C WORLDWIDE
ASSET MANAGEMENT

WHITE PAPER

GENERATIVE AI: THE ARTIFICIAL FUTURE ON STEROIDS

Generative AI: The Artificial Future on Steroids

Key Insights

- *With the move towards natural "language as a code", we are at a time in history where there's a fundamental redefinition of the human-machine interface driven by Large Language Models (LLMs).*
- *LLMs are a subset of Generative Artificial Intelligence (GenAI) models, primarily used for applications involving text. Other forms of GenAI extend these capabilities to non-textual domains like video and music.*
- *The training process involves feeding a vast amount of data into a model. Once an LLM is trained, it can be used for inference. Inference uses a trained model to make predictions based on new input data.*
- *A Mixture of Experts (MoE) model draws inspiration from the human brain's organisation. This approach employs multiple sub-models, known as "experts," each dedicated to a specific type of input or aspect and suggests a trend towards more computationally efficient models.*
- *The development of MoEs, particularly for edge devices like smartphones, mixed-reality headsets, robots, and automobiles, is an emerging area.*
- *Solely offering a basic GenAI model will not lead to structural profitability. This generic technology will be impossible for a handful of companies to dominate. Applications must be offered on top to create a profitable moat.*
- *The introduction of app-free smartphone concepts powered by an AI assistant hints at a potential overhaul of the mobile economy and challenges to current app store models.*
- *The "intelligence" of these models comes at a cost – namely, a massive increase in compute requirements. This surge in compute demand has led to significant investments in AI infrastructure, with cloud services and power management becoming increasingly crucial.*
- *While Nvidia has a firm grip on the accelerator market today, the proliferation of different compute architectures promises much more competition over the mid to longer term.*
- *Applications are becoming available to companies to improve productivity and will soon impact productivity and job market metrics.*
- *An underappreciated area is GenAI's impact on blue-collar tangible sectors. GenAI is accelerating the development of cobots and humanoid robots. As deep automation across multiple sectors takes hold, the deflationary impulse will be significant – counter-balancing the inflationary effects of shrinking labour markets.*
- *We continue to see attractive investment opportunities amongst the 'picks and shovels' of the AI revolution, proprietary data owners, and the application layer.*

Morten Springborg
Global Thematic Specialist
C WorldWide Asset Management

Introduction

The evolution of computing over the decades can be broadly characterised by several key phases, each marked by significant technological breakthroughs and paradigm shifts in how computers have been used and accessed. As new computing cycles emerge each of these new cycles has followed a similar investment pattern, starting with semiconductors, then shifting to infrastructure, and finally to software and services. As an example, in the mobile era, the initial excitement was centered around Qualcomm and ARM, owning the chip designs, then moving on to the infrastructure and devices in the form of Samsung and Apple, and eventually peaking out with a focus on software and services, creating enormous value in companies like Google, Facebook, Uber, etc. In the GenAI era, we are currently in the first phase of excitement around semiconductors. Still, as commoditisation begins, we would expect the usual pattern of value shifting towards infrastructure, apps, and services as GenAI technology matures over the coming years. It is this framework we apply when thinking about GenAI.

The early days of Neural Networks

Until recently, artificial Intelligence was primarily about neural networks and deep learning. These networks, composed of layers of interconnected nodes (mimicking neurons in the brain), could learn and make sense of complex patterns in vast amounts of data. The introduction of deep learning marked a significant leap, enabling AI to tackle more sophisticated tasks like image and speech recognition with unprecedented accuracy. We published a White Paper on this topic back in 2016 titled 'The Artificial Future'.

The rise of Generative AI

In November 2022, OpenAI released ChatGPT. Initially, this was not a particularly noteworthy event – the system's capability aligned with the work done at other AI labs, such as Google DeepMind. What made ChatGPT special was finding a way to put the technology in the hands of users everywhere. Through OpenAI's partnership with Microsoft, the company was able to scale technology aggressively, and ChatGPT set a remarkable record by reaching 100 million monthly active users just two months after its launch, making it the fastest-growing consumer application in history. This enabled Microsoft to – at least initially – leapfrog Google's early leadership in AI and to position the company as the clear early leader in GenAI applications.

