

In Silico, Part V

Investing in AI

In Silico is a multi-part series discussing artificial intelligence, its economic and financial impact, and its role as a driver of change.

- For an introduction to AI, we recommend reading: In Silico, Part I: AI's "Hello, World!" Moment
- For a review of machine learning and deep learning we recommend reading In Silico, Part II: AI and the Quiet Revolution of Machine Learning.
- For a discussion of generative AI, we recommend In Silico, Part III: The Rise of Generative AI and How It Could Change Our Future.
- For an overview of the macroeconomic impacts of AI, we recommend: In Silico, Part IV: Automation, Productivity, and Tech-Driven Deflation



Ashley Oerth, CFA®
Senior Investment
Strategy Analyst
Investment Thought
Leadership



Cyril Birks
Market Strategy Analyst
Investment Thought
Leadership

Introduction

This article is the capstone of the series, focusing on AI investment opportunities, including both investments to gain exposure to the value chain or deploying AI to increase productivity. We will also discuss the recent run-up in AI-related stocks and how we think about AI investments moving forwards. We recommend reading Parts I – IV for further context.

AI excitement reaches markets


Generative AI's sudden 2023 entry into the limelight steps in front of what we called the 'quiet revolution of machine learning'. For well over a decade before the launch of tools like ChatGPT, AI has been deployed in commercial use cases across almost every sector and industry. AI is used for cutting edge statistical analysis, content recommendation, navigation, virtual assistants, song recognition, image enhancement, signal generation, and more.¹ With the release of highly capable generative AI models, businesses have now begun to integrate automated customer service chatbots, coding assistants, and other capabilities. We believe we are only at the beginning of what we may see from AI integration.

Markets appear particularly excited about generative AI. Technology stocks have played an outsized role in supporting S&P 500 returns for much of 2023, especially in March and May of this year (**Figure 1**). As of 30 November 2023, S&P 500 info tech stocks are up 51.6% year-to-date, while the S&P 500 index excluding technology is up just 1.5% over the same period.² Meanwhile, concentration in the index is at its highest levels ever as mega cap tech names dominate market capitalization-weighted indices.³ Despite economic headwinds, tech industry layoffs, and rising interest rates, AI-related tech stocks have been remarkably resilient in 2023.

1. See In Silico, Part II: The Quiet Revolution of Machine Learning
2. Source: Bloomberg. Based on S&P 500 year-to-date price returns to 30 November 2023. S&P500 Ex-Tech measures the broad U.S. market, excluding members of the GICS® Information Technology sector. Past performance is no guarantee of future returns. An investment cannot be made directly in an index.
3. As measured by the Herfindahl-Hirschman Index, a metric of concentration where individual constituent weights (in percent) are squared then summed to generate a measure of concentration. Scores range from near 0 (infinite diversification) to 10,000 (the entire observed universe consists of one firm).

Figure 1: S&P 500 returns have been supported substantially by technology stocks

S&P 500 Price Returns, Indexed to 100 at 31 December 2019




Sources: Bloomberg and Invesco, as of 30 November 2023. Vertical line indicates the 30 November 2022 public release of ChatGPT by OpenAI.

Notes: S&P500 Ex-Tech measures the broad U.S. market, excluding members of the GICS® Information Technology sector. **Past performance is no guarantee of future returns.**

While we think we are in the early innings for generative AI, we also suspect that the recent run-up in technology stocks is due for a pause. Valuations, both on a trailing and forward basis, have climbed markedly and forward earnings expectations for major tech names have been revised upward (**Figures 2 and 3**). Valuations would likely be higher if monetary policy and credit conditions were looser. At this stage, we believe AI-related stocks must deliver on bullish earnings expectations.

Figure 2: Repricing of AI stocks has primarily been through valuations—much of the change has been in future growth expectations.


Breakdown of YTD returns of tech stocks showing contributions from earnings change, valuation change



Sources: Datastream and Invesco, as of 14 November 2023. Chart shows decomposed year-to-date returns of the Datastream World Technology Index. Earnings Growth is measured by 12-month trailing earnings per share. Multiple Change is measured by the change of the price to trailing 12-month earnings ratio.

