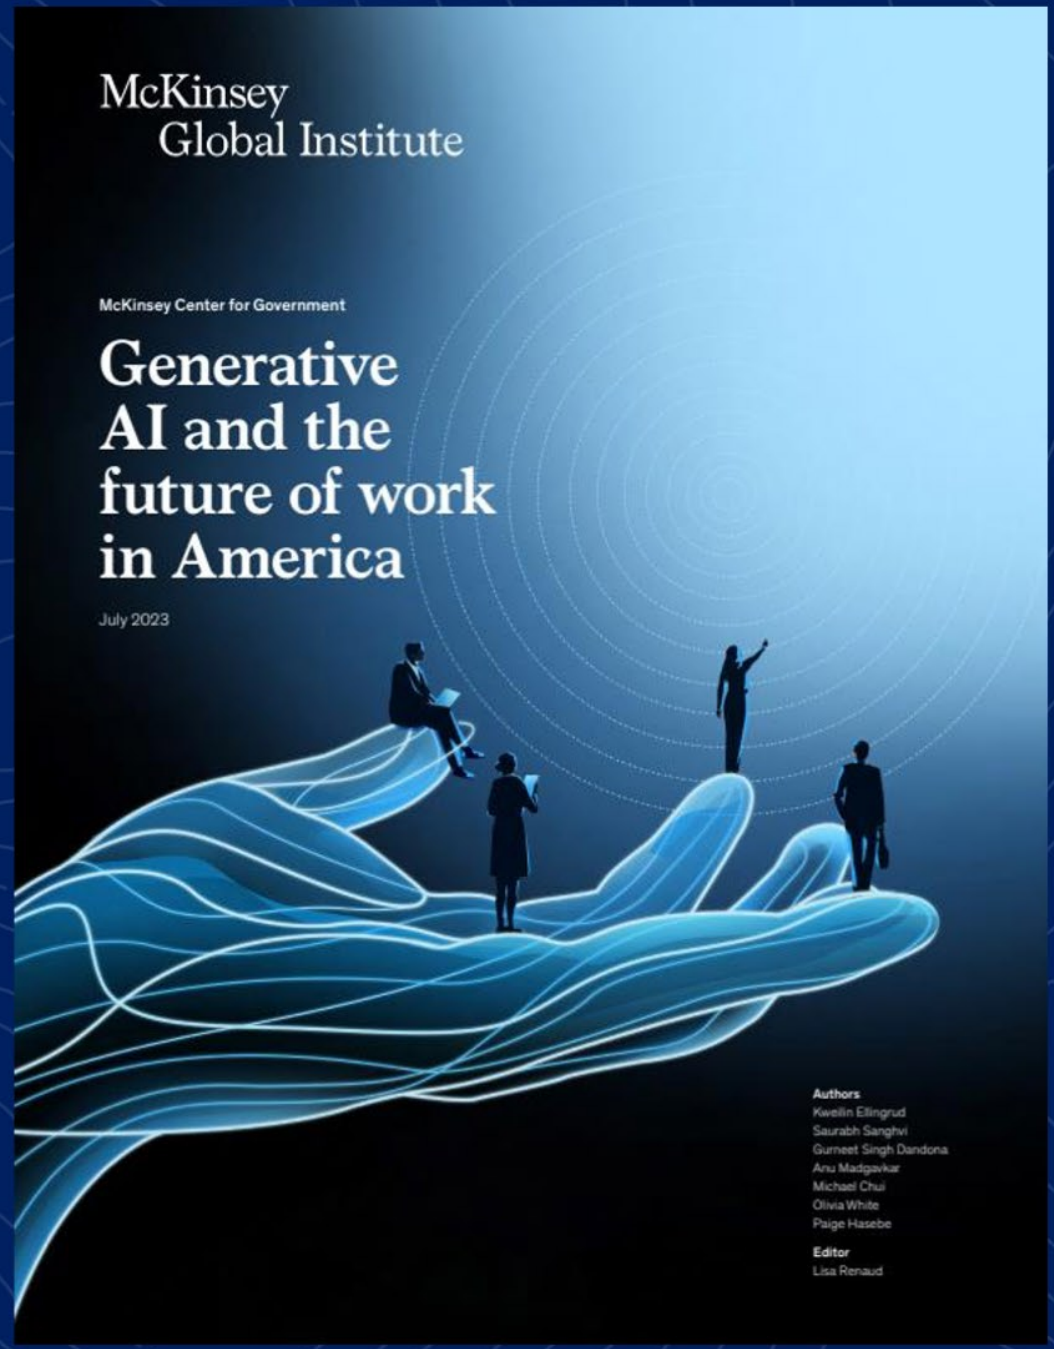

Claims that “data has no agenda” are naïve

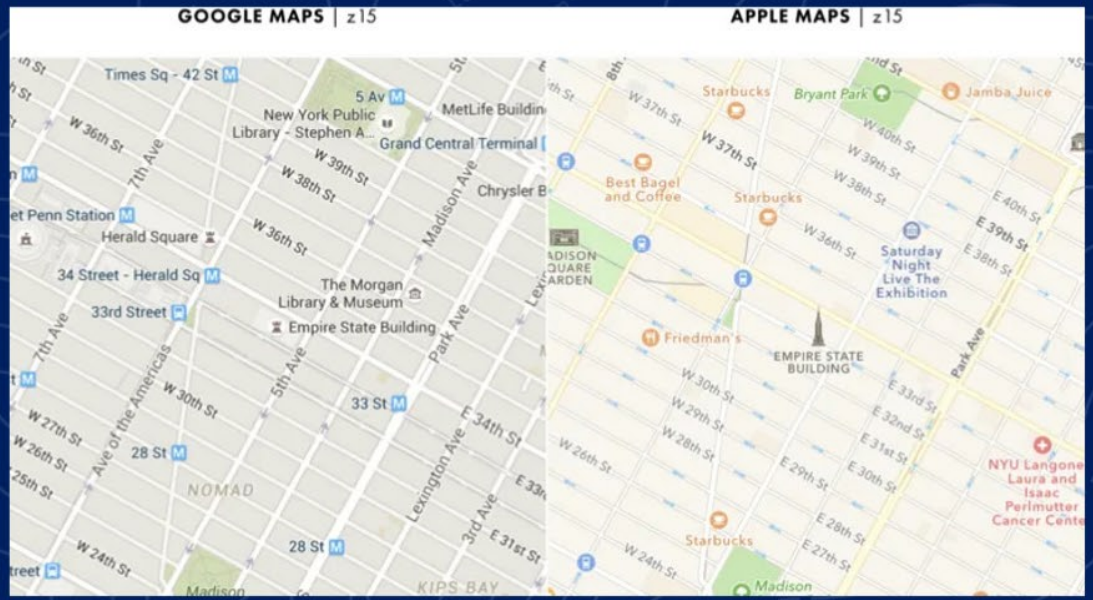

Google Maps: how to get there

Apple Maps: why to go there

“Inform” is a verb of behavior change

Apple Health doesn’t track menstruation, an omission that was quickly seized upon by many tech writers as, well, ridiculous. *The Verge* asked “is it really too much to ask that Apple treat women, and their health, with as much care as they've treated humanity’s sodium intake?”

The Atlantic

With algorithms, we don’t have an engineering breakthrough that’s making life more precise, but billions of semi-savant mini-Frankensteins, often with narrow but deep expertise that we no longer understand, spitting out answers here and there to questions we can’t judge just by numbers, all under the cloak of objectivity and science.

The New York Times

In AI, “positive feedback” *doesn’t* mean “good job!”

Predictive policing systems are increasingly used to determine how to allocate police across a city. Discovered crime data (e.g., arrest counts) are used to update the model; such systems have been empirically shown to be susceptible to runaway feedback loops, where police are repeatedly sent back to the same neighborhoods regardless of the true crime rate.

We demonstrate how to change the inputs to a predictive policing system (in a black-box manner) so the runaway feedback loop does not occur.

Runaway Feedback Loops in Predictive Policing
1st Conference on Fairness, Accountability and Transparency, 2018

Marketing AI: learn from past behavior? WCPGW*?