The foundational building blocks of GenAI

Unlike their predecessors, who were adept at analysing data and making predictions, GenAI models can create new content – text, images, music, or video that is often indistinguishable from human-generated content. This shift from understanding to creating was made possible by advances in neural network architectures, such as Generative Adversarial Networks (GANs) and Transformers, originally developed by Google. GANs brought a novel approach where two neural networks, the generator, and the discriminator, compete against each other to generate new, synthetic data that are incredibly realistic. On the other hand, transformers revolutionised how we process sequences, such as sentences in a paragraph, by capturing the context of each element about all others. This breakthrough paved the way for LLMs.

LLMs and beyond

LLMs are a subset of GenAI models and are primarily used for applications involving text, such as content creation, conversation agents, summarisation, translation, and more. Other forms of GenAI extend these capabilities to non-textual domains. For instance, GANs are used to create realistic images or videos, and there are models designed to generate music or synthetic voices.

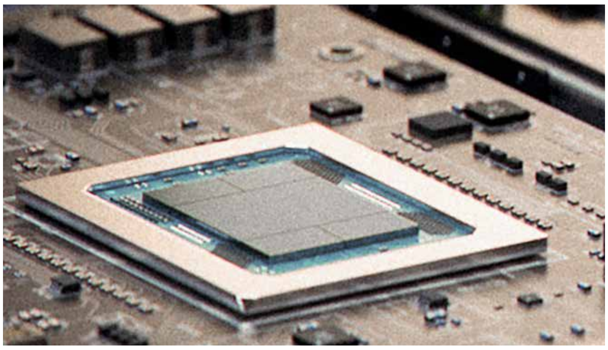

An LLM like OpenAI's GPT (Generative Pretrained Transformer) series is a type of artificial intelligence model designed to understand, generate, and interact with human language. These models are "large" in the sense that they consist of billions or even trillions of parameters, which are the parts of the model that are learned from data during the training process.

The training process involves feeding a vast amount of data into the model, basically up to all the available data on the internet, allowing it to learn patterns, relationships between words, and the structure of sentences in various languages. The objective is to adjust the model's internal parameters (weights) to minimise the difference between the predicted output and the actual output. The model learns to predict the next word in a sentence, given the words that come before it, which enables it to generate coherent and contextually relevant text sequences.

Once an LLM is trained, it can be used for inference. Inference means using a trained model to make predictions based on new input data. While less computationally intensive than training, compute demand for inference will grow massively because of the sheer number of prompts from millions or billions of users, daily.

Soaring computer requirements will put costs in focus

The "intelligence" of these models comes with a cost – a massive increase in compute requirements. Training state-of-the-art AI models require vast amounts of computational power, often involving thousands of powerful GPUs (Graphics Processing Units) working in tandem for weeks. This surge in compute demand has led to significant investments in AI infrastructure, with cloud services and specialised hardware becoming increasingly crucial. Increasingly, as the workloads transition to inference, the cost of computing will become a hugely important competitive parameter, which will drive innovation both in model design as well as the design of the chips.

"As the workloads transition to inference, the cost of computing will become a hugely important competitive parameter."

Open models and MoEs indicate more efficient models in the future

Meta, a year ago considered to be lagging in the GenAI race, has adopted an open-source approach to LLM development and has been able to reposition itself successfully. Being "open source" means the software's source code is freely available for anyone to view, modify, and distribute. Thereby companies like Meta have been able to leverage the global developer community in developing the models. This contrasts with closed or private models like GPT, where the source code is proprietary and controlled by its creators, limiting access and modification by others. Open-source models promote collaboration and transparency, whereas closed models often focus on commercial benefits, restricting access to improve or adapt the software without permission from its copyright holder. Open-source projects can accelerate innovation and accessibility in technology, providing a foundation for further advancements by the global community.

In a challenge to Meta's early mover advantage in open-sourced LLM, Google recently released new AI models for use by outside developers for free. The models, named Gemma, are based on the same technology as Google's proprietary Gemini LLMs. Gemma could benefit Google if other software developers make improvements to the open-source model, which Google could then incorporate into its proprietary model, Gemini.