Figure 3: Tech valuations have risen despite tightening financial conditions

MSCI World trailing price/earnings ratio since 2002



Sources: Bloomberg and Invesco, as of 16 November 2023. Daily data from 2 January 2003 to 16 November 2023.

Notes: The MSCI World Information Technology Index is designed to capture the large and mid cap segments across 23 Developed Markets. The MSCI World Index captures large and mid-cap representation across 23 Developed Markets.

Past performance is no guarantee of future returns. An investment cannot be made directly in an index.

We chose to begin at this time to remove some of the extreme data that took place from 2000 to 2001, when earnings of MSCI World Information Technology stocks fell to near- and below-0 levels. Note: Trailing price/earnings ratios are based on index-adjusted earnings less the effects of extraordinary items.

AI must deliver on the hype

Although technologies are often subject to accelerating improvements, ‘Hofstadter’s Law’ reminds us that change often takes longer than we expect, even when we account for that fact. We suggest that many AI technologies and use-cases are towards the peak of inflated expectations. Indeed, Gartner—the producer of the Gartner hype cycle that describes phases of inflated expectations followed by tangible progress—plotted generative AI at the zenith of inflated expectations. At this stage, it remains unclear whether real-world impacts will be far off the mark from investors’ and adopters’ expectations. The prevailing investment thesis at present, in our view, is one consisting mostly of exciting hypotheticals rather than an evidence-based approach. Therefore, we believe it is important to remember that it is possible to be right about the impact of a new technology and yet invest long-before its effects are realized.


We expect most of the AI-driven excitement is already priced in for the enablers of AI, which we define in the next section. Moving forward, we expect considerably broader market breadth compared to today’s narrow focus on Big Tech players. As such, we think it is worth considering the entire value chain for AI.

Mapping the AI value chain

We believe the long-term value of AI will be realized over the next decade, as we have laid out in previous editions of In Silico. While forecasts are subject to a high degree of uncertainty, we nevertheless see meaningful effects from AI through new products, existing product enhancements, and positive productivity changes. To capture the evolution and adoption of AI, we believe that the full value-chain is worth considering for AI-related investment.

Each step along the value chain can be characterized by the roles that companies occupy as enablers who build or power AI systems, adopters who benefit with new product offerings or increased competitiveness, and responders who evolve product solutions to adapt to the unique challenges of AI. In our assessment, the AI value chain comprises distinct yet interconnected elements, each contributing to the broader landscape of AI innovation and deployment (Figure 4). 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 is entirely possible to have an “AI” themed portfolio which is considerably more diversified than the label may suggest.

Figure 4: The AI value chain⁴



Source: Invesco. For illustrative purposes only.

At present, much of the excitement around AI has been captured by those companies involved in the “enabling infrastructure” in the value chain. These include raw commodity inputs, semiconductor manufacturers and designers, cloud computing and storage, as well as generative AI designers (especially deep-pocketed large tech companies).

4. While we expect there to be greater added value as we move up the value chain, we note that we do not anticipate every step along the way to be an equal increase in value. For example: while cloud computing and storage services are likely to have a big part to play in the AI value chain, we believe that edge-computing and locally- or internally-hosted AI tools will have a significant role to play.

Commodities: Rare earth metals and other critical raw materials

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, TPUs, and other computer hardware typically grouped together as ‘semiconductors’ along with simpler devices).⁵ These circuits are fundamental to any computing operation and have in recent years become the subject of supply chain issues and geopolitical constraints.

