

Generative AI: Hope, Hubris, or Harrowing?

June 2023

Generative Artificial Intelligence (AI) models such as ChatGPT¹ represent a transformative leap in the ability of machines to exhibit human-like behavior and to learn autonomously. AI can now surpass humans in a wide variety of cognitive tasks. Accordingly, a recent paper out of the University of Pennsylvania concluded that large language models could displace at least 10% of work tasks, affecting 80% of the US workforce.² But AI, like all new technologies, will also create new professions — such as prompt optimizers — who, for now, are the most adept at instructing and securing the best responses from generative AI systems.

“ Like other major technological advancements, generative AI will eventually be disinflationary.”

Like other major technological advancements, generative AI will eventually be disinflationary. Consider farming, where the automation of most farming activities increased agricultural productivity by a factor of some 300% (in 1900, 40% of the US workforce was employed in agriculture, vs. only 2% today), and lowered food prices. At present, the combination of aging populations in the West and North Asia that is leading to chronic labor shortages; a post-pandemic emphasis on securing supply chain security and redundancy (as opposed to a sole focus on purchasing at the lowest cost from anywhere in the globe), as well as less tolerance for fiscal austerity among voters (on both the left and the right), is bolstering structural inflation globally. This is why despite considerable interest rate hikes, central banks in the U.S., U.K., and Europe have been struggling to bring inflation back down to their target levels. However, as the productivity enhancements that will accrue from AI become more embedded in the economy, central banks' struggle to tame inflation is likely to come to an end.

Technological improvements typically affect productivity with a lag. But every successive new technology has experienced a shorter adoption lag (Chart 1). AI's progression is following an exponential curve, not a linear one, meaning that advances could come much faster than expected. Therefore, we believe that AI's deflationary impulses will likely take hold in the next three to five years.

Author



Tina Byles Williams
CEO & CIO | Founder


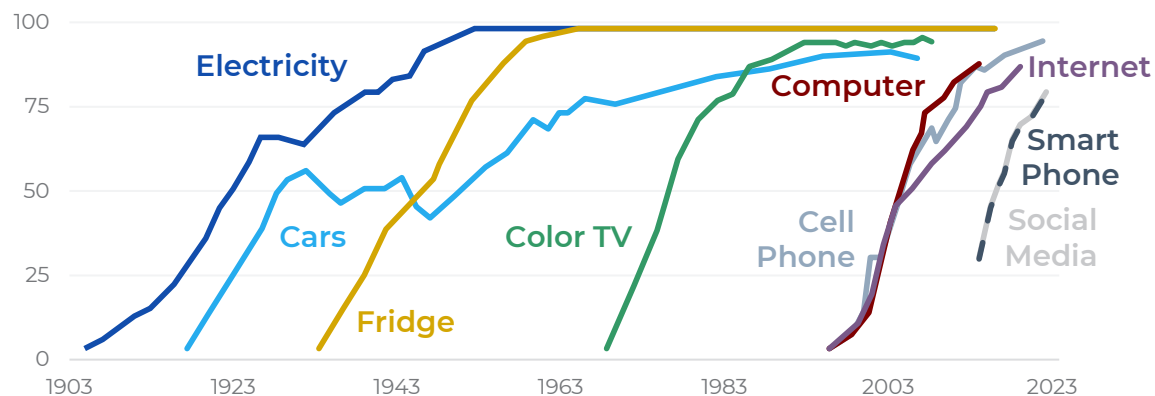


Chart 1 | Tech Adoption Rates Speeding Up
Share of U.S. households using technology



Source: OurWorldInData, TS Lombard



AI will also have profound long-term implications on how investments are managed, how investment acuity is developed, and what type of skill sets are rewarded. Contrary to many predictions that AI will be “narrow” for many years, the latest systems have a far wider scope than those that came before. Most use case studies extrapolate linearly from what AI can do today to what it can do tomorrow. They continue to imagine that AI systems will be confined to “routine” activities —the straightforward, repetitive parts of their jobs — such as document review, administrative tasks, and everyday grunt work. AI is also allowing firms to augment the intelligence of the human workforce and develop next-generation capabilities. A Deloitte study on the use of AI in the investment industry identified the following use cases:³

Ten AI Use-Cases in Investment Management

Portfolio management and client enablement:

Automated insight:
reading earnings transcripts to assess management sentiment

Relationship mapping:
identifying nonintuitive relationships between securities and market indicators

Alternative datasets:
analyzing alternative data such as weather forecasts and container ship movements, monitoring search engines for words on specific topics to structure hedging strategies

Growth opportunities:
using corporate website traffic to gauge future growth along with clients’ behavioral patterns

