



Exploring Artificial Intelligence

Assessing AI in the economy and markets

March 2024

Ashley Oerth, CFA®
Associate Global Market Strategist

Cyril Birks
Market Strategy Analyst



This marketing communication is for Professional Clients, Qualified Clients/Sophisticated Investors (as defined in the important information at the end); for Sophisticated or Professional Investors in Australia; Institutional Investors in the United States; In Canada this document is for use by investors who are (i) Accredited Investors, and (ii) Permitted Clients, as defined under National Instrument 45-106 and National Instrument 31-103, respectively; for Qualified Institutional Investors in Japan; for Professional Investors in Hong Kong, for Institutional Investors and/or Accredited Investors in Singapore, for certain specific sovereign wealth funds and/or Qualified Domestic Institutional Investors approved by local regulators only in the People's Republic of China, for certain specific Qualified Institutions and/or Sophisticated Investors only in Taiwan, for Qualified Professional Investors in Korea, for certain specific institutional investors in Brunei, for Qualified Institutional Investors and/or certain specific institutional investors in Thailand, for certain specific institutional investors in Indonesia and for qualified buyers in the Philippines for informational purposes only. It is not intended for and should not be distributed to or relied upon by the public or retail investors.



Contents

1 What is AI?

2 Why are we hearing about AI now?

3 What might AI do to our economy?

4 What are the roadblocks to AI adoption?

5 Investing in AI

Artificial intelligence is a popular phenomenon...

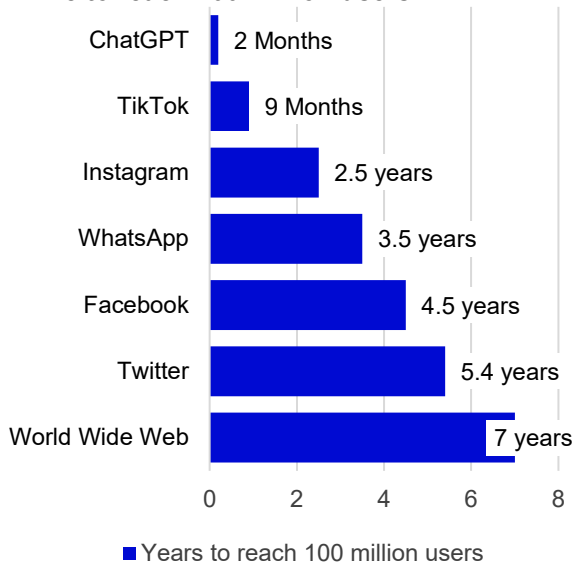
With the launch of ChatGPT, AI has captured public interest

Artificial intelligence (AI) was one of the biggest topics of 2023, as measured by Google search interest. Beginning with the release of ChatGPT on 30 November 2022, AI has rapidly taken the spotlight as the next big thing from technology, reaching 100 million users in a record period of time—just two months.

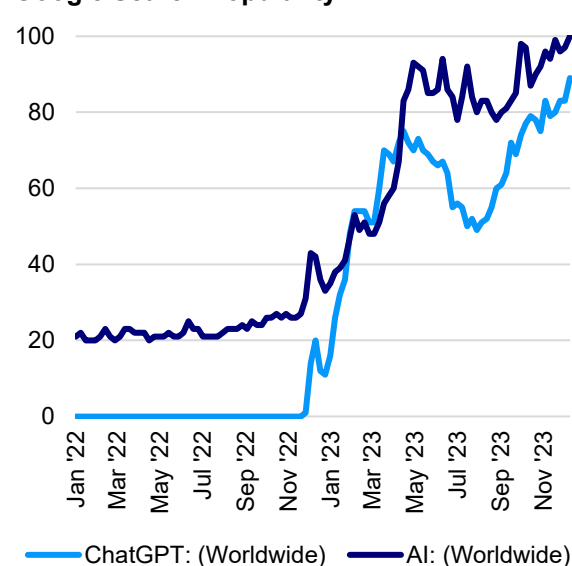
Indeed, artificial Intelligence has elicited grand ideas about how it may fundamentally alter our day-to-day interactions with the world and each other, as private citizens, as scientists and researchers, and as working professionals.

AI is increasingly being embedded in our economic systems, at every point in the supply chain from R&D to sales, from procurement to marketing to after-sale analysis. Despite only recently breaking into the public limelight, AI already has more than a decade of real-world applications.

Time to reach 100 million users



Google Search Popularity



Source: PwC Outlook Perspective report 2023, Wells Fargo Investment Institute, and Google. Charts provided for illustrative purposes only and do not constitute a recommendation.


...and it's receiving greater attention in business

From earnings calls to fundraisings, investors don't want to miss out

As interest in AI has taken off, investors and businesses alike have taken note. Across earnings calls, private fundraisings, and internal discussions, AI has taken industry by storm. Businesses and analysts are actively reviewing how the rise of AI may impact their operations, shift product offerings, and more.

In this article, we look to explore what comprises AI, how generative AI is different, why AI is making headlines now, how this technology may impact the broader economy, and how investors may be able to gain exposure to the AI value chain.


Business interest in generative AI has surged, with a noticeable increase in earning's call references



Sources: Bloomberg, as of January 19, 2024. Quarterly data from January 1, 2020 to December 31, 2023. Document search covers transcripts of earnings calls, shareholder meetings, guidance calls, investor days, and similar calls. Mentions are defined as the number of documents that contain the specified search terms. Search terms used were "generative AI", "generative artificial intelligence", "genai", "gpt", "openai", "LLM", and "large language model."

What is AI?

A brief history of AI and the rise of generative AI



'AI' means many things to many people

For many years, artificial intelligence has been the fascination of science fiction. While recent AI developments have brought new capabilities to what the technology can do, it is still far from the dreams of sci-fi.

Since the 1950s and 60s, AI referred broadly to algorithms letting computers compute and act without human aid. Such actions could be as simple as "if X is true, then do Y". Since then, the interpretation of what is "AI" has evolved.

Algorithms have become increasingly complex. Machine learning and deep learning gave rise to image recognition, algorithmic recommendations, and advanced data analysis techniques. Today, AI also includes generative tools for text, images, and video.

Whether simple or cutting edge, the implications of "AI" hinge crucially on which technologies we are referring to.

For illustrative purposes only.

"Artificial Intelligence" can carry different meanings

Artificial Intelligence

Ingesting data, interpreting patterns and relationships, and acting.



Examples include rules-based systems (essentially flow charts) and expert systems

Machine Learning

Learning from information, usually iteratively



Examples include linear regression, decision trees, and clustering

Deep Learning

Learning by abstracting data relationships.




Examples include neural networks and their applications

Today's AI achievements build on a long history of advancements

In the past decade, AI progress has accelerated, especially in Deep Learning

Number of 'notable' AI systems released each year



Notes: While we have shown OpenAI's 'GPT' versions, it must be noted that Google's BERT was first released in 2018 and multiple AI models have been released since 2018, including from OpenAI, Google, Anthropic, and Meta. Sources: Our World in Data and Invesco, as of December 31, 2023.

AI is already here

Uses of AI in today's economy

As companies scramble to show they are making use of the latest AI technology, we believe it is important to note that AI already has a substantial commercial presence in today's economy.

Voice assistants, content recommendation, spam filters, fraud prevention, facial recognition, maps routing, ridesharing optimization, and computer opponents in games are just a few examples of AI in use today.

In fact, the broad definition of what exactly comprises "AI" enables many companies to claim AI competence and use, which may contrast with consumer and analyst expectations of sophisticated AI implementations.

Examples of Artificial Intelligence in Use Today



For illustrative purposes only.

Generative AI allows models to “guess” new data

Previous popular forms of AI relied on categorization or differentiation

With the introduction of ChatGPT, attention on AI has refocused to so-called “generative AI.”

Previous forms of AI models that have been popularized in recent years have broadly focused on sorting—or “classifying”—data for the purpose of differentiation, interpretation, or understanding what actions to take based on that data. Such AI systems have broadly been labelled as “discriminative AI.”

Generative AI, on the other hand, focuses on creating (generating) new data that looks like its training data. A generative algorithm learns the relationships between pieces of data to and uses these learned patterns to “generate” output data.

For example, models trained on human language can reply to user prompts based on their understanding of previously seen data. In the right-side charts, we show how a simplistic generative model may be used to learn from an existing dataset and infer the next datapoint given an input.

Discriminative AI

Popular forms of AI in use today have primarily relied on sorting data or decisions

Input data

Model learns patterns of data to separate it

Output is a classification or decision

Generative AI

Generative AI, meanwhile, creates new data related to a prompt

Input data

Model connects input data to learned patterns from previously seen data to generate new data

Output is new data

For illustrative purposes only. The above description is a simplification of discriminative and generative models. Discriminative models essentially study the probability of a variable y given x . Generative models, on the other hand, focus on joint probabilities (e.g. the probability of x and y occurring together).

AI systems have rapidly improved in both comprehension and synthesis

Human-level performance is now here

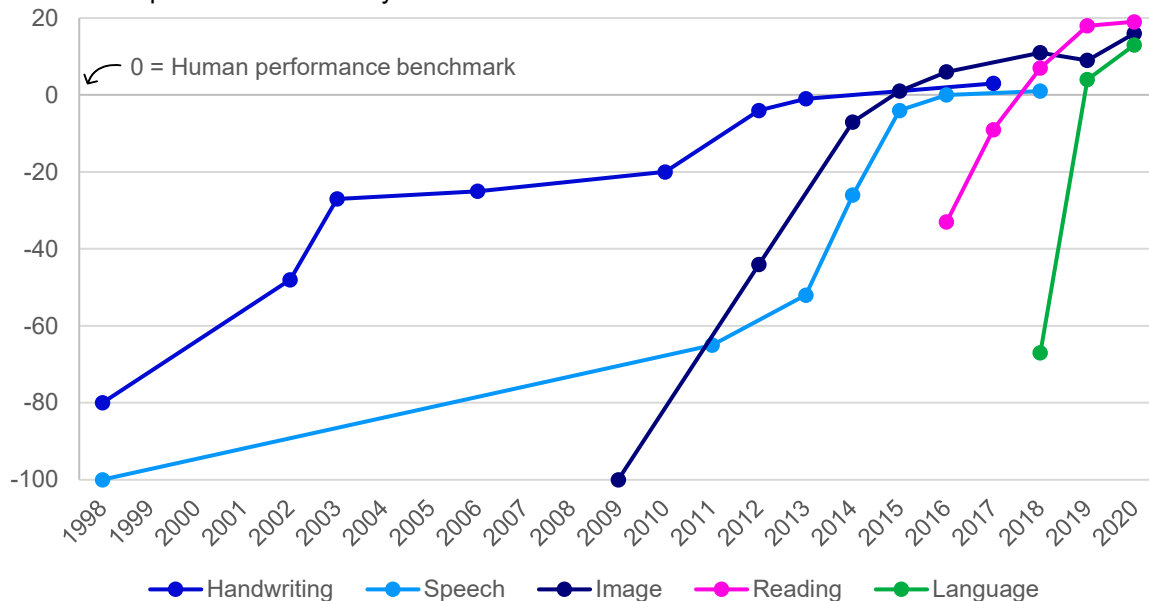
Generative AI technologies have been around for decades, with one such example in 1973 in the form of AI-generated artwork. But only recently have capabilities evolved to merit serious attention from the public.