"In the past, we used human intelligence – of mainly white males – to create advertising. Today's approach uses datasets pulled from various sources to create marketing campaigns that extend to many more media channels. And yet, those data sources are still based on past performance, created and evaluated through the lens of white males." – Larry Adams, Founder and CEO, X_Stereotype

"Marketers usually have better information about existing customers than future prospects. As a result, there can be a fair amount of risk that those algorithms are inherently more successful at finding people just like their current customers.

"Many marketers segment and target high-value customers; algorithms are typically trained mostly on data from the more common, lower-value customers. Consequently, those algorithms prove to be biased toward finding lower-value customers, hurting efforts overall." – Jake Moskowitz, Vice President of Data Strategy, Emodo Institute.

* What Could Possibly Go Wrong?

Bias distorts our data and handicaps our processes

- **Confirmation** bias? Record your beliefs *before* you begin
- **Selection** bias? Represent and randomize
- **Historical** bias? Undertake affirmative inclusivity
- **Survivorship** bias? Look hard for “what *didn’t* work”
- **Availability** bias? The middle of the road is a bad place to drive
- **Outlier** bias? Investigate and challenge distributions

www.metabase.com/blog/6-most-common-type-of-data-bias-in-data-analysis

Mitigation of bias isn't just a good idea: it's becoming the law

Enactment date: 12/11/2021
Law number: 2021/144
Title: A Local Law to amend the administrative code of the city of New York, in relation to automated employment decision tools

Summary: **This bill would require that a bias audit be conducted on an automated employment decision tool prior to the use of said tool.** The bill would also require that candidates or employees that reside in the city be notified about the use of such tools in the assessment or evaluation for hire or promotion, as well as, be notified about the job qualifications and characteristics that will be used by the automated employment decision tool. **Violations of the provisions of the bill would be subject to a civil penalty.**

Abundance of cheap data demands skeptical analytics

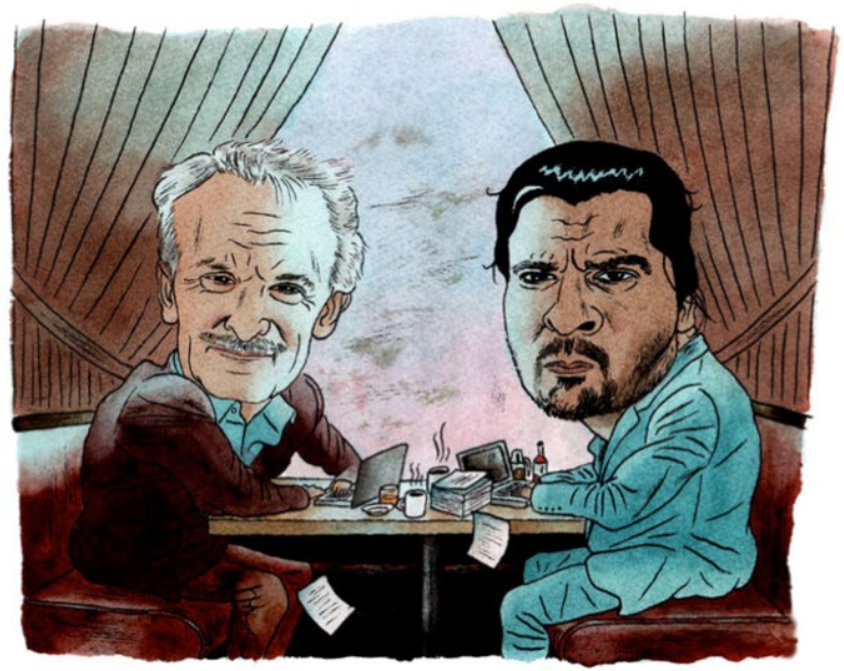

Nick Brown (left) and James Heathers (right) act as enforcers when they spot anomalies in the literature. ANDY FRIEDMAN

Meet the 'data thugs' out to expose shoddy and questionable research

"The test showed that the data in at least one of his studies–a 2012 article in the journal Preventive Medicine looking at carrot consumption among school children–appeared iffy. How so? **Running the published data through SPRITE showed that at least one child in the sample would have had to have eaten roughly 60 carrots in a single sitting.**"

www.sciencemag.org/news/2018/02/meet-data-thugs-out-expose-shoddy-and-questionable-research

OK, but...
Salesforce?
An AI company?