A semiconductor material begins its life as a pile of sand made of high purity silicon and a collection of rare earth elements with the right properties for their use in semiconductor devices, from LEDs to photovoltaic cells, to integrated circuits.⁶ 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 (see Appendix A). 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 2022.⁷

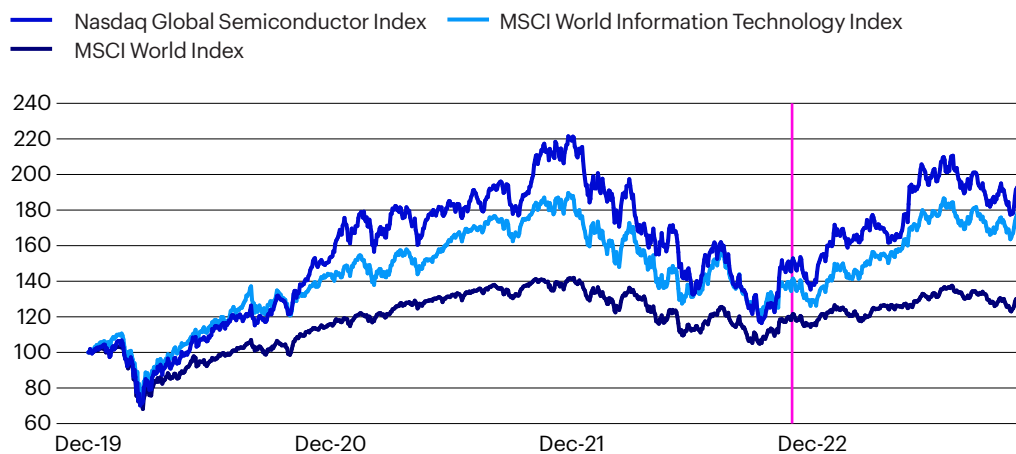
These key commodities are essential inputs to produce semiconductor devices and changes to their prices are likely to impact costs elsewhere in the AI value chain.

Hardware: Semiconductor design and manufacturing

AI systems must be “trained,” a process by which a model’s parameters are learned through a sophisticated series of steps, or algorithms, that comprise the model. To accomplish the training of recent generative AI models, AI-focused companies have built specialized supercomputers capable of training the latest high-parameter models. For perspective, training the model that underlies OpenAI’s GPT-3 on a single enterprise-grade graphics processing unit would have taken 288 years⁷ (assuming it were technically feasible). While there are many means to speed up training times and reduce training burdens, the most advanced AI models have extreme computational demands, driving AI players to invest in high-value, high-end computing hardware. Whether models as powerful as today’s offerings will be within the reach of tomorrow’s consumer-grade hardware is unknown.

After a model has been trained, it still requires additional processing power each time it is used. For instance, a prompt submitted to a generative AI tool goes on to activate the underlying model with the provided inputs, and each word entered and generated uses an amount of computing power. For example, across the ChatGPT platform, daily estimated running costs range from \$100,000-\$700,000. Recent AI demand has driven exceptional demand for processors, helping to push global semiconductor stocks 51.0% higher year-to-date (Figure 5).⁸

Figure 5: Semiconductor stocks have been among the earliest beneficiaries from the AI craze




Source: Bloomberg and Invesco, as of 30 November 2023. Total return indices are indexed to 100 at 31 December 2019. Vertical line indicates the 30 November 2022 public release of ChatGPT by OpenAI.

Notes: The Nasdaq Global Semiconductor Index is designed to measure the performance of the 80 largest semiconductor companies globally. The MSCI World Information Technology Index is designed to capture the large and mid cap segments across 23 Developed Markets. The MSCI World Index captures large and mid-cap representation across 23 Developed Markets. **Past performance is no guarantee of future returns. An investment cannot be made directly into an index.**

5. Deutsche Bank Research ‘China takes a new step in the US-China chip rivalry’ – 4 July 2023. Note: CPU = central processing unit; GPU = graphics processing unit; TPU = tensor processing unit.
6. arXiv: 2104.0447, Narayanan et al., “Efficient Large-Scale Language Model Training on GPU Clusters Using Megatron-LM” 23 August 2021.
7. Not all materials are alike. Characterizing semiconductor materials and devices is a highly specialised domain of material science and beyond our scope.
8. Source: Bloomberg. Global semiconductor stocks are measured by the Nasdaq Global Semiconductor Price Index, which measures the performance of the 80 largest semiconductor companies globally.