In the chart on the right, we show a study that benchmarked the performance of notable AI systems across a selection of tasks spanning across different modalities (formats of data, such as text, imagery, etc.). Performance improved steadily over the past several decades and accelerated dramatically beginning around 2014.

Today, benchmarks show AI systems exhibit capabilities on par with average human scores. In fact, anyone can make use of a number of free-to-access models and see these abilities first-hand.

“It’s learning!” AI systems are reaching human-level capabilities in new functions

Normalized performance of AI systems in selected tasks



Sources: Our World in Data and Kiela et al. ‘Dynabench: Rethinking Benchmarking in NLP’ (arXiv: 2104.14337)

Generative AI can treat various kinds of data

Preparing emails, coding, generating images, and speech—generative AI can do it

While ChatGPT has captured the spotlight, generative AI systems are also able to address a number of different types of data—not just text. In fact, AI systems can prepare working code, generate imagery of objects and actions, produce convincing audio of a person (existing or fictional), and more.

Better yet, generative tools can also be interactive, responding to user prompts for refinement, assisting with copyediting, and more. These AI systems can help with laborious tasks, such as literature summaries and documenting code, and also help users navigate large amounts of information.

These examples demonstrate broad potential of this technology, which may have far-reaching consequences in today’s economy.

In the graphic on the right, we summarize a number of AI use cases across different kinds of data, and we provide examples of models that already exist for these modalities. In the next page, we show two examples of generative AI in use.

For illustrative purposes only. Inspired by an Accenture graphic addressing generative AI.

Generative AI Capabilities by Data Type

Text	Code	Image	Audio	Video	Other
Note-taking	Code writing	Image Generation	Voice Synthesis	Video editing	3D modelling
Writing	Code Documentation	Media & Advertising	Automated narration	Video generation	Gaming
Customer Support	Interpretation	Design	Accessibility	Animation	Biology & chemistry
Sales & Marketing	Debugging	Creative Inspiration	Creative inspiration		
Creative Inspiration					Applications
OpenAI GPT-3 and -4	OpenAI GPT-3 and -4	OpenAI GPT-4	OpenAI Whisper	Microsoft X-CLIP	Galactica
DeepMind Gopher	Tabnine	OpenAI DALL-E 2	AudioLM	Meta Make-A-Video	Minerva
Meta LLaMA	Stability.ai	Craiyon	Jukebox	Phenaki	Alphatensor
HuggingFace Bloom	Codex	Imagen		Soundify	GATO
Cohere	Alphacode	Muse			Human Motion Diffusion Model
Anthropic					
Google PaLM					
Google LaMDA					Models

Generative AI can create images...

Prompt: "Teddy bears working on new AI research underwater with 1990s technology"



Sources: DALL-E 2 via OpenAI and Wikipedia, under public domain.

...and produce convincing text

Prompt: "Write me a short description of what it is like to relax on a beach, describing it in the style of Henry David Thoreau."



Upon the sandy expanse, where the ocean kisses the land, I recline, surrendering to the symphony of the waves.

The sun, a gentle orchestrator, casts its golden hues across the canvas of the horizon. Each grain of sand, a testament to nature's meticulous artistry, cradles my repose.

The sea whispers ancient tales, inviting contemplation and quietude, as if time itself were but an illusion.

Here, amidst the vastness of the ocean, I find solace, melding with the rhythm of nature's eternal heart.



Source: ChatGPT 3.5 via OpenAI. Produced as of September 26, 2023.

AI use cases are likely to cover greater functions

Today's out-of-the-box capabilities are just the basics, in our view

Out-of-the-box capabilities	Navigating information	Research summarization, executive summary writing, and meeting preparation
	Draft writing	A first-pass copywriter that can write emails, prepare content, and more
	Copyediting	Proofread, re-word, and re-work existing drafts of e-mails, content, etc.
Potential future state	Tailored service	Customize client experiences through more relevant content and support
	Data analysis	Leverage breadth and depth of existing proprietary data and insights.
	Compliance guidance	Identify problematic language and address common issues.
	Cross-sell	Identify opportunities for cross-selling based on interests and client needs
	Retention	Help detect when clients are feeling dissatisfied and induce human intervention
	Market research	Scan the market, better understand trends from unstructured sources, and help develop new sales propositions, products or services

Why are we hearing about AI now?

The enabling factors of artificial intelligence

2

Why the excitement now?

Recent AI advances build on successive technology and data progress

AI has been around for many decades, with even ancient Greek philosophers discussing intelligent machines. In modern history, from the 1950s onward, AI progress has gone in fits and starts. So what makes today's moment different—why is AI having its time in the sun now?

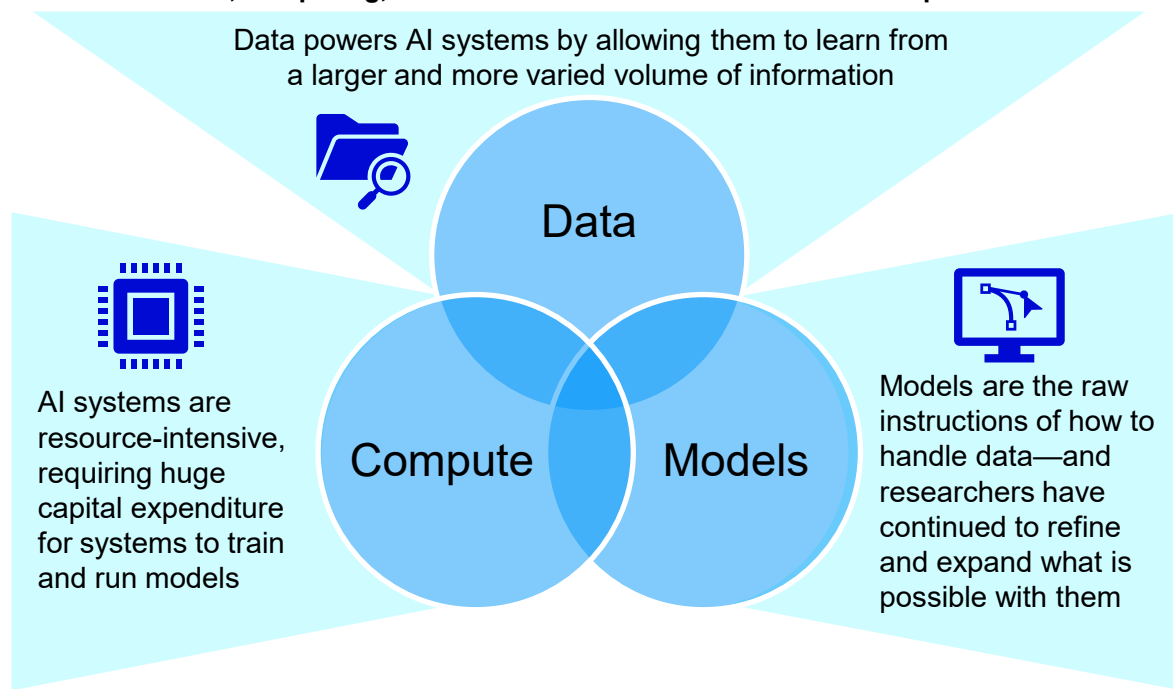
One driver is macroeconomic—as societies age and when specialized labor is expensive or hard to find, labor-saving technologies become more desirable. As businesses seek to manage costs, AI presents itself as a potential source of new efficiencies.

However, we believe progress in input technologies is a more important driver of AI advancements. We see three enabling factors, all of which are interrelated in driving AI forward: More data is captured and available today than ever before; greater and cheaper computing power; and better and more sophisticated models.

In the following slides, we explore each of these in more detail.

For illustrative purposes only.

Advances in data, computing, and models have enabled the rise of AI capabilities



Why the excitement now?


1.) More data is generated today than ever before

Global data volumes are exploding. Never before have we created and stored so much data.

In fact, 90% of global data was created in the last two years, and is expected to double every two years. More data is created per hour today than in an entire year just two decades ago.


Data is the lifeblood of AI systems. A popular rule of thumb is “garbage in, garbage out”, referring to the quality of data used to power models. As data volumes have grown, AI systems have had far more material to work with to power ever more powerful systems.

AI's Enabling Factors: Data



90% of global data was created in the last two years.

Social media is just one example of exploding data volumes



Note: Chart shows absolute increase in data created across various social media platforms between 2013 and 2022.

Source: Bank of America Global Research, Domo. NOTE: Google data is based on user queries, YouTube data is based on hours uploaded, Instagram data is based on photos shared, Twitter data is based on tweets shared, Facebook data is based on content shared, Email data is based on emails sent. Used with permission. See page 50 for BAC disclosures.

Why the excitement now?

2.) Computing power has continued to grow, with new hardware and methods


Advances in computing power, combined with more sophisticated methods of splitting computing tasks across multiple computers, have been the precondition to AI systems as powerful as today's offerings.

Creating and then using an AI system involves two resource-intensive steps—training the model, and then using it.

When AI systems are “trained,” they undergo a process by which the model essentially learns from a given set of data for it to be later applied to new data. For perspective, training the model that underlies OpenAI's GPT-3 on a single enterprise-grade graphics processing unit would have taken 288 years.


Then, after a model has been trained, it still requires additional processing power each time it is used. Across the ChatGPT platform, daily estimated running costs range from \$100,000 to \$700,000 USD, due in large part to the computational needs of the underlying model.

AI's Enabling Factors: Compute



Computing power needed to train the best-in-class AI models has increased **300,000x** in the past decade.

Moore's Law: Exponential increases in computing power



Sources: Graphic: AI and Compute, OpenAI –May 16, 2018, and OurWorldInData as of May 19, 2023. Note: Vertical axis uses exponential scale. Text: arXiv: 2104.04473, Narayanan et al., “Efficient Large-Scale Language Model Training on GPU Clusters Using Megatron-LM” August 23, 2021.


Why the excitement now?

3.) Model complexity and approaches have continued to evolve

A combination of new approaches and a growing level of model sophistication have enabled more capable AI systems.