Salesforce has been pioneering AI application in task domains since 2014

1 trillion+
predictions a week

2016
Salesforce
Einstein

SFDC AI Research Founded

Acquired RelateIQ

Acquired PredictionIO & MetaMind

Vision and Language

Bots

Paved the way for prompting with DecaNLP paper

Prediction Builder

Acquired Bonobo

Conversation Insights

Recommendation Builder

Einstein Search

CodeGen LLM

ProGe n Paper

Einstein GPT

Slack GPT

Tableau GPT

CodeTF & CodeT5+ LLMs

300 AI patents

227 AI research papers

2014 2015 2016 2017 2018 2019 2020 2021 2022 2023

"What workers have today...has **low usability, low access, and low empowerment**. It's often mostly a jumble of technology that's not aimed at a coherent employee experience...

"**We have to design a lot more of it today** and now for our emerging new world."

– Dion Hinchcliffe

An active platform enables *agile* collaboration + *focused* automation

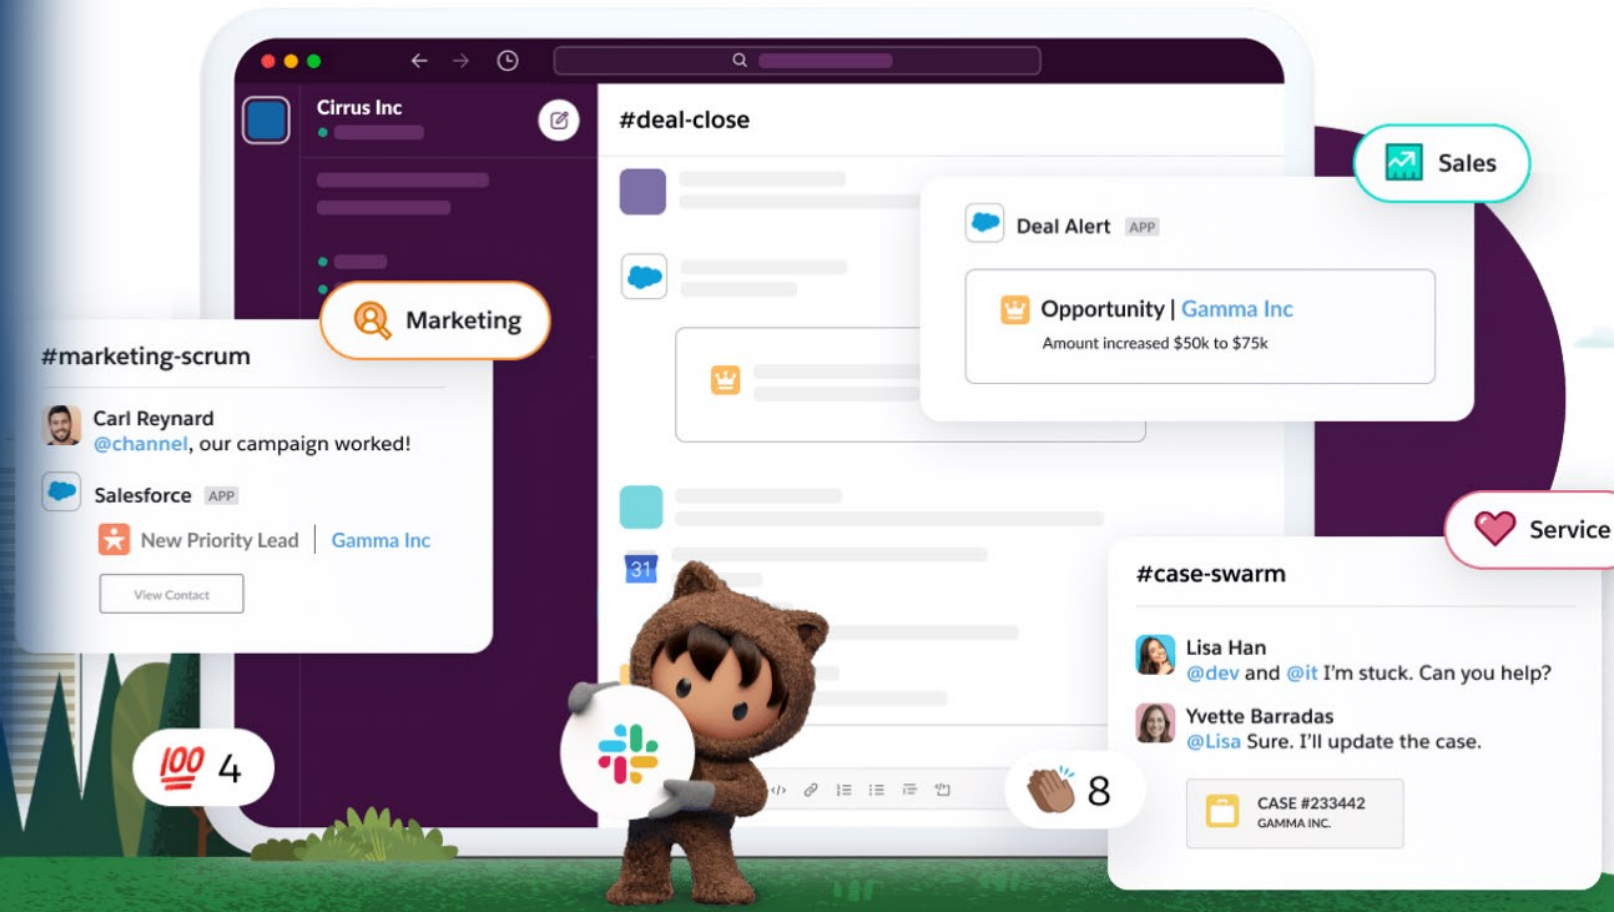

"What workers have today...has **low usability, low access, and low empowerment**. It's often mostly a jumble of technology that's not aimed at a coherent employee experience...

"**We have to design a lot more of it today** and now for our emerging new world."

– Dion Hinchcliffe

83% of users say automation solutions have provided them with the time to take on new, challenging projects

83% of users say automation solutions have provided them with the time to learn new skills

77% of users say automation solutions have provided them with more time to deepen relationships with customers and stakeholders

(Automation Survey, Salesforce Research, October 2021)

