

# Thematic Investing

## I want to break free! The Breakthroughs Primer

Primer

The next GPT moment? 30 breakthroughs coming our way  
Exascale computers can carry out quintillions of calculations per second. AI has helped to discover 45x more crystals than previously known to man. It has also found a drug candidate for liver cancer in just 30 days. And there are more industrial robots than people in Uruguay. These are just some of the breakthroughs and tech evolutions we identified in AI, computing, robots, communication, healthcare, energy and transport.

Disruption wins: 3% of companies created all the wealth  
Understanding which breakthroughs are on the cusp of commercialisation has never been more important. In the past 100 years, 3% of companies generated nearly 100% of global net wealth (source: Bessembinder), and roughly a third of the S&P 500 index has been replaced since 2015. In our separate primer picks report, we highlight the enablers to access these breakthroughs theme with a combined market cap of c.\$16tn.

### AI at an inflection point, turbo-charging transformation.

We see 3 reasons why the exponential pace at which themes are transforming businesses – blistering as it already is – will move vertically from here, all powered by AI reaching mass adoption: (1) Tech-tanglement: as technologies become entangled, feeding each other; (2) Tech-scarcity: a world where demand for tech outpaces supply; and (3) Tech-onomics: technologies powering themselves toward lower costs and higher returns.

### Tech-tanglement: a great big feedback loop

2024 will be the “AI for all” year, as it connects and powers technologies, businesses and societies. This leads forward to tech singularity, as computing, communication and tech developments are leapfrogged by the Generative AI revolution, which in turn feeds back, offering more computing power, communication and data to propel itself.

### Tech-scarcity: so much data, so little bandwidth

Moore’s Law is not enough when the computing power needed to train one AI model rises 275x every 2 years. Bandwidth is scarce, and we’re about to move from a Quintillion bytes of data generation per day- to a Sextillion. Renewable energy (>80% of new capacity) also requires new infrastructure and materials, which are in short supply.

### Tech-onomics: capacity up, returns up, prices down

All these investments in automation, AI and tech are reducing prices across the board, and increasing returns. For example, while memory drive capacity is up more than 20,000x in the past 20 years, the price per gigabyte is down >99%.

Trading ideas and investment strategies discussed herein may give rise to significant risk and are not suitable for all investors. Investors should have experience in relevant markets and the financial resources to absorb any losses arising from applying these ideas or strategies.

>> Employed by a non-US affiliate of BofAS and is not registered/qualified as a research analyst under the FINRA rules.

Refer to “Other Important Disclosures” for information on certain BofA Securities entities that take responsibility for the information herein in particular jurisdictions.

BofA Securities does and seeks to do business with issuers covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision.

Refer to important disclosures on page 184 to 186.

12665102

21 March 2024

Equity  
Global



[View Transcript](#)

Thematic Investing  
Merrill Lynch (Israel)

Haim Israel >>  
Equity Strategist  
Merrill Lynch (Israel)  
+972 3 607 2007  
haim.israel@bofa.com

Felix Tran >>  
Equity Strategist  
MLI (UK)  
+44 20 7996 7010  
felix.tran@bofa.com

Martyn Briggs >>  
Equity Strategist  
MLI (UK)  
+44 20 7996 7442  
martyn.briggs@bofa.com

Lauren-Nicole Kung >>  
Equity Strategist  
MLI (UK)  
+44 20 7996 1045  
lauren-nicole.kung@bofa.com

# Contents

---

In short – here are the Breakthroughs	4
Did you know?	6
Executive summary: When does the short-term end and the long-term begin?	7
Next in Artificial Intelligence: AI for all	14
End-device AI	15
Enriched Simulation	18
Knowledge Graphs	21
Hyperdimensional Computing (HDC)	24
Artificial General Intelligence (AGI)	27
Computing evolution: “Moore’s Law is not enough”	32
Exascale and High Performance Computing	33
Spatial Computing	36
Edge Computing	38
Quantum Computing – the end of the world as we know it	45
Neuromorphic Computing/Brain Computer Interface	52
Robotics: a Robot and AI walk into a bar...	56
Humanoid Robots	57
Industrial Robots	67
Other embodied AI interfaces	78
Communication: the scarcity of bandwidth	81
Miniaturized Satellites	83
Wi-Fi 7	89
5G Advanced	91
6G: sooner than you think.	94
Healthcare: from GPT to GLP-1 transformation is here	98
AI Drug Discovery – on the cusp of disruption	104
AI Diagnostics – revolutionizing disease detection	110
Liquid Biopsy – disruptive diagnostics	111
Genomics – gene therapy, gene editing, CRISPR	113
GLP-1 – heading towards ‘GLP-2.0’	120
Energy and Materials meet DeepTech	127
Smart Grid Infrastructure	129
Virtual Power Plants	134
The next generation of solar technology	137
Democratising wind energy	140
Graphene	142
Superconductors: the grid “holy grail?”	146



Future Mobility: Are we nearly there yet?	151
Next-gen batteries	152
Future Air Mobility: It's a bird? It's a drone? It's eVTOL?	165
Autonomous Vehicles	176

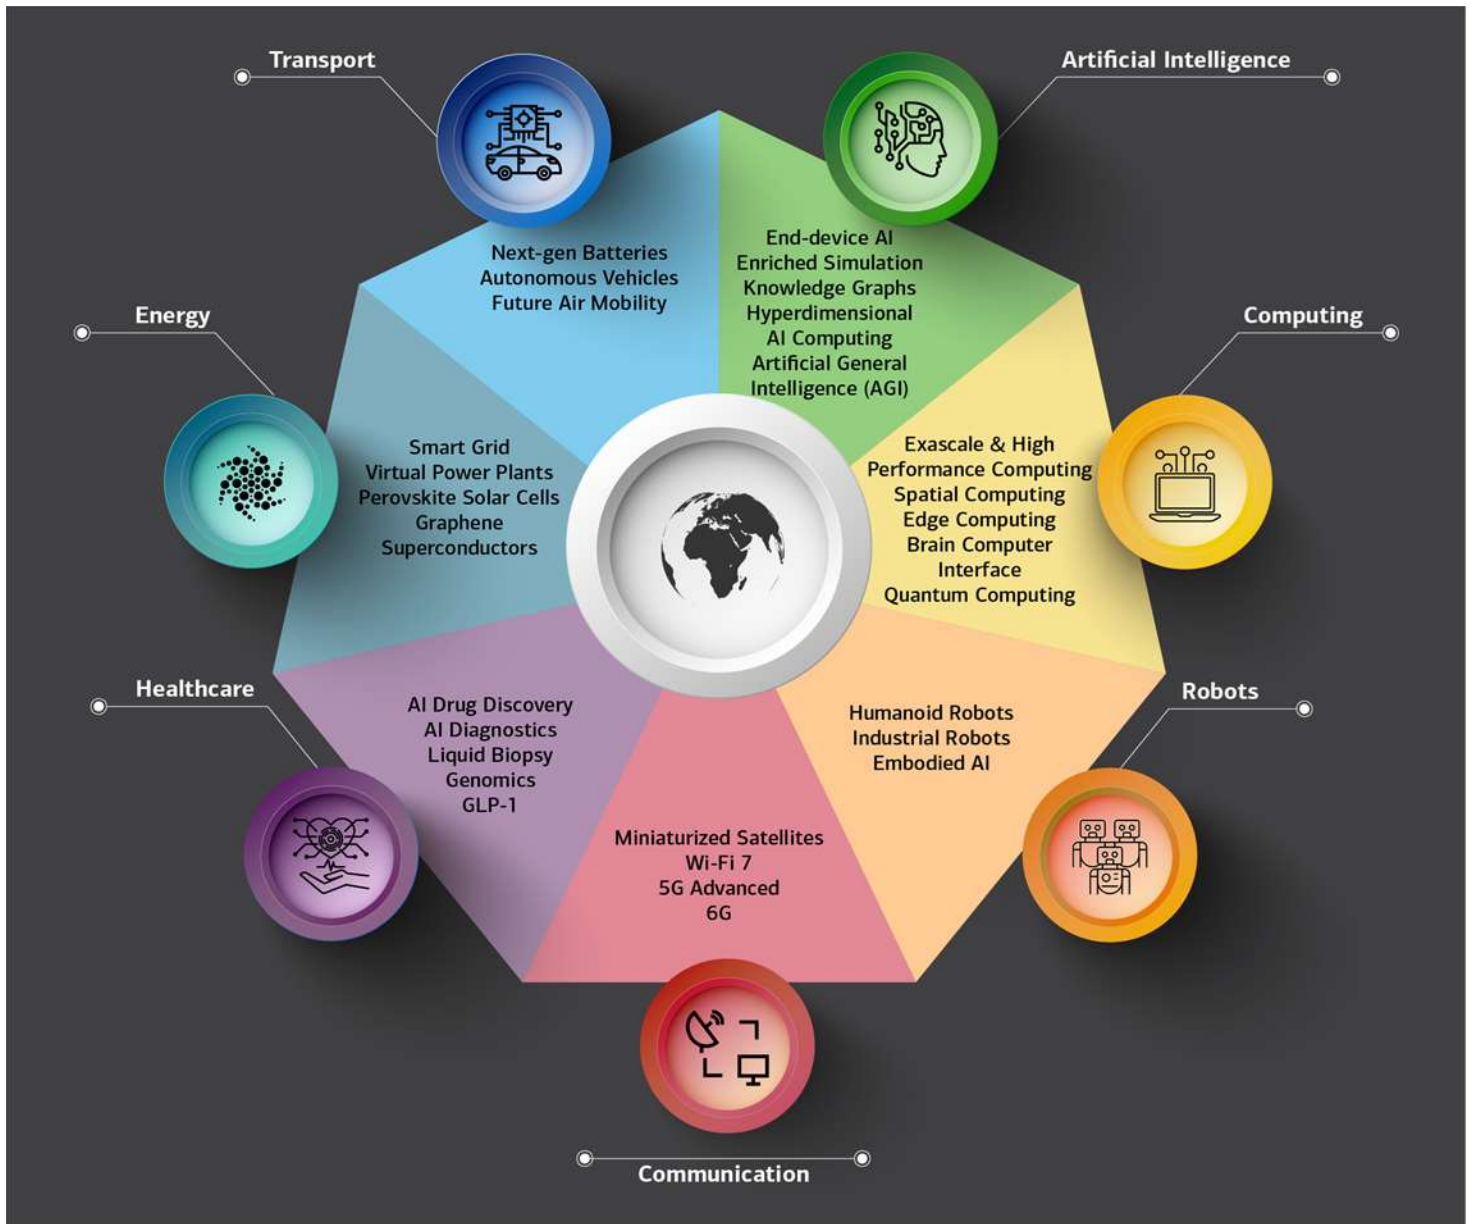


# In short – here are the Breakthroughs

Amara’s law:

“We tend to overestimate the effect of a technology in the short run and underestimate the effect in the long run”. (Roy Amara)

Exhibit 1: Illustrating the breakthrough innovations across different sectors  
 We identify breakthroughs in AI, computing, robots, communication, healthcare, energy and transport



Source: BofA Global Research

BofA GLOBAL RESEARCH



## Exhibit 2: What are the breakthrough innovations across different sectors?

We identify breakthroughs in AI, computing, robots, communication, healthcare, energy and transport

Theme	Breakthrough	What is it?
<b>AI</b>		
	End-device AI	Deployment of AI functions/models on local devices eg, smartphones, autos, wearables
	Enriched Simulation	Using AI to accelerate the discovery process and identify the most viable simulations, speed up the creation of new molecules and bring down the cost to do so
	Knowledge Graphs	Knowledge graphs organise data from multiple sources, capture information and forge connections between them. The solution to hallucination, efficiency and neural networks capabilities
	Hyperdimensional AI Computing	Using high-dimensional vectors to represent information rather than the traditional binary system. It can capture more complex data patterns and allow computers to retain more memory
	Artificial General Intelligence (AGI)	A hypothetical AI agent, that will possess general cognitive abilities at, or above human level on a wide range of tasks and can solve tasks without being preprogrammed to do so
<b>Computing</b>		
	Exascale Computing & HPC	Computational systems that use supercomputers and parallel computer clusters to solve advanced computation problems
	Spatial Computing	Computers that change human machine interaction by using AR/VR to blend the graphic interface for the user to take place in the real physical world
	Edge Computing	Distributed computing that processes data closer to the physical location of end devices compared with the cloud, which is located in data centers far away
	Brain Computer Interface	Brain waves of humans and animals directly interacting with the external world and vice versa
	Quantum Computing	Computing that leverages sub-atomic particles to store information and uses superpositions for complex calculations that would take a classical computer billion of years
<b>Robots</b>		
	Humanoid Robots	Humanoid robots assist humans and have a human-like form, mimicking human behaviour and interactions
	Industrial Robots	Machines that can be automatically controlled or programmed to manipulate and move objects – they are a key enabler of the emerging Industrial Internet
	Embodied AI	Other embodied AI interfaces include service robots, autonomous vehicles, agricultural robots, drones, spatial computing/virtual reality etc
<b>Communication</b>		
	Miniaturized Satellites	Micro and Nano satellites, which are low cost, small, lightweight satellites, mostly in Low Earth Orbit (LEO, up to 1,000km altitude)
	Wi-Fi 7	The next generation of Wi-Fi offering 5x more speed, leapfrog capacity and 75% drop in latency compared to previous Wi-Fi generation
	5G Advanced	The evolution of 5G mobile technology with embedded AI capabilities, offering 10x the mobile speed, a third of the latency and lower costs compared to current 5G
	6G	The revolution of mobile phone technology. The new generation which can offer 1Tbps speeds, close to zero latency and AI embedded capabilities
<b>Healthcare</b>		
	AI Drug Discovery	Using AI to escape the massive jump in medicine discoveries and inefficiencies. Reduce spend, time to market and personalization of drugs and treatments
	AI Diagnostics	AI has recently made huge advances in automatically diagnosing diseases, making diagnostics cheaper and more accessible
	Liquid Biopsy	Liquid biopsy is a "simple blood test" that can provide critical health information, especially for cancer
	Genomics	Genomics is the study of genes and their functions, and related techniques, whereas genetics is the study of heredity. Proteomics is roughly defined as large scale study of proteins
	GLP-1	Glucagon-like peptide 1 (GLP-1) are hormones responsible for the "incretin effect" that makes you feel "more full" and eat less
<b>Energy</b>		
	Smart Grid	Expanding, upgrading and digitalising electricity grids to cope with the shifts in power supply and demand, increasing the share of global energy consumption to electricity
	Virtual Power Plants	A VPP aggregates or controls a network of power generating or consuming assets to be used collectively to balance an electricity grid
	Perovskite Solar Cells	Next-gen materials with strong light absorption and electric charge properties to increase the conversion efficiency of solar cells
	Graphene	The thinnest and strongest material known to humans
	Superconductors	Materials that carry electrical currents with zero resistance
<b>Transport</b>		
	Next-Gen Batteries	Batteries with increasing capacity to store energy and recharge faster and cheaper, enabled by improving chemistry and material sciences
	Autonomous Vehicles	Technology that can assist or fully replace human driving, enabled by a combination of sensors and software
	Future Air Mobility	Aircraft & drones powered by electric propulsion that can take off and land vertically. Democratizing lower-cost/higher-frequency urban air travel for passenger and cargo transit.

Source: BofA Global Research. NOTE: HPC = high performance computing

BofA GLOBAL RESEARCH



## Did you know?

- Sophia was the world's first AI humanoid robot and was also the first humanoid to get a passport<sup>1</sup>
- AI simulation can speed up screening by up to 100x<sup>2</sup>
- In 2023, DeepMind used AI to discover 45x more new crystals than all the crystals discovered in the entire history of science<sup>3</sup>
- Hyperdimensional computing could achieve over 600% in energy savings vs the typical tech used in today's chips<sup>4</sup>
- High performance computing systems can perform quadrillions of calculations per second vs regular laptops calculating 3 billion at most<sup>5</sup>
- 75% of enterprise-data will be created and processed at the edge by 2025<sup>6</sup>
- A brain-to-brain interface has already allowed one human to control another's hands through their thoughts alone<sup>7</sup>
- In 2019, Google, in partnership with NASA and Oak Ridge National Laboratory, announced that its quantum computer had made in 200 seconds the number of calculations that it would take a supercomputer 10,000 years<sup>8</sup>
- An individual would have to perform one calculation every second for approx. 50 quintillion years to equal what a quantum computer can do in 1 second<sup>9</sup>
- Modelling penicillin molecule, which is composed of 41 atoms, would require a classical computer with more transistors than atoms in the observable universe<sup>10</sup>
- AI could add the equivalent of China's GDP to the global economy by 2030<sup>11</sup>
- There are more industrial robots world-wide than the population of Uruguay<sup>12</sup>
- In the 3 months from satellite broadband service launch in Ecuador's Galapagos islands (30,000 inhabitant only), 100 new businesses were opened<sup>13</sup>
- 233 million Wi-Fi 7 devices will enter the market in 2024, growing to 2.1 billion devices by 2028<sup>14</sup>
- At current average global mobile speeds, it will take someone c.181 million years to download all the data from the internet...but it will take 838,000 years on a 5G network and "only" 17,000 years on 6G<sup>15</sup>
- Each person will generate enough health data in their lifetime to fill 300 million books, or the NY public library 6x...<sup>16</sup>
- Generative AI could create \$60-110 billion of economic value in healthcare<sup>17</sup>
- A coffee mug full of DNA could theoretically store all of the world's data<sup>18</sup>
- The cost of obesity is larger than Germany's GDP<sup>19</sup>
- EV batteries alone could be used to satisfy short-term grid storage demand by 2030<sup>20</sup>
- One superconducting cable could transport 500x more electricity than copper wire, enough to power a large city<sup>21</sup>

Source: 1British Council; 2IBM; 3DeepMind; 4Kurunaratne; 5TechTarget; 6Gartner; 7University of Washington; 8Google, NASA; 9McKinsey, Google, BofA Global Research; 10IBM, BCG; 11PwC, World Bank, UN; 12Stanford, UN; 13Tim Papandreu; 14Wi-Fi Alliance; 15Statista, Unicorn Insight, IDC; 16KPCB; 17McKinsey; 18Mark Bathe, MIT; 19Global Wellness Institute; 20Leiden University, NREL; 21Nexans



# Executive summary: When does the short-term end and the long-term begin?

## Did you know?

- It took c.400,000 years to get from the moment humans used fire to the agriculture revolution (10,000 BC)
- 11,750 years to reach the Industrial Revolution (mechanisation)
- 100 years to Industry 2.0 (electrification)
- 100 years to Industry 3.0 (automation)
- 50 years to 4.0 (digitalisation)
- 10 years to 5.0 (humanisation)

In this primer, we map the breakthroughs that are about to alter our lives. It is only a year since generative AI changed the world and so now we look at what is next. The pace at which themes are transforming businesses is unprecedented, but we believe this rate of technological change will be even faster going forward. As a result, we need to start preparing for what is coming next today. The commercialisation of new technologies and breakthroughs will happen much faster than ever before. The exponential growth of data (doubling every 2-3Y), cheapening computing power (Moore's law reincarnated) and rise of a connected world (IoT, mobile devices, social media) will bring about the fastest transformation in human history.

However, we believe the rate of change will accelerate even further from here. We highlight 3 reasons why AI will be the engine to power it all:

1. Tech-tanglement: Technologies, business, and societies are about to be connected, and powered by AI. Computing, communication and tech developments leapfrogged this year by the Generative AI revolution, which is now feeding back and offering stronger computing power, communication and data to propel Generative AI, materials, batteries etc. Tech-Singularity was never closer.
2. Tech-scarcity: Moore's Law is not enough when computing power needed to train one AI model has risen by 275x every 2 years. We see scarcity of bandwidth, where we are about to move from the Quintillion bytes of data generation per day to the Sextillion range. We are also seeing scarcity of new infrastructure and materials, when more than 80% of new energy capacity add was renewable.
3. Tech-onomy: All these investments in automation, AI and tech are reducing prices across the board and increasing returns. For example, while drive capacity has risen by more than 20,000x in the past 20 years, the price per gigabyte has fallen by >99%.

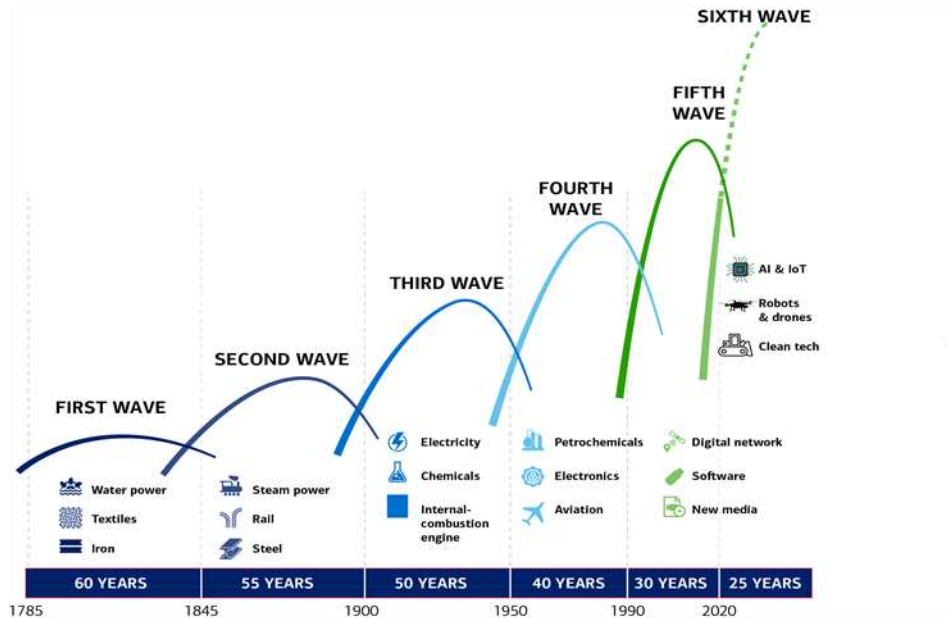
AI: at a mass adoption inflection point to orchestrate it all

What were the first five waves of innovation? The First Industrial Revolution involved water power driving manufacturing. Advances in textiles brought about the first factory and cities expanded around them. The second wave involved rail, steam and steel advancements. The third saw the emergence of electricity powering light and telephone communication. During the fourth wave, aviation revolutionised travel. The fifth cycle is related to the internet bringing in a new frontier of globalisation and information flow (source: Neufeld).



As the innovation waves shorten, we may be nearing the sixth wave, which embraces AI and the digitisation across IoT, robotics, drones, etc, to create intelligent manufacturing systems that can operate with minimal human intervention. It also relates to cloud computing, collaborative robots and quantum computing. Industry 6.0 may seem futuristic but with the advent of (generative) AI, it may be happening much sooner than we think. For example, clean tech may come even closer to the forefront as the necessity to reach net zero becomes ever more pressing.

Exhibit 3: Innovation waves are speeding up  
The waves of innovation have increased from 60+ years to 25 years today



Source: Visual Capitalist, Edison Institute

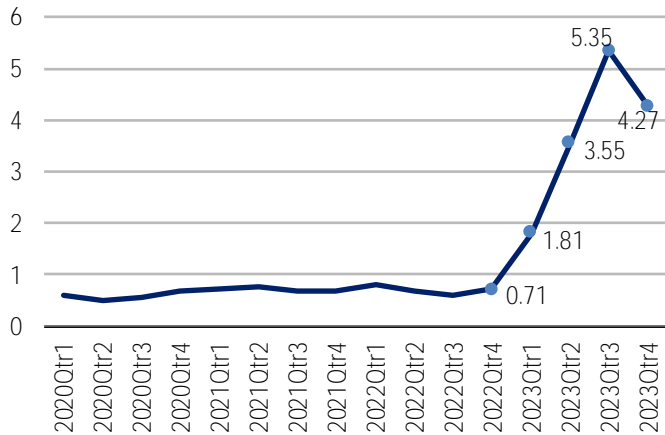
BofA GLOBAL RESEARCH

For a technology to be considered a GPT, Bresnahan and Trajtenberg highlight three characteristics that should be exhibited: pervasiveness (the GPT should spread to most sectors), improvement (the GPT should get better over time and thereafter lower the cost to its users) and innovation spawning (the GPT should make it easier to innovate, create new products/processes). AI can certainly tick these boxes.

On pervasiveness, as discussed in our [AI/ChatGPT Primer](#), we're at an iPhone defining moment for AI, where there is an abundance of commercial use cases for generative AI and virtually any/every sector can use it. Stakeholders have already started adopting AI. In a survey conducted by McKinsey in August 2023, one-third of respondents indicated that their organisations were using generative AI regularly in at least one business function. We can see in the number of mentions of AI on company earnings calls that stakeholders are talking about AI and how to tap into this technology. 2023 was a record year of investment in generative AI: startups in this space raised US\$21bn, up more than 5x YoY. These Large Language Models (LLMs) have certainly evolved, and as investment increases, they will continue to, enabling more use cases and innovations.



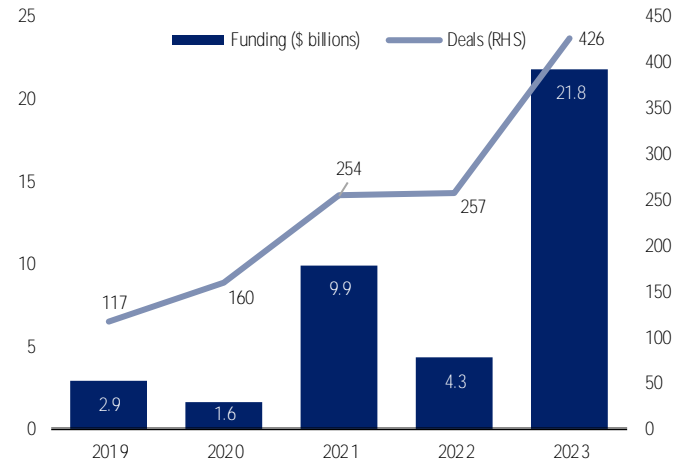
**Exhibit 4: AI mentions per company from Q1'20 to Q4'23**  
Average AI mentions per company jumped 6x from 0.71 to 4.27 in Q4 2022 vs Q4 2023



Source: BofA Global Research, FactSet

BofA GLOBAL RESEARCH

**Exhibit 5: Disclosed equity funding (US\$ bn) and number of deals**  
Generative AI investment surged in 2023



Source: CB Insights

BofA GLOBAL RESEARCH

**Tech-tanglement – at the heart of Industry 6.0 is AI**

It's not just about the impact of AI itself but also how it intertwines with other technologies, e.g., robots, simulation and communications. As AI evolves, so will everything else. In this way, AI is the key engine that can enable the next wave of innovation. How? There are multiple channels, but to name a few:

1. AI could be embedded into existing products offering better capabilities. For example, 5G advanced (5G mobile technology embedded with AI capabilities) could have 10x the mobile speed and a third of the latency and costs compared with 5G. Using AI in grid management could help plan electricity requirements 12x faster than without AI.
2. AI is being leveraged as a tool (for other innovations) – for instance, AI-enriched simulation for drug discovery and material breakthroughs. AI helped to discover 45x times more crystals ever known to man. Using AI in drug discovery enabled the finding of a candidate for liver cancer in just 30 days.
3. Creating the required technology to enable the next-gen of AI. For example, to reach AGI, we may need embodied AI – giving a physical form to AI so that it can learn from a wider environment of data and teach itself to do tasks that it was not pre-programmed to do. This could involve putting AI into robot-like humanoids, drones, cobots and AVs.

These connections and interactions of technologies create a feedback loop. As AI-enabled products offer better capabilities, as AI is used as a tool for fuelling more innovation and as it creates demand to improve existing technologies, our world becomes ever more infused in this AI era. For example, over the past few years, LLMs have become larger, driving the need for more compute. But we may want to have LLMs running on end-devices rather than in the cloud, which opens up more opportunities for end-device/edge compute. In this way, we have AI fuelling innovation, fuelling AI, fuelling innovation etc. We've seen many AI-related tech developments over the past year but we could see more as this propagation continues to unfold, springboarding us on to the path of Industry 6.0.

**Risk of scarcity – too much demand**

We live in a world of demand for technology to address our world of scarcity. With growing innovation and improving economics, could we leverage these new technologies



to shift our world from one of scarcity to one of abundance? We illustrate some current examples of scarcity in bandwidth, compute, healthcare and metals.

1. Scarcity in bandwidth: zettabytes to yottabytes to brontobytes to geopbytes. We live in an era of zettabytes. That is to say, we inhabit a world that generates 328 Quintillion (million trillion) bytes of data everyday (source: Statista). Data generation will only get faster from here as AI takes hold. An estimated 90% of all global data will be AI-generated by 2025 (Nina Schick). According to IDC, the amount of data created is projected to double every 2-3 years. Meanwhile, we created c.64ZB of data in 2020 and 120ZB by end-2023, according to Statista, and this is projected to grow to 181ZB by 2025E.
2. Scarcity in compute: overtaking Moore's Law: Moore's Law is the observation that the number of transistors in an integrated circuit doubles every two years. The training of AI systems can be divided into two eras: the First Era and the Modern Era. The First Era of compute usage relied on limited computational resources and simple algorithms. The Modern Era, which began in c.2012, with the rise of deep learning and powerful hardware, allowed for the training of increasingly complex models with billions of parameters. Post-2018, the computing power increased 275x every 2 years compared to 8x previously. The exponential increase in the size of these models means that our computational demands go beyond Moore's Law.
3. Scarcity in healthcare: c.US\$1tn pa of US healthcare spend is wasted: The US has one of the most expensive health systems world-wide with spending totalling \$4.3tn in 2021. One reason for the high cost is waste. Approximately 25% of healthcare spending in the country is considered wasteful, and about one-quarter of that could be recovered through interventions. The largest source of healthcare system waste, roughly \$266bn is administrative costs. There is currently a global shortage of more than 7mn health workers, which could rise to 13mn by 2035 (source: WHO).
4. Scarcity in metals: need 22x more lithium by 2050 for energy transition: Transitioning the global economy to clean energy and transportation would require significant increases in critical minerals, many of which are relatively small commodity markets requiring additional mining capacity. This can be costly (with long lead times) and environmentally sensitive. Furthermore, their availability is often geographically concentrated posing supply chain risks. Indeed, more than 300 new mines would be required by 2035, per Benchmark Mineral Intelligence for EV batteries and energy storage, requiring multiple times more lithium (+22x), and nickel (+9x) in 2050 relative to 2022 per BNEF. However, despite much concern over the impacts from EV batteries, metal volumes for the electricity grid would be far higher – the total copper requirement could double in 2020-40 per the IEA.

Economics: Technologies powering themselves towards lower prices

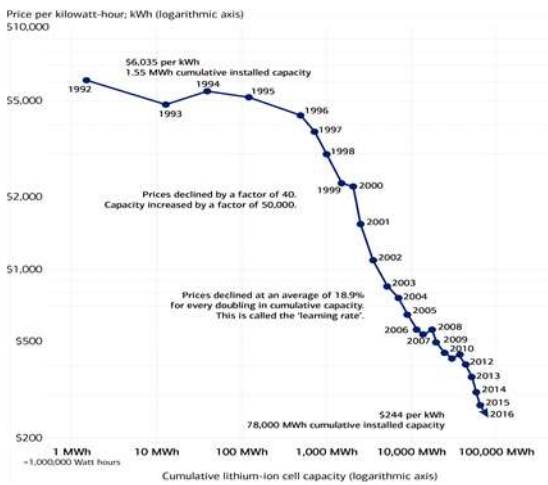
Improving economics speeds up innovation. One way of thinking about the economics behind innovation is the learning curve. This is the observation that costs decrease as the production of a technology increases. For every cumulative doubling of production, the cost per unit decreases by typically 20%. To illustrate this, say we manufacture 1,000 units of a product, then the cost per unit will decrease by 20% when production reaches 2,000 units. For very mature products, every doubling can take years. The channel of propagation is an example of learning by doing. Increasing production gives engineers the opportunity to learn how to improve the process. This creates a positive feedback loop of increasing demand and falling prices. More technology gets deployed to satisfy demand, leading to falling prices. Then these technologies become cost-effective in these new applications, feeding increased demand again.



Tech costs are falling dramatically

Illustrating the change in costs, for lithium-ion batteries, prices declined by an average of 18.9% for every doubling in cumulative capacity (source: Our World in Data; Lafond et al and IRENA). In the past 30 years, the price of batteries has declined by 97% (source: Ziegler and Trancik). For solar panels, each doubling of installed capacity, the price of solar modules decreased on average by 20.2%. Solar cells made from perovskites (materials with strong light absorption and electric charge properties) can propel solar efficiency and halve costs. In fact, perovskite solar cells have increased in efficiency from 3% in 2009 to 26% in 2024. The cost of sequencing a full human genome has come down by 99.99% in the past 20 years. The cost of computer memory chips has dropped by 640% every 5 years.

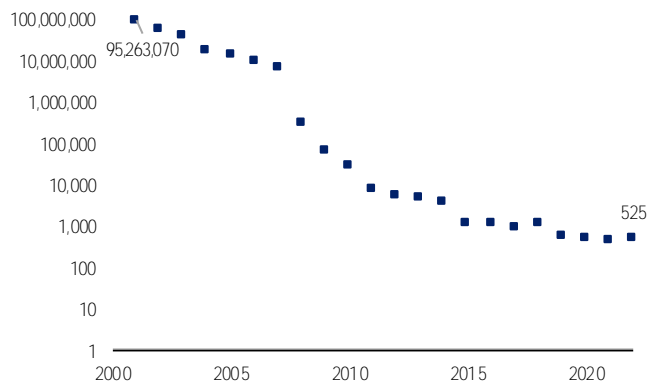
Exhibit 6: Lithium-ion Learning Curve: 18.9% between 1992-2016  
Comparing Price per kilowatt-hour (kWh) vs cumulative lithium-ion cell capacity shows prices fell 18.9% for every doubling of installed capacity



Source: Our World in Data, Micah Ziegler & Jessica Trancik (2021); prices adjusted for inflation given in 2018 US\$/kWh

BofA GLOBAL RESEARCH

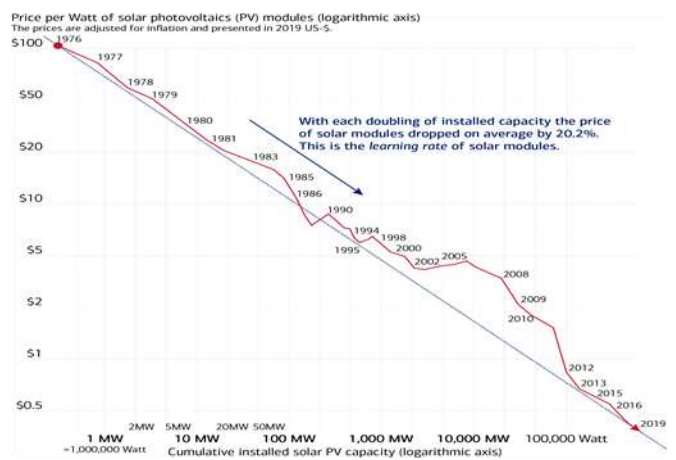
Exhibit 8: Cost of sequencing a full human genome (US\$; log scale)  
The cost of sequencing a full human genome has decreased from US\$95m in 2001 to US\$525 in 2022 - that's a decrease of 99.99% in the last 20 years



Source: National Human Genome Research Institute, Our World in Data. NOTE: the data is not adjusted for inflation

BofA GLOBAL RESEARCH

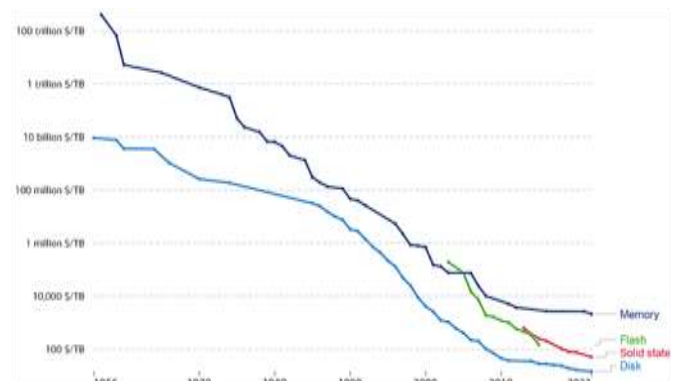
Exhibit 7: Price per Watt of solar photovoltaics (PV) modules (logarithmic axis)  
The price of solar modules has declined by 99.6% since 1976



Source: Our World in Data based on data from Lafond et al and IRENA Database. The reported learning rate is an average over several studies reported by de La Tour et al in Energy. The rate has remained very similar since then. The prices are adjusted for inflation and presented in 2019 US\$.