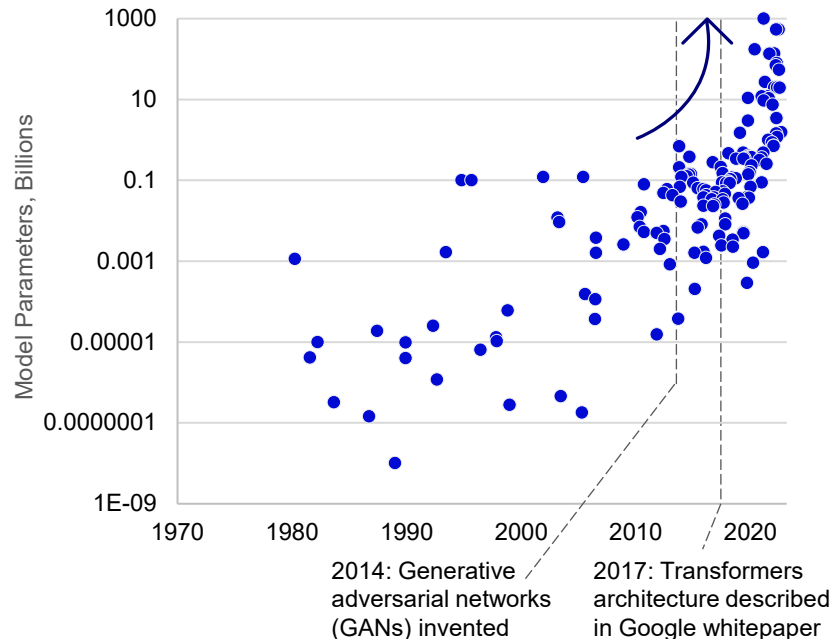
New methods—especially transformers, which underly the majority of the latest language models—have propelled model capabilities to new highs. And in the last decade, AI systems have seen an exponential growth of the number of datapoints they can simultaneously use due to advances in computing power and methods.

AI's Enabling Factors: Models



Models are ever more sophisticated –ChatGPT-4 was trained with **1 trillion parameters**.

Growth of parameter counts of AI models has accelerated



Source: Sevilla, Villalobos, Cerón, Burtell, Heim, Nanjajjar, Ho, Besiroglu, Hobbhahn, Denain, and Dudney (2023) Parameter, Compute, and Data Trends in Machine Learning via OurWorldInData. Chart shows selected models and estimated parameter counts from 1970 to July 2022. Note: Vertical axis uses exponential scale.

What might AI do to our economy?

The macroeconomic effects of artificial intelligence

3

AI in the economy

Greater possibilities with AI toolsets

We expect AI to bring about meaningful transformations to our economy by providing new toolsets that assist, automate, and augment everyday tasks.

Research also suggests that AI helps drive new outcomes by helping users with decision-making and operations, either as a new toolset or a tutor in certain tasks. For example, one study examined how chess players improved alongside the successive improvements in AI computer opponents. Players are increasingly making moves and using strategies that are closer to the actual theoretical optimum, in part because AI is better at calculating those moves.

Whether it's saving time or improving outcomes, we expect AI to increasingly be embedded in our lives, helping to enable greater economic possibilities.

Anticipated effects of AI in economic functions

Assist



From formulating emails to suggesting brand materials, AI may find its best use today as an assistive tool to everyday job activities, especially in information-heavy tasks.

Automate



Whether it's self-driving cars, automated manufacturing processes, or helping clients choose a product, AI may have far-reaching labor-saving applications.

Augment



Through enhanced product offerings, unlocking cost efficiencies, or expediting research through new toolsets, AI may boost worker outputs by providing new capabilities.

For illustrative purposes only.

The macro effects of AI

AI may enhance productivity and reduce inflation, but labor may be at risk

The future macroeconomic impacts of AI are uncertain. Expectations range wildly from a non-event to a productivity bump to a world-changing displacement of labor.

In our view, workers are likely to find use cases of AI that enable them to be more productive in their roles, increasing output per hour and helping to pull productivity growth in the developed world out of a multi-decade slumber.

As productivity increases and AI integration grows, we expect the technology to help pull inflation lower through falling costs and increased output. As part of this, we anticipate some labor displacement, especially in information-focused roles and tasks.

In the following slides, we explore each of these themes in more detail.

Anticipated macro impact areas of the rise of AI



Productivity

AI tools may increase productivity by automating laborious and monotonous tasks and providing a “co-pilot” for more sophisticated, higher value-add work



Disinflation

Continuing a long history of tech-driven deflationary forces, AI tools are likely to become more broadly implemented and cheaper over time, while also enabling an expansion of aggregate supply



Labor displacement

As AI capabilities make their way into more economic tasks, we expect a substitution effect where AI replaces some human labor, especially for information-heavy but lower value-add functions

For illustrative purposes only.

AI may make us more productive

Information and communication technology revolution of 1995-2005 may serve as guide

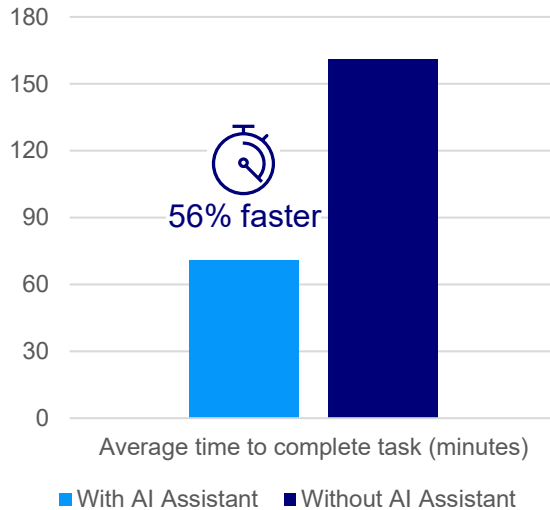
Productivity, roughly speaking, measures the output each worker produces for every hour worked. Increases in productivity have historically been associated with major technological change, such as the invention of the assembly line, automobile, personal computing, and the Internet.

AI, especially generative AI, may be the next driver of productivity increases. In one example, Github—a code collaboration and repository website—paired developers with an AI coding assistant. Those using the AI system performed the given task 56% faster than the control group—a remarkable difference that translates into more productive workers.

We believe that, as generative AI is integrated into various aspects of the economy, productivity growth may see a revival similar to what was last witnessed in the 1995 to 2005 period in the US, when the information communication technology revolution took place, which saw productivity rise at a rate three times that of recent decades.

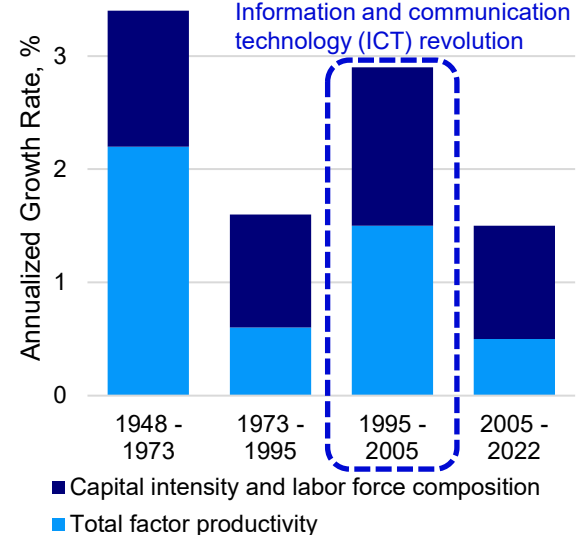
GitHub's CoPilot Coding Assistant

Time for professional programmer to complete task



US Labor Productivity Growth

Components of productivity growth, 1948 to 2022



Sources: Macrobond, Github: *Research: quantifying GitHub Copilot's impact on developer productivity and happiness* from September 7, 2022; Bureau of Labor Statistics, and Brookings Institution, as of May 19, 2023.

Productivity uplift would likely bring meaningful GDP growth

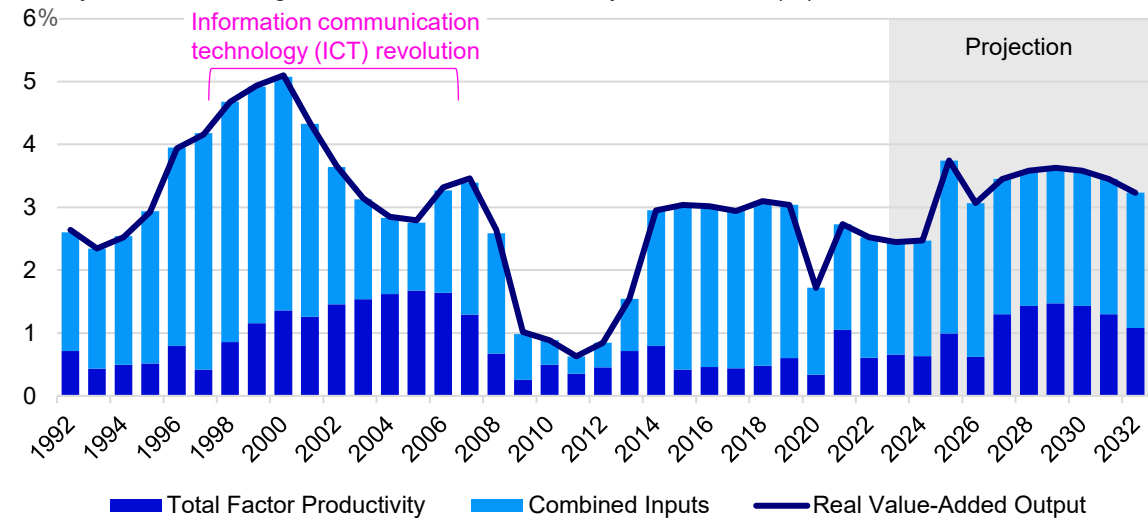
Applying the ICT revolution effect would push growth to multi-decade highs

Higher productivity growth would boost overall economic growth, in our view. In the study on the right, we apply a higher rate of productivity growth modeled after the experience of the information communication technology boom. We assume an upward glide path from 2024 that peaks in 2027, falling back to the post-2009 trend thereafter.

Based on this experience, the US economy would be 7.4% larger than without the contribution of artificial intelligence by the end of the forecast period. In our figuring, this result is from productivity growth averaging 0.7% higher per year due to the effects of AI integration across the economy.