French start-up Mistral is seen as an early leader in a new development within language models, namely the Mixture of Experts (MoE) models. Microsoft has recently teamed up with Mistral and offers the models on the Azure platform along with OpenAI models. MoEs draw inspiration from the human brain's organisation, where different regions specialise in distinct tasks. This approach employs multiple sub-models, known as "experts," each dedicated to a specific type of input or aspect. These experts operate in parallel, with their outputs integrated through a gating network based on the input data's relevance. This method aims for increased efficiency, lower computational demands, and improved performance on certain tasks by utilising smaller, specialised models. It is speculated that GPT-4 may secretly comprise eight GPT-3 style models, each specialising in a particular task. The architecture's design, mirroring the brain's specialisation, suggests a trend towards more manageable and computationally efficient models, potentially moving away from the larger, more unwieldy models seen in previous AI generations.

Smaller models will proliferate at the edge of the network

The development of MoE models and their application in GenAI, particularly for edge devices like smartphones, mixed-reality headsets, robots, and automobiles, is an emerging area since edge devices are often constrained by power, memory, and processing capabilities.

The intersection between MoE developments and GenAI edge on-edge devices signifies a shift towards more distributed, efficient, and specialised AI systems. These advancements could lead to higher-quality models that require fewer resources for training and deploying, aligning with the ongoing trend towards bringing powerful AI capabilities closer to where data is generated and consumed, at the edge of the network, while at the same time reducing the computing need at the data center.

> *"The intersection between MoE developments and GenAI edge on-edge devices signifies a shift towards more distributed, efficient, and specialised AI systems."*

This suggests a future where no single GenAI model will prevail outright. Instead, a multi-model world is emerging. Smaller, more accessible, and more compute-efficient models or MoE-style architectures are seen as increasingly practical and desirable for specific tasks, while LLMs like GPT-4 maintain their utility for broad, general-purpose applications.

On-device AI will transform the smartphone market

Google recently released Gemini Nano, a small language model designed for on-device generative AI tasks, which will debut on the Pixel 8 Pro and soon expand to more Android devices. This development, like the recently released Samsung Galaxy S24, is an example of GenAI

"Inference will be a battle of who can do it cheaper, and in many instances, smaller models on the device win simply because of response time and the low cost of computing."

moving towards on-device processing. This approach not only increases hardware needs as devices are upgraded but is also a precursor for a potentially significant market shift, with AI-capable smartphones growing significantly over the coming years. The device's ability to perform AI tasks directly on the smartphone without relying on cloud computing not only improves privacy and data security but also enables real-time processing for applications like language translation, image recognition, and personalised recommendations.

Historically, AI enhancements from companies like Qualcomm, MediaTek, and Google have improved device functionalities like imaging and battery life. However, integrating small language models requires advanced computing platforms and software innovations. Qualcomm's collaboration with Deutsche Telekom on an app-free smartphone concept, powered by an AI assistant, hints at a potential overhaul of the mobile economy and challenges to current app store models.

"This change could offer a direct challenge to the current dominance of Apple and Google's app stores."

The Qualcomm CEO highlighted in a recent Financial Times article the disruptive potential of on-device generative AI, suggesting a shift away from app store dependencies towards pre-installed GenAI models. This change could offer a direct challenge to the current dominance of Apple and Google's app stores. It signifies a major shift in computing, access, and distribution models. It implies, at a minimum, that Apple/Google is going to face more competition from Meta and Microsoft as the latter team up with OEMs to get their software in the 'device engines'.

Inference will be a battle of who can do it cheaper, and in many instances, smaller models on the device win simply because of response time and the low cost of computing.

It's still very early in the roll-out of language models. Still, events over the last year suggest that the competitive landscape will be tough and that only offering the basic model will not lead to structural profitability. This generic technology will be impossible for a handful of companies to dominate. And applications must be offered on top to create a profitable moat.

The chips doing the AI computations

The development of parallel computing has been crucial for General AI due to its ability to perform many operations simultaneously, significantly accelerating AI model training and inference. GPUs offer massive parallelism, efficiency, and speed, allowing for the handling of complex computations more effectively than traditional CPUs.

Today's massive shift towards parallel computing is primarily driven by Nvidia and its family of GPUs. Across the global AI data centers, it is estimated that Nvidia GPUs are responsible for 80% of the workloads today.