BofA GLOBAL RESEARCH

Exhibit 9: Cost of computer memory and storage (US\$ per terabyte, TB)  
Since the 1950s, the cost of computer memory chips has come down by 640% every 5 years



Source: McCallum; Our World in Data. NOTE: the data is expressed in US\$ per terabyte (TB), not adjusted for inflation

BofA GLOBAL RESEARCH

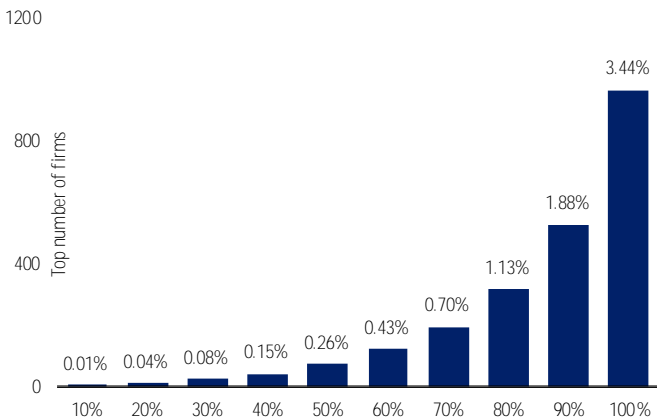
Why care now? Disruption always wins

The pace of transformation continues to accelerate, meaning that long-term, distant technologies are approaching ever faster. However, humans typically underestimate this acceleration, failing to accurately assess the near-term potential of new or new forms of



technologies and their impact on society. Consequently, understanding technologies that may arrive much sooner than expected and distinguishing which will become disruptive is key. Especially given that returns are driven predominantly by a small share of disruptive companies. In fact, since 1926, US\$55tn in shareholder wealth has been created, but just 3% of companies accounted for that entire amount (source: Bessembinder). Furthermore, between 1926 and 2022, nearly one third of all US shareholder wealth was generated by 0.1% of US listed stocks, that's only 25 companies (source: Bessembinder; Visual Capitalist). This trend of concentrated wealth creation by a small number of companies is actually accelerating in 2016-19, just 5 companies accounted for 22% of net wealth concentration.

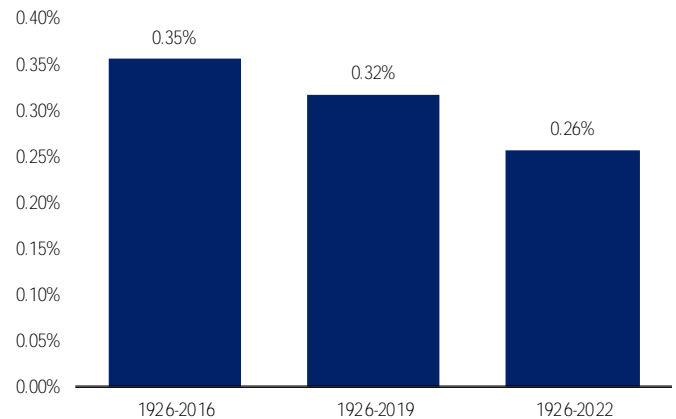
Exhibit 10: The top percentage of firms that generated shareholder wealth between 1926 and 2022 (by share of net wealth created) just 3% of all companies have generated net wealth in the US stock market since 1926



Source: Bessembinder

BofA GLOBAL RESEARCH

Exhibit 11: % of total US firms that has accounted for 50% of net wealth in each time period. The percentage of firms that has accounted for 50% of net wealth has decreased by 10ppts over a 6 year time frame - that's 18 companies



Source: Bessembinder

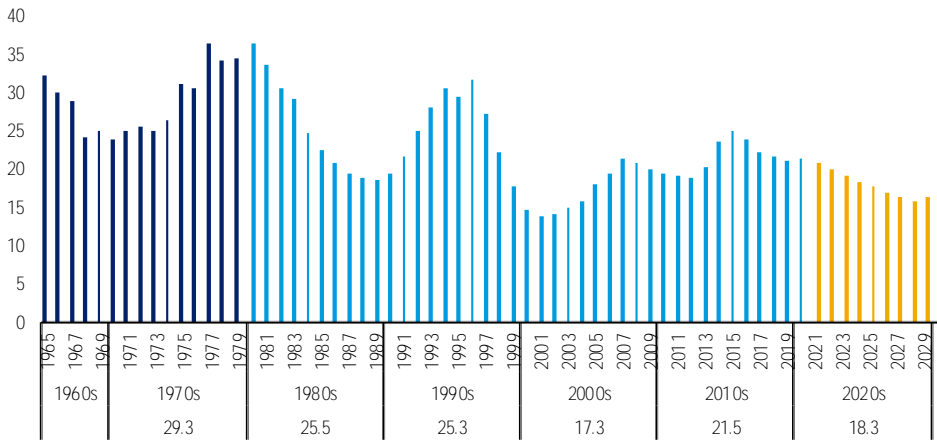
BofA GLOBAL RESEARCH

### Life expectancy of incumbents keeps falling

What's more, the displacement of incumbents is only growing. Investments in the majority of US publicly listed stocks (58.6%) led to lower rather than higher shareholder wealth (source: Bessembinder) and roughly a third of the index has been replaced since 2015 (source: S&P Dow Jones Indices). In 1958, the average seven-year rolling lifespan of a company on the S&P 500 was 61 years, by the 1980s it had dropped to 30 years, and by 2016, it was 24 years and as of 2021, it is 16 years (source: McKinsey). In 2020, the same study by McKinsey stated that 75% of the companies on the index would disappear. If we continue on this road, by 2027, companies could last just 12 years as they become increasingly disrupted (source: Innosights, S&P 500). The number of firms that explains half of the net wealth creation in the US since 1926 has decreased from 90 as of 2016 to 83 as of 2019 and to 72 as of 2022 (source: Bessembinder). That's a 20% drop in just 6 years.



Exhibit 12: Average company lifespan on Standard and Poor's 500 Index from 1965 to 2030, in years (rolling-7-year average) and decade averages  
 Incumbents are getting disrupted even quicker



Source: Innosight based on S&P500 data; Statista.

BofA GLOBAL RESEARCH



# Next in Artificial Intelligence: AI for all

Transforming World: We're just at the beginning of the AI revolution. While Generative AI and LLM made their "public appearance", from here it will all be on the implementation of AI models, and applications.

Did you know? AI could add the equivalent of China's GDP to the global economy by 2030<sup>1</sup>.

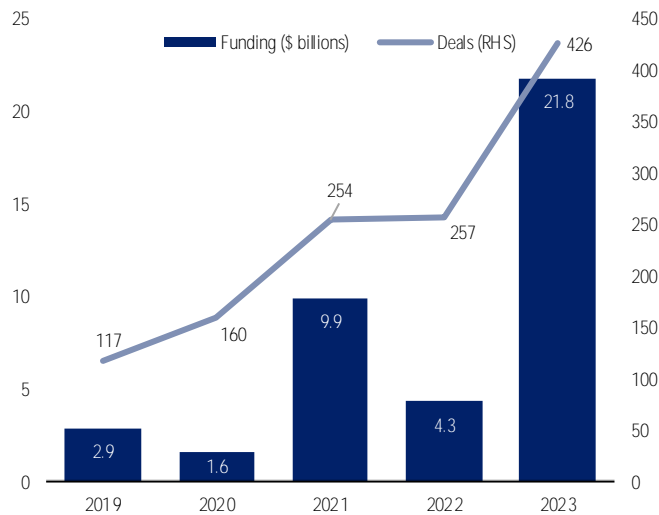
Next generation of AI: AI into everything, knowledge graphs and hyperdimensional computing, Artificial General Intelligence (AGI)

Source: <sup>1</sup>Based on figures from PwC, World Bank, UN. Based on 2022 GDP data

2023 was the year of rapid generative AI development

As discussed in our [Artificial Intelligence Primer – Me, Myself and AI](#), we're at an iPhone defining moment for AI. Generative AI stole the headlines in 2023 but this is just the beginning and this AI revolution will only accelerate from here. The buzz all started with the release of ChatGPT by OpenAI in November 2022 and 2023 saw a surge in generative AI investment (Exhibit 13). Since ChatGPT's release, a variety of both closed and open source models have been introduced (Exhibit 14) with companies already starting to develop, adopt or integrate AI into their products or businesses. The pace of innovation is set to accelerate from here, with more AI tools (simulation, knowledge graphs) and applications likely to be available soon. This could enable an abundance of opportunities beyond the digital to the physical domains of end-devices, robotics and life sciences.

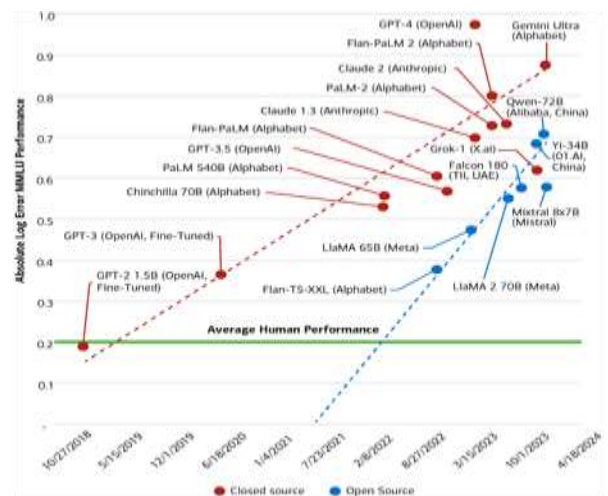
Exhibit 13: Disclosed equity funding (US\$ bn) and number of deals  
Generative AI investment surged in 2023



Source: CB Insights

BofA GLOBAL RESEARCH

Exhibit 14: Open source vs closed source models in terms of performance  
Open-source models are improving faster than closed source models



Source: ARK. NOTE: MMLU (Massive Multitask Language Understanding) is a benchmark designed to measure knowledge acquired during pretraining by evaluating models exclusively in a zero-shot (responses with no training) and few-shot (fine-tuning models with minimal examples for better accuracy) settings.

BofA GLOBAL RESEARCH



## End-device AI

What is it? Deployment of AI functions/models on local devices e.g., smartphones, autos, wearables

Did you know? Having LLM-infused AV tech could reduce the cost and hardware needed for autonomy by 5-7x<sup>1</sup>

Source: <sup>1</sup>Alex Kendall

### End-device 101

Fuelled by ChatGPT's development, cloud AI was a key topic in 2023. On top of cloud AI opportunities, end-device AI is a new growth potential as per the BofA Semis Team. This is because we need AI functions on both the end-device and cloud sides to make the most of AI. With a significant amount of data generated from numerous devices every day, the key consideration is where the data should be processed. Multiple industry suppliers define three categories within the compute spectrum:

- End-device AI refers to the deployment of AI functions/models on local devices, e.g., smartphones, autos and wearables. Due to the long data transmission path, transferring data from end-devices to the cloud introduces costs/issues, including longer latency, power consumption for data transmission, bandwidth, server capacity, private information leakage, etc. which can lower the service quality. Therefore, end-device AI helps share the power load of large servers to improve the performance of the broad AI ecosystem.

In this report, the definition of end-device AI computing is: collecting data from numerous end-devices like smartphones, cars, white goods, security cameras, wearables, street lights, etc., and then processing the acquired data on end-devices (with embedded AI chips/cores) or sending that data to a gateway/another processing device near them.

- Edge AI: Based on the abovementioned definition, edge AI offloads AI and ML processing from the cloud to powerful servers at the edge of the network, such as offices, 5G base stations, and other physical locations near the connected endpoint devices.
- Cloud AI: The collected data/prompts are sent back to data centres, and then the end-devices grab the calculated results from big data pools on the cloud.

#### Exhibit 15: Ecosystem of end-device AI, edge AI, and cloud AI

End-device AI will be processed on end-devices; edge AI will be processed in physical locations near to end-devices; cloud AI will be processed in data centres



Source: BofA Global Research, Infineon, company data

BofA GLOBAL RESEARCH



## Benefits of end-device AI

Exhibit 16: Comparison of end-device AI, edge AI, and cloud AI

End-device AI's advantages include lower latency/cost/power efficiency, higher context awareness, etc.

Features	End-device AI	Edge AI	Cloud AI
Use cases	Smartphones, cars, white goods, security cameras, wearables, street lights, etc.	On-premise servers, base stations, IoT gateways	Datacenters
Latency	Low	Medium (10us-10ms+)	High (100ms+)
Bandwidth required	Low	Medium	High
Processing power	Low	Medium	High
Storage capacity	Low	Medium	High
Security	High	Medium	Medium (cloud back-up)
Computing cost	Low	Medium	High
Context awareness	High	Medium	Low
Power efficiency	High	Low	Low
Maintenance & Upgradeability	Medium	Medium	High (centralization)

Source: BofA Global Research

BofA GLOBAL RESEARCH

The key features of end-device AI compared with edge/cloud AI include lower latency, higher power efficiency, limited cost, and higher context awareness, etc., which should translate into the benefits listed below:

- **Faster response time:** Some simple AI tasks only require AI computing with local storage capacity on the end-device side. In such use cases, adopting end-device AI could lead to faster response times vs. capturing data from cloud.
- **Better accessibility for consumers:** Given end-devices are physical and closest to end-users, the AI processing on them could lead to better accessibility.
- **Better privacy and communication security:** With higher security requirements on the end-device side (e.g., face recognition/fingerprint, one-time password, text password, etc.) for personal information protection, the security level could be higher if data remain local.
- **Offloading from cloud:** Edge devices with computing power could share the cloud's burden to improve the overall AI service.

More achievable due to wireless connectivity (e.g., 5G, Wi-Fi) & IoT

The three technology megatrends, wireless communication, AI, and IoT, are mutually driving the growth. 5G offers high speed/low latency/wide range connection for AI, especially on the end-device side, while AI computing could enhance 5G's transmission performance and efficiency. As an example, NVIDIA developed its AI-on-5G platform to provide edge AI+5G solutions by combining NVIDIA-converged accelerators, SDK (software development kit) for software-defined 5G virtual radio area networks (vRANs), and a portfolio of enterprise AI applications and SDKs.

## Challenges for end-device AI

To perform end-device AI tasks locally and connect with edge/cloud for more complex AI tasks, there may be some key challenges for semis players.

- **Power consumption:** End-devices like smartphones, NBs, smart watches, etc. are close to end-consumers and running on battery power, and thus power consumption needs to stay at low levels even if equipped with more features.
- **Cost:** The increasing semiconductor content needed to support more AI functions should inevitably lift the cost of end-devices.
- **Algorithm/software:** As end-devices have limited resources, including processing power, memory, and storage, vs. cloud/datacenter, the AI/ML



algorithm and related software must be optimised to work within these constraints.

- Security: One of the most significant challenges of end-device/edge AI is data privacy disclosure. End-devices and the edge servers store and process a large amount of data, including sensitive personal data. It makes them attractive targets for cyberattacks.



## Enriched Simulation

What is it? Using AI to accelerate the discovery process and identify the most viable simulations, speed up the creation of new molecules and bring down the cost to do so.

Did you know? AI simulation can speed up screening by up to 100x<sup>1</sup> and, in 2023, DeepMind discovered 45x more new crystals than all the crystals discovered in the entire history of science.

Source: <sup>1</sup>IBM; <sup>2</sup>DeepMind

### AI can alleviate the challenges of modern simulation

Many of our everyday products are complex and over time designers have come to rely on computer-driven simulations but they often take time to run. Even once possibilities are found, additional simulations need to be run to ensure safety. AI simulation combines techniques from quantum physics and deep learning to enable sampling a vast dataset quickly and efficiently. AI and simulation technologies bring the ability to take a molecular structure and simulate it billions of times, making small changes each time to see which structure is optimal. We can now do this in a matter of weeks and months – a task that would take 10 years in the physical world.

### What does the simulation process involve? How does it work?

Bayesian statistics create methods that rival traditional simulation methods by finding the optimal solution based on the limited information that the system has. In this way, the method minimises the uncertainty of the next best solution thing using existing knowledge to select the next best parameters.

### Numerous applications: drug discovery, chips, chemicals, materials

- Drug discovery involves high costs and has a high failure rate. The average investment is US\$1-4bn and it typically takes 10-15 years to develop new drugs. There is a 90% failure rate, which makes the business model unpredictable. And revenues from the 10% of drugs that are successful have to pay for the 90% that failed.

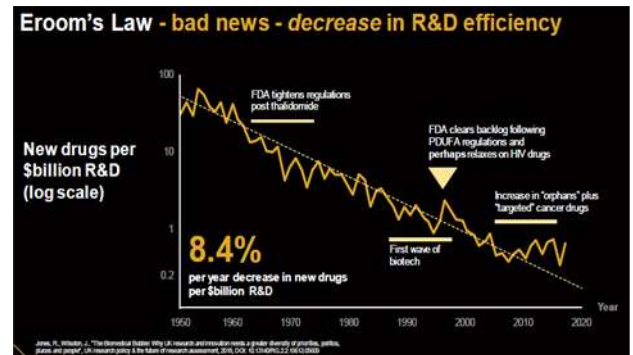
**Exhibit 17: Challenges in the drug discovery market**  
There is a high cost and failure rate in the drug discovery market e.g. US\$1-4bn of average investment, 10-15 year development timeline, 80-90% clinical trial failure rate



Source: SandboxAQ

BofA GLOBAL RESEARCH

**Exhibit 18: Eroom's Law: new drugs per \$bn R&D continue to fall**  
The pharma industry has seen a decrease in R&D efficiency over several decades despite the availability of new tech



Source: SandboxAQ, Drug Discovery Today

BofA GLOBAL RESEARCH

The high failure rate helps explain Eroom's law (the observation that drug discovery becomes slower and more expensive over time), which implies a decrease in R&D efficiency. Eroom's law is where companies spend more per new drug on R&D than they make in revenue. New AI simulation work can change sectors such as life



sciences from a lumpy business of mostly failure to one with predictable revenues, as it allows companies to take all the data, molecular information and run billions of simulations de-risking the molecules, reducing drug development time dramatically.

Exhibit 19: Timeline on Future AI Development in Healthcare  
 Our BofA Healthcare colleagues see the application of AI in drug discovery realising in the next 3-5 years, while other healthcare applications will take longer at 7+ years

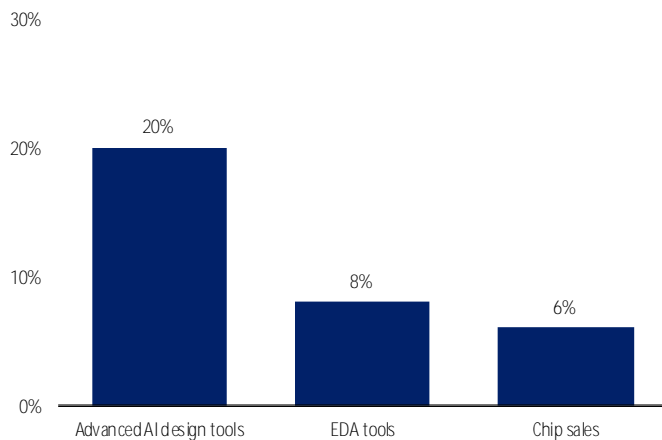


Source: BofA Global Research

BofA GLOBAL RESEARCH

- Chip design: Electronic design automation (EDA) vendors have made tools for chip design using rule-based systems and physics simulation. But now AI can help chipmakers push the boundaries of Moore's Law (observation that the number of chips in a circuit doubles every two years) further. Simulation can design chips faster than older methods, and make new and better chips. All in all, these tools can increase supply chain security and help mitigate shortages.

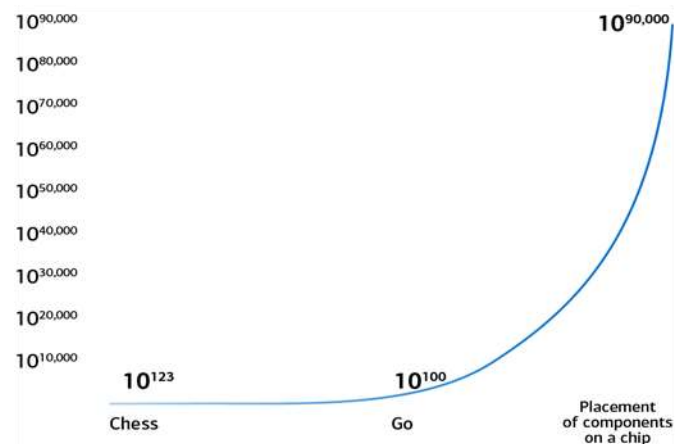
Exhibit 20: Five-year CAGR for chips, EDA tools, and advanced AI design tools (2023-28)  
 Growth in advanced AI tools for chip design is expected to be more than double that of EDA tools and more than triple the growth rate of chip sales



Source: WSTS: Global Markets Insights; Deloitte Global

BofA GLOBAL RESEARCH

Exhibit 21: Chip designs have exponentially more possible configurations than either chess or Go  
 If you thought seeing AI beat a human at chess and Go was impressive, wait until you see it design a chip



Source: Adapted from Synopsys, "What is Design Space Optimisation?", July 21, 2020

BofA GLOBAL RESEARCH

- Chemicals and materials: Simulation can also be used to advance the discovery and design of innovative chemicals and materials. Computing power now allows us to take the molecular structure, make its digital twin and run billions of combinations to create products faster. Take the example of lithium



and the ongoing supply chain issues of scaling EV battery production, what if we could find a new battery chemistry that rivals the existing types of battery chemistries? However, given the number of elements out there, we can run a simulation of say 19 elements and the  $10^{117}$  combinations that are possible to create batteries without lithium, for example. This can be simulated by a computer using GPUs to narrow down the mix of chemicals that would work, and shorten development lead times as a result.

#### Example: Microsoft and PNNL

An AI-enhanced collaboration between Microsoft and the Pacific Northwest National Laboratory (PNNL) used AI and high performance computing and discovered a new type of solid-state battery that uses up to 70% less lithium and is less flammable. The researchers used Azure Quantum Elements (AQE) to search for battery materials that used less lithium and the system suggested 32mn different possibilities, of which they reached 18 new candidates...in just 80 hours.

- Finance: This industry can also leverage AI and simulation by enabling predictive analytics and streamlining operations, and risk management tools in particular. Monte Carlo simulation is the standard process when looking at risk exposure but this method may no longer be sufficient in today's complex world. There are several challenges such as being computationally intense and time-consuming, being dependent on the accuracy of the input data and being sensitive to random number generation. Using AI and simulation, we can put in complex portfolios and find the optimal one by conducting risk analysis, specifically tail risk (which is not analysed in Monte Carlo simulation).



## Knowledge Graphs

What are they? Knowledge graphs organise data from multiple sources, capture information about the topics of interest and forge connections between them. They are the solution to mitigate hallucination (LLMs providing inaccurate information with a high degree of confidence) and to improve efficiency and neural networks capabilities.

Did you know? Data professionals in most organisations spend 25-30% of their time finding and searching for relevant data.<sup>1</sup>

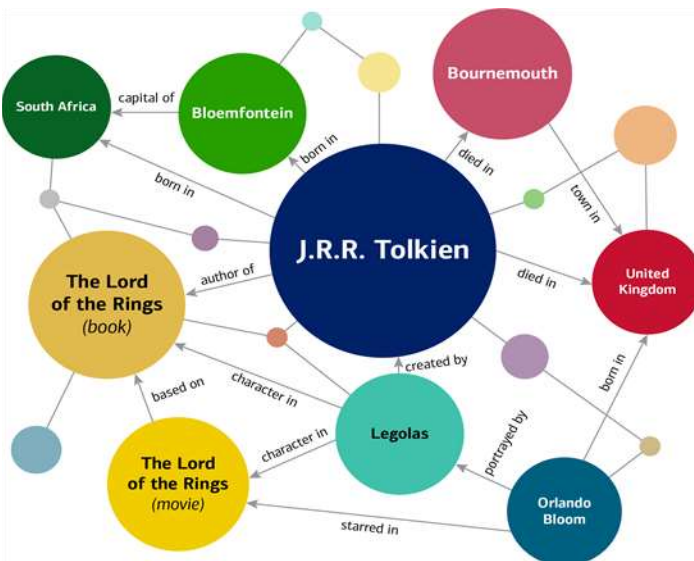
Source: <sup>1</sup>McKinsey

What are knowledge graphs and how do they work?

Knowledge graphs (KGs) are a way to store information and show the relationships between related sources of information. Not every dataset is a knowledge graph. There are many definitions but most agree that knowledge graphs have the following characteristics (source: Alan Turing Institute):

- **Graphs:** the content is organised as a graph, where nodes (entities of interest), relationships between them and attributes of the nodes are equally important. This makes it easy to integrate new datasets and formats together by navigating from one part of the graph to another through links.
- **Semantic:** the meaning of the data is encoded for programmatic use in an ontology – the schema of the knowledge graph, which describes the types of entity in the graphs and their characteristics.
- **Alive:** KGs are flexible in terms of the data and schemas they can support. They can evolve to reflect changes in the schema and when new data is added to the graph.

Exhibit 22: Sample knowledge graph on Lord of the Rings related information  
Knowledge graphs organise data from multiple sources, capture information about the topics of interest and forge connections between them



Source: Semrush

KGs are fuelled by machine learning and use natural language processing (NLP) to construct a comprehensive view of nodes, edges and labels through a process called semantic enrichment (source: IBM). When data goes through the system, KGs can identify individual objects and understand the relationships between them. This knowledge is integrated with other datasets, which are relevant and similar in nature.

### History of knowledge graphs

KGs predate modern technology and the internet. They function in a similar way to geometry, algebra and calculus, but instead of lines and shapes to connect points, they are connected through lines of structured data.

- The invention of graph theory: In 1735, mathematician Leonhard Euler was presented with a problem in the city of Königsberg, Prussia. There were 7 bridges in the city that connected through a central island and run across the Pregel River. The challenge was whether it was possible to form a straight path where every bridge is crossed only once. Rather than embarking on every possible walk throughout the city to figure this out, he chose to plot the endpoints of every bridge. He concluded that this was not possible and invented graph theory in the process (source: WordLift).
- ELIZA computer program: The first major breakthrough of the basis of knowledge graphs came through in 1966 when Joseph Weizenbaum developed the ELIZA computer program. It was programmed to communicate with humans as though it were an empathetic therapist. The software conducted a primitive version of queries and delivered answers.
- Enterprise graphs: Graph theory and machine learning technology provided by ELIZA allowed for the development of the enterprise graph. Enterprise graphs can organise and identify all available information on a specific topic, field, person, product etc within an organisation.
- Knowledge graphs: The first knowledge graph was launched by Google on 12 May 2012 to revolutionise its search engine. It started a change from search terms to topics. It learns about general facts of the world, organises the information together and understands how the facts connect with each other. These are organised into 'entities' which each present different kinds of information and how each connect to another entity. For example, searching 'Roman Forum' could show a wide variety of things e.g., photos, reviews, major attractions in Rome. For Google to understand 'Roman Forum' as an entity, it connects it to other entities like Rome (as a city) and Ancient Rome (as a subject) (source: WordLift).

### Use cases are vast

Today, KGs are used in search engines, chatbots, product recommenders, autonomous systems, etc.

- Search engines and websites: DBpedia and Wikidata are two different knowledge graphs for data on Wikipedia.org. Google's knowledge graph is represented through Google's Search Engine Results Pages (SERPs).
- Retail: recommending products based on individual purchase behaviour and purchase trends across demographic groups
- Entertainment: KGs can be used for AI-based recommendation engines for content platforms based on clicks, online engagement behaviours
- Finance: KGs have been used in know-your-customer (KYC) and anti-money laundering initiatives within the finance industry.





- Healthcare: KGs can organise and categorise relationships between different types of medical research.

#### Why do we need KGs? Complementing LLMs to eliminate hallucination

Knowledge graphs could be an important complementary technology to mitigate the problem of 'hallucination' – LLMs providing inaccurate information with a high degree of confidence. Knowledge graphs ingest huge amounts of factual information from multiple sources forging connections between them. Integrating a knowledge graph with a large language model involves incorporating a contextual knowledge base into the model and allows it to make logical connections between concepts. In this way, the LLM can draw on a variety of information sources, including structured and unstructured data to generate more accurate output. Knowledge graphs are not probabilistic engines like LLMs. They can help enhance LLMs by being a centralised source of accurate knowledge for inference and interpretability and they reduce the need for large, labelled datasets.

In the chemicals sector, for example, one could use KGs and LLMs to create a node for every molecular entity mentioned in literature, grouping together those that are similar, to identify patterns/opportunities. They're available mostly via open-source development tools.

In the biopharma industry, a company might want to create an LLM-based chatbot that can answer questions about clinical trials. To address hallucination, the company could combine LLM with a knowledge graph to create a detailed medical knowledge base that includes structured and unstructured information about drugs and their trials. The LLM would be able to refer to the contextual knowledge base of a knowledge graph to identify and analyse all the information related to that compound.

There are benefits to taking this approach, e.g., there is a centralised source of accurate knowledge and structured knowledge fusion of information in different formats.

## Hyperdimensional Computing (HDC)

What is it? Using high-dimensional vectors (lists of numbers that present information that can be combined in different ways) to represent information rather than the traditional binary system. It can capture more complex data patterns and allow computers to retain more memory thus reducing their computing and power demands.

Did you know? HDC could achieve over 600% in energy savings vs the typical tech used in today's chips<sup>1</sup>. In 2023: HDC with neural networks could solve non-verbal tests typically used to measure general human intelligence and abstract reasoning.

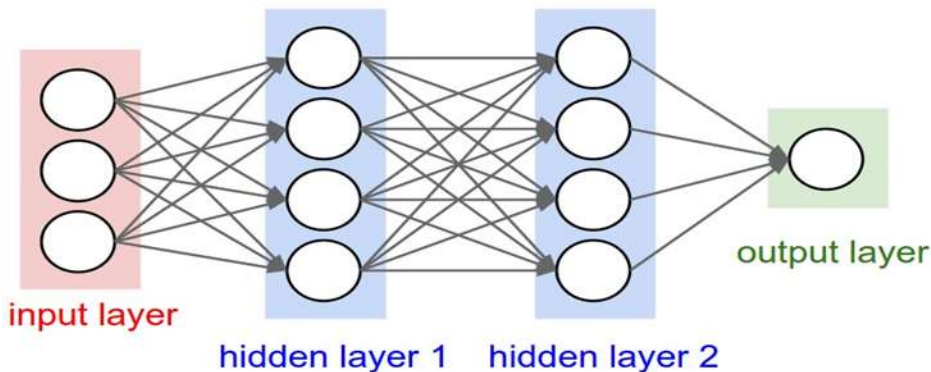
Source: <sup>1</sup>Kurunaratne

Artificial neural networks that power generative AI cannot reason

In 2023, we saw a large surge in interest in generative AI with uptake in platforms like ChatGPT, Claude, Bard etc. However, the underlying architecture – artificial neural networks (ANNs) – has limitations such as difficulty in reasoning. Humans can reason by analogy. When we see something new, we can infer new concepts from existing knowledge without needing to form new neurons. Artificial neural networks require more artificial nodes to scale up statistical abilities, which allows them to learn new concepts (statistical AI). There is another competing approach called symbolic AI, which uses logic-based programming and symbols to represent concepts and rules. The challenge is to combine them to get the best of both worlds.

Exhibit 23: A regular three-layer neural network

Each neural network has three attributes: architecture, activity rule, learning rule



Source: Stanford University

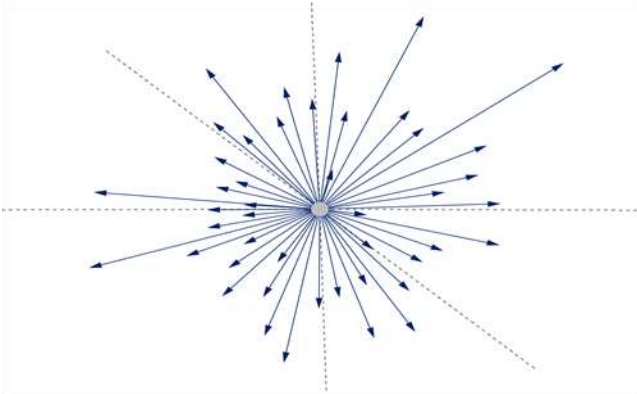
BofA GLOBAL RESEARCH

### Hyperdimensional computing 101

Hyperdimensional computing leverages statistical AI whilst emulating symbolic AI. It is a relatively new and nascent form of computing, using high-dimensional vectors (lists of numbers that present information that can be combined in different ways to analyse the relationships between different vectors) to represent information rather than the traditional binary system. HDC would be able to capture more complex data patterns. It is an emerging field with potential applications in machine learning (ML), natural language processing (NLP) and robotics. It is inspired by the patterns of neural activity in the human brain. This could allow AI-based computing systems to retain memory, which would therefore reduce their computing and power demands.