Reflecting their growing prevalence in the world economy, semiconductors were the world's second-most traded good by value in 2021, accounting for approximately 3.9% of global trade.⁹ **Figure 6** shows the global trade estimates for semiconductor devices used in computers. Semiconductor manufacturing is an enormous industry, with a range of specialised functions and players.

Figure 6: Value of global trade in computer-related semiconductors



Sources: Macrobond, Invesco, Semiconductor Industry Association and World Semiconductor Trade Statistics, and Observatory of Economic Complexity (OEC), as at 17 November 2023.

Notes: Values for 2023 and 2024 are estimates.

Notes: '% of Global Trade' calculated as a Observatory of Economic Complexity (OEC) Global World Trade / the total value of semiconductors for that year. Semiconductor Industry Association classifies semiconductors as follows: Logic and Micro, devices that perform complex functions, used in mobile phones, AI systems, personal computers, and other information and communication technologies (ICT); Analog, devices that regulate "real world" conditions. Automotives including electric vehicles, aircraft, and renewable energy technology rely on these devices; Memory, devices that store data and are commonly used in mobile phones, personal computers, and data centers; Discrete and sensors and actuators are devices perform simple functions and are heavily used in all industries; Optoelectronics, devices that emit light, are common in information communication technologies and industrial applications.

Hardware generally, and semiconductors especially, represent a fundamental bottleneck for technology of all kinds. AI needs must compete with manufacturing time for personal computers, phones, transport (from cars to aviation navigation systems), and more. Over the last decade, the growth of cloud storage and computing and cryptocurrency mining have also exacerbated demand for semiconductors.

The modern world operates on a cycle of obsolescence. Not only does hardware have a finite operational lifespan, but hardware also has a finite practical lifespan. Over time, operating systems and the software running on them has become increasingly computationally intensive, requiring better and more recent hardware to operate. To the extent our world depends on hardware, it will likely always have minimum supply needs just to maintain the status quo.

Moreover, since the birth of modern computing, the world has tended towards 'doing more with more' and greater computational power has led to greater computational needs. Consider that advances in image computation made image processing cheaper and more efficient while it made better, more computationally intensive image standards normal – digital video became high definition ("HD"), and now 4K is in the process of replacing HD. We expect this trend to continue, which should help drive continuous demand for semiconductors.

Hardware manufacturers (both foundries and fabless designers) will likely continue to be of great significance for as long as the tech sector remains on top. New demands on production, like the rapid adoption of AI, are likely to increase their importance. As a result of these technological trends and the geopolitical constraints on semiconductor trade, we expect the semiconductor market to remain tight.

9. In 2021, the world's most traded goods were crude oil (\$915b), integrated circuits (\$823b), refined petroleum (\$746b), cars (\$723b), and broadcasting equipment (\$473b), according to the Observatory of Economic Complexity – oec.world

Cloud computing and storage providers

Although some AI models can be run on a personal computer, recent AI advancements are primarily a cloud computing phenomenon. In other words, the exceptional needs of training and running AI models are typically done in dedicated computing centers (a service generally described as “cloud computing”), often physically distant from the data scientists who develop these models.

Due to these computational needs, the competitive advantage has so far gone to the tech giants – sometimes referred to as ‘hyperscalers’ – that regularly invest in deploying and updating data centers, servers, and cloud computing services. In fact, the most well-known AI start-ups like OpenAI and Anthropic are backed by Microsoft and Amazon respectively. The other two major generative AI language models are Meta and Alphabet projects, which themselves have existing access to large-scale cloud computing capacity.

As with semiconductors, data centers and cloud computing solutions are not just for AI. As the internet grows, the general need for such services has grown, too. For example, more media content than ever exists online, with music streaming revenue having overtaken physical media in 2018. As the demand for AI access increases, we expect it will compete with growing demands for other services from streaming to general web hosting needs.

Architects and Enablers: AI models and platforms

Generally speaking, AI is composed of algorithms designed by data scientists and software engineers. These algorithms comprise models which are trained on huge amounts of data that require considerable processing power. Relatively simple models can be run on a consumer-grade computing platform, while sophisticated AI models require specialized hardware with far greater performance capabilities and very large datasets. The labor, capital, and operating requirements to build, train, and maintain high-performing models mean that sophisticated AI models are the realm of just a handful of players.