GPUs are highly effective at handling parallel tasks, making them ideal for the matrix and vector computations common in machine learning and deep learning. However, TPUs (Tensor Processing Units) developed by Google are optimised to process the tensors (multi-dimensional arrays) that are fundamental to neural network computations and offer even higher efficiency for certain AI workloads compared to GPUs. Amazon AWS Tranium and Inferium chips are ASIC (Application-Specific Integrated Circuit) and custom-designed for a specific function or set of functions – in this case, optimised for machine learning training and inference tasks. Once manufactured, an ASIC's functionality is fixed and cannot be altered.

Field Programmable Gate Arrays (FPGAs), on the other hand, are integrated circuits that can be reprogrammed after manufacturing to perform a wide variety of tasks. They offer flexibility because their hardware configuration can be updated to support different functions or algorithms after deployment. AMD and Intel have, through acquisition, built positions in FPGA chips.

The choice between the different architectures depends on the specific requirements of the machine learning tasks, including the models being used, the preferred machine learning framework, budget constraints, and whether the deployment is local, or cloud-based. TPUs and Traniums/Inferiums are developed specifically by Google and AWS and are highly optimised for their respective cloud environments, and not sold to third parties. GPUs, on the other hand, offer greater versatility and support a wider range of machine learning tasks and frameworks, making them suitable for a broader range of applications.

While Nvidia, with its family of GPUs backed by its software architecture CUDA, has a firm grip on the accelerator market today, the proliferation of different compute architectures promises much more competition over the mid to longer term.

The total semiconductor spending is expected to be around US$600 billion in 2024. Of this, AI chips will be around