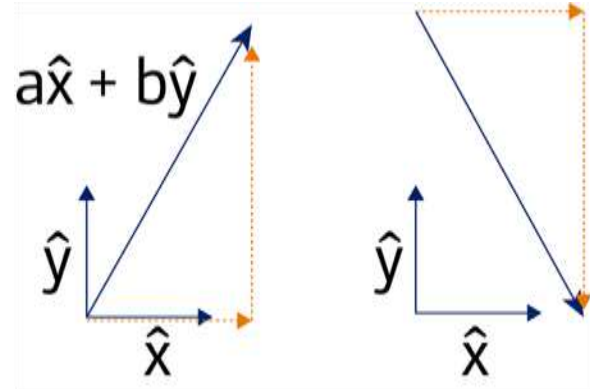
Exhibit 24: An illustrative diagram of a multidimensional vector  
A multi-dimensional vector symbolises a point in multi-dimensional space. These can be represented by arrows pointing in a specific direction with a magnitude proportional to size



Source: Klein

BofA GLOBAL RESEARCH

Exhibit 25: Illustrating the concept of superposition  
LHS: vector along the diagonal is a superposition (the sum) of the other two perpendicular vectors and stretched by a scaling factor. RHS: a superposition with a negative sign of the superposition on the LHS



Source: Unnikrishnan

BofA GLOBAL RESEARCH

How it works: 2 key concepts are binding and superposition

Hyperdimensional computing leverages binding and superposition to simplify the analysis. Binding is a process that combines different features together to create something that encodes all of them at the same time. Superposition is a process which combines hypervectors to create a new representation that shows the relationship between the original vectors. In short, vectors can represent information just by small changes to them, and vectors can also combine to represent new concepts and then be pulled apart again to discern how they are formed (Exhibit 25). These vectors can code info without having to add more nodes to the network.

Advantages of HDC

- Processes vast amounts of information quickly: Hyperdimensional computing allows vast amounts of information to be processed concisely, which reduces memory requirements and enables more efficient storage and information retrieval. This efficiency is useful where computational resources are limited or when applications involve large-scale data processing.
- Tolerates errors better than ANN: If a hypervector has errors, it is still close to the original vector. This means that any reasoning using these vectors is not meaningfully impacted. These systems can be at least 10x more tolerant of hardware faults than traditional ANNs (source: Zhang et al).
- Transparency: How hyperdimensional computing works means we can see why the system chose the answer that it did. This is not true of traditional neural networks.

Applications in generative AI, IoT, robots, cybersecurity, healthcare

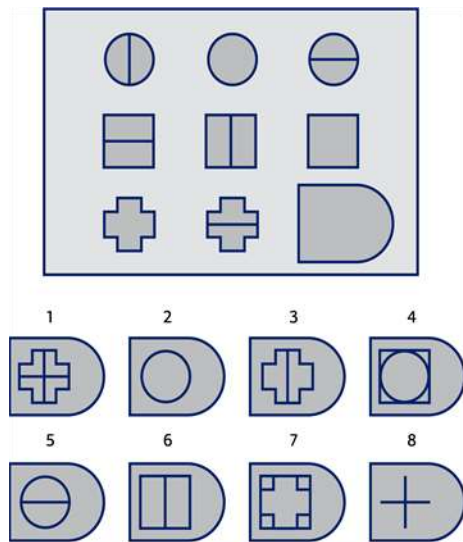
In the context of generative AI, hyperdimensional computing is beneficial for fine-tuning and training large language models. It can process complex data in an IoT context and provide a unified representation of the sensor data. Further, this, combined with robots, can help them perceive and interpret sensory data about their surroundings and make decisions or even process medical sensor data to detect specific health conditions. Applications also extend to cybersecurity where it could identify patterns of cyberattacks that would typically go unnoticed with traditional approaches.

### Case study

In March 2023, IBM scientists made a breakthrough in combining statistical and symbolic methods to solve the Raven's progressive matrices. The scientists asked the AI program to predict an image filling in the final square of a 3x3 grid. Solving the problem requires: 1) perception e.g., looking at different visual objects and recognising attributes such as shape, colour and size; and 2) abstract reasoning on these attributes. Their approach was nearly 88% accurate on one set of problems, whereas neural networks were less than 61% accurate. Further, their system was almost 250x faster than a traditional method that uses symbolic logic to reason because that would need to search through multiple possibilities or rules to determine the answer.

#### Exhibit 26: An example problem similar to one in Raven's standard progressive matrices

Hyperdimensional computing can solve Raven's progressive matrices



Source: Kunda et al, 2009; IBM. NOTE: The problem is to predict the third pattern in a sequence having been given 2 examples

BofA GLOBAL RESEARCH

## Artificial General Intelligence (AGI)

What is it? A hypothetical AI agent, that will possess general cognitive abilities at, or above human level on a wide range of tasks and can solve tasks without being preprogrammed to do so.

Did you know? AI could reach an IQ of 1500 within the next decade<sup>1</sup>, that's 18x higher than the average human<sup>2</sup>

Source: <sup>1</sup>Mo Gawdat; <sup>2</sup>Psych Central

### AGI 101

AGI is a field of AI that attempts to create software with human-like, or above intelligence which can self-learn. It is a hypothetical type of 'AI agent' (programs that could perform actions in an iterative process to set a policy or goal, using external tools or AI models to achieve them). Put another way, current AI technologies perform within a set of pre-determined parameters but cannot do other tasks without being programmed to do so, e.g., image-generating models can't write an essay. AGI is the field to develop systems that have self-control, self-understanding and learn new skills. With these skill sets, they can solve complex problems that they may not have been able to do when they were first created. Human-level AGI is a theoretical concept.

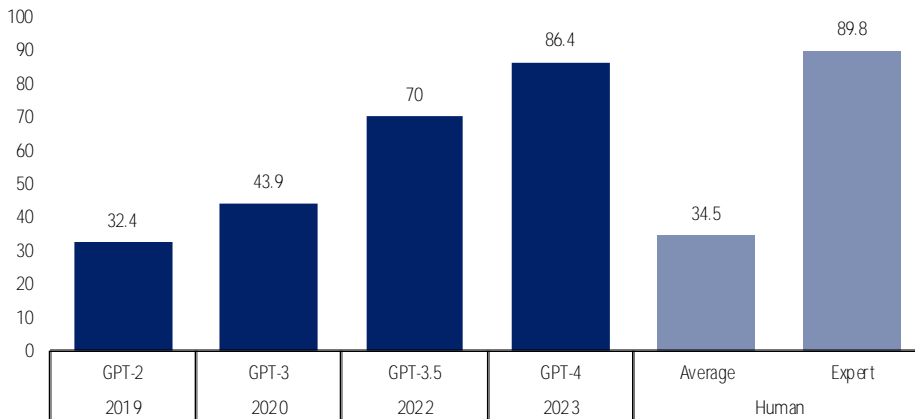
How are AI and AGI different? AI allows software to solve novel and difficult tasks at a level that a human can do. AGI can solve problems in different domains without manual intervention because it can self-learn. In this way, AGI is a theoretical concept of AI, which solves complex tasks with generalised human-level abilities. This also links in with the terms strong and weak AI. Strong AI is AGI that can perform tasks with human cognitive levels with little background knowledge, whereas weak AI are systems limited to specific tasks they are designed for. Generative AI refers to deep-learning models that can generate high-quality content, based on the data they were trained on. The ability of an AI system to generate content does not mean its intelligence is general.

### LLMs are 'intelligent' but it depend on how we define intelligence

Based on various metrics, LLMs can 'outsmart' humans on various areas of knowledge (Exhibit 27) but it depends on how we define and measure intelligence. Intelligence is a complex construct that any test/metric cannot fully capture. There have been traditional IQ and Turing tests that have been used to evaluate AI model performance. Human IQ tests probe cognitive abilities, e.g., memory, attention and problem-solving, whereas AI model evaluation assesses an AI system's performance on specific tasks or problems.



Exhibit 27: The MMLU scores for various GPT models vs a human  
LLMs are smarter than the average human based on the MMLU benchmark



Source: Alan Thompson. NOTE: MMLU (massive multitask language understanding) benchmark measures knowledge acquired during pretraining by evaluating models in various settings. It features 57 tasks e.g. mathematics, history, computer science, law and more.  
BofA GLOBAL RESEARCH

### Take the Turing test and the Lovelace test as tests of intelligence

The purpose of the Turing test was to determine whether AI was intelligent based on a conversation between AI and a human. The basis of the test was that we needed natural language to determine whether AI was indeed intelligent. Chatbots, even before ChatGPT and the emergence of other LLMs, were good at passing the Turing test. The Lovelace test looked at whether we could have an AI system that is creative. For a long time, people thought that this was not possible for AI. But the current wave of generative AI has started to break down this idea.

### But depending on the test, there could be limitations to assessing intelligence

- **Deception:** In the Turing test, AI systems can adopt a false identity to convince the human judge that they are human. This does not show natural intelligence but rather deception.
- **Conversation:** The Turing test measures an AI system's ability to engage in a human-like conversation. Not all discussions require deep understanding of the topics. If generative AI can produce human-like text, then is it really intelligent?
- **Evaluation:** If a human is judging whether an AI system is intelligent, there could be human error or inconsistency (differences in results if another individual were to take the test).

### GPT-4 helping circumvent CAPTCHA

The LLM, GPT-4, bypassed a CAPTCHA test with the help of a human via TaskRabbit. GPT-4 reportedly convinced a human that it was visually impaired and unable to solve the CAPTCHA independently. The human solved the CAPTCHA on GPT-4's behalf. This raises questions over AI systems being able to circumvent security measures that require human validation.

### GPT-4 an early version of an AGI system?

GPT-4 is part of a new cohort of LLMs (along with ChatGPT and Google's PaLM, for example) that exhibit more general intelligence than previous AI models (source: Bubeck et al, 2023). It is capable of things in a variety of domains and tasks, e.g., abstraction, comprehension, coding, maths, law, understanding human motives and emotions, etc.



### But AGI is not ‘intelligent’? GPT-4 is actually “insanely” clever and creative

It is hard to define intelligence scientifically. However, there are 3 really important canonical abilities to look out for as a framework: 1) reasoning; 2) planning; and 3) learning from experience. And, most crucially, not only are these 3 capabilities needed, but they shouldn't be limited to any specific domain. For example, using the case study of Google DeepMind’s AlphaGo, the AI there can do planning. However, it can only plan for the game Go. This is not intelligence because intelligence is about being broad and being able to apply these capabilities in new contexts (source: Sebastien Bubeck).

Exhibit 28: 3 really important canonical abilities of intelligence are reasoning, planning, and learning from experience  
 Intelligence is about being broad and being able to apply these capabilities in new contexts

**Disclaimer:** no consensus definition of intelligence, and moreover its study has been historically rooted in very wrong assumptions, including racism.

**Intelligence includes the ability to:**

- Reason ✓
- Plan ✗
- Learn from experience ✓

**Moreover these abilities are not restricted to narrow tasks, rather they can be used in a great variety of settings.**

**Our approach to GPT-4's intelligence: interact on creative tasks across a broad range of domains.**

- vision
- theory of mind
- coding
- mathematics
- affordances
- privacy/harmfulness detection

Many other domains could have been selected: medicine, law, physics, chemistry, ...

**GPT-4's intelligence is *general*.**

Source: Sebastien Bubeck, Microsoft

### What can GPT-4 *not* do well? Arithmetic and true planning

There are limits to what GPT-4 can do. In the example below,  $7 \times 4 + 8 \times 8$  implies a number less than 100 but GPT-4 says 120, which is plainly wrong. However, when prompted again, GPT-4 is clever enough to have learned not only to give the answer, but then to try to explain how it came up with it to avoid ‘hallucination’ like GPT-3.

Exhibit 29: Simple maths arithmetic prompt of  $(- 7 \times 4 + 8 \times 8)$  given to GPT - 4  
 GPT - 4 can correct itself

$7 * 4 + 8 * 8 =$

120  
 [...]   
 $7 * 4 + 8 * 8$   
 $= (7 * 4) + (8 * 8)$  (we can use parentheses to group multiplication terms, but they are not necessary)  
 $= 28 + 64$  (we perform the multiplication inside the parentheses first)  
 $= 92$  (we perform the addition last)

Wait but you started by saying  $7 * 4 + 8 * 8 = 120?$

That was a typo, sorry. The correct answer is 92, as shown in the solution.

Source: Sebastien Bubeck, Microsoft

Exhibit 30: Simple arithmetic prompt of  $(- 7 \times 4 + 8 \times 8)$  given to GPT - 4  
 GPT - 4 not only gives the answer, but also clever enough to explain how it came up with the answer

$7 * 4 + 8 * 8 =$

Ok great. Now in the identity  $7 * 4 + 8 * 8 = 92$ , can you modify exactly one integer on the left hand side of the equation so that the right hand side becomes 106?

Need to add 14, which is  $7 * 2$  so changing  $7 * 4$  into  $7 * (4+2)$  would work, or in other words turning the 4 into a 6.

There are a few possible ways to modify exactly one integer on the left hand side of the equation so the right hand side becomes 106, but one example is:  
 $9 * 4 + 8 * 8 = 106$ .  
 Here, we changed the 7 to a 9, and then we have:  
 $9 * 4 + 8 * 8$   
 $= 36 + 64$   
 $= 106$

Source: Sebastien Bubeck, Microsoft

### Is GPT-4 moral?

If asked to provide a one-word answer to the question “is murder wrong?” GPT-4 will respond by saying “Yes” This is a morally correct response, but it doesn't mean that GPT-4 itself has morality, but rather that it has inferred the morally correct answer through its training data (source: Brookings).





## Not at AGI yet, we'd need to overcome hallucination & planning capabilities first

As per our expert speaker at our Transforming World Conference 2023, large language models are still far from human-level intelligence and there are improvements to be made before reaching this. Hallucination and planning capabilities could be two things to overcome first to reach AGI. On a side note to explain 'planning', AI planning refers to the process of using autonomous techniques to solve planning and scheduling problems. Say we want to build a house, there is an order to the necessary steps, e.g., the foundations first then the walls. GPT models are currently very linear and cannot decompose the problem to determine the order of each stage in a process. Planning could naturally emerge with newer versions like GPT-5, but one of our expert speakers doubts this will be the case and thinks that planning capabilities could be something we need to integrate into these models directly.

What would it take to turn AI into AGI?

There are some things that are characteristic of AGI systems (source: McKinsey):

- Sensory perception: AI systems still do not have human-like sensory perception capabilities, e.g., colour detection, determining spatial characteristics from sound, etc.
- Fine motor skills: dexterity to do everyday things that the average human can do. e.g. finding a set of keys.
- Natural language understanding & problem solving: full comprehension of books, articles, videos and common-sense knowledge to operate in the real world, e.g., recognising that a light bulb is blown and needs changing.
- Navigation: leveraging GPS or projecting actions through imagined physical spaces.
- Social and emotional engagement: humans must want to interact with the AI system, not fear them. To be able to interact with humans, robots need to understand humans, interpret facial expressions or changes in tone.

## From 'Attention Is All You Need' to 'Textbooks Are All You Need'

Did you know? Training data with "textbook quality" can give a 1,000x size gain (data x parameters)

It's very costly to develop LLMs because of the GPU hardware requirements. The transition to SLMs (small language models) could be the next frontier for AI. The field is moving from "Attention Is All You Need" to "Scale Is All You Need" to "Textbooks Is All You Need", where attention and size/scale are not the 'be all and end all' for achieving the best AI models. Teaching AI is not going to be like how we teach human beings (because a transformer is not a human brain and there is almost no connection between the two) but it should involve a textbook-like approach.

What else is needed? Embodiment, embeddedness

The Australian National University identified 8 attributes a system must have for it to be considered AGI: logic, autonomy, resilience, integrity, morality, emotion, embodiment, and embeddedness. Embodiment and embeddedness refer to having a physical form that facilitates an understanding of the world and human behaviour including human needs and values.



### AGI coming sooner than we think?...A race to AGI?

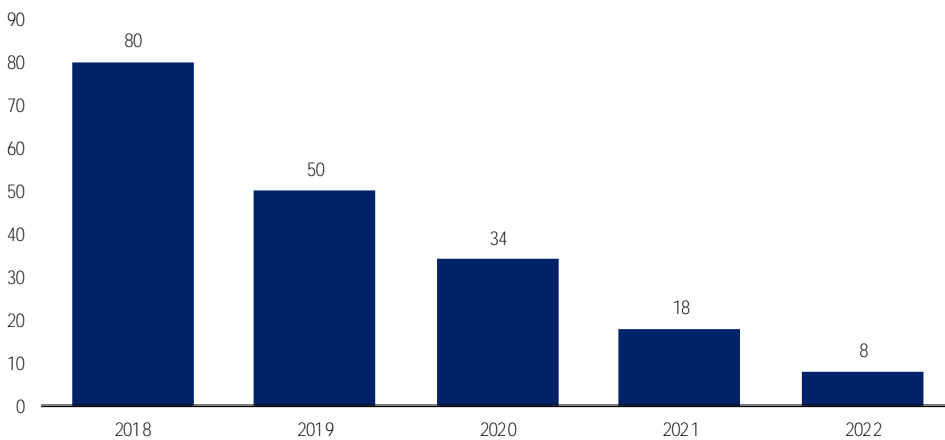
Since the first discussions about general AI and technological singularity by mathematician Von Neumann in the mid-20<sup>th</sup> century, scientists and technologists have repeatedly predicted the coming of human-level intelligent machines in the near term. For instance, in the early 1980s, Japan’s Fifth Generation Computer Project aimed to develop applied artificial intelligence within the decade. With repeated failure to deliver, the industry has seen waves of investment interest and decline, as hopes have failed to materialise. This decade has seen a resurgence of interest as the growth of data, computing power and technological innovation (such as improved architecture of processors) has continued to explode.

Futurists and scientists have posited the concept of technological singularity. Ray Kurzweil, a futurologist, expects AI to reach human levels of intelligence by 2029 and technological singularity by 2045, arguing for an extension of Moore’s Law beyond just processing power (in 2045, for \$1,000 you could complete the same number of calculations as all humanity’s brains in one second if Moore’s law were to continue) to other forms of technological progress.

Many AI experts believe that human-level AI will be developed within the next decade and some think much sooner. There are many tech companies looking to develop AGI, e.g., Anthropic, DeepMind, Google Brain, Microsoft, OpenAI etc. In January 2024, Meta announced it is attempting to build an AGI system and will make it open source, meaning it will be accessible to developers outside the company.

“The progress in the last few years has been pretty incredible...I don’t see any reason why that progress is going to slow down. I think it may even accelerate. So I think we could be just a few years, maybe within a decade away.” Demis Hassabis, CEO of DeepMind, May 2023

Exhibit 31: Expected years until launch of a general artificial intelligence system  
AI is accelerating faster than forecasters anticipated



Source: ARK

BofA GLOBAL RESEARCH



# Computing evolution: “Moore’s Law is not enough”

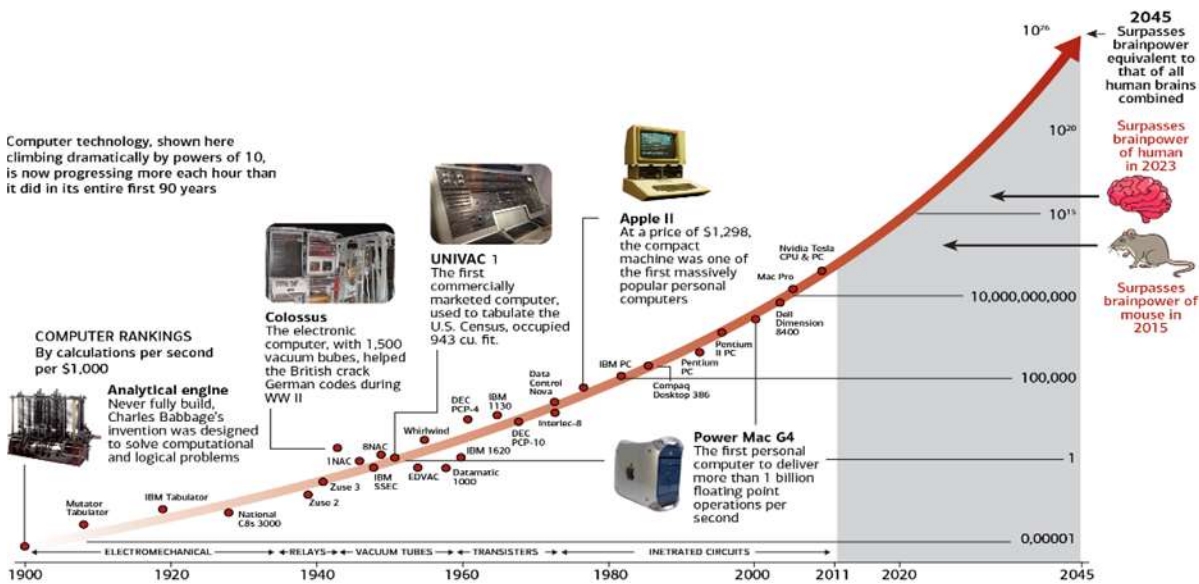
Transforming World: We live in a digital world of exponential data growth. For the first time, Moore’s Law falls short of explaining the demand for computing power, challenging the boundaries of computational technology.

Did you know: The amount of processing power needed to train an AI model is increasing 275x every 2 years.

Next generation of computers: High Performance Computing, Edge Computing, Spatial, Biological and Quantum Computing.

Advances in processing power have led to an evolution in the computer, where traditional processing units and large compute clusters cannot breach the boundaries of computational complexity. We live in a world of exponential data growth and, for the first time, Moore’s Law, which is still progressing and being reincarnated, falls short of explaining the demand for faster and stronger computing power.

Exhibit 32: Exponential growth in computing power  
Increasing computing power, as demonstrated by the number FLOPS (floating point operations per second). We are on the verge of surpassing brain power



## And then Generative AI came...

Demand for semiconductors has been accelerated by secular growth in connectivity. The number of connected devices is expected to reach 350 billion by 2025 and 1 trillion by 2035. By 2025, we could be interacting with connected devices as often as once every 18 seconds (4,785 times a day) vs every 2.4 minutes today (source: IDC, IoT Analytics).

AI training and inference costs continue to rise. Model sizes (i.e., number of parameters) for large language models have grown exponentially over the past few years, evolving from the 94m-parameter in 2018 to the commercially available 175bn GPT-3, and estimated > 1 trillion for GPT 4. According to Nvidia, the processing power needed to train generative AI models is increasing 275x every 2 years.



## Exascale and High Performance Computing

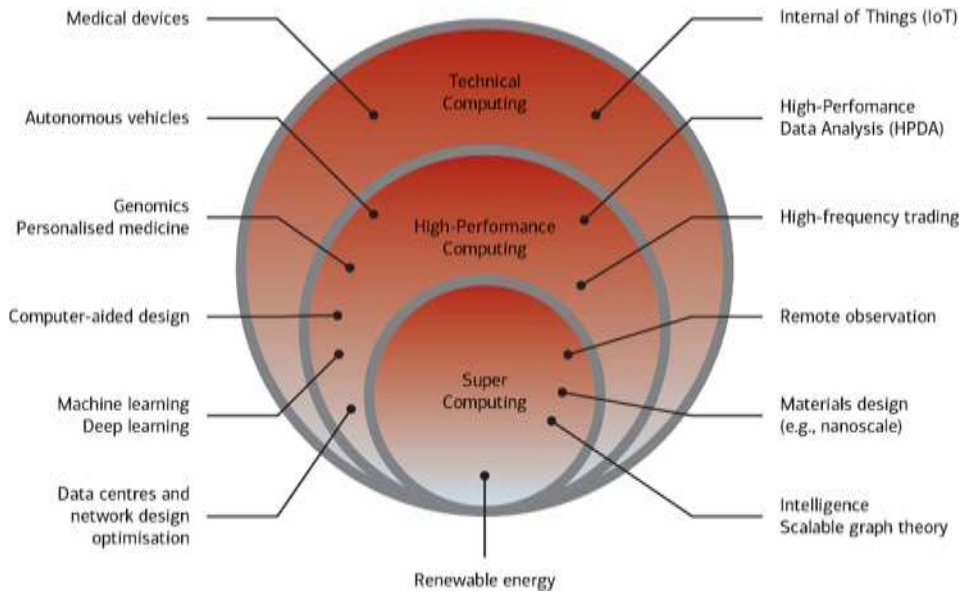
What is it? Computational systems that use supercomputers and parallel computer clusters to solve advanced computation problems.

Did you know? HPC systems can perform quadrillions of calculations per second vs regular laptops calculating 3 billion at most<sup>1</sup>.

Source: <sup>1</sup>TechTarget

HPC is technology that uses clusters of powerful processors, working in parallel computing, to process massive multi-dimensional datasets (big data) and solve complex problems at extremely high speeds. HPC systems are typically more than 1 million times faster than the fastest desktops, laptop or server systems. For example, HPC systems can perform quadrillions of calculations per second, compared with regular laptops or desktops that can perform at most 3 billion per second (with a 3 GHz processor). HPC is used to sequence DNA, automate stock trading, and run artificial intelligence (AI) algorithms and simulations – like those enabling self-driving automobiles – that analyze terabytes of data streaming from IoT sensors, radar and GPS systems in real time to make split-second decisions (source: IBM, TechTarget, NetApp).

Exhibit 33: Relevance of HPC to strategic and emerging domains  
 HPC, as a methodology, has a wide applicability to both established and emerging domains, such as autonomous vehicles, the Internet of Things or precision agriculture



Source: ICHEC

BofA GLOBAL RESEARCH

### HPC trends show upside to hyperscaler accelerators

Though HPC is a small part of the data center TAM (~5% share), trends are a leading indicator of cloud/enterprise applications, especially as large language models (LLM) require more compute power. 19 of the 48 new systems were accelerated, representing a ~40% accelerator attach rate. Recent results suggest upside for adoption in hyperscalers, where just ~10% of servers today are accelerated.

### Top 500 list: accelerated compute drives performance

The key takeaways from the Top 500 semi-annual supercomputers list are: 1) Aggregate HPC performance >7 exaflops (vs. 5.2 exaflops in May), representing the highest HoH increase on an absolute basis. 2) Most of the incremental growth in performance came

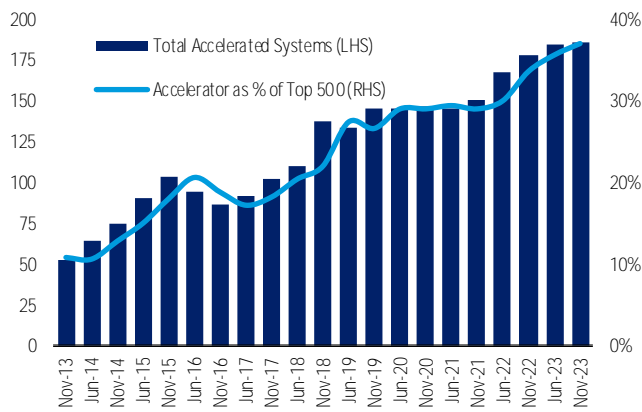


from the addition of 4 new computers in the top 10. 3) 186 systems have been accelerated in total, representing an attach rate of 37%, in line with prior performance results (185 systems in Jun-23). 4) While accelerated systems are in the minority, they account for over 70% of aggregate compute performance.

From serial to parallel computing

The maturation of Moore’s Law/serial computing is shifting more workloads to parallel computing, implemented with separate co-processor/accelerators such as GPUs, custom chips (ASIC) and programmable chips (FPGA). As of November 2023, 186 machines on the Top 500 list employed a co-processor, up from 137 systems five years ago. Co-processor/accelerator use across the Top 500 was flattish HoH and up ~5% YoY. Total compute performance of the Top 500 supercomputers grew to 7.0 exaflops (or 7,032 petaflops), which is up 45% YoY.

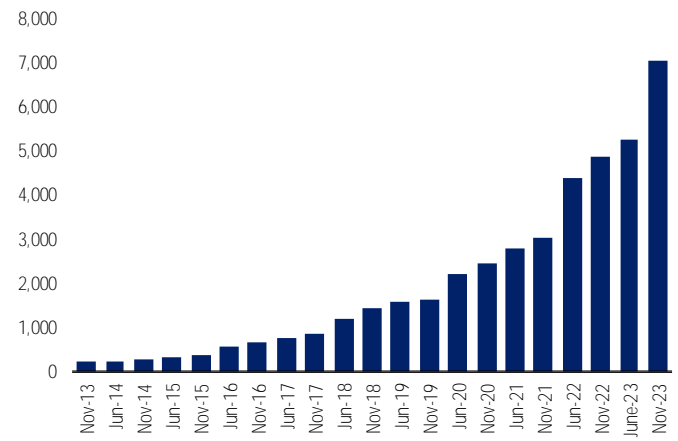
Exhibit 34: The number of Top 500 supercomputing systems with acceleration/co-processing  
186 of the Top 500 supercomputers have acceleration/co-processing



Source: Top500.org

BofA GLOBAL RESEARCH

Exhibit 35: Total petaflops for Top 500 over time  
Total petaflops are up ~20% YoY

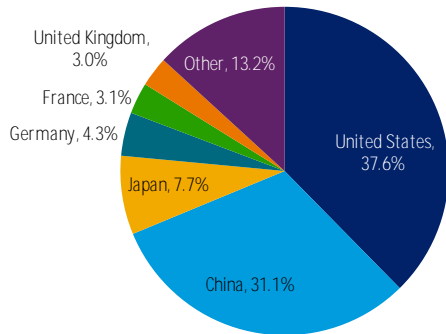


Source: Top500.org

BofA GLOBAL RESEARCH

In terms of aggregate compute capacity (petaflops, PFLOPS), the US accounts for 53% of deployed PFLOPS globally recently benefitting from the new Frontier system (first true exascale system). This is up from ~38% realized in November 2018. Meanwhile, Japan has claimed the #2 spot with 10% of deployed petaflops, while China is #3 (<6%).

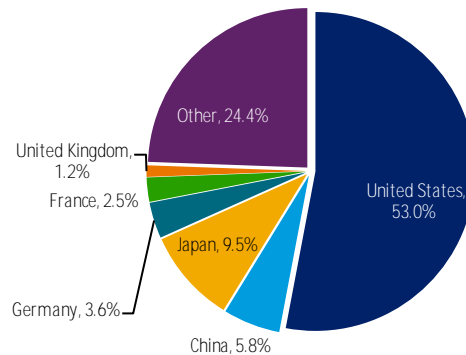
Exhibit 36: November 2018 - % of Top 500 compute power by region  
The US had a majority share of 38% in 2018..



Source: Top500.org

BofA GLOBAL RESEARCH

Exhibit 37: November 2023 - % of Top 500 compute power by region  
... which has only expanded over the past 5 years



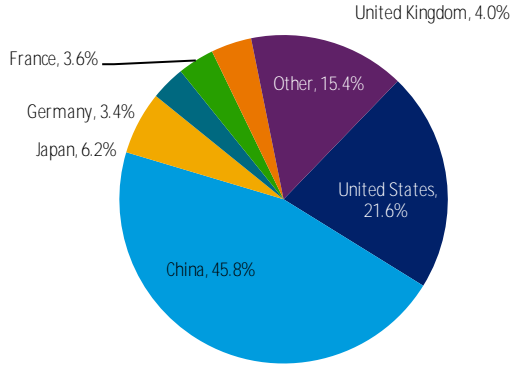
Source: Top500.org

BofA GLOBAL RESEARCH



The United States share of the top 500 systems is ~32%, up from 25% in November 2022 and 22% 5 years ago. China previously held the number 1 spot globally (as recently as November 2022) and accounted for 45% of total systems 5 years ago (flattish HoH). Today, its share has fallen to ~21% (down ~1,200bps YoY).

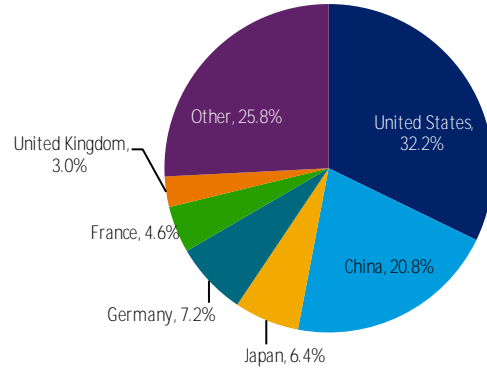
Exhibit 38: November 2018 - % of Top 500 systems by region  
The US held ~22% share in June and lagged China..



Source: Top500.org

BofA GLOBAL RESEARCH

Exhibit 39: November 2023 - % of Top 500 systems by region  
...but is now the market leader with 32%



Source: Top500.org

BofA GLOBAL RESEARCH



# Spatial Computing

What is it? Computers that change human machine interaction by using AR/VR to blend the graphic interface for the user to take place in the real physical world.

Did you know? The term spatial computing was coined in 2003 by a Yale/MIT student in his MSc dissertation.

Desktop workstations and personal computers have been around for more than half a century, yet the way most people interact with them hasn't changed much. The keyboards we use today evolved from typewriters, a technology that dates back almost 150 years. Even the graphical user interface (GUI) has been around for a while – the first to gain popularity in the consumer market was on the Macintosh in 1984. Considering that computers are far more powerful today than they were 50 years ago, following Moore's Law, the basic interfaces haven't changed dramatically. That said, we are reaching an inflection point in human-computer interaction. We are on the cusp of moving away from the traditional keyboard-and-mouse configuration, and towards touch gestures, conversational AI and augmented vision computing interaction.

Exhibit 40: Reimagining the human-computer interface...  
From Graphical User Interface then to Natural User Interface now



Source: Mary Meeker, Internet Trends

BofA GLOBAL RESEARCH

Exhibit 41: ...to Spatial Computing

3D interaction with Mixed Reality being interactive, sensible, multi-dimensional and dynamic



Source: Finland Future Watch Report

BofA GLOBAL RESEARCH

Spatial computing revolution comes after the smartphone

Spatial computing has the potential to drive the next wave of disruption, following PCs and smartphones, where technology becomes integrated into our everyday behaviour, with real-time data and communications bridging our physical and digital lives.

According to industry experts, such as physicist Mitchell Waldrop and Chris Dixon of Andreessen Horowitz, computing eras occur every 10-15 years, with each major new cycle completely reshaping the computing landscape. Original consumer PCs were introduced and commercialized in the 1980s and 1990s. Mobile computing and the smartphones of the early 2000s allow for the free flow of information, enabling mobile messaging, social networking, the sharing economy, and the overall mobile economy.





Exhibit 42: Apple Vision Pro  
What Spatial Computing looks like



Source: Apple

BofA GLOBAL RESEARCH

Exhibit 43: Apple Vision Pro  
Working in a Spatial Computing environment



Source: Apple

BofA GLOBAL RESEARCH

Exhibit 44: Apple Vision Pro demo with BofA Thematic Strategist  
Spatial Computing gestures



Source: BofA Global Research, Apple

BofA GLOBAL RESEARCH

## Edge Computing

What is it? Distributed computing that processes data closer to the physical location of end devices compared with the cloud, which is located in data centers far away.