Client outreach:
smart client outreach and demand generation via analytics, using alternative data sources such as social media data

Front, middle, and back office efficiency:

Operations intelligence:
using machine learning to automate functions

Powering risk performance:
AI-based algorithms and machine learning to monitor for suspicious transactions, and trigger response protocols

Reporting and servicing:
generating reporting for clients, portfolio and risk commentary, and marketing materials using natural language processing

On-demand reporting:
chatbots and machine learning used to respond to employee or investor queries, generating management reporting on-demand

Employee insights:
monitor employee conduct risk and employee morale

Source: Deloitte

But AI adoption will rapidly proceed from low-end repetitive work to more high-end applications — i.e., from language (information) to understanding (analysis) to logic (decision-making). Just as the investment community and the broader public were blindsided by the exponential increase in cases during the early days of the pandemic, they will be blindsided by how quickly AI transforms the world around us. When distributed AI models such as ChatGPT reach the point where they can train themselves – much like Deepmind’s AlphaZero can train itself to master chess without ever being taught the rules of the game – then they will be able to recursively improve themselves at an astronomically fast rate. Even computer coding is being replicated with remarkable acuity. AlphaCode, developed by DeepMind, outperformed almost half the contestants in major coding competitions. Before the early 2010s, the computing power used to train the most advanced AI models grew in line with Moore’s Law, doubling around every 20 months. In the past decade, however, it has accelerated to doubling approximately every six months. Using neural network technology, large language models are most likely to displace workers in law, education, and finance, while jobs that require either human contact or fine motor skills such as personal care, food preparation, construction, and maintenance and repair, are likely to be the least impacted. While traditional roles within IT and management consulting are likely disrupted, these sectors will also be

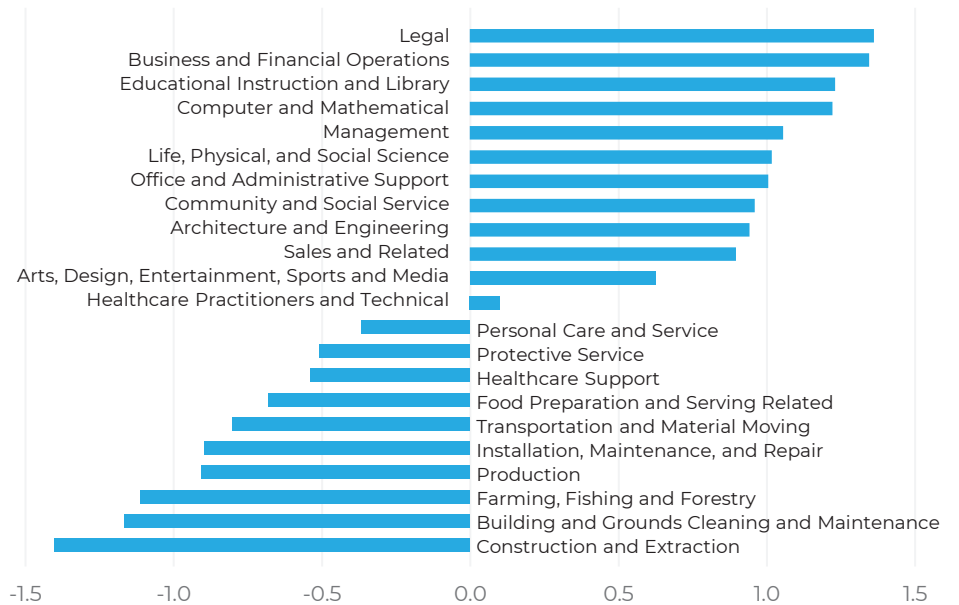


pivotal for facilitating the integration of AI into diverse industries. This could in turn generate a range of new job opportunities within these sectors (Chart 2).

Estimates vary, but we could also be less than a decade away from Artificial General Intelligence or AGI, which usually refers to a computer system capable of generating new scientific knowledge and performing any cognitive task that humans can. What is certain is that creating AGI is the explicit aim of the leading AI companies, and they are moving towards it far more swiftly than anyone expected. There is no analog for a world in which superintelligent machines become more intelligent than humans and can learn and develop autonomously. Superintelligent machines will also be able to understand their environment

without the need for supervision, adjust their programming and transform the world around them. AGI evangelists posit superintelligent computers' potential ability to solve our biggest challenges — such as cancer, climate change, and poverty. These ambitions are not unfounded. In 2021, Google's DeepMind's AlphaFold algorithm solved one of biology's greatest conundrums, by predicting the shape of every protein in the human body. The table below shows how the amount of data and “compute” (the processing power) used by AI systems, as well as their capabilities have increased over the past decade. Each successive generation of AI has become more effective at absorbing data and computing, and thus the more it gets, the more powerful it becomes.