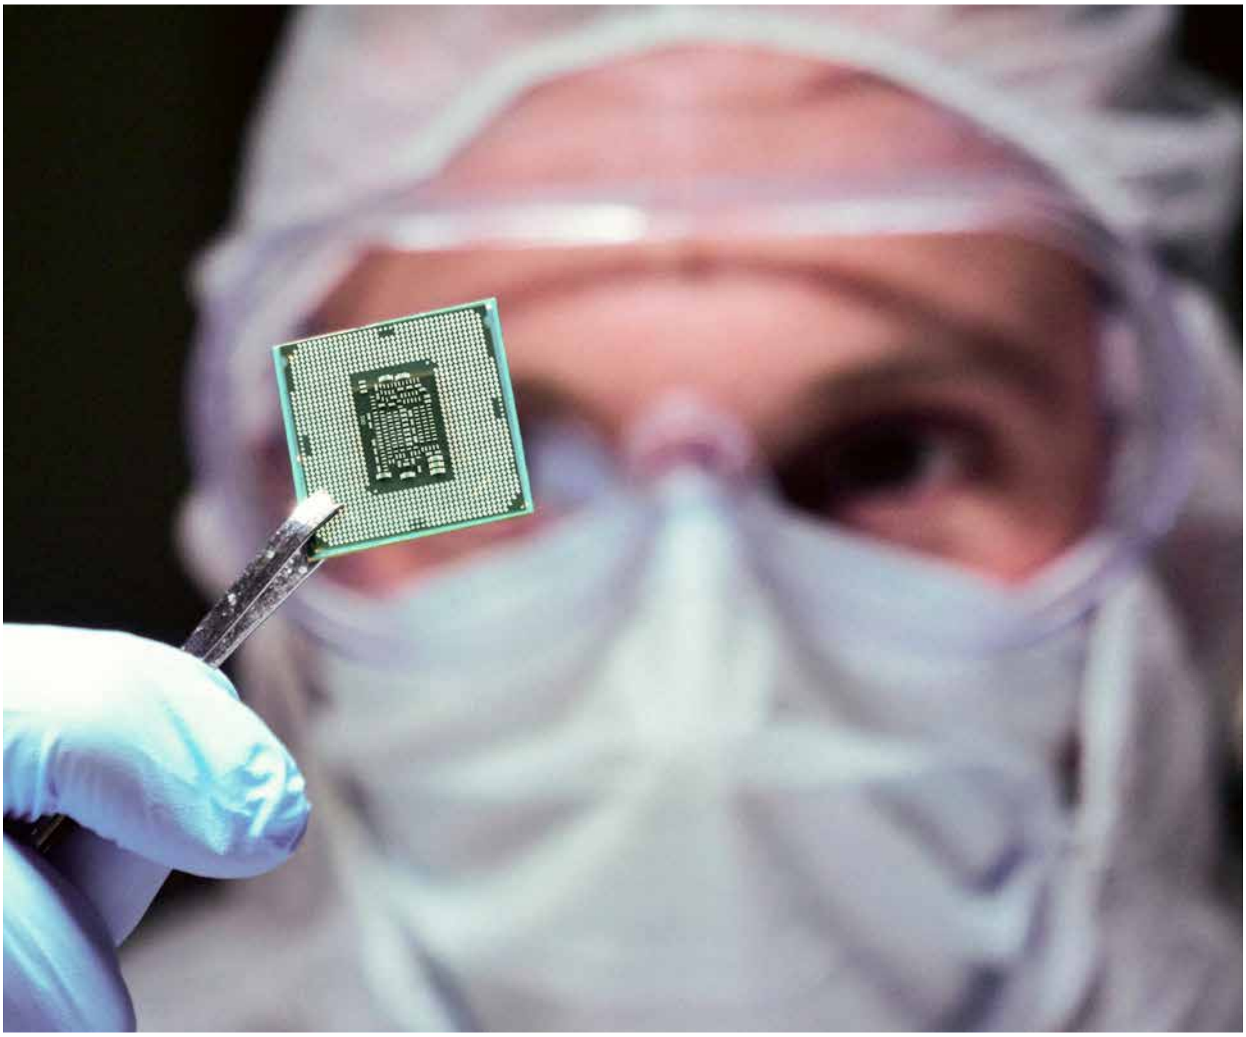

US$45 billion. AMD CEO Lisa Su recently forecasted that the accelerator market will grow to around US$400 billion by 2027, a growth rate of more than 70% p.a. in AI chips. Jensen Huang, founder and CEO of Nvidia, has estimated that accelerated computing will double the total investments in data centers from US$1 trillion in installed DC infrastructure to US$2 trillion over the next 5 years. The bottom line is that growth will be extremely strong, but so will competition as demand shifts from training to inference, increasingly at the edge, questioning the sustainability of earnings growth because of the extreme over-earing that is taking place today. As an example, in CY2019, Nvidia had a gross margin of 62.5%, while in the most recent quarter, the company reported a gross margin of 76%. We prefer to focus on the sustainability of earnings growth and invest in the growth indirectly via a picks-and-shovels approach, buying the foundries and equipment providers for semi-manufacturing, thereby being technology agnostic while benefitting from the overall massive growth in demand for AI Chips.

Deployment of GenAI has a significant effect on power demand

Data centers have immense power and cooling requirements to support the operation of high-performance computing systems for GenAI. Today's expansion of data center capacity is constrained not only by the availability of GPUs but also by energy availability and the need for significant investments in infrastructure and power management systems to meet the growing computational demands of AI models. Refurbishing of legacy CPU data centers to GenAI is complicated by the fact that the energy density of GPU racks can be a factor of 6-8 compared to CPU racks, and because the training of LLMs can run parallel on thousands of GPUs for weeks, consuming vast amounts of energy.

> *"Today's expansion of data center capacity is constrained not only by the availability of GPUs but also by energy availability."*

Morgan Stanley expects power demand from GenAI to increase from 0.2% of total US electricity consumption in 2023 to 3.6% in 2028, an 18-fold increase. The largest GPU data centers are today quoted at up to 300 MW, and the industry is preparing to source energy from dedicated nuclear power plants, exemplified by for example Microsoft, which in late 2023 started looking at next-generation nuclear reactors to power its data centers, according to job listings.

Networking within data centers also presents a bottleneck for scaling GenAI. The complexity and cost of networking infrastructure to connect 1000s of GPUs for distributed AI training tasks are significant. High-speed, low-latency networks are essential for efficient training of large AI models, and the engineering challenges in scaling this infrastructure impact the overall performance and scalability of GenAI systems.

The shift towards the inference phase which, while generally less computationally intensive than training, is expected to accelerate demand for computing substantially as AI applications become more widespread. In its latest Q4 financial results released in February 2024, Nvidia estimated that 40% of workloads were related to inference, much higher than what was estimated by industry experts, mitigating fears that Nvidia GPUs would lose relevance as workloads shifted away from training LLMs. The shift towards inference has implications for the design and optimisation of AI infrastructure, with a potential increase in the deployment of specialised hardware like the internally developed chipsets by Google and AWS for inference tasks. In the long term, the move toward GenAI Edge distributed models will accelerate the development of even more computing power in edge devices like smartphones and lead to new product cycles. And eventually, a completely new human interface will be created for computing language.