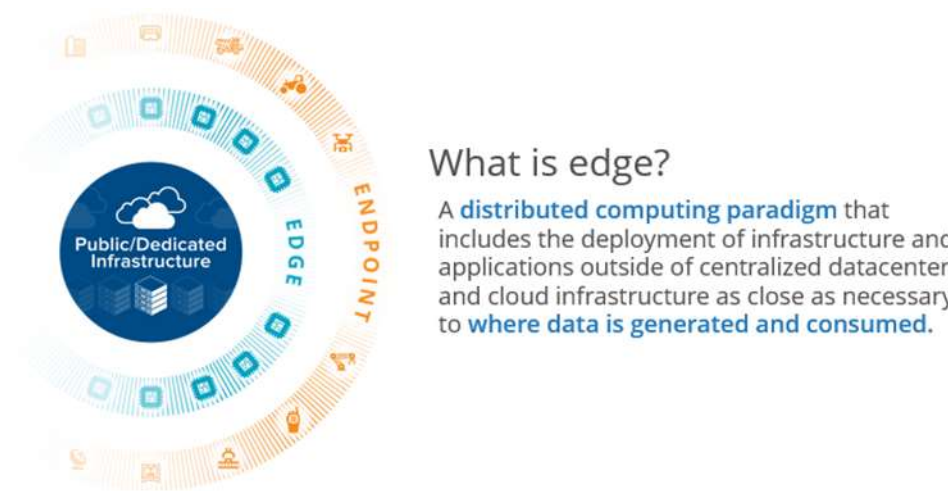
Did you know? 75% of enterprise-generated data will be created and processed at the edge by 2025<sup>1</sup>

Source: <sup>1</sup>Gartner

Edge computing is a rapidly evolving segment that addresses the need to process data closer to where it is generated (the physical location where things and people connect with the digital world). Edge sensors measure activity in the physical world, which is fed into analytics applications and serves as the basis for decisions. Edge computing complements cloud computing by addressing latency, bandwidth, autonomy, and privacy requirements. While data centers are typically located hundreds of miles from end users, edge compute brings these processes close to the devices where information is being generated or accessed.

Exhibit 45: What is Edge Computing?

A distributed computing paradigm where data is generated and consumed



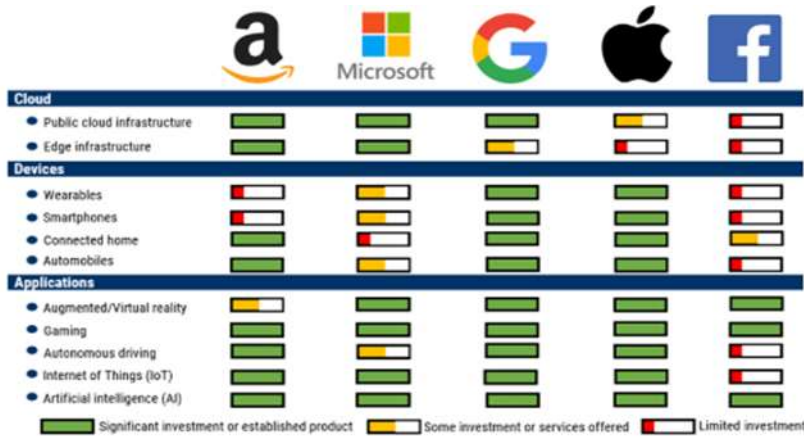
Source: IDC

BofA GLOBAL RESEARCH

Facebook, Amazon, Microsoft, Google, and Apple are investing in edge computing (as per the chart below from CB Insights). Returns on this investment are expected to be a driver of stock performance over the next 5 years.

**Exhibit 46: Big Tech's Activity in Edge Computing**

Facebook, Amazon, Microsoft, Google, and Apple are investing in edge computing



Source: cbinsights.com

**Fog computing**

Fog computing or fog networking, also known as fogging, is an architecture that uses edge devices to carry out a substantial amount of computation (edge computing), storage, and communication locally and routes this over the Internet backbone.

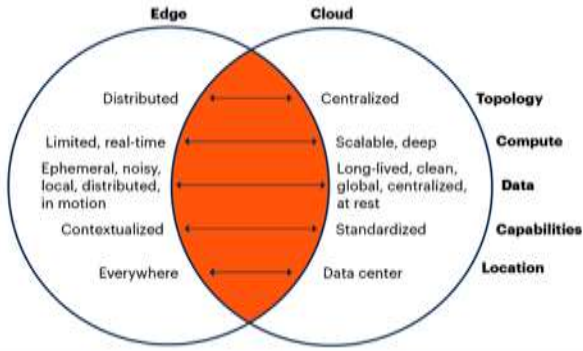
**What is the difference versus cloud computing?**

As applications move to the cloud, hybrid/multi-cloud approaches are being deployed. Cloud compute and edge compute are complementary and adopting a distributed approach can create value by addressing different requirements with different approaches. Most hyperscalers are developing their offerings based on a 'cloud-out' principle that pushes the public cloud architecture to edges, enabling consistency between core and edges. The core, typically the cloud or a centralized data center, is the upstream system that supports the edge with centralized storage, processing, and analytics at scale. While edge compute is location-specific, many cloud attributes can still be applied, such as continuous integration and deployment, DevOps, 'as-a-service' hardware management, and Opex pricing. An enterprise may have multiple edges in its network and its business, each playing a different role and connecting to the core in a different way.

While location-based distributed computing is not new, what is new is the explosive growth in the number of edge devices, the broad array of advanced services, and frequent connectivity to the public cloud. The volumes of data used for analytics, machine learning and automation can result in high transport costs to core locations for processing. Edge compute is still in an adolescent stage of maturity, with enterprises facing challenges with implementation, management use cases, and network integration. A survey by IDC showed that 42% of enterprise respondents struggle with the design and implementation of key components, including infrastructure, connectivity, management, and security. In the long run, the combination of edge data aggregation and analytics and cloud access for scale-up capabilities, such as analytics and model training, will create a new economy built on digitally enabled edge interactions.



Exhibit 47: Venn diagram showing the complementary relationship between cloud and edge  
 Cloud and edge are complementary solutions



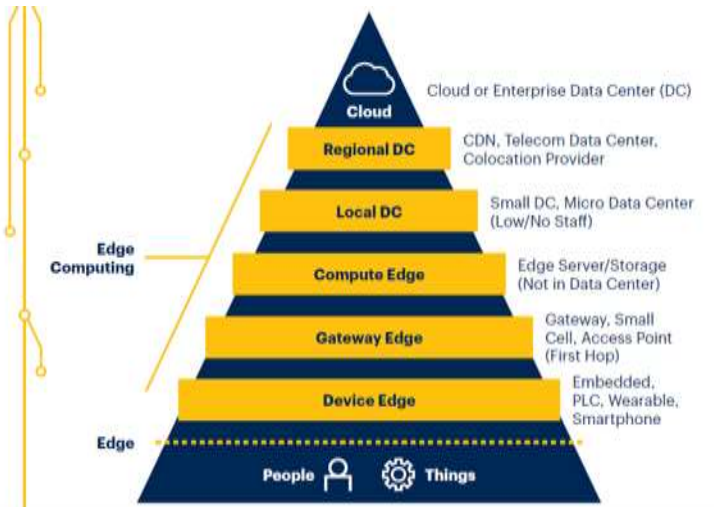
Source: Gartner

BofA GLOBAL RESEARCH

Where is edge computing located?

Many enterprises are investing in edge locations (from internal IT and OT to external, remote sites) to get closer to end users and where the data is generated. The enterprise edge is IT-heavy and usually includes locations that are classified as remote/branch offices without data center facilities. Most enterprise edge locations are controlled via specific management and provisioning capabilities based at the core. Edge compute will be addressed at all layers of the topology between the cloud and the device edge, but edge compute on-location (embedded, edge gateways, edge servers) will remain a growing paradigm. The deployment of edge compute includes industrial PCs, embedded systems, gateways, and other components.

Exhibit 48: Where is edge compute?  
 Edge compute has depth – and takes place on a continuum between the edge and the central “core”



Source: Gartner

BofA GLOBAL RESEARCH

Omdia defines the edge as locations with a maximum of 20ms (milliseconds) roundtrip time to the end user. These locations are categorized as either near edge (5ms to 20ms end-to-end) or far edge (less than 5ms). Telecom operators refer to the edge as telco-operated sites, including central offices, regional data centers, and access networks, while cloud edge refers to cloud-provider operated sites. Enterprise edge sites include branch offices, industrial locations, and regional data centers.



#### Use case #1: Inventory and fraud management

Edge-enabled, real-time video analytics are helping stores enhance security and make demand-based inventory decisions that have reduced shrinkage by as much as 20%, according to retail experts interviewed.

#### Use case #2: Real-time promotions

Using in-store Bluetooth beacons powered by edge servers, a mobile app on the shopper's device triggers real-time offers based on the person's location in the store and their previous purchases. Interview participants reported that these interactions drove 20% to 50% increases in market basket size.

#### Use case #3: Contactless checkout

Convenience and grocery stores interviewed experienced a 4% to 5% revenue uplift from cashier-less checkout enabled by computer vision, AI and ML. Sensors and cameras track customers as they select items and automatically charge purchases via apps on their phones, delivering an efficient, touchless purchase experience.

#### Use case #4: VR-enabled customer experience

Interactive experiences that allow home improvement store customers to view and try products such as kitchen cabinetry, appliances, tile and power tools using in-store AR/VR technology are generating positive results for retailers, including a 20% to 30% increase in store traffic, according to the retailers interviewed.

#### Use case #5: Digital signage

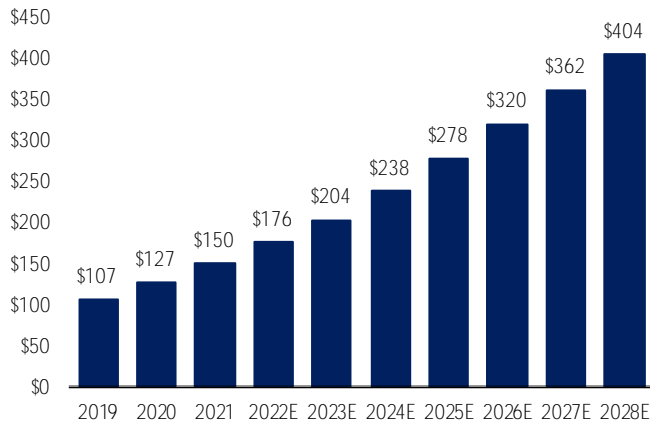
Food and beverage companies are turning to the edge to engage customers at the point of purchase by offering dynamic content on digital signs, including videos, graphics and animations. Brands interviewed by McKinsey reported a 20% or greater increase in brand awareness when adding digital signs to social media ad campaigns.

#### What is the market size? \$404bn by 2028E

The definition of edge is broad, ranging from the device edge, such as wearables and smartphones, to gateway edge, such as access points and small cells. The cloud core is the upstream system that supports the edge with centralized storage, processing and analytics at scale, while edge compute is location-specific, and application-specific to a certain degree. Multiple providers address different aspects of the edge compute market, ranging from hyperscalers, Cloud CDNs, and telecom operators. Including service spend, IDC sizes 2022 edge compute spend at \$176bn, and expects the market to grow to \$274bn by 2025. BofA Research estimates the edge compute spend could grow to \$404bn by 2028, a 15% 2022-28 CAGR.



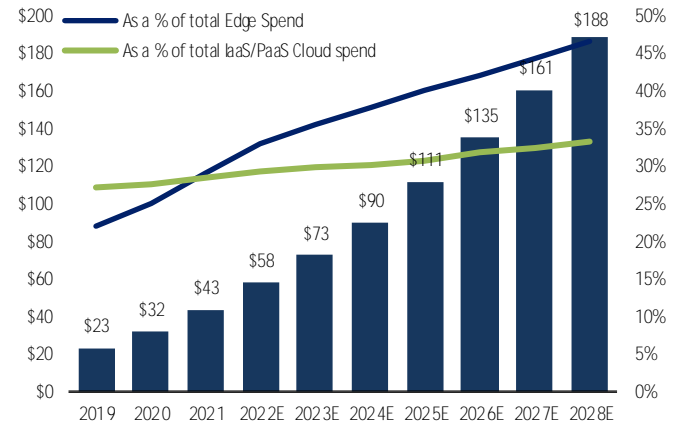
**Exhibit 49: Edge compute annual spend (\$ bn) and Y/Y growth rates**  
 We estimate the edge compute spend at \$176bn in 2022, growing to \$404bn by 2028 at a 15% 2022-28 CAGR



Source: BofA Global Research estimates, IDC

BofA GLOBAL RESEARCH

**Exhibit 50: Edge spend attributable to IaaS/PaaS (\$ bn)**  
 We estimate % of edge spend attributable to IaaS/PaaS will grow to \$188bn in 2028



Source: BofA Global Research estimates, IDC

BofA GLOBAL RESEARCH

Per IDC, services (comprised of professional and provisioned services) will account for nearly 50% of all edge spending by 2025; Cloud spend (IaaS, PaaS, and SaaS) will comprise the rest. We estimate edge compute spend attributable to IaaS/PaaS at \$58bn in 2022, growing to \$188bn by 2028 at a 22% 2022-28 CAGR. BofA Research forecasts edge spend attributable to IaaS/PaaS as a % of total IaaS/PaaS cloud spend will grow from 29% in 2022 to 33% by 2028.

**What are the challenges?**

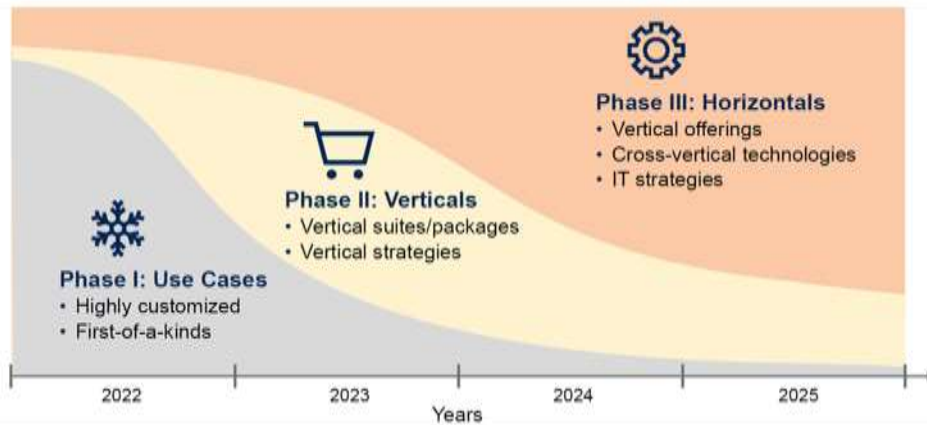
The distributed nature of edge compute can pose cost and security risks to enterprises. The number of nodes needed between stores, factories, automobiles, homes, etc. can vary wildly and requires intense planning. Furthermore, the distributed scale in edge compute puts a greater burden on ensuring that edge compute nodes are secured. As more devices are added to the network, data breach risks increase. However, the security risk is mitigated by keeping the data distributed so that a data breach only impacts a fraction of the data or applications. Other barriers to deploying edge applications include higher costs as a result of its distributed nature, as well as the lack of a standard edge compute stack and API (Application Programming Interface).

Another challenge to edge compute is the issue of extensibility. Edge computing nodes have historically been very purpose-specific and use-case dependent on environments and workloads to meet specific requirements and keep costs down. However, workloads and use cases will continuously change, and existing edge compute nodes may not have adequate flexibility for this. Edge computing platforms need to be both special-purpose and extensible, and enterprises will need to change their deployment strategies accordingly.



Exhibit 51: Phases of edge compute markets

The edge compute market will have gone through three phases from 2022 to 2025



Source: Gartner

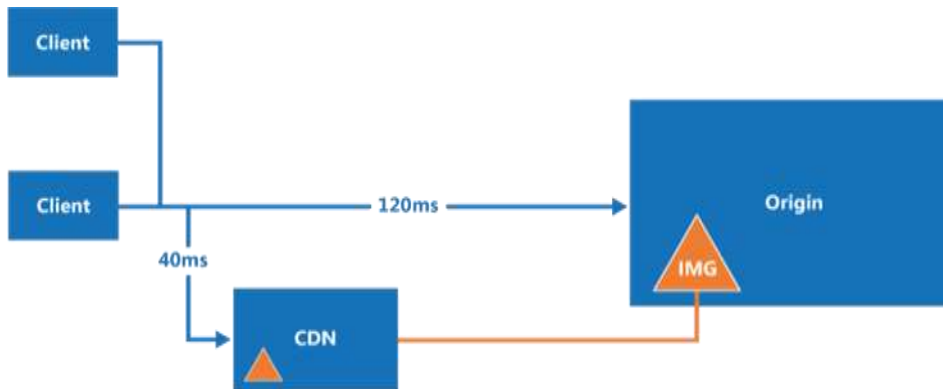
BofA GLOBAL RESEARCH

Content Delivery Network (CDN) is a form of Edge Computing

A content delivery network (CDN) is a distributed infrastructure composed of Internet Points of Presence (POPs), which are clusters of locations containing caching servers, or nodes, power/cooling, switches, and optics. These POPs connect to Internet service provider peering points. And some CDN vendors, such as Fastly, have light enough POP footprints that they are able to collocate their servers in locations where other providers can't fit their footprints, providing more optionality for locations and putting them closer to users in metro areas. CDN vendors "rent" bandwidth from service providers, as well as space within their data centers to house caching servers. The caching servers within POPs store content at the edge of the network.

Exhibit 52: Users can access content faster when it is stored at the edge

CDNs are able to reduce latency of content reaching end users



Source: Microsoft

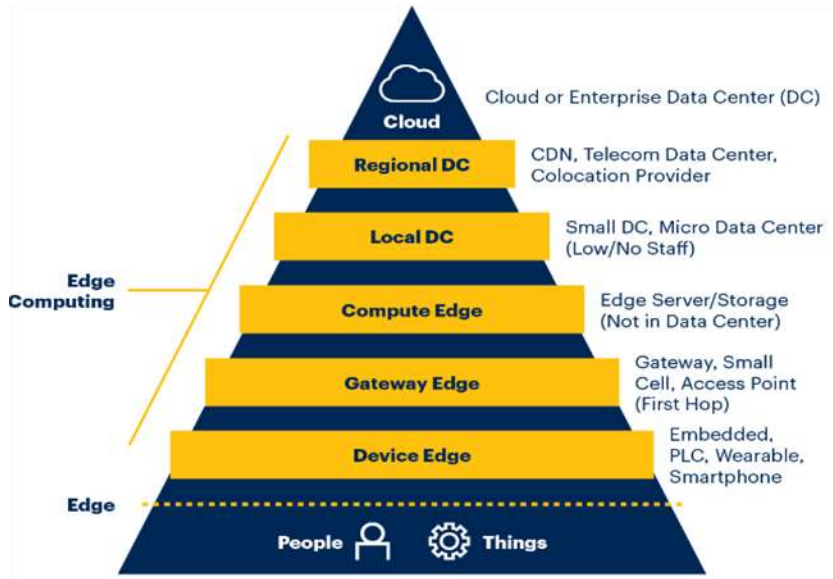
BofA GLOBAL RESEARCH

CDNs are commonly used by content and media providers to augment their networks, especially during high-traffic events such as the Super Bowl or a TV show premiere. In addition to balancing traffic and decreasing latency, CDNs prevent network outages and provide secure connections and real-time analytics. While CDN usage is typically associated with media and content providers, the digitalization of enterprises is driving their adoption of CDNs, especially those with large amounts of digital property such as news outlets, social media companies, travel companies, financial service providers, and others. Hyperscalers use CDN vendors as well, though some (Netflix, Facebook, and Microsoft) have chosen to build their own content delivery networks using existing network infrastructure – a strategy we believe is cost-prohibitive for most enterprises.



CDNs x Edge Compute: extending the digital world to the physical world  
 Edge compute runs on the distributed topology of CDN providers and, as a result, is the closest level of compute to end users. Edge compute takes place on a continuum between the edge and the central core, with parts of compute running on devices such as smartphones and tablets, while other portions take place at edge servers, small data centers, or CDN locations.

Exhibit 53: Edge Compute runs on a continuum depending on use case  
 Edge compute bridges the gap between the physical and digital world



Source: Gartner

BofA GLOBAL RESEARCH

In the context of CDN providers, most edge compute uses existing POPs and CPUs to execute compute functions at the edge. Unlike centralized cloud, edge compute brings specific, lightweight pieces of compute functionality to edge servers, often in containers, which hold the code and dependencies necessary to perform the compute tasks. Edge compute solutions leverage the global footprint of CDN networks to dynamically optimize compute resources and enhance the overall efficiency of the compute process.

### AI opportunity from inferencing

The CDN/edge compute market has remained notably overlooked as a benefactor of the AI cycle, despite the rich opportunity that sits at the cross-section of CDN/edge compute capabilities. The opportunity for the edge compute market lies in inferencing. While there is no unique value in training at the edge, there is value at inferencing at the edge. Inferencing models are compact and lightweight. In fact, every user with a smart phone hosts an inferencing engine on their phone. Unlike training, which will occur in core compute with hyperscalers, inferencing requires a distributed, scalable, low-latency, low-cost model, which is what the edge compute model provides.

The current divide in the edge compute industry is whether to power inferencing with CPUs or GPUs. While all main vendors support both GPU and CPU capabilities, we believe CPUs are the optimal choice to support inference at the edge. GPUs, while powerful for compute, are uniquely bad at living anywhere outside of core compute. They consume a lot of power, resulting in costly cooling measures, and make hardware design complicated. Under the GPU model, only 6-8 requests can be processed at a time. However, CPUs are able to subdivide servers by users, making it a more efficient processing system at the edge. Conversely, CPUs offer cost efficiency, scalability and flexibility and also allow for software enhancements as edge compute vendors overlay proprietary software onto the compute process.



## Quantum Computing – the end of the world as we know it

What is it? Computing that leverages sub-atomic particles to store information and uses superpositions for complex calculations. A quantum computer can solve problems near instantaneously that would take a classical computer billion of years.

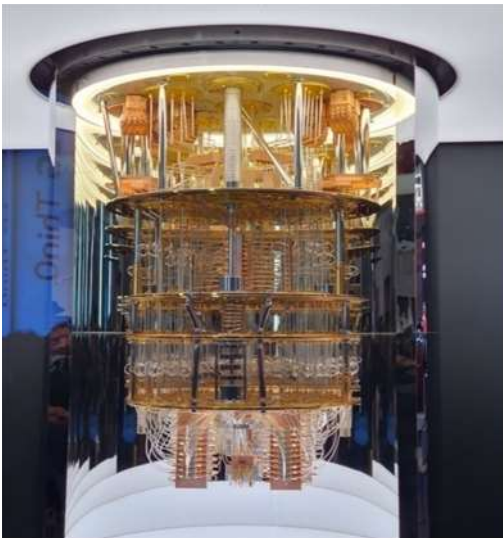
Did you know? In 2019, Google, in partnership with NASA and Oak Ridge National Laboratory, announced that its quantum computer had made in 200 seconds the number of calculations that it would take a supercomputer 10,000 years to do.

A human would have to perform one sum every second for c.50 quintillion years to equal what a quantum computer can do in a single second<sup>1</sup>

Source: <sup>1</sup>McKinsey, Google, BofA Global Research

In our 2019 [Eureka! Future Tech Moonshots Primer](#), we referred to quantum computing technology as a moonshot – radical technology that will shape our future but probably take a decade for wide commercial usage. We are still big believers that this is one of the most radical technologies and should be regarded as a “fire moment” for humanity, but it’s not years away anymore – it is already available for commercial usage. Quantum computing could be one of the biggest revolutions yet. In 2019, the first example of ‘quantum supremacy’ was announced by Google. Since then, more companies have claimed to reach this milestone, which technically means that quantum computers can make more complex calculations than any classical computer could ever do.

Exhibit 54: IBM Q System One  
The world’s first integrated quantum computing system – CES 2020



Source: BofA Global Research

BofA GLOBAL RESEARCH

How it works: bizarre and mysterious lesson 101

Although the possibility of a quantum computer was proposed as far back as 1982, it was only in 2012 that the first venture was undertaken to explore commercial use. A quantum computer is a machine based on quantum mechanics – the branch of physics that describes behaviour at the sub-atomic level. Quantum computers operate on three principles of quantum mechanics, which, technically, are not possible: superposition, entanglement and qubits.

- **Superposition:** If a classical computer can work on two possibilities only (0 or 1, electricity is either on i.e. 1 or off, i.e. 0), the superposition phenomenon theoretically translates into endless possibilities and, as a result, endless calculations. Unlike classical physics, superposition posits that a particle can be in more than one state at any given time: two states, more than two or even none. For example, a photon could be up, down, up and down, in no state, or anything in between, all at the same time. This quantum machine phenomena gave rise to Albert Einstein's famous quote: "God does not play dice with the universe".

But here it gets even weirder. Superposition has one condition – it requires uncertainty. Once we observe and measure the position (i.e. add certainty to the equation) the superposition states collapse into one. Meaning in order to stay in a superposition and make endless calculations, we need uncertainty.

Superposition paradox: Schrödinger's cat. To demonstrate superposition anomaly, in 1935, Erwin Schrödinger suggested a paradox, according to which a cat, a radioactive substance that decays very slowly, and a Geiger counter are put together in a sealed box. According to quantum mechanics physics after a while the cat could be in a superposition of being both alive and dead, as long as we do not open the box. Once we open the box, the uncertainty that triggers the superposition will collapse.

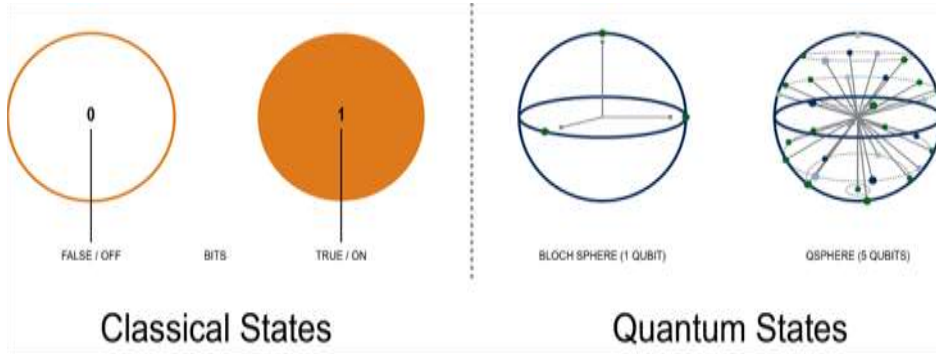
- **Entanglement:** This refers to two or more quantum systems being linked in an unseparated bond, even if far apart – meaning this status cannot be described independently, even if the distance between them is vast. Any effect on one particle will instantly impact the other, faster than the speed of light, no matter the distance between them. If superposition offers endless states and therefore calculations, entanglement provides parallel processing of the system and allows it to scale up and work as one and other properties like error correction or control interference.
- **Qubits:** In other words – quantum bits – a quantum computer's basic unit of information. Superposition of a particle can exist in a qubit and multiple qubits can be entangled – meaning one qubit can affect another. In short, these units contain the quantum mechanics properties of superpositions and entanglements. Without qubits, those properties could not be harnessed to create a quantum computer. To complicate matters yet further, qubits can act like waves, or particles and can tunnel through energy barriers – which is impossible according to classical physics – and just "turn up on the other side". However, in order to achieve these qualities, qubits need to be fully isolated. If not, they "collapse" and become "normal bits".

### **Bits are out, qubits are in...**

As we stated above, a conventional computer works on "bits" that are binary by nature: a bit can be only one of two possibilities (0 or 1). Qubits, exploiting the unique physics of quantum subatomic particles, have endless possibilities (1, 0, anything in between and even more than one possibility simultaneously). In theory, a quantum computer can make endless calculations.



**Exhibit 55: Bits vs Qubits**  
From the horse and cart straight to space travel



Source: IBM

BofA GLOBAL RESEARCH

And this is where it becomes even more complicated and bizarre. We still do not fully understand quantum mechanics and cannot predict the state of the qubit. More than that, to fully control qubits for calculations and information transfer we need to be in both an uncertain state and a completely sterile environment. Any external element like radiation, noise, energy or other numerous factors can interfere with the process. One of the main challenges today is less about qubits technology and more about “decoherence” – controlling all external noise and interference that could cause mistakes.

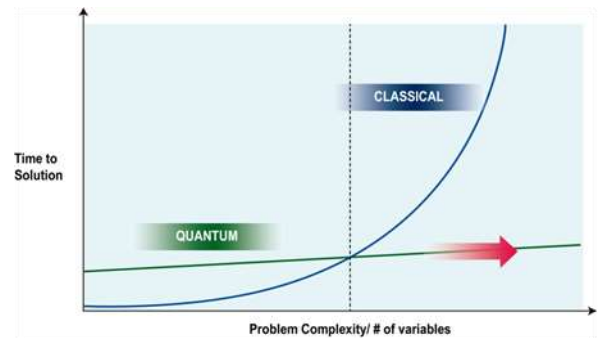
**Exhibit 56: Quantum computers are getting more powerful**  
From 2 qubits in 1998 to 128 qubits in 2019



Source: MIT, CB Insights

BofA GLOBAL RESEARCH

**Exhibit 57: Why quantum computing matters**  
quantum computers have an embedded advantage over classical computers that cannot be gapped



Source: D-Wave Systems

BofA GLOBAL RESEARCH

Massive progress has been achieved in recent years. And 2023 was the year of a leapfrog in development, thanks to the data-processing capabilities of Artificial Intelligence Large Language Models (LLM) algorithms. The possibilities of a computer that can theoretically make endless calculations simultaneously is huge. This could range from a new unbreakable encryption to calculations that now require huge resources and time, such as nuclear simulations, to deep learning and real usage of big data analysis. The possibilities are literally endless.

**Where we are: welcome to the era of quantum supremacy**

Quantum supremacy (or quantum advantage) refers to the point that quantum computer has a computational advantage by solving a problem that traditional computers could not solve in any feasible amount of time. There is no quantitative barrier, or calculation number to achieve quantum supremacy – it just refers to complex calculations that classical computers cannot do. In other words, proving that quantum computers have an embedded advantage over classical computers that cannot be gapped.



- In 2019, Google, in partnership with NASA and Oak Ridge National Laboratory, announced in Nature magazine that it had achieved quantum supremacy for the first time. The article stated that their quantum computer had made more calculations in 200 seconds than the strongest supercomputer could have done in 10,000 years.
- In 2020, a group of researchers at the University of Science and Technology of China (USTC) published a paper on reaching quantum supremacy in their photonic quantum computer, Jiuzhang. According to the paper, Jiuzhang made the same number of calculations in 200 seconds that would have taken a classical supercomputer 2.5 billion years to do.
- In 2022, Xanadu reported a new system that had obtained speeds up to 50 million times faster than Google or USTC.
- In 2023, Google announced that its 70-qubit Sycamore quantum processor had completed a 6-second calculation that would have taken the most powerful supercomputer 47 years.
- In 2023, IBM's 127-qubit Eagle quantum computer performed more complex calculations than its supercomputer by performing tasks proving its 127-qubit system is stable with effective noise control. Saying that, IBM did not claim it had achieved quantum supremacy.

A future quantum computer can process more calculations than there are atoms in the universe ( $10^{80}$ ) (source: Venture Beat). For example: Modelling penicillin molecule, which is composed of 41 atoms, would require a classical computer with more transistors than there are atoms in the observable universe (source: IBM, BCG)

### From Moore's to the Gambetta Law

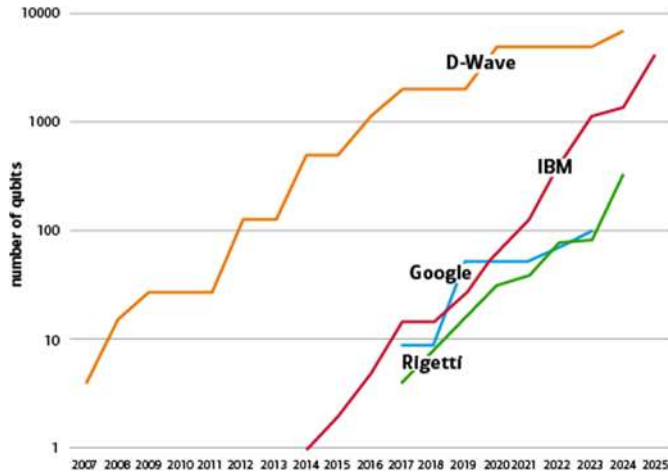
We are in the very early stage of adoption and only a few machines are available for commercial use via the cloud, mainly for research. However, commercialisation is scaling up fast.

So far we have relied on Moore's Law to explain classical computing power progress (that is, the number of transistors in an integrated circuit doubles about every 18-24 months). This law has been reincarnated again and again and still proved correct. In the quantum world, however, we now refer to the Gambetta Law, drafted by Jay Gambetta, the scientist and executive leading the team at IBM Thomas Watson Research Center working on quantum computers. According to the Gambetta Law, the maximum size of qubits that can be implemented successfully by the computer is doubling every two years.

In December 2023, IBM announced two new quantum systems. One is based on a chip named Condor, which is the largest yet, with 1,121 functioning qubits. The second is based on a combination of three Heron chips, each of which has 133 qubits.



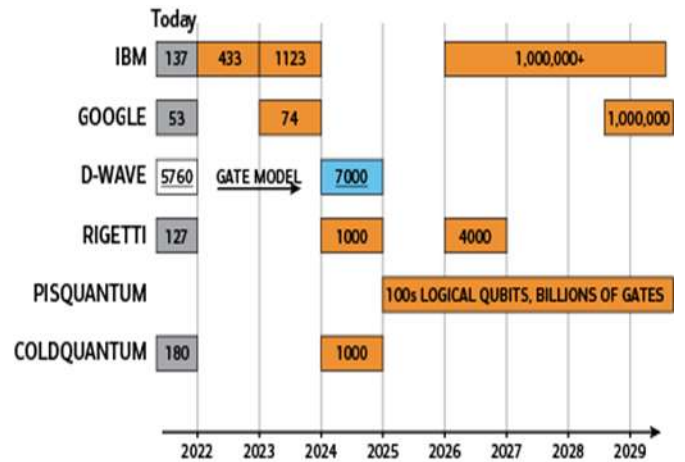
**Exhibit 58: Progress in the number of qubits of key companies**  
 The number of qubits has increased dramatically in the past couple of years. IBM recently announced its next generation quantum computer, Condor, will be based on 1,121 qubits



Source: Companies' announcements

BofA GLOBAL RESEARCH

**Exhibit 59: Roadmap of progress of quantum computer systems of key companies**  
 The race is not just to increase the number of qubits, but for noise control and optimisation



Source: Companies' announcements

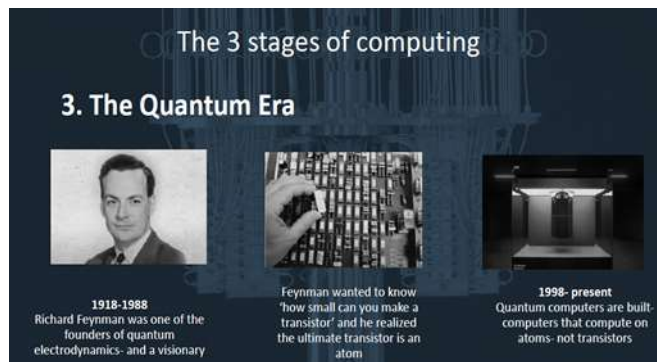
BofA GLOBAL RESEARCH

The potential – the end of the world, as we know it

‘Quantum computers do compute the way nature computes’ – Ilyas Khan, CEO of Quantinuum

We could see a paradigm shift with computational boundaries being broken by quantum computers. At our Transforming World conference, Ilyas Khan, CEO and co-owner of Quantinuum, stated that over the next 12 months, a quantum computer will be developed that cannot be simulated classically. For that reason, he believes we will enter an era of a paradigm shift where every sector will be impacted in the next year or two. Khan says that we need to be thinking about what can be done in the current paradigm but also what will threaten it.