Generative AI may revive productivity growth—and boost GDP

US 5-year annualized growth of real value-added by contribution (%), 1992 to 2032



Notes: annual data from 1992 to 2032. Values from 2023 to 2032 are projections, where combined inputs grow at the average growth rate observed from 2010 to 2022 and total factor productivity grows at an annualized growth rate of 1.6% (the growth rate observed from 2010 to 2022 plus an additional 1.1% that is assumed to derive from the effect of artificial intelligence). The rate of total productivity growth glides upward to 1.6%, peaking in 2027 and then falling back to the 2010-2022 trend thereafter. Because of this glide path, the total productivity effect is assumed to be lower than the ICT revolution experience from 1995-2005. As of December 31, 2023. Source: Macrobond, US Bureau of Labor Statistics Productivity Database and Invesco Global Market Strategy Office

Technology has a long history of deflationary effects

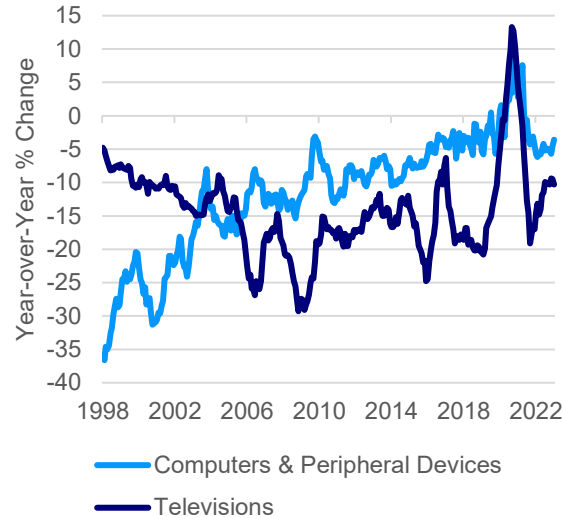
Expansions of aggregate supply should be deflationary, too

Technology has a long history of deflationary effects, especially when valuable labor hours are freed up to be allocated to other tasks, increasing aggregate supply in the process. As supply expands and workers become more productive, they tend to also become richer in the process.

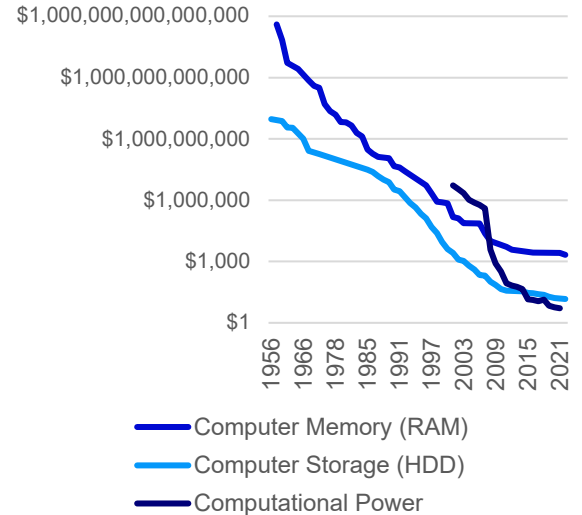
Moreover, hardware involved in information technology are increasingly being produced more cheaply and at greater scale. Recent technological developments that make use of computers, data storage, and connectivity have accelerated deflationary forces as prices have fallen.

Following in the footsteps of past technology developments, AI has the potential to exert deflationary forces via both increased aggregate supply and falling costs of the required hardware.

Ex-pandemic, tech goods have fallen in price
US CPI sub-indices on selected technology goods



Computation & data costs have fallen exponentially
Cost of selected computing-related factors



Sources: Macrobond and Our World In Data, as of December 31, 2023. Right-side chart: Note: Scale is logarithmic. Some unreported or missing values have been interpolated for illustrative purposes. 'Computer Memory (RAM)' and 'Computer Storage (HDD)' are in USD per terabyte. RAM = Random Access Memory, HDD = Hard Disk Drive. 'Computational Power' is a measure of the cost of sequencing one billion base pairs of DNA in USD per gigabase.

AI may displace labor, but history suggests it will persevere

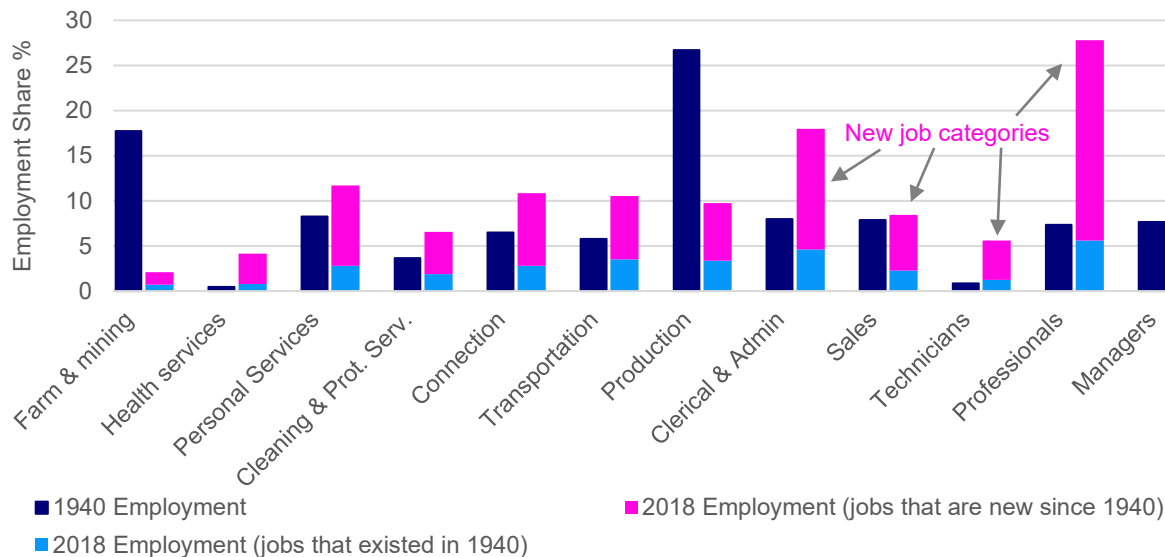
Technology tends to evolve the role of labor, not remove it

History is full of examples of creative destruction, where technological advancements create an ever-changing business landscape and, usually, economic growth. For example, electricity, cars, computers, the Internet, and countless other examples have transformed the ways we work.

As technology changes, so does the role of labor. In 1870, 46% of US occupations were in agriculture, while about 14% of jobs were in services. By 1940, the share shifted to just 17% of occupations in agriculture and 41% in services. Today, the US is an information and services economy, where services make up a whopping 86% of jobs.

Generative AI may begin a new cycle of creative destruction, displacing labor but maintaining the need for a “human in the loop” across most functions. Indeed, even today, countless examples of AI have still left us with unemployment rates at multi-decade lows (for example, official end-2023 rates are 3.7% in the US and 5.9% in the Eurozone).

History suggests technology tends to change labor allocation to higher-value-add functions
US employment by type of job, 1940 versus 2018



Note: “Cleaning & Prot Serv.” is Cleaning and Protective Services. For illustrative purposes only, based on Autor, D. “The Labor Market Impacts of Technological Change; From unbridled enthusiasm to qualified optimism to vast uncertainty” NBER Working Paper.

What are the roadblocks to AI adoption?

Challenges range from fundamental capabilities to ethical and legal issues

4

Challenges to AI adoption

AI abilities, legal challenges, and familiarity pose friction

Trust in AI models



Hallucinations

AI systems are known to fabricate datapoints, literature, facts, and more, leading professionals to question their abilities.



Safeguards

Large language models and image generation tools may be used for malicious purposes, which forces AI engineers to build in safeguards to prevent unwanted behavior.

Legal & regulatory challenges



Misinformation

Image and voice generation enable propagation of fake or misleading stories, images, videos, voiceovers, and more, which may drive restrictions on AI deployments.

Barriers to entry



Upfront & ongoing running costs

AI requires substantial investment for designing, training, and specializing software. Afterward, running costs are incurred every time a pretrained model is used.



Bias

Models are only as good as the data that we feed them. Datasets may have ingrained biases. Users of such tools must be vigilant to their shortcomings.



Intellectual Property & Rights

Since AI is built on previous content, it can use work from others—oftentimes without authorization.



Worker Sentiment

Workers may be reluctant to use AI for fear of their jobs becoming automated and ultimately displaced. Workers may also distrust AI tools.



Technical expertise & know-how

Business leaders may wish to make use of AI but not have the technical capabilities to integrate it into products or/and workstreams.

For illustrative purposes only.

Hallucinations

AI systems can fabricate or misinterpret information

AI hallucinations refer to instances where an AI system has misinterpreted or fabricated information.

Hallucinations may be caused by insufficient or unrepresentative training data, overly complex models that see patterns in the data where there are none, or may simply be a byproduct of how AI models are designed.

Regardless of the cause of hallucinations, the end result is an erosion of trust—users may put too much confidence in model outputs and later be proven wrong, or users may be reluctant to make use of the tools at all.

Hallucinations are difficult to solve, with the usual solution being mitigation and reduction of their frequency. For example, developers may try to use broader and more comprehensive datasets for training, and to apply robust validation techniques to ensure the model generalizes well to new data.

Nevertheless, for users of AI output, proofreading and factchecking outputs remains essential.

AI tools are by no means infallible

April 2023
Non-existent Citations

Researchers find in a study that almost half of 178 generated research references were non-existent or could not be found.

May 2023
Fake Legal Precedents

A US lawyer submits fake case precedents to the Southern District of New York, generated by a chatbot research aid.

2021
Blind Vehicles

Researchers discover that real world objects can be rendered 'invisible' to image recognition AI used in self-driving vehicles.

Sources: Gnanasambandam et al. Optical Adversarial Attacks ([arxiv.org /2108.06247](https://arxiv.org/abs/2108.06247)); Anthaluri et al. 'Exploring the Boundaries of Reality: Investigating the Phenomenon of Artificial Intelligence Hallucination in Scientific Writing Through ChatGPT References', April 2023; 'Here's What Happens When Your Lawyer Uses ChatGPT' New York Times, May 27, 2023.

Intellectual property concerns remain

Can AI models be trained on unlicensed content?

Artists spend a lifetime studying and being inspired by great works. Some would suggest this is no different to AI relying on libraries of text and image data. Others contend that AI is a tool, developed by humans, profiting off the unfair, unethical, unauthorised use of intellectual property. Another camp believes the issues of copyright and intellectual property lie not in training AIs, only in the use of their outputs. Additionally, the global nature of AI complicates matters further, as copyright and IP laws vary by country.

These fundamental questions – and others like them -- are now the subject of court cases, like NYT v OpenAI and Getty Images v Stable Diffusion.

Companies providing AI models are keen to suggest that current copyright legislation is sufficiently robust and that, notwithstanding safeguarding procedures, responsibility ultimately lies with the user. The evolution of this debate will likely have important implications for how AI is developed going forward.