Looking ahead, applications will drive mass adoption of GenAI

We are at a time in history, where there's a fundamental redefinition of the human-machine interface more significant than the mainframe to PC or PC to smartphone shifts. One can view it as a huge democratisation of technology in that the barrier to access is being lowered tremendously as we move towards natural "language as a code". It will be

“We are at a time in history, where there’s a fundamental redefinition of the human-machine interface more significant than the mainframe to PC or PC to smartphone shifts.”

by far the most profound shift ever in the engagement between humans and machines if the AI assistants do achieve ‘human level’ intelligence and engage with us daily in natural language conversation. We are not there yet, but as Sam Altman, founder of OpenAI, said in a conversation with Bill Gates recently, “these current LLMs are the stupidest the models will ever be....this is the fastest technology revolution by far”, and we should expect these models to become extremely powerful and effective over the coming few years. Mind you, we are only 18 months into the GenAI age.

We are very early in the integration of AI into workflows. However, straightforward applications are becoming available to companies to improve productivity. At the forefront is Microsoft Copilot. It will not be long before the AI application layer starts to impact productivity and the job market metrics, from overall wage growth to hiring dynamics. That evidence will begin in early adopter sectors, notably tech itself – where demand for software engineers is beginning to be affected as “language as a code” tools get integrated. Software coding assistants are clear early examples of the power of GenAI given their tangible productivity gains (e.g., 50% improvement in developer code writing velocity, according to Microsoft) and large return on investment (GitHub Copilot Enterprise costs only $19 per month per user).

Another example is interactive chatbots, which empower greater personalisation and allow for real-time customer interactions. Swedish ‘buy now, pay later’ company Klarna recently said its AI assistant is doing the work of 700 full-time agents and that its chatbot has had two-thirds – or 2.3 million customer service chats in 35 languages – of Klarna’s customer service chats in a month.

GenAI is a generic technology, and anyone can, with the help of IT consultants and cloud service providers, create a ‘GPT‘bespoke AI model, assuming they have the data. Therefore, one should assume there will be a wide variety of these corporate ‘AI apps’ to automate repetitive business functions. Practically, it will take time as corporations need to develop data policies and have “clean” and consolidated data of high quality from different systems, and that takes time and effort.

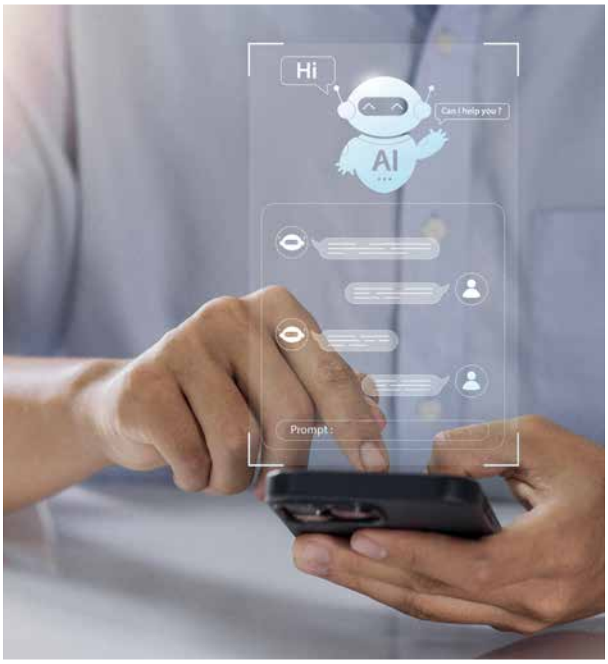

“These current LLMs are the stupidest the models will ever be....this is the fastest technology revolution by far.”

On the analytics front, GenAI technology enables a broader range of people to use natural language queries to search and ask questions about enterprise data. Beyond these applications, we see GenAI systems targeting manual writing tasks such as emails, outlines, summaries, and presentations – with the promise of saving workers a significant amount of time.

From a vertical industry perspective, we see the low-hanging fruit for GenAI coming from industries (e.g., healthcare, legal, and media) that produce large quantities of mostly unstructured data that are difficult/time-consuming to process or summarise manually.