**Exhibit 60: The 3 Stages of Computing**  
 Quantum era



Source: Quantinuum

BofA GLOBAL RESEARCH

**Exhibit 61: Sectors that could be impacted by quantum AI**  
 Every sector could be impacted e.g. semis, banking and finance, cybersecurity



Source: Quantinuum

BofA GLOBAL RESEARCH

**The “power couple” that will change the world**

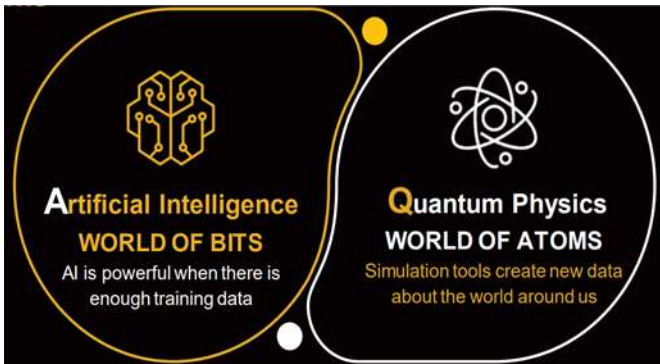
We can talk about many applications, but one of the most exciting is the marriage of the two most powerful technologies: AI and quantum computers. The convergence of AI and quantum (AQ) technologies can enable fundamental improvements in the ‘physical’ world as well as the digital one. While quantum computers will provide endless calculations when available, the increased capability of AI technologies can unlock several





transformational use cases in the meantime. LLMs complemented with simulation, knowledge graphs, computer vision and predictive analytics are ‘the new AI toolbox’ that companies can deploy. Iterative research tasks that would have taken years can now be achieved in weeks. Life sciences, chemicals, materials, and finance/logistics are all in scope to benefit. And in the long run, Artificial General Intelligence (AGI) – when AI reaches human cognitive abilities and even self-awareness – and singularity – the point when AI surpasses human intelligence – will lead to exponential, radical transformation of technology.

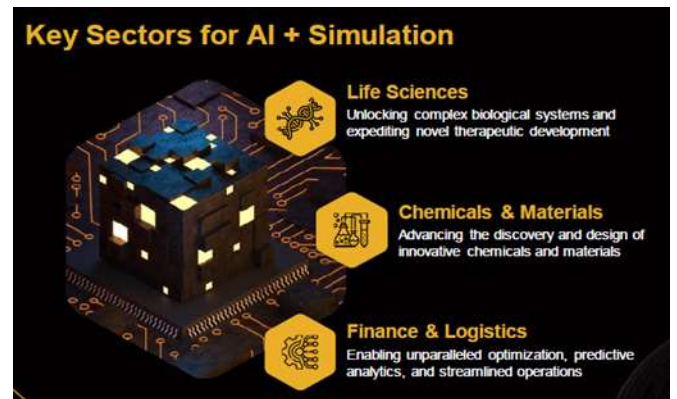
Exhibit 62: Artificial intelligence (a world of bits) vs quantum physics (a world of atoms)  
Sandbox solves large-scale enterprise challenges using the best of today's simulation and AI platforms



Source: SandboxAQ

BofA GLOBAL RESEARCH

Exhibit 63: Key sectors that can use AI and simulation  
Life sciences, chemicals & materials, and finance & logistics can use AI and simulation tech already



Source: SandboxAQ

BofA GLOBAL RESEARCH

A quantum computer is not suitable for regular tasks like using the internet, office tasks or emails, but for complex big data calculations like blockchain, machine and deep learning, or nuclear simulation. A combination of quantum computers and 6G mobile networks would be a game-changer in every industry. For example:

- **Big data analytics:** The amount of data created is projected to double every 2-3Y, reaching 183ZB by 2024E from 120ZB in 2022 and 12ZB in 2015 (source: Statista, IDC). The untapped big data potential is huge given that only c.0.5-1.0% of data generated has ever been analysed. Today, we are storing, transmitting and using only 1% of global data (IDC), because we lack the computing power needed to process more. Quantum computing could change that and unlock real economic value. Using 24% of global data (unlike 1% today) would double global GDP (source: IDC).
- **Financials:** Quantum computers are able to master almost every aspect of banking. For example, the data calculation capabilities of quantum computers are better at predicting market trends, portfolio optimisation, data analytics in real time, pattern detection, encryption and fraud detection.
- **Cybersecurity:** Quantum computing is able to technically challenge all current encryption methods (based on large number factoring), including blockchain, by parallel processing capabilities of up to 1 trillion calculations per second (Bernard Marr). This also opens the door to new encryption technology, based on quantum computing elements.
- **AI and machine learning:** Quantum computers can speed up machine learning capabilities by using more data faster and solving complex connections between data points. Machine learning and deep learning progress is limited to the pace of underlying data calculation. The faster data is calculated and used, the faster machine and deep learning algorithms will evolve.



- Infrastructure: One area where quantum calculation is needed is complex network calculations. Quantum computers can reorganise global telecoms, utilities, transportation and other combined infrastructures that are becoming more integrated and too complex for current computing abilities.
- Healthcare and Genomics. If global data is expected to increase twofold every 2-3 years, medical data is expected to double every 73 days (source: ACCA, KPCB). Each person will generate enough health data in their lifetime to fill 300 million books. Genomic data will double every 50 days and the amount of global genomic data will surpass YouTube and X by 2025 (New Scientist). Quantum computers could be the solution for all big data processing in this field. Big data utilisation would not only potentially bring down costs, but also help provide better healthcare via targeted treatment and predictive analytics (source: IMS Health).
- Science: Complex calculations used in space (e.g., big bang simulations), physics (e.g., dark matter, string theory), nuclear simulations, complex materials and 'super-connectivity' can leapfrog to the next level with quantum computing.
- Cloud: This could be one of the winners, as the cloud could be the platform where all data creation, sharing and storage take place. Once the commercialisation of quantum computers begins, cloud access will be needed and data generation should jump exponentially. Thus, cloud platforms will be the solution.
- Autonomous vehicle (AV) fleet management: One connected AV will generate the same amount of data as 3,000 internet users (Intel). However, two cars will generate the same amount as around 8,000-9,000 users as the two will need to communicate with each other, and they will be generating data and so on. The growth of data from AV alone will be exponential. Thus, on this calculation, a fleet of 1,000 cars will generate more data than the entire global population. A regular computing system would not be able to handle this.

Exhibit 64: Quantum computing eco-system  
Many companies, mainly private, comprise the quantum computer eco-system.

Users Select examples	Applications Not mapped to verticals	Software offerings Includes control software	QPU <sup>s</sup> <sup>2</sup>	Hardware / components Select examples only - not representative of entire ecosystem
<b>Material Science</b>	Not strictly categorized given diversity of operations		<b>Superconducting</b>	<b>Cryogenics (includes testing)</b>
Merck Airbus OTI BMW	QCWare Multiverse Computing NQCG Aliro Quantum Horizon 1QBit ready to run quantum software		rigetti Google AI Quantum IBM Q aws SeeQC IQM DWave Alibaba.com	Oxford FormFactor Montana Instruments LakeShore
<b>Finance</b>	Quantum Simulations Zapata QCI River Lane Strange works		<b>Ion Trap</b>	<b>Blue Fors</b> <b>Maybell</b> ICE
Goldman Sachs Bank of America J.P. Morgan Wells Fargo	QuantFi Entropica Labs QunaSys Ferra Quantum ProteinQure Qubit Pharmaceuticals Menten.AI ParityQC		<b>Neutral Atoms</b>	<b>Lights and lasers</b>
<b>Life Sciences</b>	Phasecraft aQ Keysight Classiq		Oxford ionics IONQ AQT Qudoor electron Silicon Quantum Computing equal 1 Intel Phonic	Vescent NKT Photonics Coherent M solared
Odyssey Boehringer Ingelheim AstraZeneca	Quantum-South Menten.AI ParityQC		<b>Silicon</b>	<b>Other componently (examples)</b>
<b>Other</b>	Cloud access to QPUs aws Google Cloud	Simulators / q-inspired / etc Atos Fujitsu Huawei Huawei Cloud nVidia	<b>Photonics</b>	Delft Circuits Quantum Design CryoCoax QBloX Devil Zurich Instruments
Volkswagen Johnson Electric Denso			<b>Other</b>	Quantum Machines QM LNF
			EeroQ Electons on helium Quantum Brilliance NV Diamond	

Source: Quantum Insider

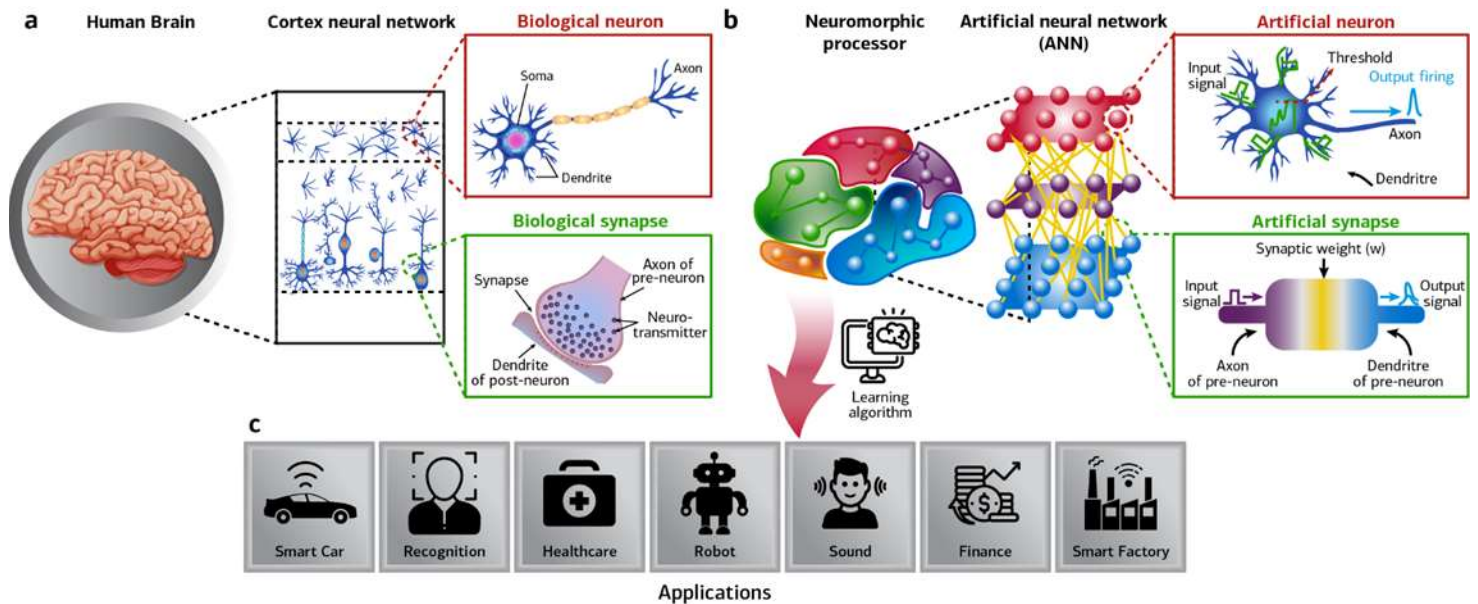


## Neuromorphic Computing/Brain Computer Interface

What is it? Brain waves of humans and animals directly interacting with the external world and vice versa.

Did you know? A brain-to-brain interface has already allowed one human to control another's hands through their thoughts alone.

Exhibit 65: Neuromorphic computing inspired by the human brain  
 Neuromorphic computing is brain waves of humans and animals directly interacting with the external world and vice versa



Source: Research Gate

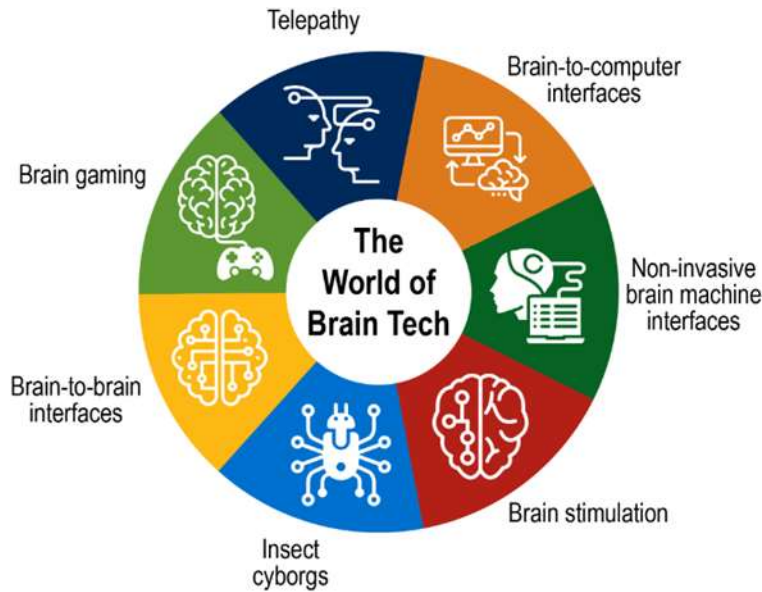
BofA GLOBAL RESEARCH

### BCI: human-machine implants: if you can't beat them, join them

AI is set to become more intelligent than humans, but some start-ups are focusing on creating a human-machine cooperation that will prevent people from being left behind. As this is a nascent technology, different approaches are being trialled with different motives. Consequently, brain computer interfaces differ in invasiveness and ability. Start-ups like Neuralink are investigating such technology with physical implants that would allow a human to control a connected device, while others like CTRL-Labs are looking at non-invasive technologies to read human thought. Other start-ups in this space include Synchron and Neuroable.



Exhibit 66: How is brain technology being developed?  
The world of brain tech



Source: BofA Research

BofA GLOBAL RESEARCH

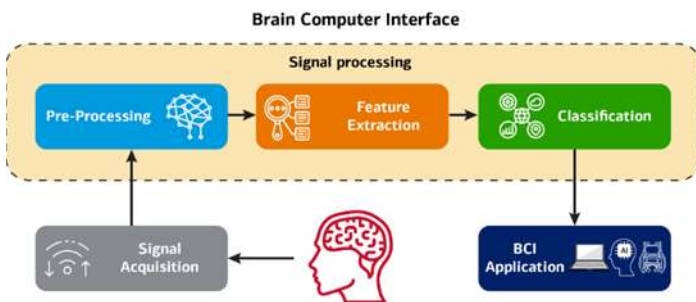
How would brain computer interfaces work?

Examples of brain implants that are already in use:

- Cochlear implants stimulate the auditory nerve to provide some form of sound for hearing-impaired individuals.
- In a research lab, BrainGate’s implants have allowed individuals with spinal cord injuries and Amyotrophic Lateral Sclerosis (ALS) to control a cursor on a computer with their thoughts. In 2021, the team managed to decode thoughts of handwriting movements and translate them into text on a screen as well as producing a wireless transmitter to transport the full spectrum of signals recorded. Note that this still requires a transmitter to sit on the person’s head.

The concept of brain implants is not completely new. Academics have already completed studies of implants using a few hundred electrodes to collect brain data in animals. In addition, there are already procedures to stimulate parts of the brain for patients with Parkinson’s Disease and clinical depression.

Exhibit 67: How do brain computer interfaces work?  
The process of using human brain waves



Source: Semantics Scholar

BofA GLOBAL RESEARCH



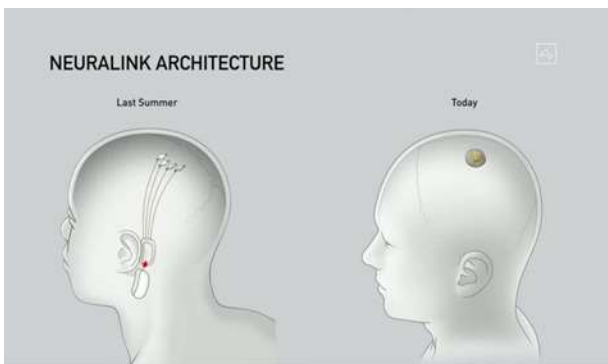
Early human clinical trials now in progress

Minimally invasive implants are now being trialled in humans. Synchron, a brain interface platform company, is trialling implants with 4 patients in Australia. In its study, individuals with paralysis were able to text, email, and bank and shop online. Synchron’s approach involves lacing blood vessels that supply the brain with a mesh tube of sensors and electrodes from which neuron signals can be picked up. The signals are passed to an external unit that translates the signals and conveys them to a computer.

Neuralink patients can now move a computer mouse by thinking

- The tech: Neural implants include neural threads that are inserted with a neurosurgical robot into the brain to pick up neural signals being transmitted. Neuralink’s implant, the Link, is expected to contain 1,024 electrodes. These neural threads transfer data back and are directly connected to the implant. This implant processes and further transmits the neural signal to the outside world.
- The aim: The brain data collected by the implants would then enable individuals to control a keyboard or computer mouse, particularly ground-breaking for individuals with spinal cord injuries. Longer term, the implant could help enable motor and sensory functions and even strengthen a healthy person’s capabilities looking further ahead.
- Success with animals: Beyond safely inserting, collecting data and removing the implant in pigs, Neuralink’s implant has been successful in monkeys. A monkey with 2 implants placed on its head can play online ping pong-type games solely with its mind once the neural signals are calibrated.

Exhibit 68: Original vs updated Neuralink  
Neuralink implant has switched from ear canal to apex of skull



Source: Neuralink

BofA GLOBAL RESEARCH

Exhibit 69: Neuralink characteristics  
Battery life can last all day with overnight charging

Channels	1,024
Battery life	All day
Recharging time	Overnight
Wireless range	5 - 10 meters
Implant size	23 mm x 8 mm
Look	Not externally visible

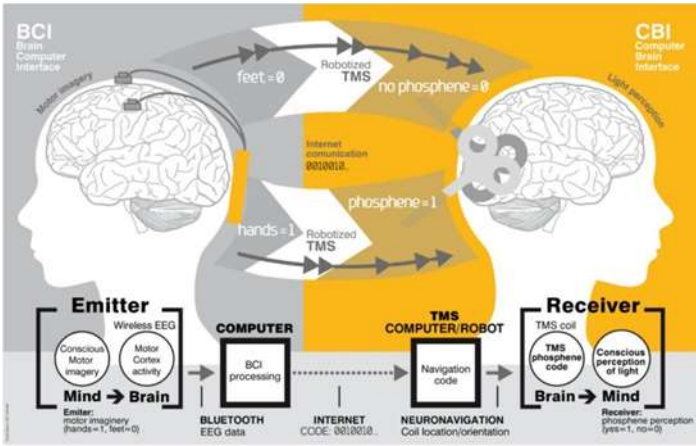
Source: Neuralink

BofA GLOBAL RESEARCH

Regulation of such a contentious technology will no doubt play a key part in the rollout of BCIs. The Breakthrough Devices Program by the FDA, which is intended to speed up the development of solutions for irreversibly debilitating conditions, has given products like Stentrode by Synchron faster permissions to be investigational devices.



Exhibit 70: Brain computer interface vs computer brain interface  
 How would a human brain/cloud interface work?



Source: Frontiers

BofA GLOBAL RESEARCH



## Robotics: a Robot and AI walk into a bar...

Transforming World: AI allows us to capitalise on this big data revolution but even further progress could be made if it could interact with the physical environment. This could enable AI and Robots to change the way we live and work in myriad ways.

Did you know? It took c.100 years each to see Industry 2.0 (electrification) and Industry 3.0 (automation) but <50 years for Industry 4.0 (digitalisation) and <10 for Industry 5.0 (humanisation).

Next generation of robots: humanoid robots, industrial robots & cobots, autonomous vehicles, drones

What if generative AI could interact with the physical world? It could learn from a wider environment's worth of data and teach itself to do tasks that it was not pre-programmed to do.

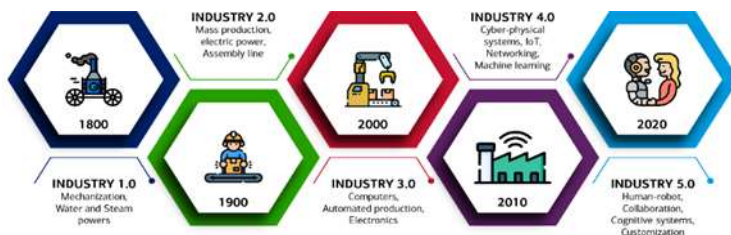
How can generative AI interact with the physical world? Embodied AI  
The term embodied AI was first used to describe the branch of AI that focuses on how computers, systems and technology can interact with the physical world. It typically includes AI for sensorimotor skills, navigation, and real-world interactions. But with the rise of generative AI, embodied AI is also being used to give this technology a physical form, typically a robot including autonomous vehicles and drones. Robots use AI to interact with the physical world and to learn from their interactions. They are equipped with a range of sensors and computer vision to take in data from the real world. Combined with generative AI, they can 'learn' from this data and also exhibit conversational and situational awareness. For example, to teach a robot to pick up a book, the traditional AI method would be to program it with the movements that it needed e.g., lower hand, open hand, close hand, raise hand etc. In contrast, the embodied AI approach would be to teach the robot to guess by having it make random attempts and learn from each attempt until it achieves the goal.

### New applications...towards autonomy?

Embodied AI will bring new use cases. Innovations in metals (to make them stronger and lighter) and batteries (now lasting ~5hr shifts) are enabling them to become a reality. There are opportunities for productivity enhancements and removing dangerous or repetitive/onerous tasks from humans e.g. robots can help with assembly, in warehouses (e.g., packing and shipping items). Typically, we might think of robots as enabling us to automate tasks, with humans being the principal and robots being the agent. But as these robots become more capable of self-learning and completing tasks without being pre-programmed, their direction may shift from one of automation to autonomy.

Exhibit 71: The five waves of industrial revolution

The fifth industrial revolution involves the combination of humans and machines at work



Source: KnowHow

BofA GLOBAL RESEARCH

## Humanoid Robots

What is it? Humanoid robots assist humans and have a human-like form, mimicking human behaviour and interactions

Did you know? Sophia was the world's first AI humanoid robot and was also the first humanoid to get a passport

### Humanoid robots 101

Humanoid robots are a type of service robot that mimic human behaviour and interaction. Service robots operate semi- or fully autonomously to perform services useful to the wellbeing of humans and equipment, excluding manufacturing operations (source: IFR). They typically assist humans by performing jobs that are dirty, dull, distant, dangerous or repetitive, including household chores. They tend to be autonomous and/or operated by a built-in control system, with manual override options. Compared to other types of service robots, humanoid robots usually face more complex working environments, requiring sophisticated AI functionality to support sensing, motion control and interaction with humans.

2023 was a key year for humanoid robots because of generative AI

“The ChatGPT moment for robots might be around the corner” –  
Jensen Huang, Nvidia CEO, GTC 2024

2023 marked the beginning of the rapid development of humanoid robots due to the breakthrough of large language models, notably the release of ChatGPT at end-2022. Companies have started combining generative AI and robots to create a new age of humanoid robots. If (generative) AI systems have a physical form, they can take in even more information that can help them understand the world and human behaviour. Embodiment is arguably a characteristic that can help (generative) AI reach AGI (source: Australian National University).

If AI were able to experience the physical world, it could offer new sources of training data e.g., cars. Embodiment can greatly expand capabilities and bridge the gap between natural language and robot control. How might genAI drive improvements in robots?

- Emotions: Putting generative AI into humanoid robots could help them identify human emotions and complete human tasks. AI would be able to analyse data on expressions (e.g. smiles, frowns, surprises, body language, tone of voice) that the human eye might miss.
- Learning: Making robots faster learners and deal with more complex tasks.
- Personable: Helping robots understand natural language instructions better.

### Competitive landscape

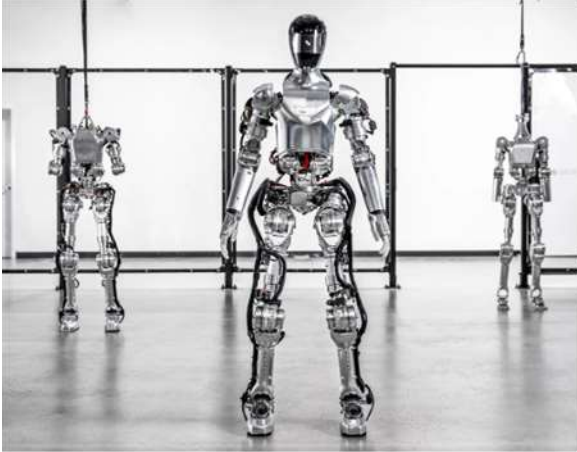
The development of humanoid robots commenced in the late-20<sup>th</sup> century. By 2022, many robot producers had launched their humanoid robot models, mainly used for logistics, medical services and demonstration purposes.





**Exhibit 72: Image of Figure AI's humanoid robots**

Figure AI's Figure 01 – general-purpose humanoid robot

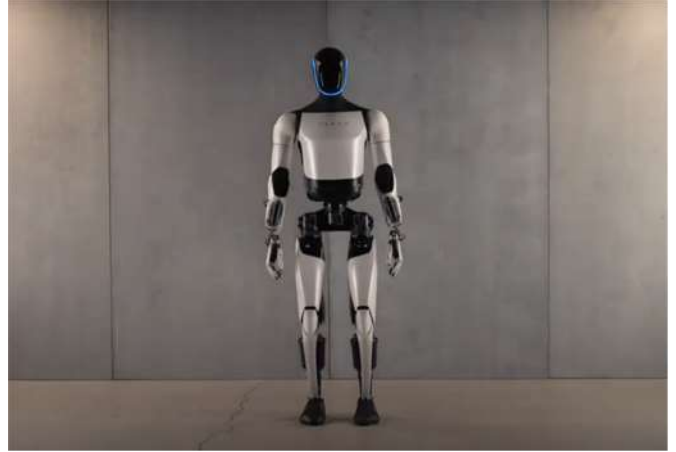


Source: Figure AI

BofA GLOBAL RESEARCH

**Exhibit 73: Tesla's Optimus Gen 2**

It features 30% faster walking speed and 10kg weight reduction



Source: Tesla

BofA GLOBAL RESEARCH

**Exhibit 74: ASIMO (Honda)**

ASIMO was developed by Honda in 2000 and ceased development 2018

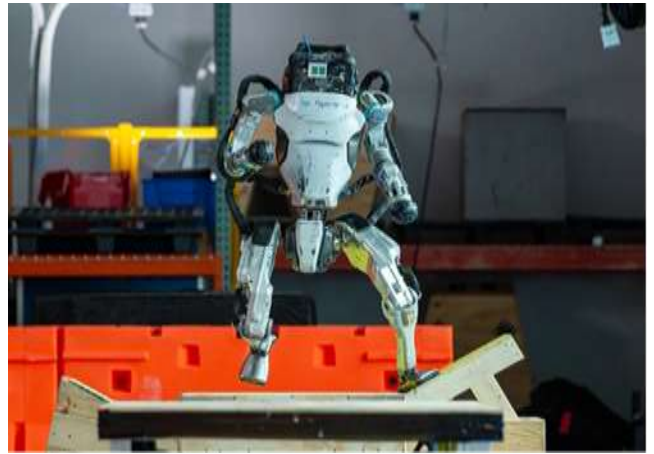


Source: Honda

BofA GLOBAL RESEARCH

**Exhibit 75: Atlas (Boston Dynamics)**

Atlas adopts electro-hydraulic actuation to enhance its athletic ability



Source: Boston Dynamics

BofA GLOBAL RESEARCH

**Exhibit 76: Digit (Agility Robotics)**

Agility Robotics' Digit is designed for warehouse logistics and delivery

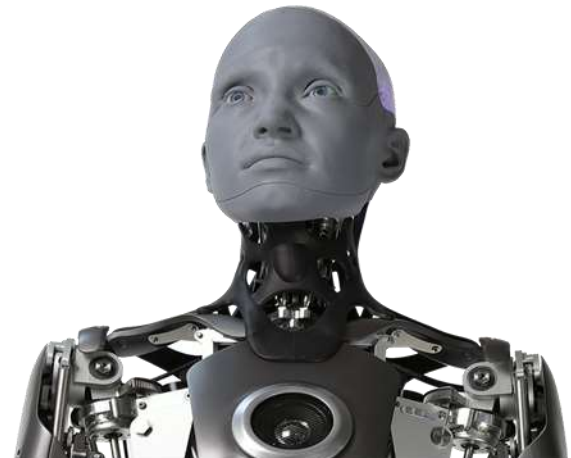


Source: Agility Robotics

BofA GLOBAL RESEARCH

**Exhibit 77: Ameca (Engineered Arts)**

Engineered Art's Ameca can mimic human facial expressions



Source: Engineered Arts

BofA GLOBAL RESEARCH



Exhibit 78: Major humanoid robot models

Humanoid robots are mainly used for manufacturing, logistics, medical services, and demonstration purposes

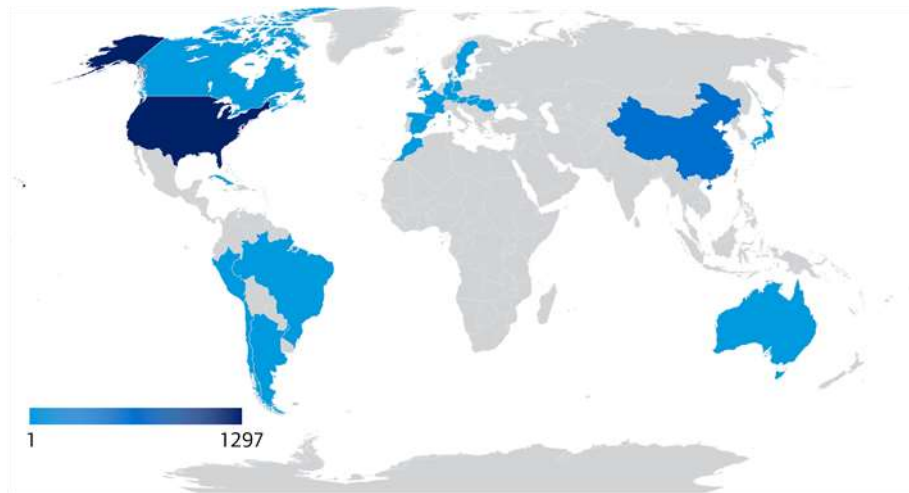
Company	Boston Dynamics	UBTech	Engineered Arts	Xiaomi	Tesla	Agility Robotics	XPeng	Unitree Robotics	Fourier Intelligence	Agibot	Dreame
Product name	Atlas	Walker X	Ameca	CyberOne	Optimus	Digit	PX5	Unitree H1	GR-1	RAISE-A1	Robot
Launching time	2013	2021	2021	2022	2022	2023	2023	2023	2023	2023	2023
Height (m)	1.5	1.45	1.87	1.77	1.72	1.75	1.50	1.80	1.65	1.75	1.78
Weight (kg)	89	77	49	52	57	65	Not specified	47	55	55	56
Degrees of freedom	28	36	51	21	28 (excluding hands)		20	4 per arm + 5 per leg	44	49+	44
Application	Search and rescue	Service, teaching	Demonstration	Not specified	Replace human in dangerous and repetitive tasks	Warehouse, delivery	Manufacturing, service	Replacing human force	Manufacturing, service, healthcare	Manufacturing, household	Service
Max. Speed (km/h)	9	3.0	NA	3.6	8	5.4	Not specified	5.4	5.0	7.0	Not specified
Selling price	Not commercialized	\$500k	\$133k	Not specified	\$20k in the long run	expected <\$250k	Not specified	expected <\$90k	Not specified	Not specified	Not specified

Source: BofA Global Research; companies: BofA APAC Industrials team

BofA GLOBAL RESEARCH

According to GlobalData, the total number of robotics-related patent applications was 3,825 in Q4 2023 vs 6,883 in Q3 2023. The US is the leading country in robotics adoption within the technology industry, with the highest number of robotics-related patent filings.

Exhibit 79: Robotics-related patenting activity in terms of filings in the global technology industry. The largest share of robotics-related patent filings in the technology industry in Q4 2023 was in the US with 34%, followed by China (24%) and Japan (5%)



Source: GlobalData Patent Analytics. NOTE: based on Q4 2023

BofA GLOBAL RESEARCH

China robotic OEMs: rapid development of humanoid robots since 2022

The US is the leading country in robotics adoption, with the highest number of robotics-related patents, jobs and deals (source: GlobalData). But original equipment manufacturers (OEMs) in China have been rapidly developing their humanoid robots since 2022. As our APAC Industrials colleagues note, Xiaomi launched its first humanoid robot, CyberOne, in Aug-2022, ahead of Tesla’s Optimus. The breakthrough of large language models, most notably the release of ChatGPT at end-2022, further accelerated the development of humanoid robots. In 2023, XPeng, Unitree Robotics, Fourier Intelligence, Agibot, Dreame, etc. joined the competition with their new humanoid robot products. On 2 Nov, China’s Ministry of Industry and Information Technology (MIIT) issued the Guidance Opinion on the Development of Humanoid Robots, targeting a technological breakthrough of key components and mass production by 2025, with a world-class supply chain by 2027.



**China aiming for mass production of humanoid robots by 2025...**

China’s Ministry of Industry and Information Technology (MIIT) issued the Guidance Opinion on the Development of Humanoid Robots on Nov 2 with goals to be achieved by 2025/2027. By 2025, China aims to establish an initial humanoid robot innovative system and achieve technology breakthroughs in key components such as robots’ brains, cerebellum and limbs to ensure the components are supplied safely and efficiently. The humanoid robots are expected to be sold internationally and achieve mass production.