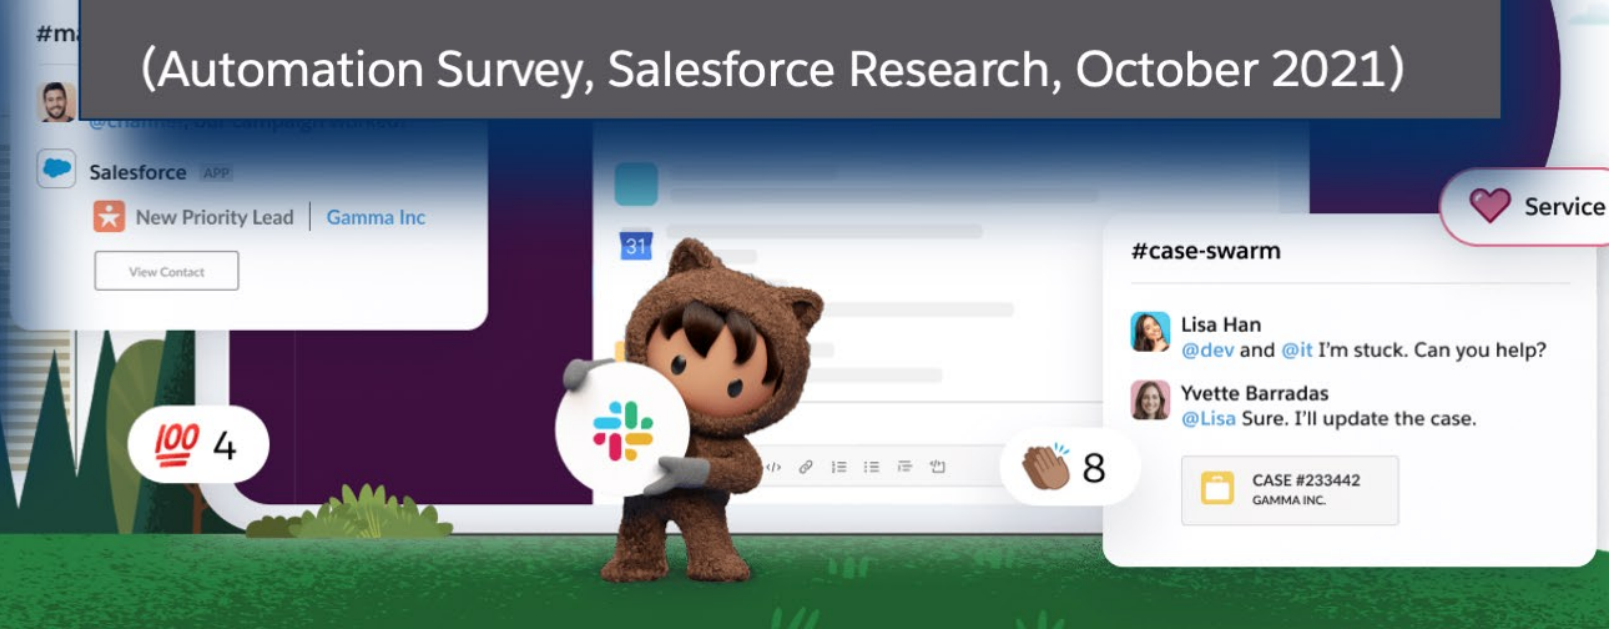

Peter Coffee
@petercoffee

"Socher & colleagues at @Salesforce Research have devised a method of merging commonsense reasoning...with neural network models...in a paper titled 'Explain Yourself! Leveraging Language Models for Commonsense Reasoning.'"

Common Sense Makes Progress with Deep Learning

We've witnessed incredible progress in the capability of deep learning models to not only understand text, but to generate it too. While the ...

datanami.com

10:55 AM · Jul 9, 2019 · Twitter Web Client

Generative AI: 5 Guidelines for Responsible Development

FEBRUARY 7, 2023

Paula Goldman

CHIEF ETHICAL AND HUMANE USE OFFICER, SALESFORCE

Kathy Baxter

PRINCIPAL ARCHITECT, ETHICAL AI PRACTICE

1. **Accuracy**: We need to deliver verifiable results that balance accuracy, precision, and recall in the models by enabling customers to train models on their own data. We should communicate when there is uncertainty about the veracity of the AI's response and enable users to validate these responses. This can be done by citing sources, explainability of why the AI gave the responses it did (e.g., chain-of-thought prompts), highlighting areas to double-check (e.g., statistics, recommendations, dates), and creating guardrails that prevent some tasks from being fully automated (e.g., launch code into a production environment without a human review).
2. **Safety**: As with all of our AI models, we should make every effort to mitigate bias, toxicity, and harmful output by conducting bias, explainability, and robustness assessments, and red teaming. We must also protect the privacy of any personally identifying information (PII) present in the data used for training and create guardrails to prevent additional harm (e.g., force publishing code to a sandbox rather than automatically pushing to production).
3. **Honesty**: When collecting data to train and evaluate our models, we need to respect data provenance and ensure that we have consent to use data (e.g., open-source, user-provided). We must also be transparent that an AI has created content when it is autonomously delivered (e.g., chatbot response to a consumer, use of watermarks).
4. **Empowerment**: There are some cases where it is best to fully automate processes but there are other cases where AI should play a supporting role to the human – or where human judgment is required. We need to identify the appropriate balance to "supercharge" human capabilities and make these solutions accessible to all (e.g., generate ALT text to accompany images).
5. **Sustainability**: As we strive to create more accurate models, we should develop right-sized models where possible to reduce our carbon footprint. When it comes to AI models, larger doesn't always mean better: In some instances, smaller, better-trained models outperform larger, more sparsely trained models.