Traditional Content Creators

“

By OpenAI's own admission, **high-quality content, including content from The Times, was more important and valuable for training the GPT models as compared to content taken from other, lower-quality sources.** ... Using the valuable intellectual property of others in these ways without paying for it has been extremely lucrative... ”



— New York Times case filing in lawsuit against OpenAI for alleged intellectual property infringement

AI Developers

“

It is the user who determines whether the output implicates the exclusive rights of a copyright owner... ... evaluating whether an output infringes the exclusive rights of a copyright owner turns on how the output is ultimately used, something that is context-specific and falls beyond the control of an AI model developer. ”



— Fred von Lohman, Associate General Counsel, Copyright, OpenAI

Sources: Left-side quote from Case 1:23-cv-11195 filed in the United States District Court, Southern District of New York, filed December 27, 2023. Right-side quote from US Copyright; regulations.gov, 'Docket Document (COLC-2023-0006-0001)' comment ID COLC-2023-0006-8906, November 1, 2023. Bolding is our own.

AI price tag remains high

Trailblazing in AI comes with substantial talent, services, and hardware costs

Using the latest and greatest AI technology often means substantial costs, especially if capabilities are scaled or purpose-built. For example, developing in-house solutions requires substantial investment for designing, training, and specializing software.

Even after model training, running costs are incurred every time a model is used. In many cases, these costs are charged on a pay-as-you-go basis that can add up quickly.

These high costs are likely to dissuade significant, specialized AI use cases in the near-term. Over time, however, we expect costs to fall, with today's cutting edge becoming more widely and cheaply available.

Costs to play in the AI space remain uncomfortably high

Selected costs for AI-related hardware and services



\$288 million per day: Rough estimate of cost to process all daily Google search queries with generative AI chatbots¹



\$100 million: Cost for OpenAI to train GPT-4²



\$10,000: Cost for a single Nvidia A100, a specialized GPU used for AI computing³



\$162,000 per year: Average base salary of a machine learning engineer⁴



\$200,000 per month: Cost for a start-up to process several million generative AI model queries per day⁵

1 = Assuming 4 words in the average search, 0.75 tokens on average per word, 8.5 billion searches per day, and a cost per 1,000 tokens of \$0.03 to \$0.06, this would cost \$1.1 million per day just for processing searches, not including search results. Assuming each search result is 1,000 words, the total cost rises to \$288 million per day. Figures sourced from <https://openai.com/pricing>. 2 = Sam Altman quote from Wired article. April 18, 2023. 3 = CNBC article, "Meet the \$10,000 Nvidia chip powering the race for AI". February 23, 2023. 4 = Indeed.com. 5 = CNBC article, "ChatGPT and generative AI are booming, but the costs can be extraordinary" detailing Latitude's usage of OpenAI's early GPT software and associated running costs in late 2021, from March 13, 2023.

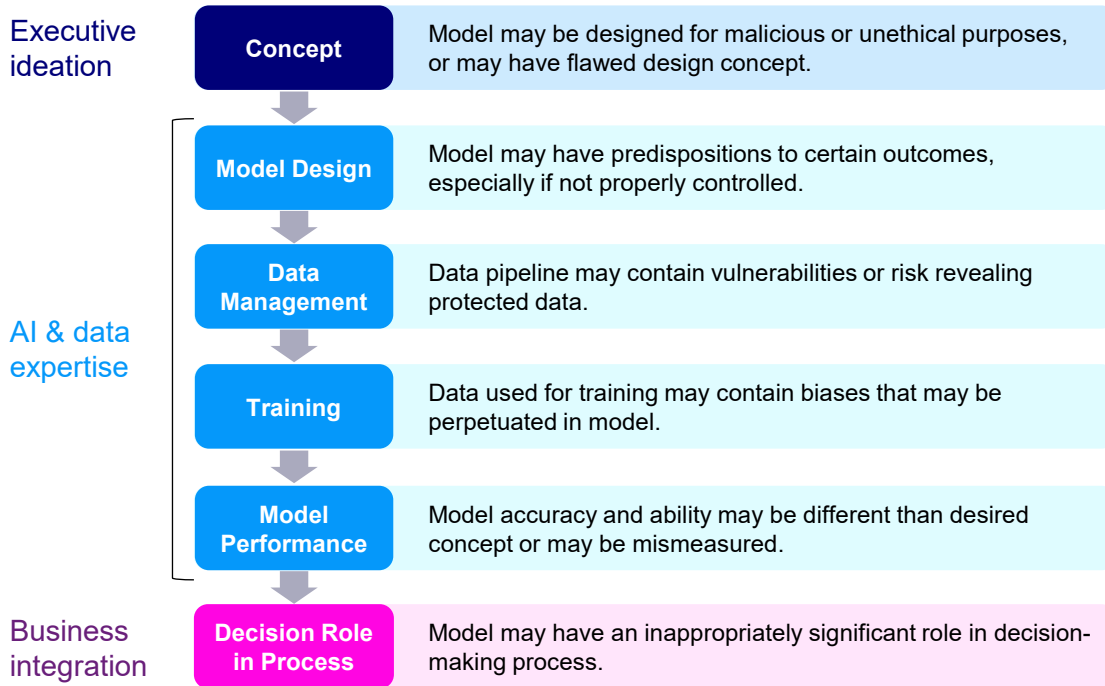
In-house AI expertise appears key to implementation

AI integration into business comes with unique challenges

New tools afford new opportunities, but new possibilities may not be attained if tools are misused. AI tools might be easy to interact with, but they perform best in the hands of expert users who know which tool is right for the job and when to use it.

On the right-hand side, we have outlined some of the risks involved with AI development and the importance of domain specific expertise in both AI systems and the arena they will operate in. Ultimately, we expect a collaborative approach between builders and adopters will be best, unless adopters have those skills in-house.

Source: Invesco. Inspired by McKinsey study on confronting the risks of artificial intelligence. For illustrative purposes only.



Regulation looms as AI presents new threats, challenges

Regulators around the world are facing up to the challenges of a new technology. They are being asked to balance the trade-off between creation and destruction in real-time; tasked with mitigating possible harms without stifling innovation. There are two primary sets of harms: Social harms, and technological harms.

Social harms include society-wide issues, like sudden shocks to labor and the nature of working. Technological harms are potentially vast, spanning from deepfakes to using AI models to generate new biologics with ill-effects.

Lawmakers and regulators have been playing catch-up, consulting industry experts and academia to plot a way forward. Broadly, regulators look to enable continued progress while providing safeguards to social and technological harms.

On the right side, we highlight some commentary from recent developments in the US, European Union, and UK.

Sources: United States White House, European Commission (Shaping Europe's digital future, AI Act website), and UK Parliament (Communications and Digital Committee, "Large language models and generative AI 1st Report of Session 2023-24").

United States

Working through Executive action

[AI] has the potential to ... **make our world more prosperous, productive, innovative, and secure**.... [but] irresponsible use could exacerbate societal harms ... displace and disempower workers; stifle competition; and pose risks to national security. [Harnessing AI for good] demands a **society-wide effort** that includes government, the private sector, academia, and civil society.



— Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, October 2023

European Union

AI Act brings regulatory framework

While most AI systems pose limited to no risk and can contribute to solving many societal challenges, **certain AI systems create risks that we must address to avoid undesirable outcomes.**



— European Commission on the AI Act, February 2024

United Kingdom

Government: Existing rules sufficient for now

In the right hands, LLMs may drive major boosts in productivity and deliver ground-breaking scientific insights. **In the wrong hands they make malicious activities easier and may lay the groundwork for qualitatively new risks.**



— UK Parliament Communications and Digital Committee, February 2024

AI regulation in five pillars

Government working groups center on common themes

Security



Regulatory approaches emphasize the importance of national security, including:

- Managing internal threats and limiting dangerous functionality
- Monitoring adversarial use
- Maintaining economic stability (for example: protecting workers).

Accountability



States appear eager to ensure that there is a balance of responsibility between creators' and users' responsibilities. Concerns include:

- Addressing misinformation
- Managing bias
- Protecting copyright and IP
- Assigning appropriate liability
- Continued monitoring and awareness

Values



AI systems come with baked-in social, political, and moral values based on training data and design choices. There is growing interest to ensure systems reflect national values. For example:

- Promoting democracy and the rule of law
- Refusing, or at least not encouraging 'unethical' tasks
- Designing AI for positive use cases

Transparency



AI systems' can be a 'black box', making it challenging to understand how they arrive at their results or how to reliably predict how they will behave. Transparency means focusing on:

- Being clear about what information developers need to make available
- Educating the public about the uses and limits of AI.
- Being open about what data is collected, when it is collected, how it is used, and how to opt out.

Innovation



Countries want to promote innovation and encourage endogenous solutions to regulatory concerns, rather than dictating terms:

- Emphasizing light-touch regulation
- Supporting national-innovation
- Unlocking AI's potential to revolutionize scientific innovation and new ways of working.
- Maintaining an international advantage
- Creating solutions to transparency and security challenges

Sources: Invesco's 2024 Global Policy Outlook.

Overcoming hurdles

“A problem well-defined is a problem half solved.” — John Dewey

The roadblocks to short- and long-term adoption of AI are numerous, but all are able to be overcome, in our view. In fact, dedicated attention to these issues is precisely what may accelerate AI adoption.

Despite risks, AI tools have been widely adopted as co-pilots and productivity aids where possible. In fact, these tools have proven sufficiently exciting to kickstart a wave of generation, with companies building their own platforms and environments to safely access current tools and mitigate risks.

Meanwhile, regulators have chosen not to stamp out or slow down innovation in AI. While safeguards and barriers are being put in place, the spirit of regulation nevertheless looks to foster innovation.

Silver linings



Risks are known & addressable

AI tools are imperfect; they can hallucinate, produce biased results, and may even infringe intellectual property in some cases. However, these are *known* issues which can be overcome.



Regulators are supportive

In response to immediate and long-term concerns, from IP and bias to national security and labor displacement, new regulation is emerging. However, regulators are clearly keen to embrace AI technology.



“Human-in-the-loop” approaches foster adoption

While we expect AI to replace some human labor, especially for information-heavy but lower value-add functions, we have also seen AI embraced as a co-piloting tool and welcomed as a productivity aid. AI fears may be allayed by keeping a “human-in-the-loop.”

For illustrative purposes only.

Investing in AI

Exploring the AI value chain and its impact on markets

5

Recent run in US stocks has been primarily a technology phenomenon


Magnificent 7 are also AI leaders

Markets appear particularly excited about generative AI. Indeed, technology stocks played an outsized role in supporting S&P 500 returns in 2023.


2023 marked one of the most concentrated periods of equity returns in US history, but it may have been justified based on fundamentals. The “Magnificent 7”—a basket of mega-cap companies that are also major players in the AI theme—accounted for roughly 60% of the S&P 500 Index’s ~26% calendar year total return, while the remaining 493 stocks comprised the other 40%.

However, the index’s earnings per share was largely supported by the Magnificent 7, whose aggregate earnings growth meaningfully outpaced the rest of the market.