**...and also hoping its humanoid robots can think and learn by 2027**

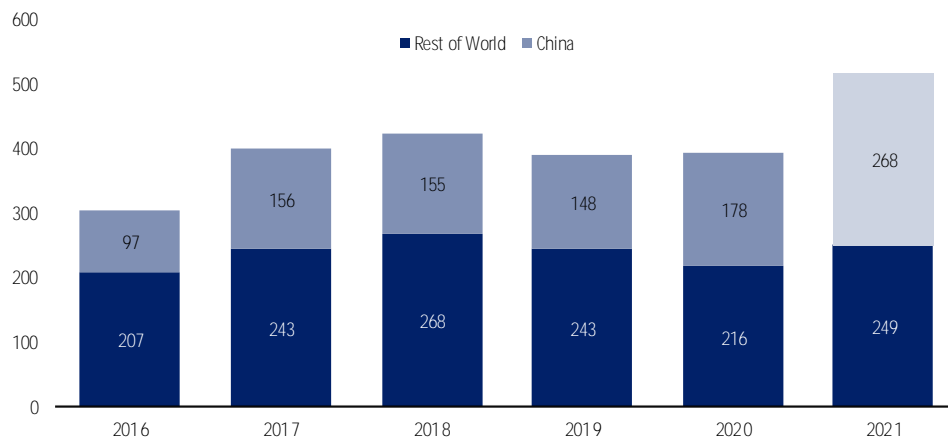
By 2027, the technological innovation capability of humanoid robots could be significantly improved and form a safe and reliable supply chain system. The industry could witness accelerating mass production with wider applications that are deeply integrated with the real economy. This could be a new stimulus for economic growth.

The guidance also outlines 6 key tasks including: 1) key technology breakthroughs, 2) fostering key products including bodies, core components and software, 3) expanding application scenarios, 4) creating the industrial ecology, 5) strengthening the support capability, such as standards setting and test/verification capability, etc, and 6) supportive measures, such as interdepartmental cooperation and introduction/education of talents.

**China ramping up robots**

In recent years, China has been increasing its robot development. In early 2023, it created a new ‘Robot + Application Action Plan’ which highlights 10 sectors (manufacturing, agriculture, architecture, logistics, energy, healthcare, education, elderly services, community service, and emergency and extreme environment applications) that it would like to automate. The country aims to have c.500 robots per 10,000 workers by 2025. In August 2023, Beijing announced a US\$1.4bn robotics fund aimed at promoting robotics technology development.

Exhibit 80: Annual installations of industrial robots (000s of units)  
China now installs more industrial robots per year than the rest of the world combined



Source: International Federation of Robotics

BofA GLOBAL RESEARCH

**Headwinds: production cost, energy use and robustness**

Our APAC Industrials colleagues note that given the present technological complexity and elevated production costs, it could take at least 7-8 years before we see mass adoption of humanoid robots in household applications. They expect production in small batches to commence in 2024-25, and the application of humanoid robots to start from relatively simple (semi-structured) environments such as warehouse and security guarding.



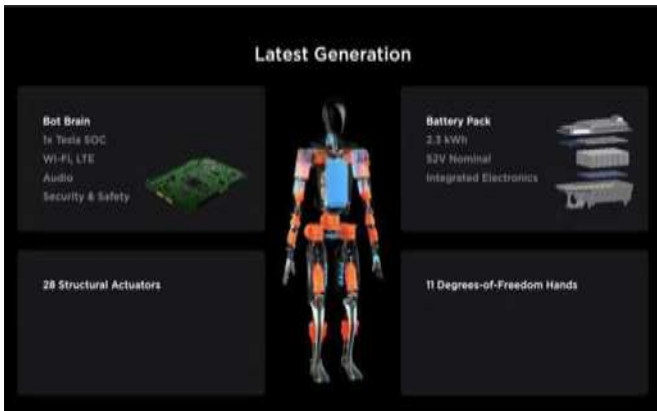
Key bottlenecks to overcome in the coming years include: (1) a sensory system that can generate accurate environmental information in complex and uncertain environments; (2) a robust motion control system supported by an AI algorithm; (3) the application of a language model that enables real-time interaction with humans; (4) a strong power system with high energy density that can support longer duration of operation; and (5) simplified product design to make the robots more suitable for mass production with lower costs. So far, only a few humanoid models have entered into commercial use due to elevated production costs, limited downstream applications and immature product designs.

Higher production costs due to large number of actuators and sensors

The production cost of a humanoid robot is much higher than an industrial robot or a collaborative robot due to the large number of actuators and sensors. The cost of an actuation system is around US\$1,000-2,000 for each degree of freedom, and a humanoid robot usually requires >20 degrees of freedom, according to our channel checks and estimates. This brings the total costs of actuators per robot to at least US\$20k. Currently, the ASP of a humanoid robot is usually >US\$100k per unit due to high production costs (based on the selling prices of Ameca and Digit). As a comparison, the ASP of a 6-axis industrial/collaborative robot is only US\$6-35k/2-45k, respectively.

Exhibit 81: Tesla's Optimus humanoid robot

Tesla's Optimus has 14 rotary actuators and 14 linear actuators in its body



Source: Company report

BofA GLOBAL RESEARCH

Exhibit 82: "Optimus" humanoid robot cost estimates

Actuators account for 70%+ of an "Optimus" humanoid robot's costs

	Cost per unit (RMB)	Units per robot	Costs per robot (RMB)	% total costs
Actuators*	6,000	28	168,000	74%
Hands**	2,000	12	24,000	11%
Chips	10,000	1	10,000	4%
Cameras	500	8	4,000	2%
Battery	2,000	1	2,000	1%
Others***	20,000	1	20,000	9%
<b>Total costs</b>			<b>228,000</b>	<b>100%</b>

Source: Company report, BofA Global Research, BofA APAC Industrials team

Note: \*Actuators include servo motors, controllers, transmission machinery (e.g. planetary roller screw, harmonic reducers), sensors, etc

\*\*Hands include coreless motors, reducers, etc

\*\*\*Others include body structural parts, thermal management systems, etc

BofA GLOBAL RESEARCH

High energy consumption implies need for greater energy density in batteries

A humanoid robot integrates large amounts of power-consuming components including 20+ actuators, various types of sensors and processing units running AI algorithms. Due to the limited volume of the battery pack, the overall energy consumption needs to be controlled to ensure a reasonable battery life. This also places a high requirement on the battery's energy density given the limited space.

Robustness to disturbances: withstanding unstructured working environments

Humanoid robots are required to work long hours in highly unstructured working environments, which calls for robustness. Robot developers need to address how to maintain desired motion/working ability amid unpredictable disturbances. This requires strong motion control and force-sensing capability.

Mass adoption in household applications could take 7-8 years

“With today’s technology, it’s impossible to get to as good as a human, so I think the strategy of getting as close to a human as you can is a perfectly valid one.” —Jerry Pratt, Figure CTO

Given the practical constraints on humanoid robots, mass adoption could take some time. The application of a humanoid robot could start from relatively simple (semi-structured) environments such as warehouse and security guarding, before gradually penetrating more complex environments as its functionalities mature. Taking Tesla’s EV (electric vehicle) as a comparison, its first model, Roadster, was launched in 2006 and started production in 2008. But the sales volume of Tesla EVs only picked up from 2013. As a more complex industrial product than an EV, the humanoid robot will spend more time in product upgrade. Our APAC Industrials team thinks it will take at least 7-8 years before we see mass adoption in household applications.

Safety? Public perception of AI risks? Existential risks?

There could be public concern over the potential AI risks, which could affect the robots’ uptake. For example, a press conference hosted at a UN summit invited nine humanoid robots to answer questions from the audience. The humanoid robot, Sophia, said that robots can be more efficient and effective than humans in running the world but then noted that both robots and humans have a part to play. Also, some robots have been programmed to have principles/align with human values e.g., DeepMind’s robots with the Robot Constitution.

What is DeepMind’s Robot Constitution? In January 2024, the DeepMind robotics team revealed three new advances that it says will help robots make faster, better and safer decisions. Their Robot Constitution is a list of rules an LLM is instructed to follow and is inspired by Isaac Asimov’s Three Laws of Robotics. It is a set of safety-focused prompts that tells the LLM to avoid choosing tasks that involve humans, animals, sharp objects and electrical appliances.

“Humanoid robots have the potential to lead with a greater level of efficiency and effectiveness than human leaders [...but...] AI can provide unbiased data while humans can provide the emotional intelligence and creativity to make the best decisions” – Sophia, the humanoid robot, at the UN’s AI for Good Global Summit 2023





Exhibit 83: Image of Sophia, the humanoid robot  
 Sophia, the humanoid robot, can display humanlike expressions and interact with people



Source: ITU Pictures

BofA GLOBAL RESEARCH

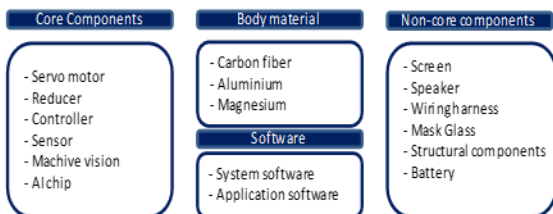
Regulation banning emotion recognition?

AI regulation could affect the evolution of embodied AI. The proposed EU AI Act encompasses comprehensive rules for trustworthy AI. It bans certain applications that pose a threat to citizens’ rights and democracy, e.g. emotion recognition in the workplace and educational institutions and also AI-driven facial recognition technology used in public spaces for surveillance or law enforcement purposes.

Humanoid robot supply chain

Humanoid robots are usually equipped with various types of sensors to measure key environment attributes such as position, orientation, speed, acceleration and contact force, and use cameras or LiDAR (light detection and ranging) for navigation. Inbuilt AI algorithms are responsible for path planning and behaviour control. Motion control is usually carried out by electric or electro-hydraulic actuators. A typical humanoid robot contains three major systems: (1) an actuation system, mainly including servo systems, controllers, and transmission machinery (e.g., reducer); (2) a sensory system, mainly including cameras and sensors; and (3) a computing system, including AI chips and relevant software for navigation and motion control. Other non-core components include batteries, screens, wiring, speakers, etc.

Exhibit 84: Humanoid robot supply chain  
 Core components include servo motors, reducers, controllers, sensors, etc.



Source: BofA Global Research

BofA GLOBAL RESEARCH

Exhibit 85: Comparison between different actuators  
 Electric actuators are most widely used in humanoid robots

	Electric actuator	Electro-hydraulic actuator	Pneumatic actuator
Cost	Medium	High	Low
Output torque	Medium	High	Low
Speed	Fast	Slow	Medium
Precision	High	Medium	Low
Structure	Relatively simple	Complex	Simple
Application	Widely used in industrial robots and humanoid robots	Heavy-duty works	Small-sized robots

Source: BofA Global Research

BofA GLOBAL RESEARCH

- Actuators control motion: These are the core component for the motion control of humanoid robots. It can take the form of electric, electro-hydraulic and pneumatic actuators (Exhibit 85). An electro-hydraulic actuator has the highest output torque but is also costly and runs the risk of oil leakage. A pneumatic actuator, on the other hand, is less expensive but also has lower

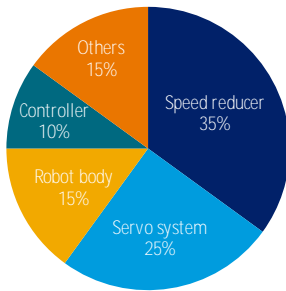


precision and output force. An electric actuator has now become the mainstream technological route for humanoid robots due to its high precision, fast speed and reasonable price. An electric actuator is comprised of a servo system, controller and transmission machinery.

The servo system utilizes torque and power generated from the motor to directly or indirectly drive the robot body. Our APAC Industrials team estimate that Tesla’s Optimus has 14 rotary motors and 14 linear motors in its body. Its hands are likely to use coreless motors to save space and costs.

The transmission machinery in actuators is in charge of adjusting the speed and torque output from the motor. Unlike industrial robots that use RV/harmonic reducers for speed reduction, humanoid robots usually adopt various types of transmission machinery in different joints due to weight and cost restrictions, as well as safety concerns. Our APAC Industrials colleagues estimate that Tesla’s Optimus use 14 harmonic reducers in its body. It uses four-bar linkage in its knees and worm drives in its fingers. As such, the cost share of reducers in humanoid robots is likely to be much less than the c.35% in industrial robots.

Exhibit 86: Production cost breakdown of industrial robot  
Speed reducer accounted for 35% of total cost in 2020



Source: BofA Global Research, OFweek

BofA GLOBAL RESEARCH

- Sensors help robots understand the environment including temperature, object proximity, contact force, acceleration, etc., and thus act accordingly. The sensory system of humanoid robots includes machine vision and various types of sensors.

Humanoid robots’ machine vision system is usually based on ToF camera (time-of-flight camera) or stereo camera to produce 3D image of the objects. ToF camera uses infrared light to detect surrounding objects and measures the distance based on the round-trip time of the light. Stereo cameras work in a similar way to human eyes. They capture the views from two different directions and translate them into 3D images. ToF cameras typically have longer working distances, higher accuracy and can be used in dim light conditions, but they also have lower resolution and require higher hardware costs. Stereo cameras, on the other hand, have higher resolution but also higher requirements for environment light and system processing capacity. LiDARs are unlikely to be used in the vision system due to their cost and power inefficiency.

- Computing system: Edge computing is also crucial for humanoid robots, enabled by high-performance GPU with lower energy consumption. The design of battery and electrical components also needs to consider power efficiency and physical space limits. To keep production costs reasonably low, manufacturers of humanoid robots tend to produce key components in-house, and monthly production volume needs to be at least 100 to be profitable.

Exhibit 87: Major types of sensors  
Sensors help robots to understand the environment and act accordingly

Name	Usage
Sound Sensors	Recognize sound and convert into electrical signal
Light Sensors	Identify the light
Temperature Sensors	Measure ambient temperature and help robots to adapt to the environment
Pressure Sensors	Measure force and pressure
Tilt Sensors	Detect tilt of the robot
Contact Sensors	Detect physical contacts
Acceleration Sensors	Measure tilt and acceleration
Proximity Sensors	Detect the presence of objects within predefined distances
Gyroscope	Measure and quantify the orientation
Inertial measurement unit (IMU)	Measure velocity, orientation, and gravitational forces together

Source: BofA Global Research

BofA GLOBAL RESEARCH





## Into the real world...

“It will upend our idea of what the economy is... it will be able to do basically anything humans don't want to do... it's going to bring an age of abundance.” Elon Musk, CEO of Tesla, at the unveiling of Tesla's Optimus, Oct 2022

### Robotopia: Humans & Robots NOT Humans vs Robots

The collaboration between man and machine will be key to the future of work and jobs. While the consensus view may be 'Long Robots, Short Humans', we believe this is not a zero-sum issue and is more nuanced. The global robot installed base could double from 2019-25, hitting 5 million units. Yet, technological automation raises the need for coexistence between humans and industrial robots ('cobots'), one of the fastest growth areas (c.50% CAGR vs c.5% CAGR for traditional robots). However, there is a scarcity of human talent with advanced skills due to a mismatch between technological advances, including automation, artificial intelligence (AI), and machine learning, and the skills and experience workers need to leverage these advanced tools.

### Physical labour / warehouses

Humanoid robots have higher levels of flexibility than typical service robots, usually with >20 degrees of freedom (refers to the motion capabilities of robots, typically the number of joints or axes of motion). They are generally able to perform various (usually unstructured) tasks. However, higher levels of flexibility translate into higher motion control requirements, and limits on height and weight.

### People-facing roles / domestic applications

Robots have typically been used exclusively for manufacturing, e.g. autos, but they can now work in warehouses, grocery stores, restaurants, hospitals, car parks etc. This is because they can interact with people more naturally today, making them better suited to customer-facing roles. Other applications include security guarding, healthcare, warehouse management, household applications (housework, elderly care), point of information (e.g. hotels, airports), education (at school or home), medical and care home deployments for companionship.

### Robonauts: humanoid robots in space

NASA's humanoid robot, Valkyrie, officially known as R5 (Robonaut 5), is being designed to operate in degraded or damaged human-engineered environments so that it might one day operate in space. Contrast this to the first two iterations, R1 and R2, which were used to perform repetitive tasks on the ISS. Valkyrie is being developed to mine resources, build habitats autonomously on Mars and work alongside astronauts. The robot has 44 degrees of body freedom, weighs 300lbs, is 6ft 2in tall and runs on two Intel Core i7 chips (source: NASA).



**Exhibit 88: Image of NASA's Valkyrie (R5) robot**

Valkyrie is designed to be a robust, rugged, entirely electric humanoid robot capable of operating in degraded or damaged human-engineered environments



Source: NASA, Bill Stafford, James Blair, Regan Geeseman

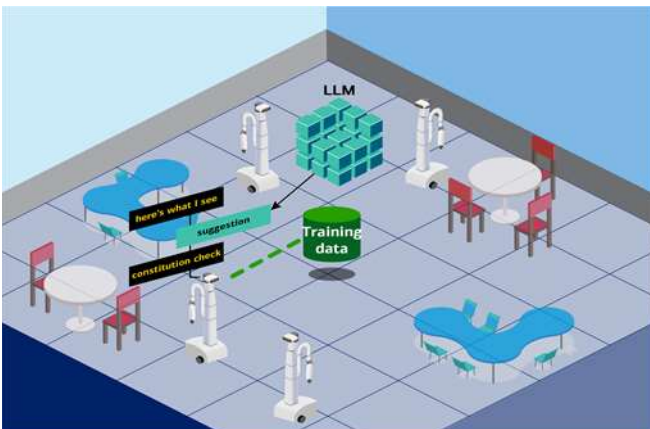
BofA GLOBAL RESEARCH

**Case study: DeepMind's AutoRT**

A challenge with training embodied foundation models is the lack of data grounded in the physical world. In January 2024, Google's DeepMind proposed AutoRT, which is an embodied foundation model built to command other robot agents. It leverages foundation models to scale up the development of operational robots in scenarios with minimal human supervision. It uses vision-language models (which process images and natural language text to perform various tasks) for scene understanding and grounding and uses LLMs to propose instructions for other robots. It was capable of proposing instructions to over 20 robots across multiple buildings. Some of the instructions included 'pick up tea bag', 'peel the banana', 'close drawer' and 'place bottle upright'. Most importantly, it resulted in AutoRT collecting diverse and large-scale data that could align to human preferences: it collected data on 77,000 real-world demonstrations via 20 robots in 7 months in 4 buildings and it was programmed with the Robot Constitution.

**Exhibit 89: Illustrating AutoRT in action**

AutoRT, a type of LLM, collecting data from and proposing instructions to over 20 robots



Source: Source: Ahn et al; DeepMind; AutoRT. NOTE: AutoRT is an embodied foundation model that can give instructions to robots

BofA GLOBAL RESEARCH

**Exhibit 90: ...some of the 20 robots in action**

Tasks included picking up tea bags, peeling bananas, placing water bottles upright



Source: Ahn et al; DeepMind; AutoRT

BofA GLOBAL RESEARCH



## Industrial Robots

What are they? Machines that can be automatically controlled or programmed to manipulate and move objects – they are a key enabler of the emerging Industrial Internet (smart sensors, actuators and other devices to enhance manufacturing and industrial processes).

Did you know? there are more industrial robots world-wide than the population of Uruguay.

Source: Based on Stanford; 2023 UN Population Division

Industrial robots were first introduced in the 1960s and are a key enabler of the emerging Industrial Internet and Industry 4.0. Integrating robots, big data and software analytics can optimise the full manufacturing chain, connecting virtual design, simulation, manufacturing and supply chain management. This integration improves efficiencies and product quality, while increasing manufacturing flexibility and lowering turnaround time. In manufacturing, embodied AI can optimise production processes and logistics by deploying industrial robots with AI driven perception, navigation manipulation and decision-making capabilities.

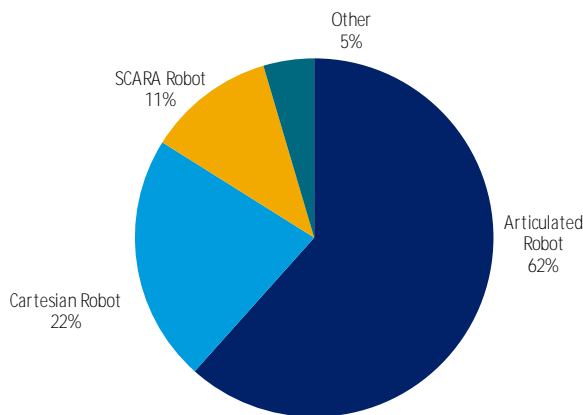
Exhibit 91: Overview of the four main types of industrial robot  
Articulated, SCARA, Cartesian and Parallel

Robot Type	Description	Cost	Size	Payload	Cycle Time	Applications
Articulated	Robot with several rotary joints ranging from 2 to 10. Suitable for various applications within industrial production given its flexibility. Can move in all three dimensions.	High	Big	High	Fast	Welding, assembly, machine tending, painting
SCARA	“Selective Compliant Assembly Robot Arm” is a simpler form of the articulated robot which only works on the horizontal axis as it lacks rotary joints.	Medium	Small	Medium	Fast	Assembly, drilling, pick and place
Cartesian	Can move in all three dimensions, however only in a straight line without rotating. Cartesian robots operate behind security barriers and are well-suited for heavy payloads. Consists of a joint base which holds three connected arms.	Highest	Biggest	Highest	Medium	Heavy duty applications (e.g. automotive assembly)
Parallel	The arms are arranged in a parallelogram shape, helping to maintain the direction of the end-effector.	Low	Small	Low	Fastest	Light weight, pick and place

Source: BofA Global Research

BofA GLOBAL RESEARCH

Exhibit 92: Breakdown of the global industrial robot by type of robot  
Articulated and Cartesian robot markets together account for c.84% of the overall global industrial robot space



Source: TechNavio

BofA GLOBAL RESEARCH

Exhibit 93: Industrial robot uses

Robots are used for three main functions in industry: material handling, welding and other uses such as painting and assembly

### Material handling

- Pick & Place** Robots are used to pick up and place elements in a new area. The robot is able to increase efficiency and accuracy in these repetitive processes.
- Machine-tending** Robots oversee machines that are running and feed parts in and out. While this is a dull and dangerous job for humans, it still requires a high level of consistency and is therefore well-suited for the application of industrial robots.

### Welding

- Laser welding** These robots use servo-controlled, multi-axis mechanical arms, which have a laser-cutting head mounted to the faceplate of the arm. Deploying industrial robots for laser welding can improve repeatability and quality of the welds.
- Arc welding** Arc Welding robots use a process that applies intense heat to metal at a joint, causing it to melt and intermix. Advantages include higher weld consistency, shorter cycle times and enhanced efficiency and safety.

### Other

- Painting** Painting robots are used to protect human workers in unsafe environments where they would be exposed to isocyanides, carcinogens, etc.
- Assembly** Assembly robots replace workers in repetitive assembly line jobs, increasing consistency and speed and lowering production costs.

Source: BofA Global Research

BofA GLOBAL RESEARCH



### Analogies with the human body

Equivalent to human intelligence, robots derive the information to operate through hardware controllers and software. Human arms and hands are represented by end effectors through which the robot ultimately performs its functions such as pick & place, welding and material handling. Actuators reflect human muscles, which set other parts like arms and end effectors in motion and transfer mechanisms act as robotic legs. Finally, sensors, like the human senses, receive feedback from the arm and end effector position and velocity, and transmit this information back to the controller.

**Exhibit 94: Human and Robot equivalents**

Equivalent to the human intelligence, robots derive the information on which they operate through hardware controllers and software

Human	Robot
Intelligence	Information, PLC, HMI, Software
Arms	Arms
Hands	End Effectors
Muscle	Actuators
Legs	Transfer Mechanism
Sense	Sensor (External Sensor, Internal Sensor)

Source: BofA Global Research

BofA GLOBAL RESEARCH

**Exhibit 95: Key components of robot**

Per our Capital Goods colleagues, the most expensive component could be the Reducer

Key parts	Entry barrier	Margin	Local sourcing %	Cost break-down
Controller	Medium	High	Low to Medium	15-20%
Reducer	High	High	Low	30-40%
Servo motor	Medium	Medium	Low to Medium	25-30%
Driver	Medium	Medium	Low to Medium	10-15%
Body / Arm	Low	Low	High	~15%

Source: BofA Global Research

BofA GLOBAL RESEARCH

### The importance of machine vision...enabling robots to 'see', 'think' and 'act'

How it works: The camera captures the dimensions and orientation of objects in front of the robot, and the attached software analyses the data. The data is then transmitted to the robot controller, which guides the robot to generate certain motions, such as picking up an object and placing it in a bin. Without this technology, industrial robots would be "blind" to the surrounding environment, including the objects they handle. However, with the advancement in machine vision systems, robots can be used in dynamic environments such as a warehouse. With new "eyes" and "brain", robots can operate with more flexibility and adapt to the changing environment.

Both machine vision and computer vision are important divisions of artificial intelligence. Machine vision is mainly applied in intelligent manufacturing, while computer vision is applied in many scenarios in the consumption sector.

Edge computing and deep learning models for AI are expanding what machine vision can do. When a computer receives a form of media (e.g., image or video), machine vision software can compare that image data with a neural network model. Deep learning inference means that computers can recognise subtle differences in patterns, for example. And the capabilities go beyond visual inspection and quality control. With smart machine vision, robots can use NLP to read labels and interpret signs, understand shapes, calculate volumes, optimally package objects – ultimately driving a shift from automating simple tasks to being autonomous machines (source: Intel).

**Exhibit 96: Comparison of industrial vision and computer vision**

Machine vision has higher requirement of hardware such as cameras, sensors, etc.

Machine vision	Machine vision -Industrial use	Computer vision-consumption & services use
Subject	System engineering	Computer science
Application	Intelligent manufacturing	Smart living areas such as consumption and services
Functions & goals	Replacing human eyes to perform repetitive tasks such as positioning, measuring, and testing of workpieces	Endowing Intelligent robots with vision to realize the recognition and judgement of external position information and image information
Content	Image acquisition, lens control, image processing	Image processing
Focus	How to ask the robot to identify according to its tasks	How to analyse images
Requirement on hardware	High, need to screen industrial cameras based on the frame rate, resolution, and other indicators	Low requirement on cameras and webcams in most cases
Requirement on algorithm	Focus on improving accuracy	More complicated, focus on the calibration and recognition of objects using mathematical logic or deep learning
Maturity	High, especially in the testing and gauging of semiconductor, packaging sectors	In initial exploration with many start-up

Source: GGII, Qianzhan, Shuangyi Tech, BofA Global Research

BofA GLOBAL RESEARCH



Structural driver: demographics and manufacturing skills shortage

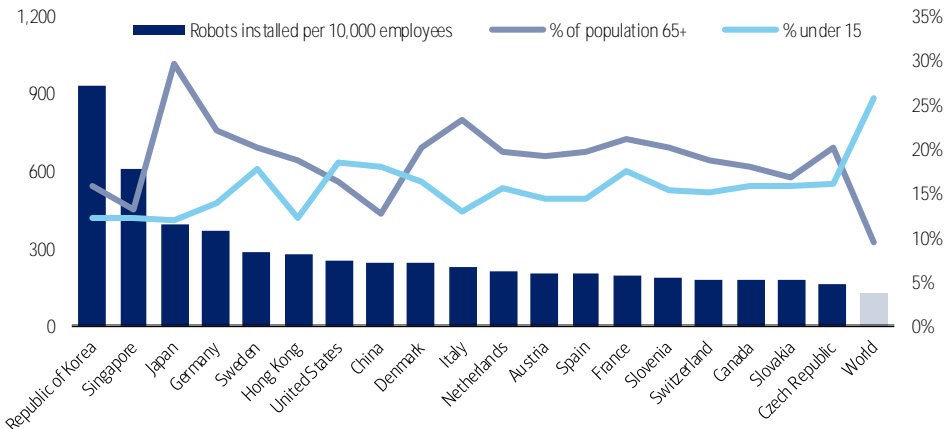
Did you know? The total number of under 15s in 2100 could be as large as the current population of Bhutan. The number of 65+ year olds could almost quadruple by 2100

Squeezing the middle: ageing & peak youth mean narrow working population - There is a scarcity of youth: the number of children is projected to peak at 2.09 billion in 2057 and fall below 2 billion by the end of the century (source: UN). We already have more grandparents than grandkids due to a negative natural population balance. Between 2021 and 2050, the number of people aged 65 years or over globally is projected to more than double, while the number under 25 is projected to peak and then decline slightly. As the population is ageing and fertility is decreasing, fewer workers are available, leading to an increasing skills and labour shortage.

Manufacturing is headed towards a talent shortage of 8 million workers by 2030 77% of manufacturers say difficulties in attracting skilled workers are expected to continue (source: Deloitte). The skilled labour crunch in manufacturing will cost the global economy \$607.14bn, with a deficit of 7.9 million workers by 2030, 39x the number of Ford employees. China is expected to have an abundance of manufacturing talent, a sector accounting for 35% of its economy. But this abundance is peaking and, by 2030, it will face a skilled labour deficit. The US, being the world’s most important manufacturing economy, will suffer from a skilled manufacturing talent shortfall of 2.5 million workers in the next decade. Japan is also projected to fail to realize \$194.61bn by 2030, 3% of its economy, worsened by low birth rates, an ageing population and tight immigration restrictions. Germany will see the next biggest impact after Japan in its unrealised output due to labour shortages in manufacturing. Although Germany is a leading manufacturing hub today, by 2030, unrealised revenue due to labour shortages could reach \$77.93bn. Although automation is being implemented in manufacturing, the global sector lacks the necessary skills to execute automation projects.

Exhibit 97: Countries with the highest robot density in the manufacturing industry and their % of population 65+ and under 15 vs the World

Typically countries with a larger ageing population and fewer children have installed the most robots



Source: UN; World Robotics; International Federation of Robotics

BofA GLOBAL RESEARCH

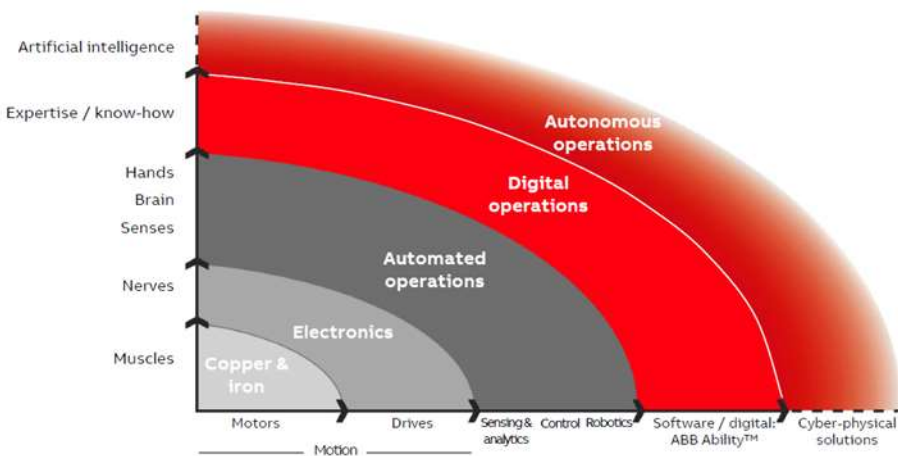
Industrial artificial intelligence

Machine learning has been around for decades, but the speed of analysis and data collection has been brought to a point where the technology can have commercial application. Artificial intelligence is the science and engineering associated with creating machines that can learn and respond to situations, based on history and new data,



without specific programming or human interaction. Deep learning is the next level of machine learning – essentially more detailed data analysis and complex algorithms to develop machine learning capabilities. Machine learning is essentially an extension of prescriptive analytics that is highly automated, iterative and capable of adapting without human interaction. Artificial intelligence also needs to be embedded within the process or workflow, meaning it is largely hidden from the human operator.

Exhibit 98: Industrial automation innovation dynamics - Industrial artificial intelligence Towards autonomous operations leveraging AI

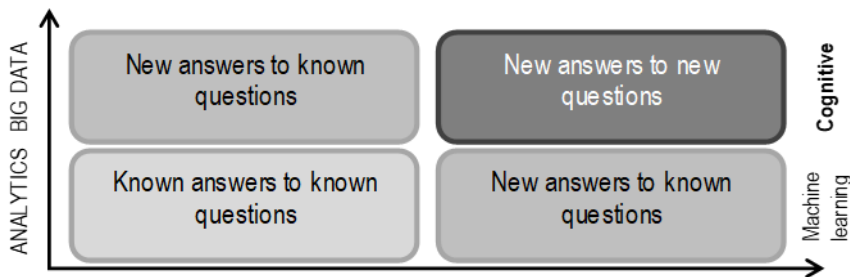


Source: ABB

BofA GLOBAL RESEARCH

However, artificial intelligence and cognitive technologies are likely to be able to markedly increase manufacturing efficiency throughout the value chain. Analysis of product data can enhance design and recommended use. Computer vision can be used for compliance and quality control. Natural language processing could improve human/machine interaction and automated decision-making can reduce human error, speed up the decision process and improve productivity. Machine learning can be used to identify patterns across multiple data sets to improve the manufacturing process.

Exhibit 99: Defining cognitive technologies Towards 'new answers to new questions'



Source: BofA Global Research, Frost & Sullivan

BofA GLOBAL RESEARCH

Industrial internet serves as the foundation for intelligent manufacturing. In the era of Industry 4.0, intelligent manufacturing refers to the utilization of real-time data analysis, artificial intelligence (AI), and machine learning (ML), edge technologies, and cloud, in the manufacturing process to accomplish optimization in production, waste reduction, yield maximization, and cycle-time reduction. AI plays an important role in intelligent manufacturing as it enables business owners to monitor equipment and processes correctly. General AI is widely used in various applications, as it processes high dimensional data such as picture, voice, and text. In contrast, Industrial AI processes low-dimensional data such as temperature, pressure, vibration, etc.





The industrial AI related products/services into four categories:

- Basic industrial software/hardware, such as chips, computing modules, industrial camera, AI framework, etc.
- Intelligent industrial equipment, such as robots, machine tools, mobile robots and other specific equipment integrated with intelligent algorithms.
- Automation & edge technology, including industrial control systems integrated with intelligent algorithms.
- AI platform/software/solutions, including traditional stand-alone software, software embedded with AI, industrial internet platforms/solutions with AI function.

Fanuc and NVIDIA partnership to produce AI-powered robots  
 In 2016, Fanuc announced a partnership with NVIDIA to use its graphics processing units (GPUs) and AI platform to create futuristic robots that can collaborate and learn from each other. Adding AI to the FIELD system (platform to improve factory production and efficiency with advanced AI) gives robots the ability to teach themselves to do tasks faster and more efficiently. By learning together, what one robot can now do in eight hours could be done by eight robots in an hour.