"The irony of the above focus on AI's impact on "white collar" and intangible knowledge sectors is that the biggest impact longer term could be in the historically most susceptible to automation – blue-collar tangible sectors."

The rise of the robots

The irony of the above focus on AI's impact on "white collar" and intangible knowledge sectors is that the biggest impact longer term could be in the historically most susceptible to automation – blue-collar tangible sectors. This is because the rise of GenAI will have a massive impact on the development of cobots and humanoid robots.

GenAI is propelling advancements in robotics and we are seeing significant contributions from companies like META, Tesla, Google, Nvidia, and Boston Dynamics. Advancements in 3D mapping, large foundation models such as LLMs and Visual Language Models (VLM), where the robot learns by observation, and cheaper hardware sensors and actuators are lowering the investment thresholds for robotics across various sectors, from healthcare to hospitality and manufacturing. This evolution is poised to cost-effectively replace a significant percentage of low-wage/skill jobs, potentially signaling a shift from cost optimisation via offshoring to "botshoring", reducing costs by deploying (humanoid) robots.

As an example, META has developed the AI simulation environment, Habitat, a platform that is teaching robots to interact with the physical world, enhancing AI's ability to understand and navigate complex environments. Tesla has made significant strides with its humanoid robot, Optimus, which according to the company, is close to mass deployment, first in Tesla's factories, later intended to be sold to 3rd parties. Nvidia and Google DeepMind are also deeply involved in the development of robot systems, with a focus on allowing robots to make decisions faster, understand their environments better, and navigate more efficiently. These systems combine foundation models with robot control models, enabling robots to perform diverse tasks in novel environments.

In summary, these companies are at the forefront of integrating GenAI with robotics, contributing unique technologies that advance the capabilities of robots, cobots, and humanoids. AI's role in enhancing robot autonomy, decision-making, and physical interaction is becoming increasingly sophisticated and potentially opens new end-markets for the respective companies in industrial automation in the years to come.

The social aspects of AI

While a deeper discussion of the social aspects of AI is not the purpose of this paper, the technological shift will probably be a challenge for labour markets and, in the long term, could lead to reconsideration of social welfare and tax systems to maintain social cohesion amidst rapid change. However, one must also recognize that this technological acceleration happens at a particularly fortunate time, when the world is beginning to see accelerating demographic aging and shrinking labour pools. It will be a balancing act, but as deep automation across multiple sectors takes hold by the end of the decade, the deflationary impulse will be significant –

"This technological acceleration happens at a particularly fortunate time, when the world is beginning to see accelerating demographic aging and shrinking labour pools."

counterbalancing the inflationary effects of shrinking labour markets.

We are at the dawn of a new era not only in the IT industry but in the global economy, and just like the initial internet revolution, many of the eventual user cases for the new technology have not even been conjured up yet. Many industries will be upended, and many more will be created in the coming years; breakthrough technologies will be developed on the back of GenAI that will accelerate productivity growth and value creation in equity markets.

Investment conclusions

Across our different investment strategies and funds, the investment focus today is on infrastructure providers such as Microsoft, Amazon, and Alphabet, which offer the compute power and GenAI models. We are also focused on the manufacturers of semiconductors such as Samsung and TSMC, both logic as well as memory, as well as companies supplying the machines and other equipment going into producing semiconductors. Examples are companies like ASML, Hoya, Atlas Copco, and Schneider.

Furthermore, proprietary data, which might not be monetised today, has the potential to become a unique asset as GenAI reveals patterns and uncovers knowledge from the data, which can be monetised. We are, therefore, also focused on companies such as MSCI, TransUnion, Verisk, SAP, Deutsche Börse, Relix, Aon, and S&P Global, which own proprietary data. Finally, the application layer is emerging with early leaders like Microsoft with Copilot and Adobe with Firefly. We expect many more investment opportunities in this space as we get deeper into the GenAI age.