AI leaders—particularly the Magnificent 7—have dominated US large-cap stock performance
S&P 500 stocks versus the Magnificent 7



Earnings growth (YoY)



Source: Bloomberg. Daily return data as of January 31, 2024. Earnings data shows latest Q4 2023 (with 91.6% of companies reporting) versus Q4 2022 data, as at February 28, 2024. The “Magnificent 7” includes Microsoft, Apple, NVIDIA, Amazon, Meta, Tesla, Alphabet, with performance calculations based on their respective weight within the S&P 500 Index. **Past performance is no guarantee of future returns.** An investment cannot be made directly into an index.

Investing in AI is more than just tech companies

Value chain includes infrastructure, models and ultimate use cases

To capture the evolution and adoption of AI, we believe that the full value-chain is worth considering for AI-related investment.

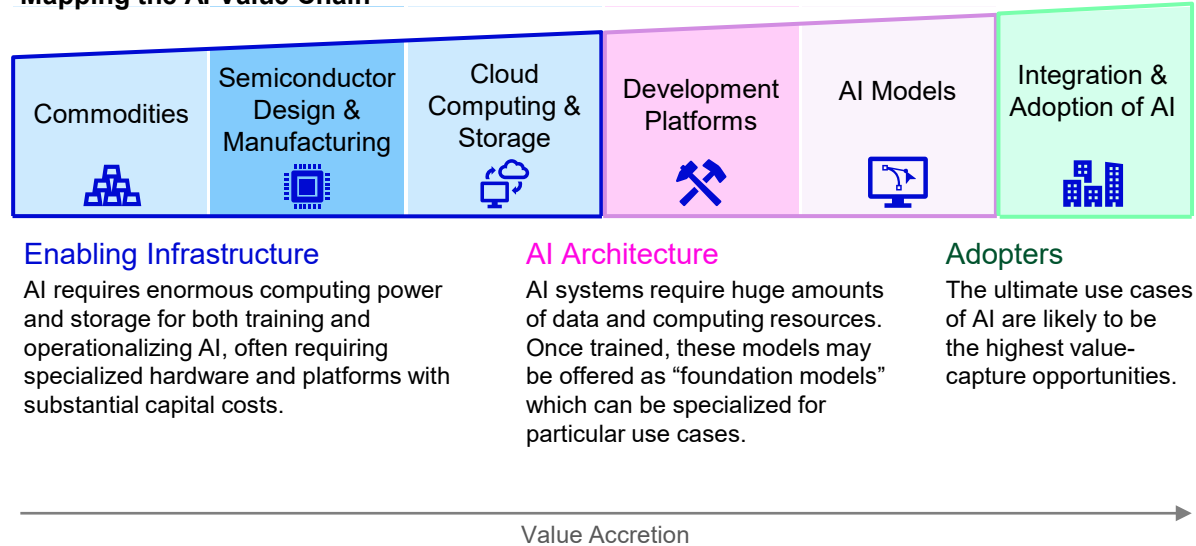
In our assessment, the AI value chain comprises distinct yet interconnected elements, building from raw commodities, to sophisticated computing equipment deployed at scale, to broad-purpose AI systems, to specialized use cases.

In our analysis, few AI-related investments are pure plays, with most companies likely to focus on more than just AI-related products and services. Reflecting this, we note that it is entirely possible to have an “AI” themed portfolio which is considerably more diversified than the label may suggest. We also note that our value chain assessment leaves off important contributors to the AI theme, including AI-specific consulting and services, AI-driven cybersecurity solutions, and more.

In the following slides, we explore each of the themes shown on the right in more detail.

For illustrative purposes only.

Mapping the AI Value Chain



Commodities: Key rare earths in semiconductor manufacturing

Semiconductor raw materials tend to have concentrated suppliers

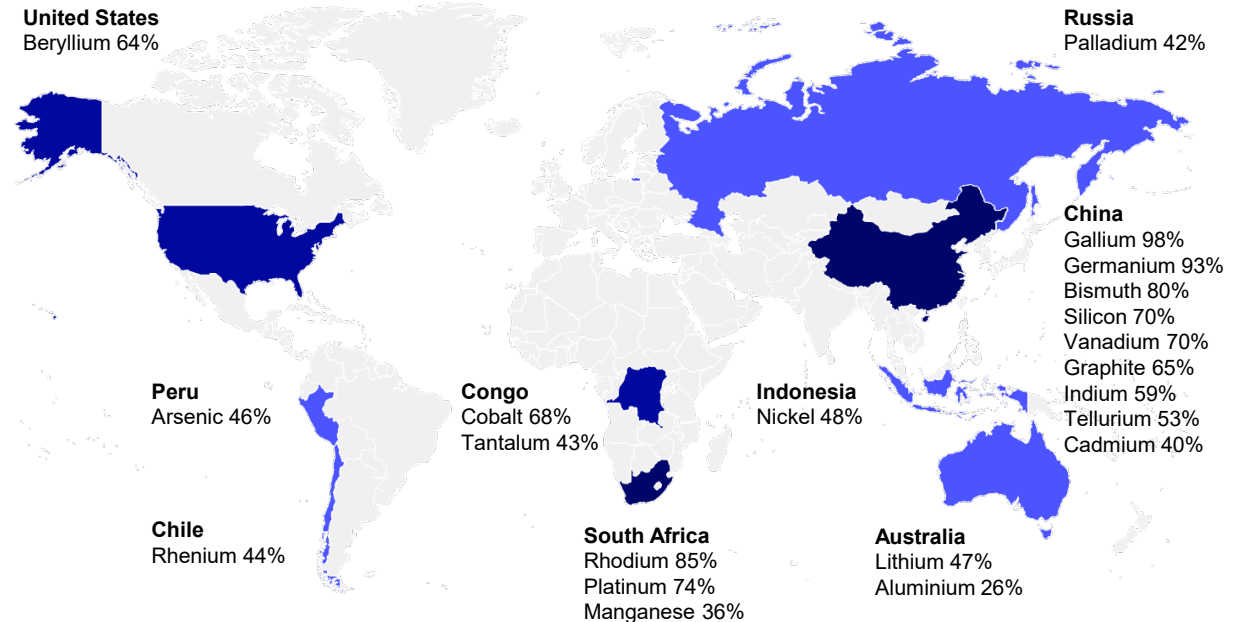


Critical raw materials play a foundational role in technology, being essential to produce over 200 kinds of vital commercial products including integrated circuits (CPUs, GPUs, and other computer hardware typically grouped together as 'semiconductors' along with simpler devices).

The United States Geological Survey (USGS) lists a collection of important elements for manufacturing associated with AI, whether directly in semiconductor materials or indirectly in other computer hardware and telecommunications. Based on an inventory of these inputs, China dominates world production, producing 70% of the world's rare earth metals in 2022. Indeed, the US and EU imported 74% and 98% of their rare earths from China in 2022.

These key commodities may see sustained demand as AI becomes ever more advanced and demands more hardware to enable continued advancements.

Single-largest producers of critical minerals (share of world total supply, %)



Sources: US Geological Survey, 'Mineral Commodity Summaries 2023' and World Mining Data 2021.

Note: World Mining Data and US Geological Survey data do not always perfectly align.

Semiconductors: Computing power is key AI ingredient

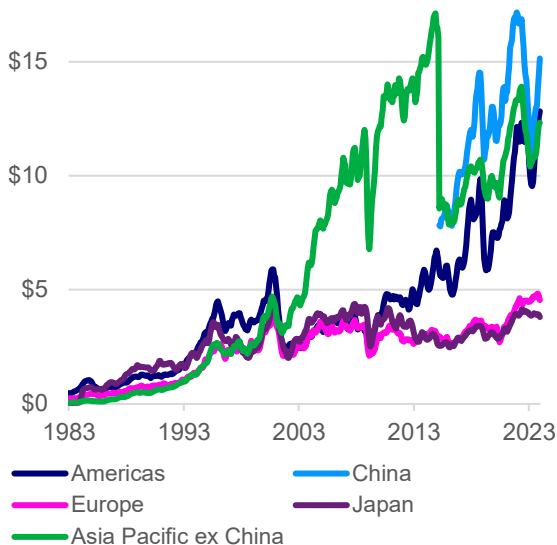
Advanced chip design & manufacturing look well-placed for AI theme

The digital world runs on semiconductors, which are necessary for everything from advanced computing, to navigation systems, to solar power. AI adds to global demand at a time when supply is already strained, in our view, setting a positive outlook for semiconductor-related equities.


Some companies have radically outperformed others—primarily the most advanced producers. Cutting edge computing – including AI – demands cutting edge manufacturing and design. A handful of global semiconductor names have captured the lion's share of recent AI-related gains, reflecting the growing need for high-end, specialized hardware.

Three-month moving averages of semiconductor sales show regional concentration in Asia and North America, where many of the world's leading foundries are based. These locations are also historically considered manufacturing bases for many leading tech companies.


Semiconductor sales
(3-month moving average, USD billions)



Global market share by geography
(by market capitalization of publicly listed equities)



% of total value-added by production segment



Note: "Packaging" includes assembly, packaging, and testing. Mfg = Manufacturing.

Sources: Left: Semiconductor Industry Association Factbook 2023, based on end-2022 data. Right-side, bottom: Strengthening the Global Semiconductor Supply Chain in an Uncertain Era, April 2021 from Boston Consulting Group and Semiconductor Industry Association.

Cloud Computing: Computing at scale is necessary for broad AI use

Training and running models requires at-scale resources



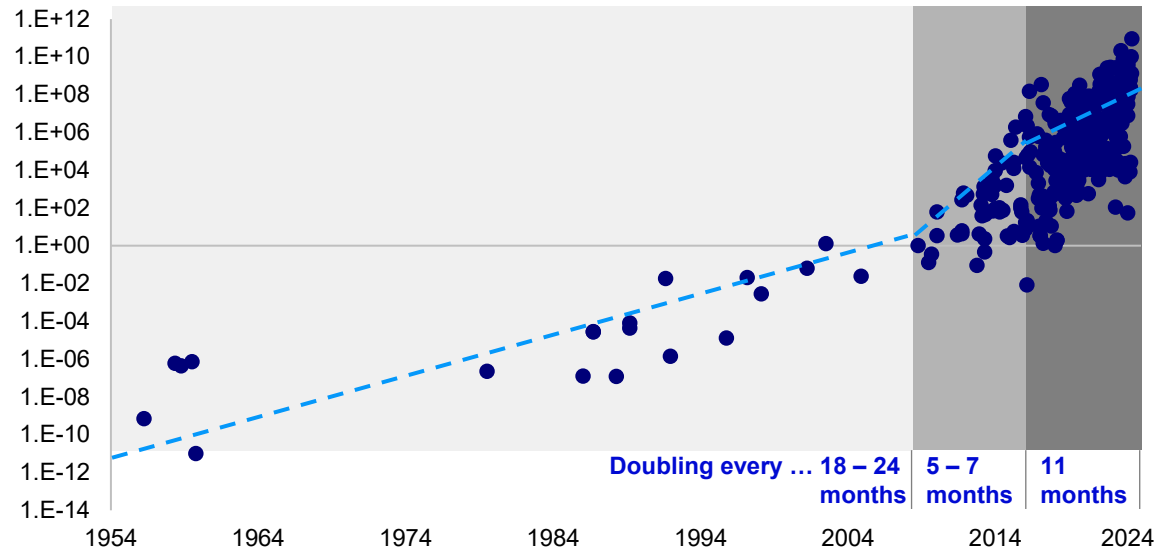
Cutting edge AI models require cutting edge technology. Training the model that underlies OpenAI's GPT-3 on a single enterprise-grade graphics processing unit would have taken **288 years**.¹ AI is powered by sophisticated, high-end, dedicated hardware

Cutting edge technology is almost prohibitively expensive to develop, deploy, and maintain. This means that R&D is primarily in the hands of well-established players with deep expertise and even deeper pockets.