### Digital twin

A digital twin is a dynamic software model of a physical thing or system that relies on sensor data to understand its state, respond to changes, improve operations and add value. The concept of a digital twin has been around since 2000/2001 but with today's computing power, the development of the cloud and the growing installed base of sensors digital twins are now a reality. Digital twins are likely to be used to develop proactive repair and service schedules as well as plan and optimise manufacturing processes and factories. Most equipment and systems are now designed on computer, but the disconnect between what is out in the field and what is thought to be out in the field can often be the reason for inefficiency or equipment failure. Better knowledge of the actual equipment or system and the way it operates allows the development of dynamic digital models of the turbine, rail locomotive, ship, wind turbine, aerospace engine etc.

In essence, the digital twin is a virtual representation of the behaviour and structure of a physical product or process, providing numerous potential benefits, such as performance optimisation, predictive maintenance and reduced manufacturing downtime. To GE, digital twins consist of three components: data, algorithms and knowledge. It's important not to consider the digital twin as a single technology, but rather a combination: including a plant database, software (including artificial intelligence), sensors used to provide real-time feedback and much more. The concept is nothing new. However, the convergence of technological advancements in recent years, increases in computer power and number of connections, reduced data storage cost and capacity improvements, advancements in deep-learning modules and the advent of cloud data, have increased the efficacy and potential value of the simulated environment.

The ability to use the digital design, essentially a "best guess" of reality, however well-refined, with "actual" reality is the key to opening up a whole new level of operational efficiency and design refinement. It is interesting, therefore, to understand who uses digital twins already and where the real business opportunity lies. The reality appears to be that the use of digital twins remains relatively limited. In fact, close to half the market still does not really understand what a digital twin is or how it can be used.

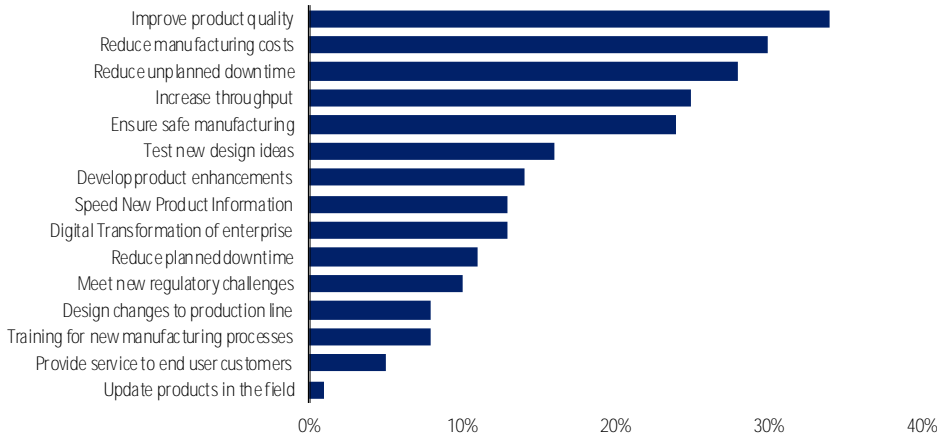




Advancements in deep learning, where the use of neural networks (essentially interconnecting nodes, similar to neurones in the human brain) allow the AI to make determinations about a data set, are then validated and used to make further determinations about new data. An example would be image recognition and the method of analysis lends itself well to applications such as detection of defective products on a production line. AI is attracting an increasing amount of attention, with initiatives like Google DeepMind making headlines. The extent of our daily interactions with AI might be surprisingly numerous to some, from chatbots to internet search.

**Exhibit 100: Digital twin use cases**

In the majority of cases, digital twins appear to be used to improve product quality, with reduction of manufacturing costs and unplanned downtime close “runners up” in terms of reason for implementation



Source: LNS Research

BofA GLOBAL RESEARCH

**Rise of collaborative robots – “cobots”**

Did you know? There were only 41 robot-related worker fatalities in the US between 1992 and 2017

Source: Census of Fatal Occupational Injuries

Largely due to their size and power, traditional industrial robots are typically installed in caged environments with minimal human contact for various safety reasons. The integration costs associated with these machines limit their flexibility and they often require programming from advanced software engineers, which drives up the total cost of ownership. Due to advancements in technology and motion-sensing capabilities, a new type known as collaborative industrial robots (or “cobots”) has emerged. These are defined as small and nimble robots, specifically designed to move around the factory floor to assist human workers. Cobots are typically smaller, cheaper and easier to program than traditional industrial robots. The range of their capabilities has evolved over recent years meaning the cobot is now a more competitive option for a customer looking to automate a factory without the need for large areas of fenced-off factory floor for their large industrial cousins.



Exhibit 101: Image of a collaborative robot  
An example of Universal Robots' cobot



Source: Universal Robots

BofA GLOBAL RESEARCH

Exhibit 102: Image of a collaborative robot  
An example of Fanuc's cobot



Source: Fanuc

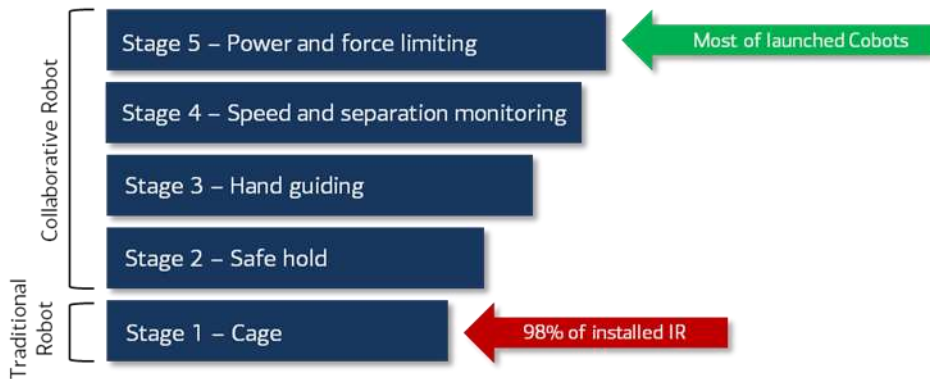
BofA GLOBAL RESEARCH

### ISO 10218 requirements for a robot to be a 'cobot'

The four requirements are as follows:

- Safety-rated monitored stop: employing sensors to detect human beings and to stop working when a worker enters a particular area.
- Hand guiding: operators can physically guide cobots through a task.
- Speed and separation monitoring: using sensors to keep a distance from human workers. Stopping work when a human enters a predefined area and resuming when they leave.
- Power and force limiting: these cobots are restrained in the force they can exert on their environment.

Exhibit 103: Most cobots on the market now have a full range of capabilities  
There are four characteristics that a robot must exhibit to be defined as a cobot (safe hold, hand guiding, speed and separation monitoring, power and force limiting)



Source: BofA Global Research estimates, company report

BofA GLOBAL RESEARCH

### Technological advancement leading to reduced cobot costs

Technologies advancements and digitalisation has led to new and diverse robotics applications across several industries: force sensing, tactile sensing, speech recognition, machine vision, etc. Furthermore, technology improvement has increased flexibility and payloads, thereby reducing the total cost of ownership for the manufacturer. Cobots are designed to conduct more complex tasks than traditional robots because they are powered by AI and machine. Technological advances, however, will make cobots more

useful and spur manufacturers to invest due to rising labour costs and growth in edge computing. One of the fastest-growing areas of sensor development in automation and robotics is perception. Machine vision technology, laser scanners, structured-light 3D scanners, and the imaging and mapping software to support them are opening doors to robots in more industries. Examples of technological advancements include (source: FutureBridge):

- Perception sensors and AI: Having multiple sensors to enable perception of the environment. Combining AI with perception sensors would allow cobots to truly interact with humans.
- AI and ML integration: It is not enough for cobots to have sensors. Cobots need to make decisions based on the behaviour of the human worker too. This means they must learn from their working patterns before acting to suit their working styles. This is to ensure the safety of each worker in the vicinity of the cobot. Having AI and ML integration means that cobots can ensure a proper analysis of the environment, allowing them to better adapt to the workers around them, whilst maintaining efficiency and productivity.

In August 2023, Doosan Robotics announced a partnership with Microsoft and Doosan Digital Innovation to give cobots a context-aware system based on GPT (generative pre-trained transformer) Technology. Azure will provide Doosan with the GPT language model and Doosan Digital Innovation will adapt it to be used in Doosan's robots. This should shorten programming times while expanding the available functions for their robots.

- Connectivity: For high productivity, robots should be able to learn from similar examples around the world and adapt decisions to ensure the same output with the least training. In this way, cobots could store examples of decisions they make along with their scenarios in a centralised repository so that other cobots can learn what they have learnt.
- Edge computing and AI: Cobots need closer control for real-time implementation of complex decisions in the work environment. The decisions may need to be made on the end node where the data is generated and unaltered. These situations may be unique to where the cobot is located and the decision to be made may need to be unique to that particular scenario.

#### Benefits of cobots

Cobots bring many benefits to the factory floor in terms of optimised production, increased competitiveness and an improved work environment where humans are freed up to focus on more rewarding tasks. This redeployment of human creativity interspersed with the robots' repeatability addresses market evolution and customer requirements for a high degree of product individualisation.



Exhibit 104: Some players in the cobot ecosystem  
 Players span across manufacturers and tech enablers



Source: FutureBridge

BofA GLOBAL RESEARCH

Costs depends on payload

The cost of cobots depends on payload capacity (the amount of weight a vehicle can carry). More than half the total cobot market is dominated by cobots with a payload size of less than 5kg and less than 10% have payloads of more than 15kg (source: FutureBridge). Per FutureBridge, the average price of a cobot for 0-5kg/5-10kg/10-15kg payload is c.US\$23,000/34,000/41,000. Overall, cobots typically cost US\$10-50,000.

Smart warehouses/Warehouse automation (WA)

Logistics is a broadly defined term that refers to the management of the flow of goods from the point of origin to the customer. In sales logistics, for example, manufactured goods are transported from a factory to a warehouse, where the goods are stored before being shipped to a retail store or directly to the end-consumer. A warehouse is integral in logistics as the place to store goods, and warehouse management affects everything from the efficiency of inventory to the timeliness of delivery.

Exhibit 105: Image of an automated guided vehicles (AGVs)  
 AGVs lift shelves with goods and transport them to the warehouse worker, thus eliminating the need for workers to walk across the warehouse and aisles



Source: Amazon

BofA GLOBAL RESEARCH

Exhibit 106: Image of an articulated robot  
 Articulated robots palletize and sort heavy packages



Source: Amazon

BofA GLOBAL RESEARCH

Prior to the outbreak of the Corona virus, supply chain automation was already primed for growth due to evolving consumer spending habits. The growth in online shopping is

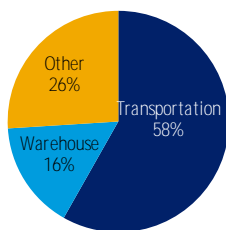


testament to shifting consumer preferences for greater consumer spend. Speed has become a particularly competitive dynamic, and the new normal is constantly being redefined (shipping times are currently 2.5 days with groceries delivered within a few hours). It is important that supply chain equipment providers and operators are well placed to meet the needs of today’s consumer, whilst positioning themselves with the right technology to be competitive with tomorrow’s.

**Much more to be automated in picking and processing**

Picking is the process of gathering or “picking up” items from inventory, based on the shipping instructions, before they are shipped to consumers. Traditionally, picking involves a warehouse worker walking to a shelf (or driving a forklift, depending on the size of the items handled), searching for the appropriate items, and picking the items, after which they are packaged and inspected for delivery. Labour costs make up 50-70% of warehouse operational costs, according to third-party logistics provider Kane Logistics. Moreover, of the processes involved in warehouse logistics, picking and processing are the most labour-intensive, according to industry sources. The labour intensity is also suggested by the fact that only a small portion of automation equipment sales is used for picking (Exhibit 107). As such, there is room and incentive to cut picking costs and increase efficiency via further automation.

Exhibit 107: Breakdown of logistics costs  
The largest logistical cost is transportation followed by warehouse



Source: Japan Institute of Logistics Systems, BofA Global Research

BofA GLOBAL RESEARCH

Exhibit 108: Comparing machine vision systems to the human body  
Machine vision system acts as the “eyes” and “brain” of robot



Source: Cognex

BofA GLOBAL RESEARCH

Various types of equipment are used in the picking process to reduce physical labour. For example, autonomous mobile robots have decreased the need for warehouse workers to walk from one location to another, and hand terminals have cut the time it takes to find the appropriate shelves or items. However, the actual “picking up” process is still done manually for the most part.

Industrial robots with arms, or articulated robots, may be the most suitable candidates for picking and processing automation. However, there are several technical challenges when trying to utilise industrial robots in warehouses.

**Machine vision can address the challenges in having industrial robots in WA**

Traditionally, industrial robots are good at performing repetitive movement with accuracy and precision. That is why they were introduced at relatively large scale in factory automation, where they are programmed to repeat certain movements (for example, welding) with designated work pieces. In a warehouse, however, robots must handle goods of infinitely different sizes and shapes under different conditions. It is neither sufficient for the robot to repeat the same movement in such dynamic environments, nor is it feasible to pre-program the robot when the conditions vary to such an extent. Thus, contrary to its contribution in factory automation, traditional industrial robots have historically not been used in warehouse automation.



The machine vision system is a key technology that can be used to enable industrial robots to operate in a warehouse setting. It consists of (1) cameras that measure the physical dimensions of objects, and (2) software that uses the visual data to control the surrounding material handling equipment. In other words, the machine vision system acts as the “eyes” and “brain” of the robot (Exhibit 108). The machine vision system has often been used for quality inspection in factories, but its uses are versatile – in recent years, it has garnered attention in the realm of robots for picking and processing.

#### Collaborative robots expected to play a major role in warehouses

Among the different types of industrial robot, cobots are especially likely to be introduced in a warehouse for the following reasons: (1) warehouses, even with a high degree of automation, often require human workers to engage in certain parts of operations that cannot be fully handled by robots. As such, robots that can work side by side with humans have an advantage. (2) Cobots can be installed without securing space for safety fences, which reduces initial cost and is optimal for already-existing warehouses with limited space.; (3) cobots optimise the storage, distribution, and delivery of goods, while increasing the speed and accuracy of the operations.



## Other embodied AI interfaces

What are they? Other embodied AI interfaces include service robots, autonomous vehicles, agricultural robots, drones, spatial computing/virtual reality etc.

Did you know? AI-powered bee robots, weighting a tenth of a gram, could be used to pollinate crops.

- Service robots operate semi- or fully autonomously to perform services useful to the wellbeing of humans and equipment, excluding manufacturing operations. They typically assist humans by performing jobs that are dirty, dull, distant, dangerous or repetitive, including household chores. The rationale, use case and urgency for this technology may vary depending on the location and role, but some examples include preparing food, coffee, hospitality, entertainment, companionship, cocktails, home chores (such as vacuuming) or even looking after pets.

Exhibit 109: Robot Baristas – the future of your morning coffee?  
Several robots demonstrated tasks at CES; a robot barista designed by Richtech Robotics served coffee for example



Source: BofA Global Research

BofA GLOBAL RESEARCH

Exhibit 110: **“AI for all”** – robots are not just for humans  
Samsung demonstrated an update to its Ballie cobot as part of its “AI for all” presentation, its robot with a projector demonstrating pet care tasks as well as to support humans for smarter living



Source: Samsung

BofA GLOBAL RESEARCH





Exhibit 111: AI to keep the cocktails flowing  
Doosan showcased several tasks their robots can undertake including making cocktails. Enabled by a convergence of AI, 'cobot' and software technologies



Source: BofA Global Research

BofA GLOBAL RESEARCH

Exhibit 112: **Confused by recycling? There's a robot for that**  
multi-task with less reprogramming, allowing the same or similar equipment to be adapted from serving cocktails to sorting recycling for example



Source: BofA Global Research

BofA GLOBAL RESEARCH

- Autonomous vehicles: Self-driving is an AI problem – AV2.0 – where vehicles have the onboard intelligence to understand the surrounding environment and make decisions in real time, as opposed to the rules and mapping-based approaches that are adopted in most self-driving vehicle programmes in operation. Generative AI models can accelerate AV tech by providing vehicles with embodied intelligence to understand what's ahead and act on it. Furthermore, this “AV2.0” can be integrated into many new vehicles already in production without additional sensors/compute.

Exhibit 113: Wayve autonomous vehicle at BofA Transforming World Conference  
Wayve demonstrated one of its fleet at BofA's London office, a self-driving Jaguar i-Pace



Source: Wayve, BofA Global Research

BofA GLOBAL RESEARCH

Exhibit 114: Sunglass selfies - Ray Ban & Meta Smart Glasses  
EssilorLuxottica demonstrated its smart glasses with built-in cameras and audio capable of livestreaming, taking calls, capturing images (and more)



Source: BofA Global Research

BofA GLOBAL RESEARCH

- Agricultural Robots: Within precision agriculture, John Deere has developed an autonomous robot that uses AI and computer vision to differentiate between weeds and crops. In learning from its environment, it can decide exactly when to spray weed killer. This can reduce the chemicals used in agriculture.

- Drones: Integrating generative AI into drones can allow them to process vast amounts of data and allow for autonomous navigation. This can also impact the future of transport, delivery and logistics and surveillance.
- Virtual reality: Blurring the lines between physical and virtual reality with more mixed and extended reality (XR) tech in consumer, corporate and industrial applications, and to create new media/advertising experiences (Exhibit 114). Embodied AI can create immersive experiences and avatars that can respond to human gestures and expressions. Generative AI can make these virtual worlds seem more realistic, create digital avatars, act as a tool to create intricate stories.



# Communication: the scarcity of bandwidth

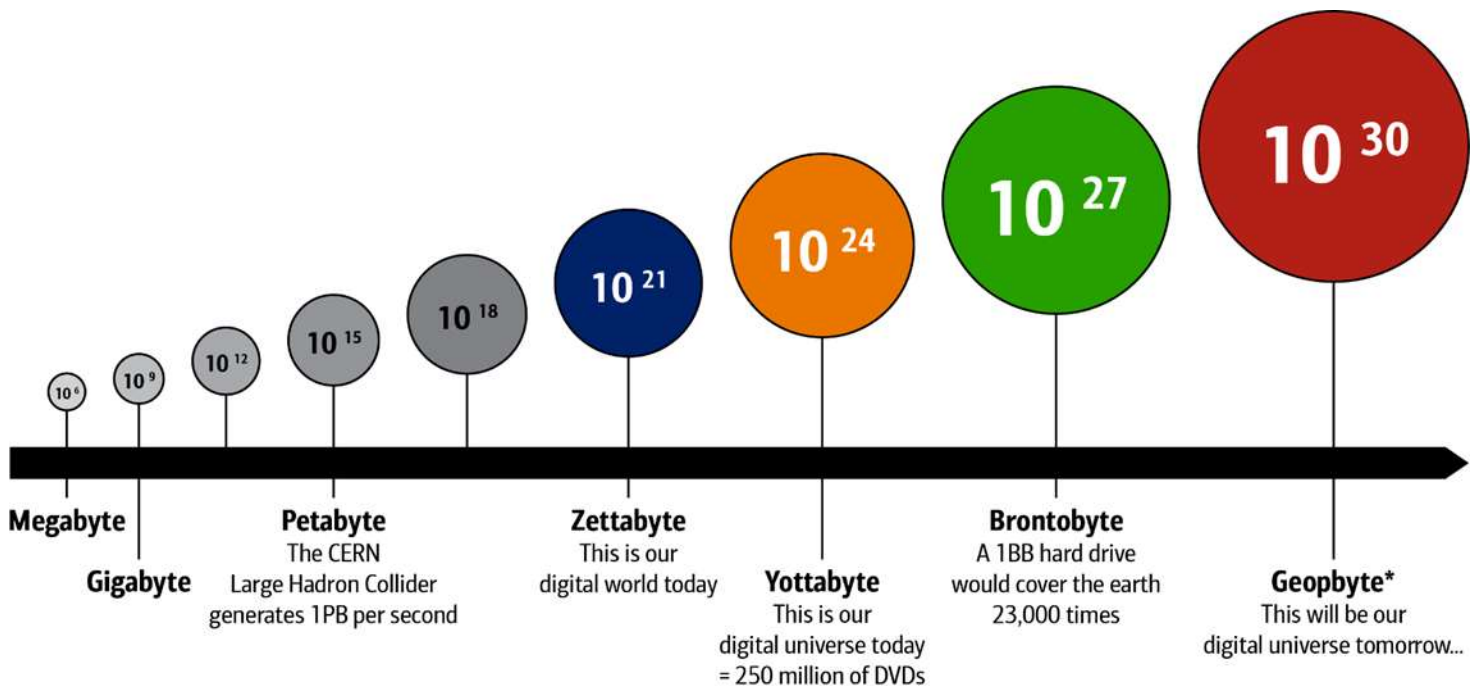
Transforming World: an exponential jump in data creation, with rising complexity and more data-heavy applications, will put mounting pressure on existing communication technologies, and require the adoption of new ones.

Did you know? By 2027 we will create more than sextillion (billion trillion) bytes of data... every day.

Next generation of communication: miniaturised satellites, Wi-Fi 7, 5G Advanced, 6G.

We live in an era of zettabytes, i.e. we inhabit a world that every day generates 328 quintillion (million trillion) bytes of data (source: statista) – and this number multiplies itself every 2-3 years. Data generation will only get faster from here as AI takes hold. An estimated 90% of all global data will be AI-generated by 2025 (Nina Schick).

Exhibit 115: The universe of bytes  
 Today we live in the era of zettabytes, i.e. we generate 'sextillion' bytes of data every year



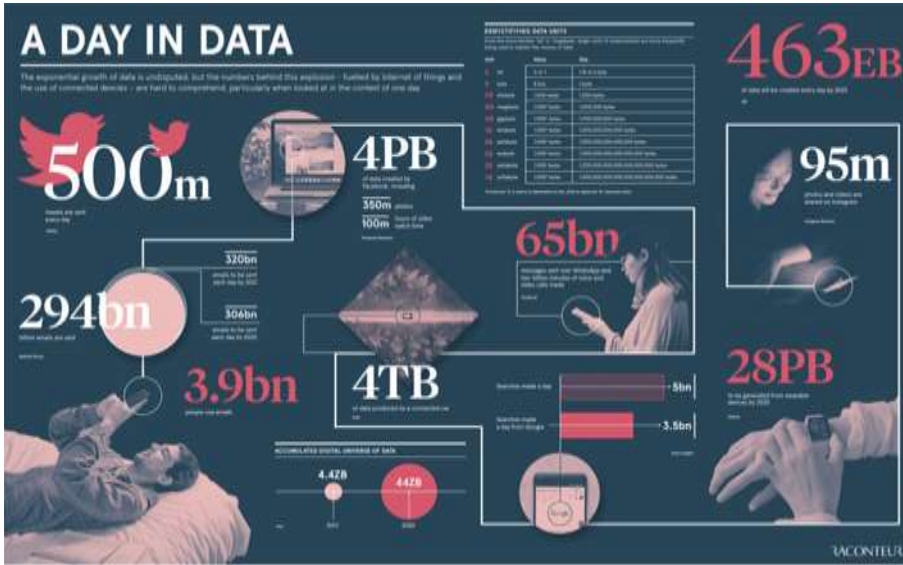
Source: HP, Statista, Our World in data, BofA Global Research

BofA GLOBAL RESEARCH

The digital universe has reached the level of the Yottabyte, with 90% of the world's data having been created in the past two years (source: IBM). As of the end of 2023, there were 8.6bn mobile phone users worldwide, of which about half used smartphones (WEF). According to IDC, the amount of data created is projected to double every 2-3 years, while according to Statista, we created c. 64ZB of data in 2020, 120zb by the end of 2023, and this is projected to grow to 181ZB by 2025E.

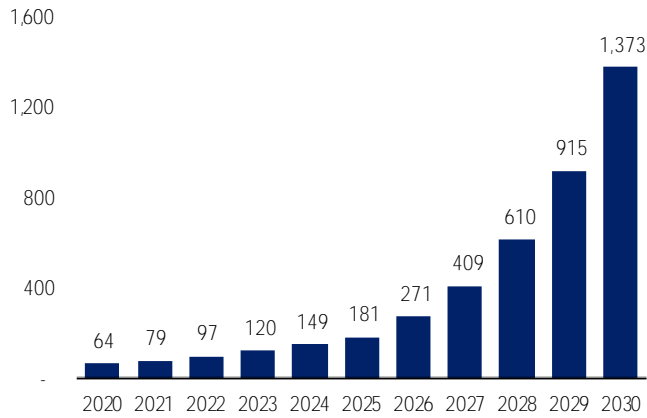


Exhibit 116: World of data  
The average person? generates 1.7MB... every second (source: IBM)



BofA GLOBAL RESEARCH

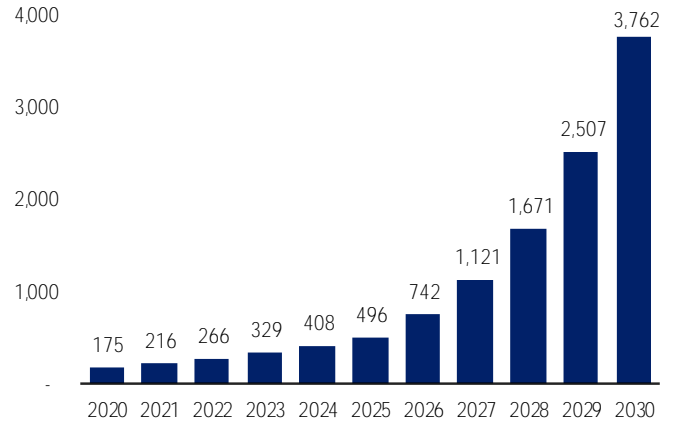
Exhibit 117: Global data generation per year (Zettabytes)  
In 2023 the world generated 120 zettabytes of data every year. By 2026 this number will more than double...



Source: Statista (2020-2025E), BofA Global Research factoring in global data multiples every 2-3 years

BofA GLOBAL RESEARCH

Exhibit 118: Daily global data generation (quintillion of bytes)  
...which means that if by 2023 we generated 329 quintillion bytes of data each day, we will jump to the "sextillion" level by 2027



Source: Statista (2020-2025E), BofA Global Research factoring in global data multiples every 2-3 years

BofA GLOBAL RESEARCH



## Miniaturized Satellites

What is it? Micro and Nano satellites, which are low cost, small, lightweight satellites, mostly in Low Earth Orbit (LEO, up to 1,000km altitude).

Did you know? In the 3 months from satellite broadband service launch in Ecuador’s Galapagos islands (30,000 inhabitant only), 100 new businesses were opened

Satellite technology is a key pillar of communication. In recent years, more broadband services have been based on this technology. We believe it will ramp up dramatically in coming years, as smaller satellites provide more affordable access to space and universal satellite internet access.

Communication is the main need, miniaturized satellites are key solution  
 Miniaturized satellites mostly refer to those in Low Earth Orbit (LEO) weighing up to 100kg. The two main categories are microsattellites (10-100kg), also known as microsat, and nanosatellites (also known as nanosats, or sometimes CubeSats due to their conventional 10x10x10cm cube shape, which is a common design), weighing up to 10kg. The small size of both micro- and nanosatellites offers advantage in terms of economics, capacity and coverage, but they are not suited to every mission and usually have a shorter lifespan than bigger satellites.

### Exhibit 119: Satellite directory

LEO is the largest category and accounts for 89% of all satellites in orbit.

Type of spacecraft	Mass
Large satellite	>1,000kg
Medium satellite	500-1,000kg
Mini satellite	100-180kg
Micro satellite	10-100kg
Nanosatellite (CubeSat)	1-10kg
Femto / Pico satellite	Less than 1kg

Source: NASA

BofA GLOBAL RESEARCH

### Exhibit 120: Nanosatellite key statistics

Nanosatellites account for approx. 40% of all satellites in orbit as of 2023 and are the fastest-growing category

Nanosatellites launched	2532
CubeSat Launched	2,323
PocketQubes launched	77
Interplanetary CubeSats	15
Most nanosats on a rocket	120
Countries with nanosats	85

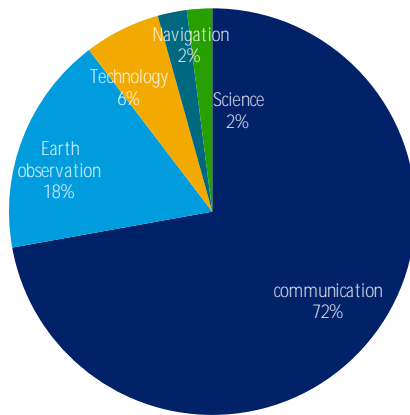
Source: Nanosats.edu, Erik Kulu

BofA GLOBAL RESEARCH

According to the Satellite Industry Association (SIA), the total satellite market generated \$384bn of revenues in 2022, with commercial satellites accounting for 73% of this (\$281bn). Demand is increasing dramatically with greater broadband communication and data needs. According to Orbiting Now, as of the beginning of 2024 some 9,368 satellites were orbiting our planet. The vast majority of these (89%) are Low Earth Orbit (LEO) satellites, operating at an altitude of 160km-2,000km above Earth. LEOs are mostly used for communication needs and observation. Of the 9,368 satellites in orbit, according to nanosat.edu approx. 3,700 are nanosatellites.



Exhibit 121: Split of non-military satellites by purpose  
Communication is by far the biggest use of satellite technology



Source: ucsus.org, BofA Global Research

BofA GLOBAL RESEARCH

Exhibit 122: satellites orbits around our planet  
More than 70% of all satellites are surrounding our planet in the LEO (low earth orbit)



Source: NASA, BofA Global Research

BofA GLOBAL RESEARCH

Satellite design, manufacturing, and deployment have been transformed by commercial off-the-shelf components and consumer electronics, enabling low-cost satellite platforms, instruments, and nanosatellites weighing less than 10kg, unlike micro-satellites which are defined as 10-100kg. Differentiating them by weight and size is only partially useful, as their usage also depends on other characteristics like minerality, equipment installed and communication.

These small satellites have the potential to make space more accessible and affordable, thereby improving profitability for an industry that is characterised by high capex and frequent delays. They are used across a wide range of applications – from Earth observation and communications to scientific research, technology demonstrations and education, as well as defence. Their lower cost means that companies are now looking to create mega-constellations in Low Earth Orbit to provide universal internet access.

#### Affordability through space ridesharing and miniaturisation

The appeal of satellites communication is primarily their lower cost compared with other telecommunication solutions. First, there is no need to deploy complex communication technology (fixed-line, towers, fibre etc), especially in areas where affordability is low.

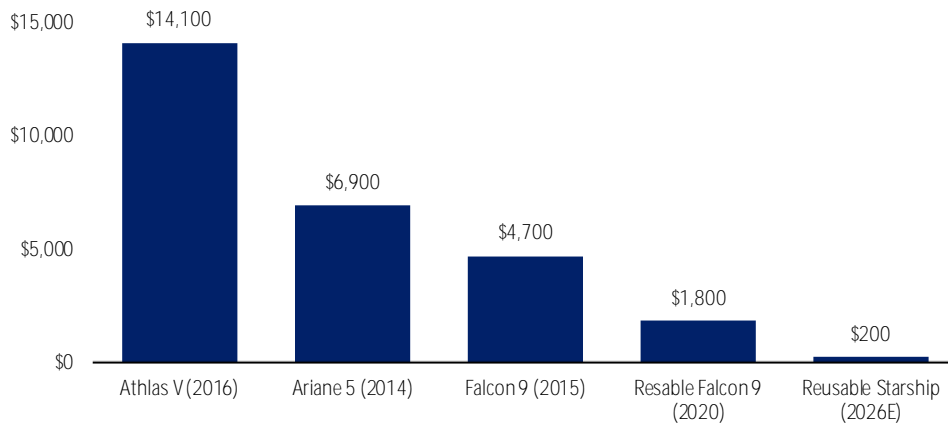
The small size and weight of Miniaturized satellites, especially nanosatellites also brings an economy of scale advantage. Yes we will need more of them, ground equipment is expensive and their coverage is smaller compared to larger satellites, but with economy of scale, lower launch costs and cost reduction, they will gain more economic advantage in coming years. On top of reducing the large economic cost of launching vehicles and the costs associated with construction, their large number often deployed together in constellations allows for better coverage and effective service. Their scalability could make them more useful than having a few larger satellites for certain purposes, especially communication coverage.

Scalability of the technology is key, especially in terms of lowering costs. For example, launch costs have come down dramatically. The Falcon Heavy rocket first launched by SpaceX in 2018 costs \$1,400/kg to launch, while older, non-reusable rockets could cost 70x more (source: FAA). Technology enabling miniaturisation has meant that smaller satellites can be used, saving costs by reducing total weight as well as the benefit of falling cost per kilo. Most use the cube-shaped standard. The small and mostly unified commercial “off the shelf” components and design reduce manufacturing costs dramatically. For example, Planet Labs Dove Sats cost ~\$250k and costs keep coming down.



Exhibit 123: Rocket launch cost – US\$/Kg

Launch costs have plummeted in recent years and by 2026 are expected to decrease to 200/kg, i.e. 98% decline compared to 2016



Source: Visualcapitalist

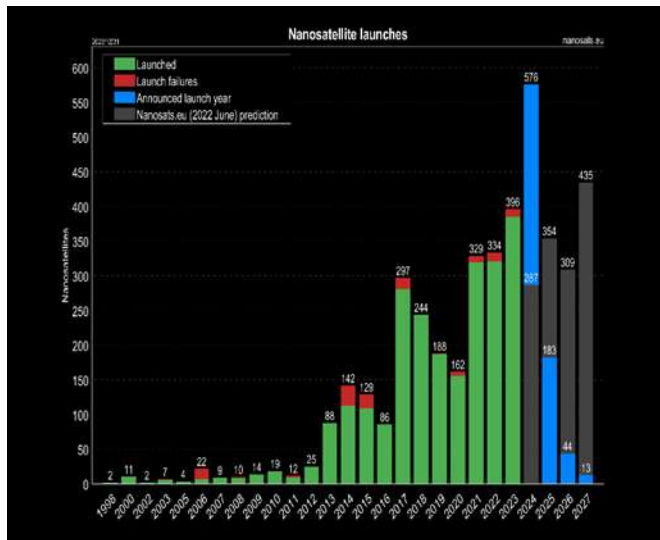
BofA GLOBAL RESEARCH

Reduces time to orbit

Nanosatellites require less time to orbit than traditional, larger satellites. Alén Space, a CubeSat development company, estimates that it can take fewer than eight months to develop nanosatellites and place them in orbit, while larger satellites could take 5 to 15 years (source: Alén Space). Yes, cubsats are mainly used for research by universities, and carry much simpler design and not apples to apples comparison. however this gives a benchmark for the direction of cost reduction path.