Since early 2023, there has been an explosion in the number of publicly available generative AI models for personal and commercial use. The landscape includes major players and tools, like OpenAI’s ChatGPT and DALL-E, Meta’s LLaMa, Google’s Bard, and Anthropic, as well as community-backed projects like Hugging Face. These players are united by their access to world class data scientists and software engineers and their ability to access large-scale computing resources.

Being some of the earliest and largest names in this emerging industry, 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.¹⁰ Furthermore, these large players have gone far further than most to demonstrate product capabilities and product-market fit. Indeed, the future of widely adopted, commercial AI may depend on AI training platforms and rentable hardware/server time that only the incumbent tech giants can supply. Even if new players become the dominant force in AI, it seems highly likely to us that the biggest tech players will still get a slice of the pie because they are responsible for the infrastructure on which much of the Internet runs.

Given the considerable barriers to entry, we believe that AI models are likely to be packaged for broad purposes, such as question answering, autocomplete suggestions, image generation and auto-fill, and other functions. Specialization on these tools will likely be conducted and refined by external players, perhaps led by internal expertise of the adopters of these models. Indeed, Big Tech names are developing so-called “foundation models” for later specialized refinement and use and thus helping secure their place in the AI value chain.

10. Microsoft, Meta, Google, and Anthropic have formed an industry body known as the Frontier Model Forum, which seeks to provide oversight and governance on ‘frontier models’ – those multi-modal and general-purpose AI models currently in development which will be more powerful than today’s offerings. This forum appears to be working proactively in conjunction with governments and regulators.

Early adopters and long-term beneficiaries

We anticipate that AI will see gradual and widespread adoption for the rest of this decade. In Appendix B, we have provided an overview of some job tasks we think are most likely to be affected, with examples of how they may be positively impacted.

Beyond AI companies, early adopters are adapting their business models or repositioning their offerings in response to widespread AI interest. Early examples include research and writing aids, search engine integration, software engineering support tools, web navigation assistants, and more. Desktop and mobile operating system designers are also integrating generative AI capabilities into consumer-grade use cases, including video and image editing, navigating interfaces, and integration with voice assistants. Research use cases are also being explored, including applications in biopharmaceuticals where models are applied to protein synthesis, allowing researchers to quickly propose and examine millions of potentially viable novel proteins and synthetic equivalents.

In our view, generative AI tools should help power three themes in business: revenue generation through new product offerings and insights, greater customer focus through more tailored and refined customer experiences and enhanced operational productivity by equipping workers with tools that can streamline or optimize existing workflows.

Moving forward, we expect that such use cases will continue to be explored and developed. Perhaps most importantly, we anticipate that adoption will grow as AI tools become cheaper, more readily scalable, and more familiar to end-users. In our view, the greatest value-add from AI is likely to be found in those adopters that can effectively integrate AI into today's economy.

Conclusion

The very cutting edge of AI research is defined by those large players who can deploy sufficient capital to do more with more, scrappy innovators who have worked out how to do more with less (often through more focused product offerings), and those adventurous few who seek out new niches or use-cases altogether typically as part of collaborations between academia and industry.

As AI continues to develop, we expect its revenue impact to be realized more broadly. Use cases and capabilities must still be identified and rolled out. Until then, most of the AI excitement is concentrated in the architects and enablers of AI, with few use cases meaningfully deployed and adopted. In other words, we believe the economy must catch up to these latest innovations and help bring the hype cycle from expectations to reality and, through this process, enable sustainable revenue growth.