Because of its computational requirements, we believe that AI will largely be accessed via cloud-hosted platforms relaying data between end users and dedicated servers. However, we also believe that AI may push innovation for personal computing and mobile phones, as comparatively lightweight, locally hosted AI models emerge. Indeed, flagship phones and enthusiast-grade consumer GPUs and CPUs already have dedicated AI cores.

Training AI systems has grown exponentially more compute-intensive

Computation used to train AI systems (petaFLOPs²)



Note: 1. Nvidia V100 Tensor Core GPU; Trendlines are illustrative.

2. A FLOP is a floating-point operation, a metric used in describing the size of computing tasks. Peta = one quadrillion, or 1,000,000,000,000,000.

Sources: Our World in Data and arXiv: 2104.0447, Narayanan et al., "Efficient Large-Scale Language Model Training on GPU Clusters Using Megatron-LM" August 23, 2021.



Development Platforms: Toolsets for producing AI software

Platforms help power higher-quality models

Machine learning development involves a stack of technologies that includes computing and data resources, development environments, and cloud services for model training and deployment. Each of these tools plays an important role in speeding up model development.

Companies that provide these services are usually well-resourced, with significant cloud computing resource access and expertise in hardware acceleration and data management.

Every one of today's major AI tools on market makes use of development platforms of some form, representing a critical link in the AI value chain.

Key value-add components of AI development platforms

Data management



AI platforms provide assistive tools for managing large datasets, including cloud-based sources

Data labelling
services

Data library
management

Datacenters
streamlined
for AI
development

Training, evaluation, and fine-tuning



From managing hardware resources to tweaking model specifications, AI platforms provide a range of tools for producing quality AI systems.

Hardware
acceleration tools

Hyperparameter
tuning

Performance
assessment

Distributed
training

Deployment and monitoring solutions



Toolsets for operationalizing AI models include cloud uploads and management, streamlined access to system functions and outputs, and monitoring tools for ongoing performance updates.

Monitoring suite for
resource use

API
tools

Cloud-ready model activation

For illustrative purposes only. API = Application Programming Interface, a simplified toolset for developers to make use of key software functionalities.



AI Models: Specialized or broad-purpose AI tools

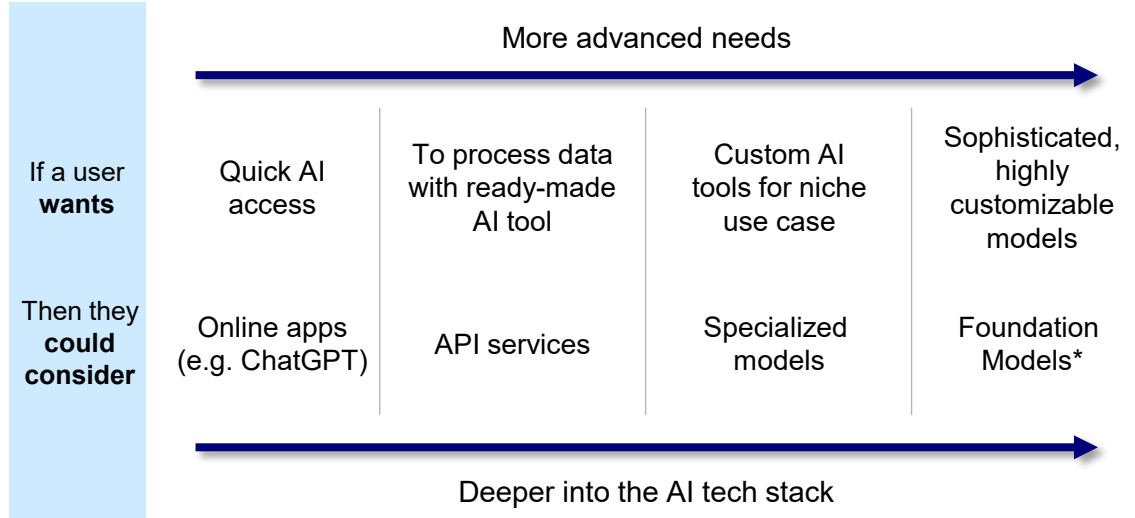
Development expertise is key to building bespoke AI solutions

Artificial intelligence systems can range from task-specific tools to large-scale models capable of being applied to a myriad of functions. In the latter case, these models can be offered as so-called “foundation models”, capable of being applied to as many contexts as you can imagine. Specialization of these tools is typically expected to be carried out by external players.

Being some of the earliest and largest names, large tech players have a history of AI development and currently dominate the market. They may further entrench their position by placing themselves ahead of regulatory changes, suggesting that players in this space may remain fairly narrow.

We expect two themes to emerge in this space: large, established players are likely to maintain their leadership positions, while small, nimble companies with highly specialized refinements of the technology can become trailblazers.

A user perspective on AI offerings



* Proprietary, closed-access models can also be developed for large datasets. For illustrative purposes only.















AI Integration & Adoption: Uses of AI are key to broader value-add effects

The ultimate value-add of AI comes from its implementation in economic use cases. In our view, generative AI tools should power three themes in business, with a broad range of potential applications:

1. Revenue generation through new product offerings and insights.
2. Greater customer focus through more tailored and refined experiences.
3. Enhanced operational productivity by equipping workers with tools to streamline or optimize workflows.

Selected examples of AI implementations

 <p>Multimedia Design</p> <ul style="list-style-type: none"> • Modify content with natural text prompts, auto-fill and extend backgrounds 	 <p>Entertainment</p> <ul style="list-style-type: none"> • Auto-populate scripts, descriptions, directing aids, and visual effects 	 <p>Software Engineering</p> <ul style="list-style-type: none"> • Software development aids that can auto-complete code and assist in designing subroutines
 <p>Construction</p> <ul style="list-style-type: none"> • Auto-generate design iterations and construction plans 	 <p>Healthcare</p> <ul style="list-style-type: none"> • Personalized and rapid healthcare advice and monitoring 	 <p>Financial Services</p> <ul style="list-style-type: none"> • Personalized financial literacy and advice aids, and research securities through natural language
 <p>Education</p> <ul style="list-style-type: none"> • Personalized learning aids 	 <p>Materials Sciences</p> <ul style="list-style-type: none"> • Rapidly create and test novel material combinations for strength, durability, conductivity, etc. 	 <p>Disability Aids</p> <ul style="list-style-type: none"> • Produce on-the-fly interpretations of content for hearing and visually impaired users
 <p>Information Management</p> <ul style="list-style-type: none"> • Natural language search capabilities with ease of refinement 	 <p>Data Visualization</p> <ul style="list-style-type: none"> • Create and modify graphical aids via natural language prompts 	 <p>Gaming</p> <ul style="list-style-type: none"> • Generate content-rich virtual environments, including interactive elements in procedural generation

For illustrative purposes only. For further examples, please see the Appendix: AI Tasks and Use Cases.

Appendix

Appendix: AI Tasks and Use Cases

Multimodal – Combining multiple types of data

Task	Description	Example use case
Document Question Answering	Answering questions about documents	Searching for key information
Feature Extraction	Identifying important characteristics in data sets and/or transforming data into numerical values while preserving the original information	Extracting keywords from a document or the most relevant features of a dataset
Graph Machine Learning	Learning patterns and relationships from graph-based data structures	Network analysis (e.g.: mapping social media relationships; financial crime and transaction monitoring)
Image-to-Text	Converting images to text descriptions	Accessibility aids, like describing onscreen information for the visually impaired
Text-to-Image	Generating images from text	Creating art on demand for multimedia use cases
Text-to-Video	Generating videos from text	Storyboarding or producing simple explanatory animations
Visual Question Answering	Answering questions about images	Interactive exhibits, accessibility aids, searching for scenes in videos, and analyzing large image databases

For illustrative purposes only.

Appendix: AI Tasks and Use Cases (cont.)

Vision and Robotics – Interpreting images and understanding spatial dimensions

Task	Description	Example use case
Depth Estimation	Estimating the distance between objects and the camera	Autonomous driving; robotic hands, including surgical aids; LiDAR assistance; diving, mining, and spelunking aids; 3D images from 2D and volume estimation
Image Classification	Categorizing and classifying images	Identifying animals, plants, and people; generating keywords for image libraries;
Image Segmentation	Dividing an image into distinct sections	Separating foregrounds and background for lens blur and other bokeh effects
Image-to-Image	Transforming and manipulating images	Colorizing images or creating art in other styles; up-sampling and increasing resolution
Object Detection	Locating discrete objects in an image	Detecting objects (e.g.: self-driving or missing items)
Unconditional Image Generation	Generating brand new images, based on a training set but not within the training set, without specifying an initial condition	Creating novel or original artwork; automating product design that is not part of existing intellectual property
Video Classification	Categorizing and classifying videos	Classifying video according to content guidelines or monitoring for illegal imagery
Zero-Shot Image Classification	Novel classification/classifying images without seeing relevant prior examples	Identifying rarely seen or new objects; action recognition
Robotics	Building and programming physical machines/hardware to perform tasks	Training robots to walk, pick up objects, or perform other physical tasks; lights-out manufacturing

Appendix: AI Tasks and Use Cases (cont.)