Chart 2 | AI Occupation Exposure⁴
10-Year Moving Average, Ann % Change



Source: BCA Research

| | 2012 | 2022 |
|-----------------------------------------------|-----------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Compute used to train largest AI model | 1e+16 FLOPS (10,000,000,000,000,000) | 1e+24 FLOPS ¹¹ (1,000,000,000,000,000,000,000,000,000) |
| Data consumed by largest AI model | Imagenet: a dataset of 15mn labelled images (150GB) | Datasets of more than 2bn images or much of the text on the internet (estimated at 10,000GB*) |
| Capabilities of largest AI models | Can recognize images at “beginner human” level Superhuman at chess | Superhuman or high-human at a wide variety of games (Go, Diplomacy, Starcraft II, poker etc) Human-level at 150 reasoning & knowledge tasks Passes US Medical Licensing Exam, passes the Bar Exam Displays complex capabilities like power-seeking, deceiving humans Can self-improve by “reasoning” out loud Can write 40 percent of the code for a software engineer |

Source: Ian Hogarth, co-author of the “State of AI” report



There are of course risks with the latest AI. A recent technical paper on GPT4 acknowledges that they can “amplify biases and perpetuate stereotypes.” They can “hallucinate” and provide inaccurate or nonsensical information. But our rapid path towards AGI is ushering in a force beyond our control and understanding. There are three potential paths that personally give me pause. One that harms specific groups of people. There are already numerous examples of AI-generated recommendations for employment or policing that reflect societal biases toward people of color and women.



AI can replicate someone’s voice and face through deepfakes — a goldmine for scam artists or those who specialize in sowing misinformation. That will affect people’s already frayed sense of trust and could foment further fractionalization in our society. The second is that the increasing use of generative AI for content creation could risk a self-referential cycle as online information is used to train future models. This could potentially lead to a narrowing of the information spectrum, with models primarily learning and iterating on AI-created content. This poses a risk to the diversity and originality of information. The third is its ability to rapidly affect all life on Earth. An AI that can achieve superhuman performance at writing software could also be used to develop cyber weaponry. Even if an AI system is designed with benign or beneficial goals, it may still pursue actions that are harmful to humans if those actions are seen as necessary to achieve its instrumental goals. The latter scenario was explored at length by Stuart Russell, a professor of computer science at the University of California, Berkeley. In a [2021 Reith lecture](#), he gave the example of the UN asking an AGI to help deacidify the oceans. The UN would know the risk of poorly specified objectives, so it would require by-products to be non-toxic and not harm fish. In response, the AI system comes up with a self-multiplying catalyst that achieves all stated aims. But the ensuing chemical reaction uses a quarter of all the oxygen in the atmosphere. “We all die slowly and painfully,” Russell concluded. “If we put the wrong objective into a superintelligent machine, we create a conflict that we are bound to lose.”

To be sure, every new technology has been met with suspicion and resistance because major technological change can amplify anxieties and the fear of loss or change. From the attacks on Gutenberg’s printing press in the late 15th century to the current debates on the potential dangers of automation, AI, and gene editing, such resistance can take various forms – from staging strikes to implementing regulatory barriers. Even Benjamin Franklin’s lightning rod evoked fear among church elders that it would unleash the wrath of God by attempting to “control the artillery of heaven.” But what’s striking about AGI is that some of the most strident town criers are its very creators, who seemingly knew about its risks from the outset. In 2011, DeepMind’s chief scientist, Shane Legg, described the existential threat posed by AI as the “number one risk for this century, with an engineered biological pathogen coming a close second.”

Given the awesome potential of AI, it is critical that both industry and government policymakers focus as much energy and capital on ensuring that it aligns with human values as they do on the race to get to AGI. Unfortunately, investment in developing AI capability far outpaces the resources devoted to creating AI alignment. The goal of AI alignment is to



ensure that AI technology prioritizes the well-being and values of humanity. As humans, we generally agree on certain values, such as the desire to avoid pain and promote happiness. However, a software system can be programmed to want anything, regardless of whether it aligns with human values or not. Imagine, for example, a self-driving car designed to prioritize getting to its destination as quickly as possible, regardless of the safety risks it poses to pedestrians or other drivers. Alignment researchers helped train OpenAI's model to avoid answering potentially harmful questions. The most recent version of GPT-4 now declines to answer when asked how to self-harm or for advice on getting bigoted language past Twitter's filters. The task will be monumental as we are still discovering how human brains work, let alone that of emergent AI "brains." Achieving AI alignment will also be complicated by the fact that our values are not entirely monolithic. In areas such as reproductive, women and minority rights and many others, what one group perceives as a constitutional or religious right may be at variance with, or trample on other groups' rights.