Appendices

Table 1: Curated list of USGS Critical Minerals¹¹

Critical mineral	Applications	Lead Producer	% World Total
Aluminium	Metallurgy and many sectors of the economy.	Australia	26%
Arsenic	Semiconductors.	Peru	46%
Beryllium	Semiconductors, aerospace, and defence.	United States	64%
Bismuth	Semiconductors, medical, metallurgy, and atomic research.	China	80%
Cadmium	Semiconductors.	China**	40%
Cobalt	Batteries and metallurgy.	Congo	68%
Dysprosium	Data storage devices, lasers, and permanent magnets.		
Erbium	Fiber optics, glass colorant, lasers, and optical amplifiers.		
Gallium	Integrated circuits and optical devices.	China	98%
Germanium	Defence and fibre optics.	China**	93%
Graphite	Batteries, fuel cells, and lubricants.	China	65%
Indium	Semiconductors and LCDs.	China	59%
Lanthanum	Batteries, catalysts, ceramics, glass, and metallurgy.		
Lithium	Batteries.	Australia	47%
Manganese	Batteries and metallurgy.	South Africa	36%
Nickel	Batteries and metallurgy.	Indonesia	48%
Platinum Group Metals*	Catalytic converters, catalysts, integrated circuits, and other electrical components.	N/A	N/A
Palladium		Russia	42%
Platinum		South Africa	74%
Rhodium		South Africa**	85%
Rhenium	Semiconductors, chemical catalysts, high-temperature components.	Chile**	44%
Silicon	Semiconductors.	China	70%
Tantalum	Capacitors and metallurgy.	Congo	43%
Tellurium	Metallurgy, solar cells, and thermoelectric devices.	China	53%
Terbium	Fibre optics, lasers, permanent magnets, and solid-state devices.		
Vanadium	Batteries, catalysts, and metallurgy.	China	70%

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

Note: % of World Total refers to the production of the country displayed in the adjacent column.

*Platinum Group Metals include iridium, osmium, palladium, platinum, rhodium, and ruthenium
 **Figures from 2021 World Mining Data. These figures can sometimes conflict significantly with 2023 USGS figures.

Appendix B: AI Use Cases and Sector or Industry Impact

As described throughout this series, AI refers to many different technologies and use cases. In our view, discussing the anticipated impact of AI across industries is best done using a ‘job tasks’ approach, the same approach we took to understand the impact of AI on the labor market. See **Table 2** for a non-exhaustive list of job tasks AI can perform.

In broad strokes, we expect that all industries and sectors engaged in ‘Basic’ and ‘Translational’ science will benefit from the advanced data modelling and imaging capabilities of machine learning and deep learning AI technologies. Applied sciences may benefit from these capabilities as powerful AI tools become more efficient to use in a wide number of live settings, including continual monitoring and adaptation use cases. Collectively, these use cases are extensions and compliments to human capabilities, not freeing up labor per se but broadening its scope and practice areas.

On the other hand, generative AI (and, separately, emerging multimodal robotics, like self-driving cars, robotic arms, drones, etc.) are more direct competitors to human labor, removing a human from the production or services loop altogether. Complex, delicate, and repetitive procedures, from surgery to manufacturing, will all be automated in part or entirety, across industries and sectors. Content generation, media, and other digital services may also be increasingly aided or outsourced to AI, notwithstanding labor agreements.

Table 2

AI Tasks ¹²	Description	Example Use Case
Multimodal	Combining multiple types of data (e.g.: text, images, and audio)	
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
Computer Vision	Interpreting images and "seeing"	
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

12. list from: Hugging Face.

Table 2

Natural Language Processing	Processing and analyzing human languages	
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
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
Translation	Converting text from language to another	Translating articles papers, or books into new languages
Zero-Shot Classification	Classifying text without seeing relevant prior examples for classification	Generating alternative classifications for thematic analysis
Audio	Processing or generating audio data	
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
Tabular	Using data structured as a table	
Tabular Classification	Assigning predefined categories to data in a table	Customer categorization
Tabular Regression	Predicting numerical values based on data in a table	Predictive analytics and forecasting e.g.: customer turnover, revenue, expected returns
Reinforcement Learning	Learning through trial and error using a penalty and reward system	
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

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.

Important Information

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.

Issued in 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.

© Copyright of this document is owned by Invesco. You may only reproduce, circulate and use this document (or any part of it) with the consent of Invesco.