Natural Language Processing – Processing, analyzing, or translating into human language


Task	Description	Example use case
Conversational	Participating in conversation that feels human-like	Chatbots for customer support or recreation
Fill-Mask	Filling in the blanks in a piece of text	Predicting missing words e.g.: restoring historical texts
Question Answering	Providing a response to a question	Automated assistants
Sentence Similarity	Rating the similarity of sentences based on predefined criteria	Identifying plagiarism or reducing the pool of potential anonymous authors
Summarization	Extracting key information from a larger body of text	Writing abstracts or conducting literature reviews
Table Question Answering	Extracting information queried from tabular data	Finding specific information in a large table of data
Text Classification	Assigning predefined categories to a body of text	Filtering text, including priority emails or spam and scams; sentiment analysis based on affect ratings for words and phrases
Text Generation	Producing new text	Copywriting
Text-to-Text Generation	Transforming a text input into a new text output, including converting text from language to another	Rewriting text for different audiences, translating articles papers, or books into new languages
Token Classification	Labelling individual words in text	Documenting and filing key information by recognizing and storing the data of named-entities like people, places, important dates, activities
Zero-Shot Classification	Classifying text without seeing relevant prior examples for classification	Generating alternative classifications for thematic analysis

Appendix: AI Tasks and Use Cases (cont.)

Audio – Processing or generating audio data

Task	Description	Example use case
Audio Classification	Categorizing and classifying audio data	Identifying instruments or classifying music genres
Audio-to-Audio	Transforming audio data to produce new audio data	Voice modification for actors and singers
Automatic Speech Recognition	Converting spoken languages into text data	Real-time transcription
Text-to-Speech	Converting text into spoken language audio data	Reading text aloud
Voice Activity Detection	Recognizing a human voice in audio data	Voice isolation for phone calls; voice activation

For illustrative purposes only.



In Silico Paper Series on AI

Part I: Understanding AI's “Hello, World!” Moment

Examines AI's explosion of interest in 2023, including a historical overview.

Part II: AI and the quiet revolution of machine learning

Explains the differences between statistics, machine learning, and deep learning, including commercial applications.

Part III: The rise of Generative AI and how it could change our future

Discusses the economic impacts of generative AI and its variants.

Part IV: AI and its impact on labor, productivity, and tech-driven deflation

Explores the macro impact of AI across labor, productivity, and deflation.

Part V: Investing in AI

Reviews the AI value chain and approaches to investing in 2024 and beyond.

Disclosures for Page 16 Graphic

Reprinted by permission. Copyright © 2024 Bank of America Corporation (“BAC”). The use of the above in no way implies that BAC or any of its affiliates endorses the views or interpretation or the use of such information or acts as any endorsement of the use of such information. The information is provided “as is” and none of BAC or any of its affiliates warrants the accuracy or completeness of the information.

Under no circumstances shall BofA Securities or affiliates be liable to you or any third party for any damages (including but not limited to direct, indirect, special and consequential damages), losses, expenses, fees, or other liabilities that directly or indirectly arise from this license, the Report or the Content or your use of the materials. You hereby waive and release BofA Securities and affiliates from any claims for damages, losses, expenses, fees, liabilities, causes of action, judgments and claims arising out of or related to your use of the Report or the Content, whether now existing or arising in the future.

You recognize that information contained in the Content or Report may become outdated and that BofA Securities and affiliates are under no obligation to update the Content or Report or notify you of any changes to the Content or Report. The Report and Content are provided “AS IS,” and none of BofA Securities and affiliates make any warranty (express or implied) with respect to the Report or any content including the Content, including, without limitation, any warranty of ownership, validity, enforceability or non-infringement, the accuracy, timeliness, completeness, adequacy, merchantability, fitness for a particular purpose, or suitability of the material for any intended audience.

Disclosures

Investment risks

The value of investments and any income will fluctuate (this may partly be the result of exchange rate fluctuations) and investors may not get back the full amount invested.

Artificial Intelligence (AI) - Companies engaged in AI typically face intense competition and potentially rapid product obsolescence. These companies are also heavily dependent on intellectual property rights and may be adversely affected by loss or impairment of those rights. Additionally, AI companies typically invest significant amounts of spending on research and development, and there is no guarantee that the products or services produced by these companies will be successful. Companies that are capitalizing on Innovation and developing technologies to displace older technologies or create new markets may not be successful.

Disclosures (cont.)

Important information

This marketing communication is for Professional Clients and Financial Advisers in Continental Europe (as defined below), for Qualified Clients/Sophisticated Investors in Israel; for Professional Clients in Dubai, Ireland, the Isle of Man, Jersey, Guernsey and the UK; for Sophisticated or Professional Investors in Australia; Institutional Investors in the United States; In Canada this document is for use by investors who are (i) Accredited Investors, and (ii) Permitted Clients, as defined under National Instrument 45-106 and National Instrument 31-103, respectively; for Qualified Institutional Investors in Japan; for Professional Investors in Hong Kong, for Institutional Investors and/ or Accredited Investors in Singapore, for certain specific sovereign wealth funds and/or Qualified Domestic Institutional Investors approved by local regulators only in the People's Republic of China, for certain specific Qualified Institutions and/or Sophisticated Investors only in Taiwan, for Qualified Professional Investors in Korea, for certain specific institutional investors in Brunei, for Qualified Institutional Investors and/or certain specific institutional investors in Thailand, for certain specific institutional investors in Indonesia and for qualified buyers in Philippines for informational purposes only. It is not intended to for and should not be distributed to, or relied upon, by the public or retail investors. For the purposes of this document, Continental Europe is defined as Austria, Belgium, Finland, France, Germany, Greece, Italy, Liechtenstein, Luxembourg, Netherlands, Norway, Portugal, Spain, Sweden and Switzerland.

By accepting this document, you consent to communicate with us in English, unless you inform us otherwise.

These materials may contain statements that are not purely historical in nature but are "forward-looking statements." These include, among other things, projections, forecasts, estimates of income, yield or return or future performance targets. These forward-looking statements are based upon certain assumptions, some of which are described herein.

Actual events are difficult to predict and may substantially differ from those assumed. All forward-looking statements included herein are based on information available on the date hereof and Invesco assumes no duty to update any forward-looking statement. Accordingly, there can be no assurance that estimated returns or projections can be realized, that forward-looking statements will materialize or that actual returns or results will not be materially lower than those presented.

All material presented is compiled from sources believed to be reliable and current, but accuracy cannot be guaranteed. This is not to be construed as an offer to buy or sell any financial instruments and should not be relied upon as the sole factor in an investment making decision. As with all investments there are associated inherent risks. This should not be considered a recommendation to purchase any investment product. Investors should consult a financial professional before making any investment decisions if they are uncertain whether an investment is suitable for them. Please obtain and review all financial material carefully before investing. The opinions expressed here are based on current market conditions and are subject to change without notice. These opinions may differ from those of other Invesco investment professionals.

This document may not be reproduced or used for any other purpose, nor be furnished to any other person other than those to whom copies have been sent. Nothing in this document should be considered investment advice or investment marketing as defined in the Regulation of Investment Advice, Investment Marketing and Portfolio Management Law, 1995 ("the Investment Advice Law"). Investors are encouraged to seek competent investment advice from a locally licensed investment advisor prior to making any investment. Neither Invesco Ltd. nor its subsidiaries are licensed under the Investment Advice Law, nor does it carry the insurance as required of a licensee thereunder.

All information as of February 29, 2024, in USD, unless stated otherwise.

Disclosures (cont.)

This document is issued in:

Australia: This document has been prepared only for those persons to whom Invesco has provided it. It should not be relied upon by anyone else. Information contained in this document may not have been prepared or tailored for an Australian audience and does not constitute an offer of a financial product in Australia. You may only reproduce, circulate and use this document (or any part of it) with the consent of Invesco. The information in this document has been prepared without taking into account any investor's investment objectives, financial situation or particular needs. Before acting on the information the investor should consider its appropriateness having regard to their investment objectives, financial situation and needs. You should note that this information: may contain references to dollar amounts which are not Australian dollars; may contain financial information which is not prepared in accordance with Australian law or practices; may not address risks associated with investment in foreign currency denominated investments; and does not address Australian tax issues.

- **Australia** by Invesco Australia Limited (ABN 48 001 693 232), Level 26, 333 Collins Street, Melbourne, Victoria, 3000, Australia which holds an Australian Financial Services Licence number 239916.
- **Austria** and **Germany** by Invesco Asset Management Deutschland GmbH An der Welle 5, 60322 Frankfurt am Main, Germany.
- **Belgium, Finland, France, Italy, Ireland, Luxembourg, Greece, Portugal, Netherlands, Norway, Spain** and **Sweden** by Invesco Management S.A., President Building, 37A Avenue JF Kennedy, L-1855 Luxembourg, regulated by the Commission de Surveillance du Secteur Financier, Luxembourg.
- **Canada** this document is for use by investors who are (i) Accredited Investors, and (ii) Permitted Clients, as defined under National Instrument 45-106 and National Instrument 31-103, respectively. Canada by Invesco Canada Ltd., 120 Bloor Street East, Suite 700, Toronto, Ontario, M4W 1B7.

- **Dubai** by Invesco Asset Management Limited, Index Tower Level 6 - Unit 616, P.O. Box 506599, Al Mustaqbal Street, DIFC, Dubai, United Arab Emirates. Regulated by the Dubai Financial Services Authority.

- **Hong Kong** by Invesco Hong Kong Limited 景順投資 管理有限公司, 45/F, Jardine House, 1 Connaught Place, Central, Hong Kong.

- **Japan** by Invesco Asset Management (Japan) Limited, Roppongi Hills Mori Tower 14F, 6-10-1 Roppongi, Minato-ku, Tokyo 106-6114; Registration Number: The Director-General of Kanto Local Finance Bureau (Kin-sho) 306; Member of the Investment Trusts Association, Japan and the Japan Investment Advisers Association.

- **Singapore** by Invesco Asset Management Singapore Ltd, 9 Raffles Place, #18-01 Republic Plaza, Singapore 048619.

- **Switzerland** and **Liechtenstein** by Invesco Asset Management (Schweiz) AG, Talacker 34, 8001 Zurich, Switzerland.

- **Taiwan** by Invesco Taiwan Limited, 22F, No.1, Songzhi Road, Taipei 11047, Taiwan (0800-045-066). Invesco Taiwan Limited is operated and managed independently.

- **UK, Israel, Jersey, Guernsey** and the **Isle of Man** by Invesco Asset Management Limited, Perpetual Park, Perpetual Park Drive, Henley-on-Thames, Oxfordshire, RG9 1HH, United Kingdom. Authorised and regulated by the Financial Conduct Authority.

Israel: This document may not be reproduced or used for any other purpose, nor be furnished to any other person other than those to whom copies have been sent. Nothing in this document should be considered investment advice or investment marketing as defined in the Regulation of Investment Advice, Investment Marketing and Portfolio Management Law, 1995 ("Investment Advice Law"). Neither Invesco Ltd. nor its subsidiaries are licensed under the Investment Advice Law, nor does it carry the insurance as required of a licensee thereunder.

- **US** by Invesco Advisers, Inc., 1331 Spring Street NW, Suite 2500, Atlanta, GA 30309.