Disclaimers

EU: This is marketing material. This publication is prepared by C WorldWide Asset Management Fondsmaeglerselskab A/S. It is provided for information purposes only and does not constitute, and shall not be considered as, an offer, solicitation or invitation to engage in investment operations, as investment advice or as investment research. The publication has thus not been prepared in accordance with legal requirements designed to promote the independence of investment research, and it is not subject to any prohibition on dealing ahead of the dissemination of investment research. Opinions expressed are current opinions only as of the date of the publication. The publication has been prepared from sources C WorldWide Asset Management Fondsmaeglerselskab A/S believes to be reliable and all reasonable precautions have been taken to ensure the correctness and accuracy of the information. However, the correctness and accuracy is not guaranteed and C WorldWide Asset Management Fondsmaeglerselskab A/S accepts no liability for any errors or omissions. The publication may not be reproduced or distributed, in whole or in part, without the prior written consent of C WorldWide Asset Management Fondsmaeglerselskab A/S. Past performance does not indicate future performance.

US: This is marketing material. This publication has been prepared by C WorldWide Asset Management Fondsmaeglerselskab A/S (CWW AM). CWW AM is a registered Danish investment firm located at Dampfaergevej 26, DK-2100 Copenhagen, Denmark. CWW AM's Danish company registration no is 78420510. CWW AM is registered with the SEC as an investment adviser with CRD no 173234. This publication is provided for information purposes only and does not constitute, and shall not be considered as, an offer, solicitation or invitation to engage in investment operations, as investment advice or as investment research. The publication has thus not been prepared in accordance with legal requirements designed to promote the independence of investment research, and it is not subject to any prohibition on dealing ahead of the dissemination of investment research. Opinions expressed are current opinions only as of the date of the publication. The publication has been prepared from sources CWW AM believes to be reliable and all reasonable precautions have been taken to ensure the correctness and accuracy of the information. However, the correctness and accuracy is not guaranteed and CWW AM accepts no liability for any errors or omissions. The publication may not be reproduced or distributed, in whole or in part, without the prior written consent of CWW AM. All figures are based on past performance. Past performance does not indicate future performance. The return may increase or decrease as a result of currency fluctuations.

UK: This document has been prepared by C WorldWide Asset Management Fondsmaeglerselskab A/S(CWW AM). CWW AM is a focused asset manager registered as an investment firm with the Danish FSA. CWW AM is located at Dampfaergevej 26, DK-2100 Copenhagen, Denmark, CVR registration number 7842 0510. This document is directed at persons having professional experience of participating in unregulated schemes (investment professionals) and high net worth companies (as defined under art. 14 and 22 of the Financial Services and Markets Act 2000 (Promotion of Collective Investment Schemes) (Exemptions) Order 2001). The document and any investment or investment activity to which it relates is available only to such persons and will be engaged in only with such persons. Any other person should not rely or act on the statements made in this document. The content of this presentation is confidential, and redistribution or reproduction is prohibited. The presentation is provided for information purposes only and does not constitute, and shall not be considered as, an offer, solicitation or invitation to engage in investment operations as investment advice or as investment research. CWW AM is not responsible for the suitability of the information in the presentation. Opinions expressed are current opinions only as of the date of the presentation. The presentation has been prepared from resources CWW AM believes to be reliable and all reasonable precautions have been taken to ensure the correctness and accuracy of the information. However, the correctness and accuracy are not guaranteed and CWW AM accepts no liability for any errors or omissions. This presentation contains information on investments which does not constitute independent research. Accordingly, it is not subject to the protections afforded to independent research and CWW AM, and its employees may have dealt in the investments concerned. It is emphasised that past performance is no reliable indicator of future performance ant that the return on investments may vary as a result of currency fluctuations.

Active Investments

C WorldWide Asset Management Fondsmaeglerselskab A/S is a focused asset manager. Our objective is to deliver consistent, long term capital growth for clients through active investments in listed equities on global stock exchanges.

Our clients are primarily institutional investors and external distribution channels. Our product range includes discretionary asset management services and commingled fund products.

The combination of a unique investment philosophy based on careful stock picking and long-term global trends coupled with a stable team of experienced portfolio managers, has since 1986 resulted in world-class investment performance.

Please find more of our White Papers on cworldwide.com

This publication was printed on MultiArt® Silk, which is FSC and EU Ecolabel-certified.

C WORLDWIDE ASSET MANAGEMENT FONDSMAEGLERSELSKAB A/S

Dampfaergevej 26 · DK-2100 Copenhagen

Tel: +45 35 46 35 00 · VAT 78 42 05 10 · cworldwide.com · info@cworldwide.com

Q2 2024