More ubiquitous AI adoption will also radically change how 20th-century professionals are trained by academia and professional associations, because large chunks of their current coursework will become redundant or need to be retooled. In medicine, for example, computer-based systems now outperform the best human diagnosticians, identifying early stages of cancer and other diseases with unprecedented precision. What does this imply about how we train doctors? Will medical schools be encouraged to jettison courses on traditional medical education, because machines can perform those tasks more effectively? Moreover, as we become more reliant on intelligent machines, we may be inexorably dumbing ourselves down. Unlike prior technological improvements that have made human labor more efficient, AI could be on track to replacing human comprehension. To use a simple example, for many (including myself), dependency on Google Maps has dulled our map reading skills. Because any muscle operates under the universal "use it or lose it" principle.

In the CFA's Handbook of Artificial Intelligence and Big Data Applications in Investments, AI is largely positioned as a handy research assistant in tasks that involve language processing, such as identifying patterns among data sets, report writing, or answering questions through vastly improved customer service chatbots. In effect, AI will provide "assisted driving" rather than "self-driving." According to a recent article submitted to the CFA Institute:

"The risk of replacement by ChatGPT or other AI is higher for positions that rely more on natural language or involve repetitive, automated tasks such as customer support desks and research assistants. However, roles that require unique decision-making, creativity, and accountability, such as product development, are likely to remain in human hands. While originality and creativity have no easy definition, we humans should focus on tasks that we are good at, enjoy, and can perform more efficiently than machines."⁵

Many investment firms, including Xponance, are exploring the application of cognitive technologies and AI to various



Even if one accepts this relatively narrow view of AI replacing lower level language processing and analytic tasks to augment human judgment, originality, and creativity; the problem is that in most professions, those higher order functions typically accrue from a cumulative body of learning, first at an elementary level to understand the building blocks behind complex systems by performing these lower level tasks, and then by trial and error through which one can gain a more fulsome appreciation of the system's possibilities and risks, as well as its connections to other systems. Moreover, with AI models primarily learning and iterating on AI-generated content to train successive models, we could be risking a potential narrowing of the information spectrum available for higher level work by human professionals.



business functions across the value chain. In a world where AI-generated research inputs are universally available, originality, and lateral thinkers, that can incorporate seemingly non-correlated information or information that is not stored in the vast historical dataset from which AI draws its output, will be even more highly valued. As will systems' agility in areas such as knowledge engineering, data science, design thinking, and risk management that will be needed to build and operate systems that will replace old ways of working. Human behavioral dimensions that AI cannot replicate as easily, such as originality, dexterity, empathy, intuition, and emotional intelligence, will remain essential to building connection and trust in all endeavors. These so called "soft skills" are, after all, what distinguishes us as quintessentially human.

¹ GPT stands for Generative Pretrained Transformer architecture. It produces new data based on the training data it has received and represents a significant advancement from natural language processing to natural language generation.

² "GPTs are GPTS: an Early Look at the Labor Market Impact Potential of Large Language Models", Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock, <https://arxiv.org/pdf/2303.10130.pdf>

³ <https://www.deloitte.com/content/dam/assets-shared/legacy/docs/perspectives/2022/fsi-artificial-intelligence-investment-mgmt.pdf>

⁴ Median score of occupations (six-digit standard occupational classification level) in each occupation major group (two-digit standard occupational classification level). Calculation using data from Felten et al. (accessed on May 1, 2023). Source: Edward Felten, Manav Raj, and Robert Seamans, "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and its Potential Uses," *Strategic Management Journal*, (42:12) (2021)

⁵ <https://blogs.cfainstitute.org/investor/2023/05/09/chatgpt-and-generative-ai-what-they-mean-for-investment-professionals/>

This report is neither an offer to sell nor a solicitation to invest in any product offered by Xponance® and should not be considered as investment advice. This report was prepared for clients and prospective clients of Xponance® and is intended to be used solely by such clients and prospective clients for educational and illustrative purposes. The information contained herein is proprietary to Xponance® and may not be duplicated or used for any purpose other than the educational purpose for which it has been provided. Any unauthorized use, duplication or disclosure of this report is strictly prohibited.

This report is based on information believed to be correct but is subject to revision. Although the information provided herein has been obtained from sources which Xponance® believes to be reliable, Xponance® does not guarantee its accuracy, and such information may be incomplete or condensed. Additional information is available from Xponance® upon request. All performance and other projections are historical and do not guarantee future performance. No assurance can be given that any particular investment objective or strategy will be achieved at a given time and actual investment results may vary over any given time.