www.salesforce.com/news/stories/generative-ai-guidelines/

Generative AI: 5 Guidelines for Responsible Development

"As businesses race to bring this technology to market, **it's critical that we do so inclusively and intentionally**. It's not enough to deliver the technological capabilities of generative AI, we must **prioritize responsible innovation** to help guide how this transformative technology can and should be used – and **ensure that our employees, partners, and customers have the tools they need to develop and use these technologies safely, accurately, and ethically**."

www.salesforce.com/news/stories/generative-ai-guidelines/

Accuracy: deliver verifiable results…communicate when there is uncertainty about the veracity of the AI's response

Safety: make every effort to mitigate bias, toxicity, and harmful output…protect the privacy of any personally identifying information

Honesty: Respect data provenance and ensure that we have consent to use data…be transparent that an AI has created content when it is autonomously delivered

Empowerment: identify the appropriate balance to "supercharge" human capabilities

Sustainability: develop right-sized models where possible to reduce our carbon footprint. In some instances, smaller, better-trained models outperform larger, more sparsely trained models.

Don’t expect a path of least resistance

Innovation is an unnatural act

Opportunity & energy? Insufficient. **People are inclined to *refine***

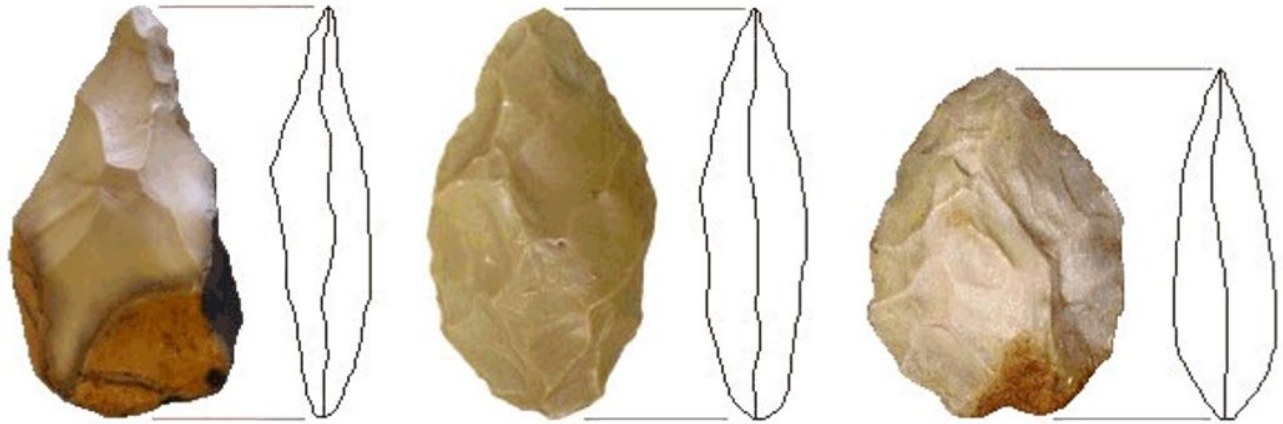

Ten thousand generations of cave-dwellers' effort…
refined their hand axes from 5.5 to 4.5 centimeters

If *everyone* is OK with it – **probably *not* a "revolution"**

Merriam-Webster

revolution noun

rev·o·lu·tion | \ˌre-və-ˈlü-shən\

- a sudden, radical, or complete change
- a fundamental change in the way of thinking about or visualizing something
- a change of paradigm – "the Copernican *revolution*"

Invention ≠ innovation; **innovation ≠ revolution**

"Revolution is a science for the few who are competent to practice it. It depends on correct organization and above all, on communications."

- R. A. Heinlein, 1966

Are you ready to be a "revolutionary"?

- Prepared to incur people's anger?
- Accepting of the burden of communication?

We'd like to partner with you on what comes next

"**Make no little plans**; they have no magic to stir people's blood and probably themselves will not be realized.

"Make big plans; **aim high in hope and work**, remembering that a noble, logical diagram once recorded will never die, but long after we are gone will be a living thing, asserting itself with ever-growing insistency."

- Daniel Burnham, via *Chicago Record-Herald*, Oct 15, 1910

Thank You


Salesforce Strategic Research
We Don't Predict. We Reveal.