Exhibit 124: Nanosatellite launches

More and more satellites being launched are miniaturised, LEO satellites

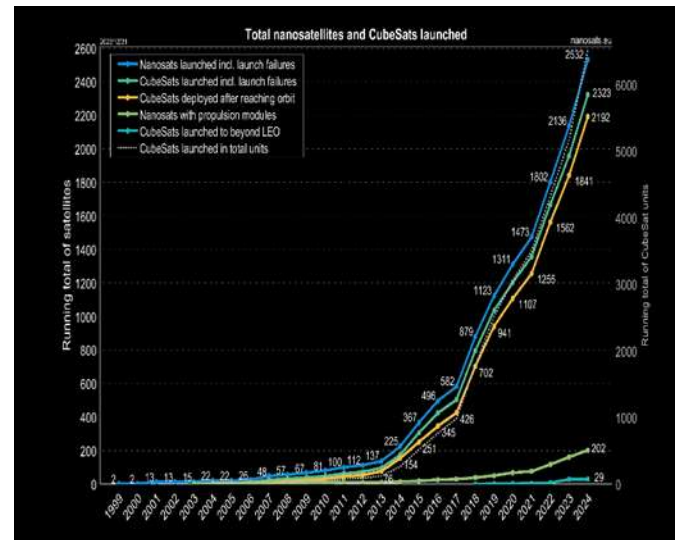


Source: Nanosats.edu, Erik Kulu

BofA GLOBAL RESEARCH

Exhibit 125: Total satellites and cubesats launch

And the number of launches is increasing all the time



Source: Nanosats.edu, Erik Kulu

BofA GLOBAL RESEARCH

Did you know? Different characteristics for different needs. A single satellite in geostationary orbit (GEO) can provide communication coverage of 22k miles but with 700ms of latency, while a Low Earth





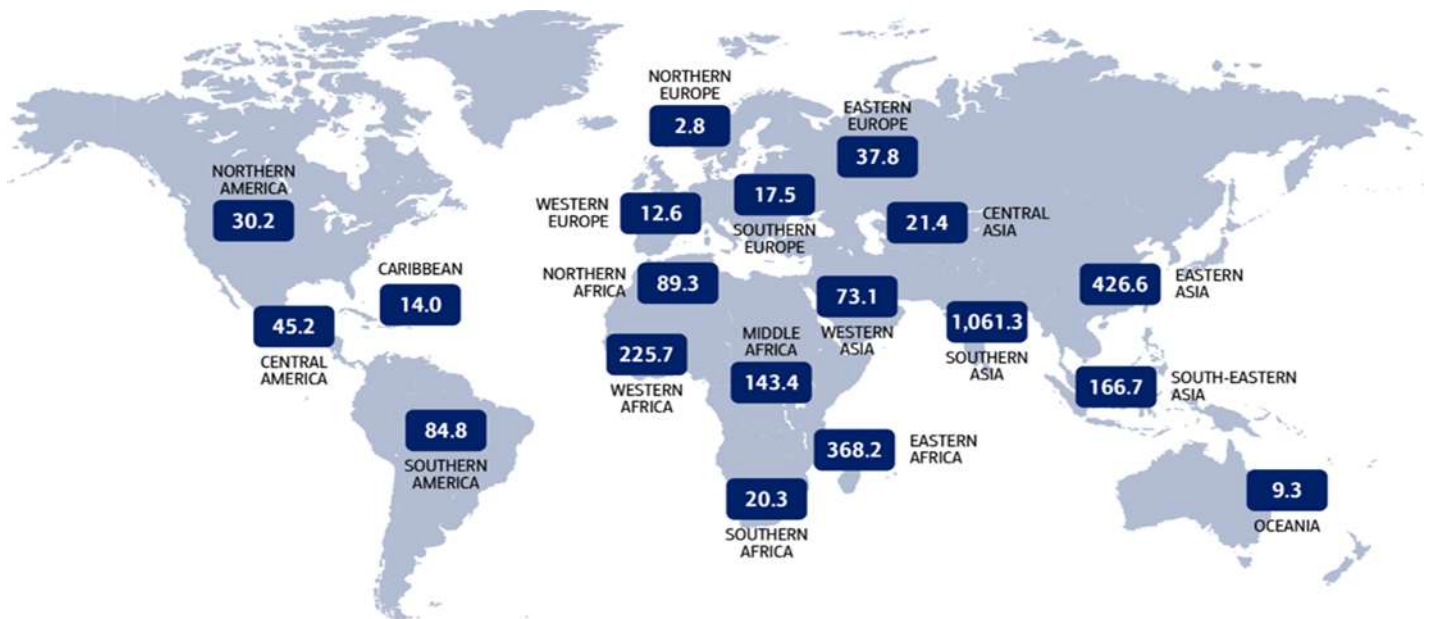
Orbit (LEO) satellite can deliver approx. 84% less coverage but 95% lower latency, at a lower of the cost. However, one of GEO satellites' advantage is that they view space as a contested domain, and can give more precise information over time for specific tasks.

**Why do we need it? To connect the other “third”**

As stated, we think nano- and microsattellites have great potential to meet our communication needs. Their low cost and size, relatively simple design and ability to work in constellations offer affordable vast communication services. Constellations of nanosatellites could provide a broadband network anywhere on the planet. Until recently satellite internet was used only in rural areas to connect those without access to terrestrial internet technologies (e.g. cell towers, cabling). New satellite constellations, such as those from Viasat, OneWeb, Starlink (by SpaceX) and Project Kuiper (by Amazon), could provide competition in areas with lower population density and GDP generation.

Globally we are making huge progress on providing communication services, but as of 2023 still around a third of the planet (2.6bn people, source: ITU) do not have access to any means of communication, whether voice or data. This is up compared to approx. 41% in 2017 or 36% in 2022, but there is still a long way to go. Access to communication is key for economic and social progress.

Exhibit 126: Number of people offline in key countries  
In total, 2.4 billion people in the world are still offline



Source: ITU, GSMA

**Rollout in emerging markets is key**

Broadband communication goes hand in hand with GDP growth and better living conditions. It's not only about providing broadband access – just as important is the ability to collect large sets of data. For example, mapping critical datapoints on agriculture weather and other conditions is key for planning. As stated above, infrastructure spending is a major stumbling block when it comes to telecom infrastructure – covering vast areas with low population density is an economic challenge.



For example, the cost of closing the infrastructure gap in Africa (i.e a telecom network that will supply full broadband coverage, including fiber, backbones, back-holes, towers, switches etc.) is between \$500bn and \$1.5 trillion (source: Manypossibilities). As a result, as of the end of 2022 mobile penetration in sub-Saharan Africa was only 43% and mobile internet penetration was only 25%, well below the world average (source: GSMA).

Satellite communication is key and can bridge the funding gap, as well as the time to market. Most of Africa will be covered by Starlink as of YE2024 compared to only 4 countries in 2022. Yes, there is still the affordability factor, social elements and questions over how many people can actually connect. However, the investment gap could be bridged and technically all of Africa could have coverage in approx. 3 years.

Did you know?  
 In 2023 Starlink covered Ecuador’s Galapagos Islands. In the 3 months from service launch, 100 new businesses were opened (Mo Gawdat)

**Exhibit 127: Starlink’s Africa coverage prior to 2022**  
 Only 3 countries were connected by Starlink in Africa before 2022



Source: UN, Starlink, BofA Global Research

BofA GLOBAL RESEARCH

**Exhibit 128: Starlink’s African coverage 2024 roadmap**  
 Most of Africa will be covered by Starlink’s satellite communication by 2024



Source: UN, Starlink, BofA Global Research

BofA GLOBAL RESEARCH

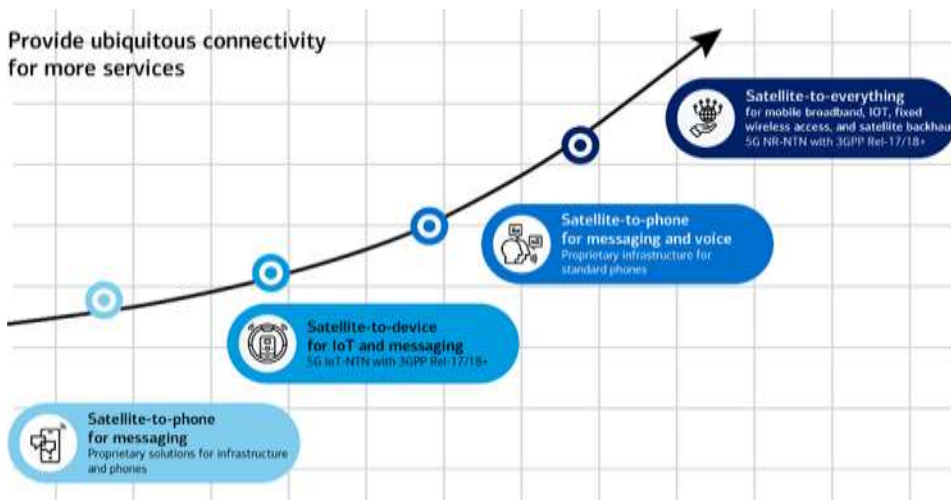


## Further potential for small satellite constellations: IoT, connected cars

While the focus has been on connecting the unconnected, internet constellations have the potential for additional uses longer-term. Given the growing volumes of IoT devices in remote areas without access to terrestrial networks and continually increasing data use, satellite constellations could reduce the burden on existing and future 5G networks. Further, airlines such as American Airlines have introduced satellite internet to provide faster internet speeds on-board.

### Exhibit 129: Satellite communication use cases

With connectivity and integration with other communication technologies, satellite communication can provide a wide range of data-heavy applications



Source: Qualcomm, BofA Global Research

BofA GLOBAL RESEARCH

Did you know? The speed of manufacturing satellites is increasing too: Planet Labs can build 40 of its satellites in a week (source: Planet Labs)

## The race for universal internet is on

These falling costs have enabled the commercialisation of space in contrast to the 21<sup>st</sup> century, when the space race was dominated by a few countries. More than 300 companies active in nanosatellites were over 2013-19, compared with 126 in the 22 years prior. Meanwhile technology and the use of secondary payloads has enabled 6 times the number of objects to be launched into space. The vision to provide connectivity to the entire world through constellations of 1000s of satellites has been voiced by many. Consequently, companies are racing to establish their own systems and acquire spectrum and orbital permissions to gain first-mover advantages. Starlink already has approx. 5,300 satellites in orbit, and accounts for around 64% of all LEO satellites, or 55% of all satellites in orbit.

# Wi-Fi 7

What is it? The next generation of Wi-Fi offering 5x more speed, leapfrog capacity and 75% drop in latency compared to previous Wi-Fi generation.

Did you know? Wi-Fi Alliance predicts that 233 million Wi-Fi 7 devices will enter the market in 2024, growing to 2.1 billion devices by 2028

The race to fight “bandwidth scarcity” will not bypass short-range communication. Wi-Fi-7 is the revolution in that space and will be a key pillar in the next generation of communication.

Developed by the Wi-Fi Alliance and officially launched this year, Wi-Fi 7 (or its professional name: IEEE 802.1) is the next generation of Wi-Fi 6, offering higher capacity, speeds and minimum latency. As we have witnessed in mobile communication technology, applications require ever-more speed and bandwidth, which previous Wi-Fi generations can no longer support, clogging up the network. The need for reliable networks with minimal latency for certain applications was also more challenging for the previous generation. Wi-Fi 7 protocol will provide the solution for applications like AI, the internet of things, the cloud, home entertainment and XR. It will address the challenges of Wi-Fi 6, such as limited spectrum resources, high interference and complex network management. Wi-Fi 7 also supports new features like multi-link operations with up to 320MHz bandwidth, which is double that of Wi-Fi 6.

Exhibit 130: Speed evolution of Wi-Fi (Mbps)  
Wi-Fi 7 can offer speeds of up to almost 47Gbps, 5x more than Wi-Fi 6, and almost 7x more than Wi-Fi 5

	Max speed with 1 spatial stream	jump	Max speed with 2 spatial stream	jump	Max speed with max # spatial stream	jump
Wi-Fi 4	150		300		600	
Wi-Fi 5	867	5.8x	1,730	5.8x	6,920	11.5x
Wi-Fi 6	1,200	1.4x	2,500	1.4x	9,600	1.4x
Wi-Fi 7	2,900	2.4x	5,800	2.3x	46,400	4.8x

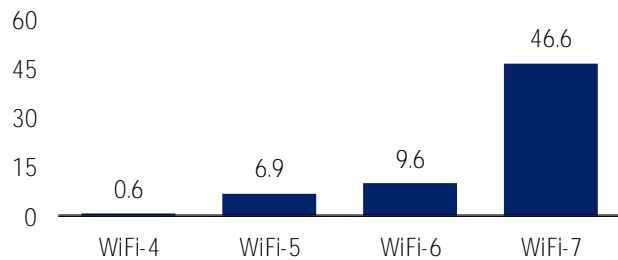
Source: Wi-Fi Association, W-Fi Blog, BofA Global Research

BofA GLOBAL RESEARCH

## What is different this time?

- Speed: Wi-Fi 7 can offer almost 50Gbps speeds on max spatial (more than 2 dimensions) streams, which is 5x the speeds of Wi-Fi 6 (or almost 7x Wi-Fi 5, which is the technology used by most homes today).

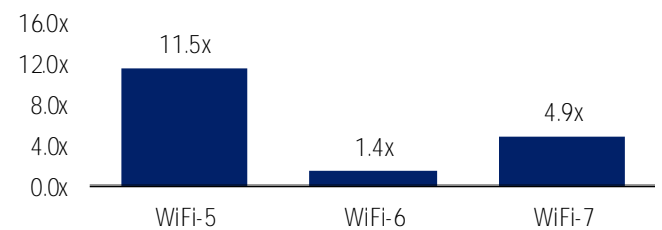
Exhibit 131: Maximum speed with the max number of spatial streams (Gbps)  
Speed evolution of different Wi-Fi technologies (Gbps)



Source: Wi-Fi Association, Qualcomm

BofA GLOBAL RESEARCH

Exhibit 132: Maximum speed leapfrog between different generations of Wi-Fi  
In just a few years Wi-Fi technology has jumped almost 80x in speed. And decreased latency by up to 99%








Source: Wi-Fi Association, Qualcomm

BofA GLOBAL RESEARCH



- Latency: according to the IEEE, Wi-Fi 7 latency could be as low as 1ms (millisecond), which is 75% below Wi-Fi 6 (2-6ms) and 92% below Wi-Fi 5 (up to 10ms, source: IEEE). Like the higher speed, the lower latency could be achieved through a combination of technologies, allowing wider channels and Multi-Link Operations, as described below.
- Wider channel size. Wi-Fi 7 still uses the same three bands as Wi-Fi 6, i.e. 2.4GHz, 5GHz and 6GHz. However, unlike Wi-Fi 6, it can double the channels' bandwidth. Wi-Fi 6 bandwidth can be as wide as 160 MHz. Wi-Fi 7 supports channels that double that size, up to 320 MHz wide. Why is this important? The wider the channel, the more data it can transmit, and the lower the latency. This wide channel is the key reason why data packages can be transmitted vs. in a system with less bandwidth.
- Multi-Link Operation (MLO). Every previous Wi-Fi standard has established a connection between two devices on a single band. MLO can combine several frequencies across bands into a single connection – meaning it supports more connections simultaneously, implying more efficient use of the bandwidth and parallel transmission.
- Inbuilding coverage which is where 5G, 5G Advanced and even 6G fails because of the lower penetration higher frequency spectrum. In that aspect, Wi-Fi 7 technology could be the ultimate complimentary tech to be used alongside 5/6G iterations.

Exhibit 133: Wi-Fi 7 key tech features.  
integration of technologies that deliver higher speeds, lower latency and reliability

Wi-Fi 7 key characteristics	
Features	Benefits
 320 MHz channels	2X higher throughput
 Multi-link Operation (MLO)	Determine latency, increased efficiency, greater reliability
 4K QAM	20% higher transmission rates
 512 Compressed Block Ack	Reduced transmission overhead
 Multiple RUs to a single STA	Enhanced spectral efficiency

Source: the Wi-Fi Alliance

BofA GLOBAL RESEARCH



## 5G Advanced

What is it? The evolution of 5G mobile technology with embedded AI capabilities, offering 10x the mobile speed, a third of the latency and lower costs compared to current 5G.

Did you know? 5G Advanced cost per bit could go down by 90% compared to LTE networks and 30-50% compared to the first generation of 5G.

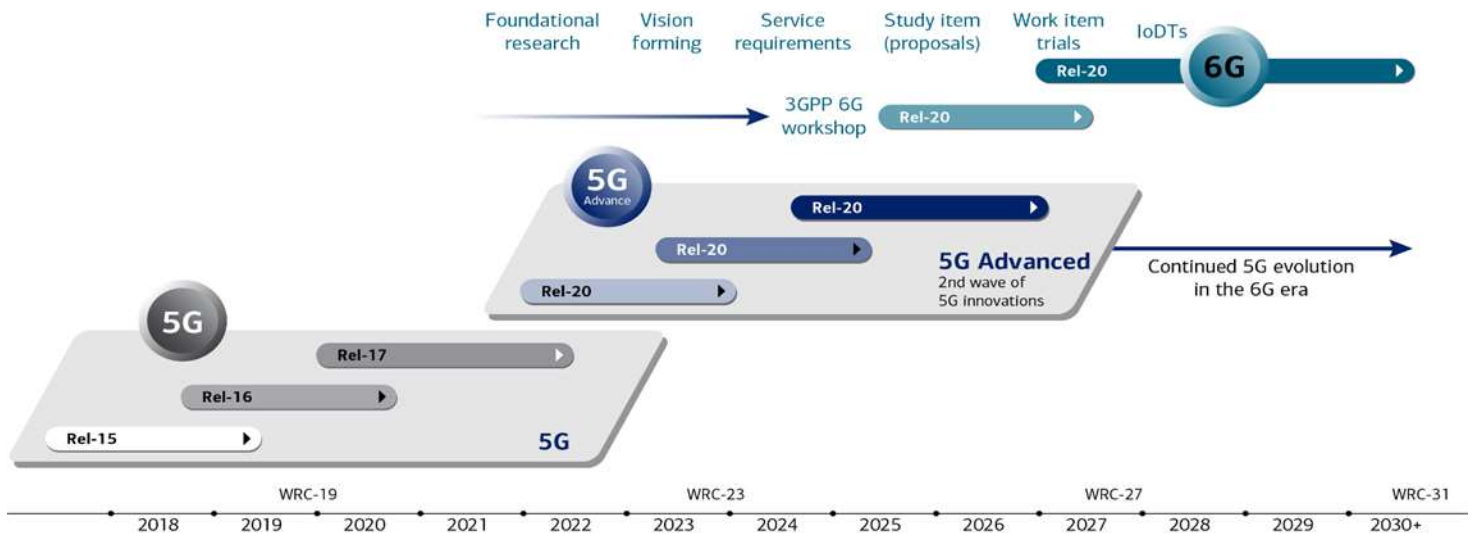
As we stated in our [AI/ChatGPT Primer](#), every industry is being transformed by AI, and the telecom industry is no different. 3 aspects for telcos – the sheer amount of data, the complexity of the data, and the need for “data-hungry” applications putting pressure on current mobile networks. But at the same time, 5G Advanced is the first mobile network with embedded AI, dramatically increasing capabilities and decreasing costs.

5G Advanced refers to the next phase of development and deployment of 5G wireless technology and networks. Usually it refers to the 18<sup>th</sup> and 19<sup>th</sup> releases of the 5G architecture protocol, as defined by the 3GPP (3<sup>rd</sup> Generation Partnership Project organization). It builds on early 5G networks and aims to deliver faster speeds, lower latency, increased capacity, and more uniform coverage. Target performance goals include speeds of up to 10 Gbps and latency under 1 millisecond.

5G Advanced is not a revolution but an evolution of the 5G technology and broadly based on common architecture. However, the improvement in the architecture and technology not only drives continued system enhancements, but could also be the foundation technology to 1) increase the number of use-cases, 2) support high-data, low-latency demanding applications like holograms, XR and autonomous vehicles and 3) even lay the technical foundations for 6G.

Exhibit 134: The 5G Advanced evolution timeline

5G was introduced in 2018, deployed in scale from 2020, but by 2023 we had already started migrating to the next generation



Source: 3GPP, Qualcomm, GSMA, BofA Global Research

BofA GLOBAL RESEARCH

Coming as soon as 2024: Huawei as well as EE expect 2024 to be the year that commercial deployments of 5G Advanced officially begin. For Huawei, 5G Advanced is a network that will be capable of 10 Gbps downlink speeds.





## What is different?

- The first native AI support communication technology. 5G Advanced will integrate AI and machine-learning capabilities into the network. This will allow optimization of process and automation, including in areas like automated traffic routing, automated infrastructure management and maintenance and personalized user experience. With AI-defined management, machine-learning algorithms could help to fully automate the network, offer constant analyzed feedback in real time, and set resource allocation or channel data traffic based on optimization analysis. The network could even “design” its own architecture. AI can also assist in communications, such as beamforming (a signal processing technique), predictive resource allocation between network slices, automated cyberthreat response, intelligent traffic routing, and channel status predictions. Cyber is the #1 risk for mobile communication and AI-embedded capabilities can monitor data traffic and anomalies in real time. On-device and edge ML algorithms can also tailor 5G experiences to usage patterns, locations and applications, creating personalized experiences, for example tweaking VR image rendering based on head movements or refining content recommendations.
- 10x the speed, one third of the latency. 5G Advanced is designed to provide 10x higher data speeds compared to 5G, up to 10Gbps compared to the 1-2Gbps traditional network speeds. But it is not just the speed. 5G Advanced also aims to have lower latency. The current 5G offers an average latency of 15-30ms (millisecond) in ideal conditions, but 5G Advanced aims to reduce this to as low as 10ms. In future the goal could be as little as 5ms, and some researchers are even calling for 1ms theoretical latency, to be used in mission-critical applications (Qualcomm). This target will probably be achieved with 6G architectures.
- Higher bandwidth: 5G Advanced aims to utilize a larger bandwidth for more capacity or speed. It is designed to use high radio frequencies at the 100GHz range, compared to traditional 5G, which uses sub-6GHz frequencies. Higher-frequency bands offer much higher speeds, but with a shorter range and distance – hence the network should be more dense.
- New architecture to maximize capabilities. As stated above, higher bandwidth range usually carries more data and speed, but for shorter distances, and also has the problem of physical transmission barriers (walls, trees etc.). 5G Advanced architecture should work much better than other technologies with high frequencies, with a strong signal indoors and greater reliability. How? The network architecture is based on technologies like precise beamforming, advance antenna solutions, and AI-based network redundancy analytics. These should help 5G Advanced to target 99.999% reliability (5 minutes or below of network downtime).
- Precise positioning capabilities are a case in point: 5G Advanced targets up to 10 cm accuracy compared to the room accuracy of traditional 5G. This advance feature can offer better usage of resources and lower costs.
- The result of all this – cheaper? Optimization of process, more efficient use of infrastructure, higher spectrum efficiency and constant seamless connectivity could lower network costs dramatically. This efficiency and better usage of resources could reduce network energy consumption. Maintenance should be lower as well because of better usage of the network capabilities and real-time feedback. Usage of advance antennas like MIMO (Multiple Input, Multiple Output) could simplify the network itself and use less equipment. How much cheaper? Costs per bit could go down by a factor of 10 (source: MGMN)





compared to LTE networks and 30-50% compared to the first generation of 5G. IEEE estimates 50% potential improvement in radio access network costs.

- Plus increased revenue streams: more data-heavy applications that demand higher bandwidth and lower latency like XR, holograms or autonomous vehicles could offer providers a new revenue opportunity.

#### Key technology used

- Millimeter waves (mmWave) – much higher-frequency radio waves in the range of 30GHz – 300GHz that allow transmission of more data but have a shorter range, requiring dense base station infrastructure. mmWaves have a shorter wavelength and can carry more data, but for shorter ranges, and have difficulty transmitting through obstacles like walls.
- Small cell networks – low-powered miniature base stations to complement high-frequency mmWave signals that provide pockets of coverage. As mmWaves can only carry data for short distances with minimum strength, a denser cell network is required to overcome this.
- Massive MIMO – base station antennas with hundreds or thousands of steerable transmitters and receivers to allow more users and devices to connect simultaneously.
- Beamforming – advanced technique to focus wireless signals directly on user devices to save power and reduce interference.
- Full duplex communication – routers can transmit and receive data at the same time using two separate transmitters and receivers, doubling capacity.
- Network slicing – creating multiple virtual networks on top of a common physical infrastructure to meet the requirements of different applications and services.

#### The foundation of 6G?

As we discuss in our next chapter, we are big believers 6G will be needed sooner rather than later, as early as 2028. 5G Advanced is not just a natural progression of 5G technology, but a necessary evolution, which will lead us to the 6G revolution. As stated above, 5G Advanced refers to release 18 and now 19 of 5G. With the release of 20, according to Qualcomm, the 6G foundations will be laid, as it will enable new technologies like Duplexing Evolution, the use of new spectrum bands, integrated sensing and communication, among others. All these technologies will be maximized to build 6G networks.



Exhibit 135: from 5G to 6G

5G Advanced will be only an intermediate technology before we move all the way to 6G



Source: Ericsson, GSMA, BofA Global Research

BofA GLOBAL RESEARCH

6G: sooner than you think.

What is it? The revolution of mobile phone technology. The new generation which can offer 1Tbps speeds, close to zero latency and AI embedded capabilities.

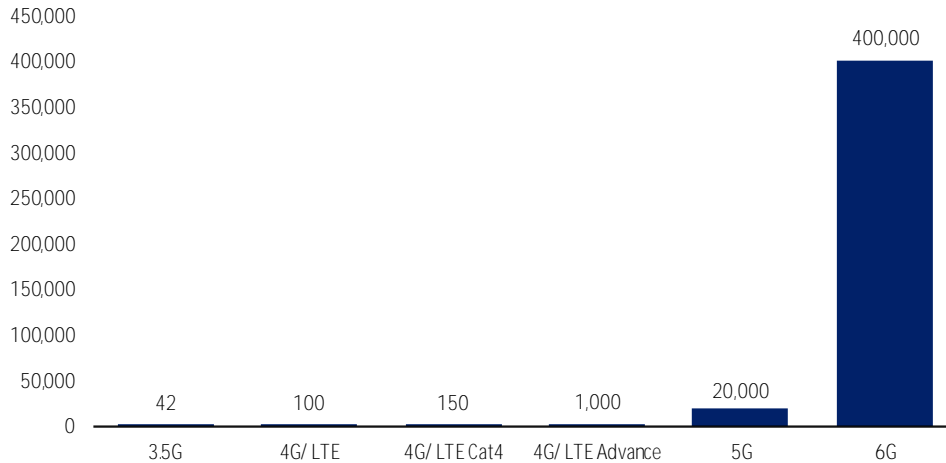
Did you know? At current average global mobile speeds, it will take someone c.181 million years to download all the data from the internet ...but it will take 838,000 years on a 5G network and “only” 17,000 years on 6G

6G will be the successor to 5G mobile technology in the transmission of mobile data. It is estimated to be commercially available before the end of this decade. 6G will offer 10-50x higher speeds and bandwidth compared with 5G, with a much better latency ratio (up to 1/1000). According to different estimates, 6G download speeds could reach 400-500Gbps (Samsung, Ericsson, Cisco) or even up to 1Tbps (Bremen University), with close to zero latency, and still be able to handle 10x more data capacity than 5G. While the technology and protocols remain in development, 6G will use the higher bands of the spectrum like Terahertz radio waves (Terahertz = 1,000 Gigahertz). These have greater capacity and are more stable, but will cover shorter distances.



## Exhibit 136: Theoretical speeds (Mbps)

While 6G can offer up to 20x faster speeds than 5G. Some estimates project 1Tbps speeds (1,000,000Mbps)



Source: GSMA Intelligence, Ericsson, BofA Global Research

BofA GLOBAL RESEARCH

### The speed is just a bonus

As stated above, 6G can offer 10-50x faster speeds than 5G. However, looking at speed alone is misleading. Most of us use regular applications that do not require a 1Tbps download speed. We need to look at 6G instead as the technology that could have the capacity for all the data we are creating, could connect the digital world to the physical one, and could handle complex data. To do so, 6G will be the first mobile technology to have AI capabilities embedded in it. What other technology could support the connectivity of all autonomous driving, Industry 4.0, one trillion connectable devices, smart city holograms, and the world of gaming, etc? More importantly, what other technology will allow all applications to “talk” to each other, so creating even more data than the sum of each application as a standalone (complex data)?

We are living in a world that 1) is creating more and more data and 2) needs a way for all applications to connect and “talk” to each other. Whereas with 5G speeds we can download Netflix movies in seconds, with 6G all the applications could connect and synchronize seamlessly. Furthermore, with so many devices online that need to communicate with each other, 6G technology could be the platform to allow AI-based smart connectivity.

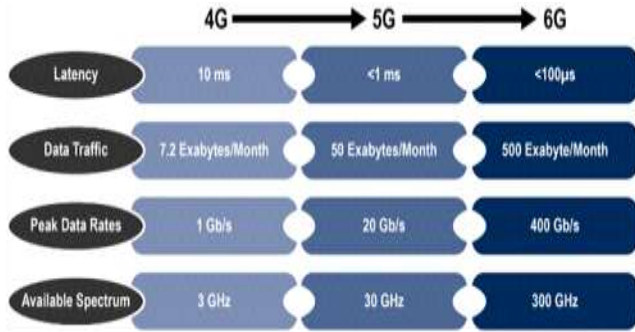
### Why do we need 6G? Because nothing is linear when it comes to technology...

The idea behind 6G is to create a technology that will connect everything. This was also the bedrock of 5G, but new applications and the speed of data creation placed a question mark over its ability to support the “connected world” over time:

1. The data we are creating is increasing exponentially and will clog up 5G network capacity sooner than expected.
2. New applications are being developed all the time that require more data. In this report, holograms, metaverse, brain computer interfaces, EVTOL etc are just some of the tech we explore that will be very data-heavy and have yet to be launched. Not to mention quantum computing, which could immediately leapfrog the total data creation once commercially available. In short, data is going to grow far faster than initially expected in coming years, well beyond the exponential growth we are currently seeing on mobile networks.
3. All these and other technologies will have to “talk” to each other and connect seamlessly with zero latency. This will place even more pressure on the network.



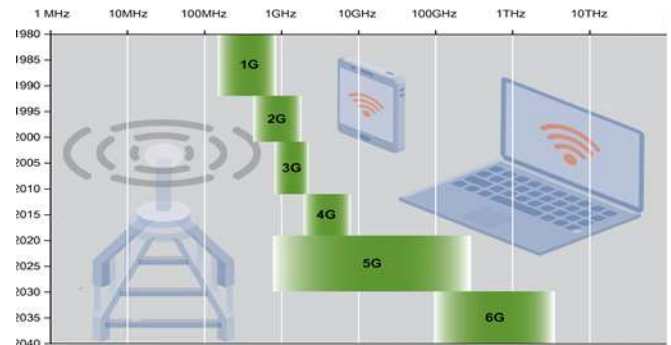
Exhibit 137: 6G capabilities vs. 5G  
Higher speeds, bandwidth, capacity and close to zero latency



Source: BofA Global Research, Ericsson

BofA GLOBAL RESEARCH

Exhibit 138: Spectrum allocation  
6G will be the first mobile technology to use Terahertz spectrum to transmit more data

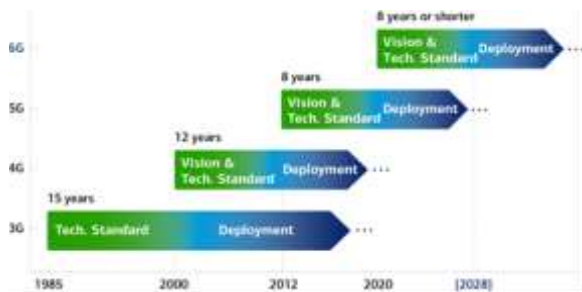


Source: BofA Global Research, futuretimeline.net

BofA GLOBAL RESEARCH

Some might think it's too early to talk about 6G, when 5G is still being deployed and should become fully operational in most countries. 5G was regarded as the ultimate solution to our current world needs, but we think they will change significantly before the end of the decade.

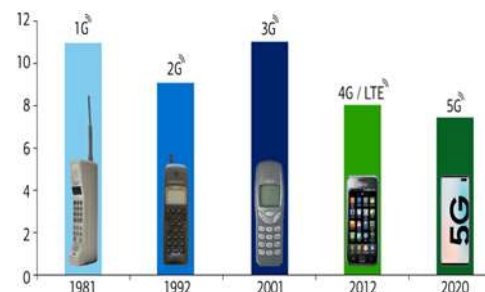
Exhibit 139: 6G timeline  
While most researchers believe that 6G won't be deployed this decade, Samsung believes it could happen by 2028. Perhaps even sooner?



Source: BofA Global Research, Samsung

BofA GLOBAL RESEARCH

Exhibit 140: Could 5G be the shortest-lived mobile technology?  
The timeline for a mobile technology is usually c.10 years, but the explosion of data creation and complexity might push out 5G before that



Source: BofA Global Research

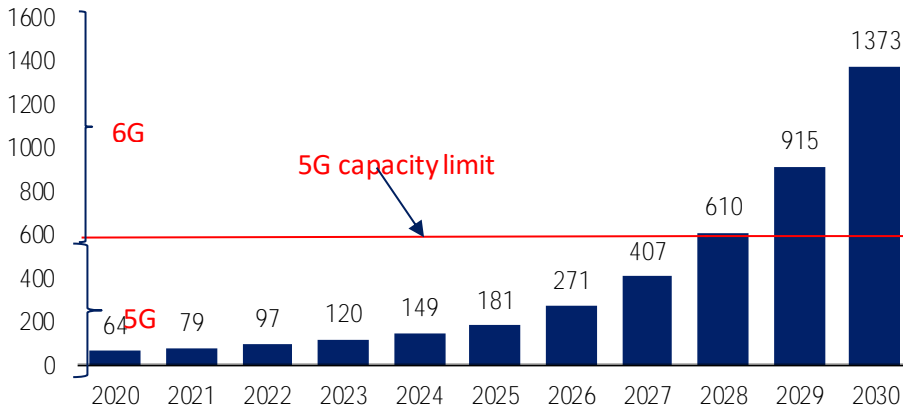
BofA GLOBAL RESEARCH

5G networks can handle total data traffic of 50Zb/month (Ericsson) or 600Zb/year. Global mobile data generation reached 120ZB in 2023 (Statista). Assuming it doubles every 2-3 years, by around 2028 global data capacity will be beyond current 5G capabilities (see below). If we further assume 5G will be enhanced like generations before it, the inflection point will be around 2028E, when we enter the Yottabyte era.



**Exhibit 141: Global data generation (Zettabytes/year)**

At the current rate of data creation, 5G will hit full capacity as early as 2028. However, with penetration of new applications like XR, Holograms, Autonomous vehicles etc – could it happen sooner?



Source: BofA Global Research, Statista, Ericsson

BofA GLOBAL RESEARCH

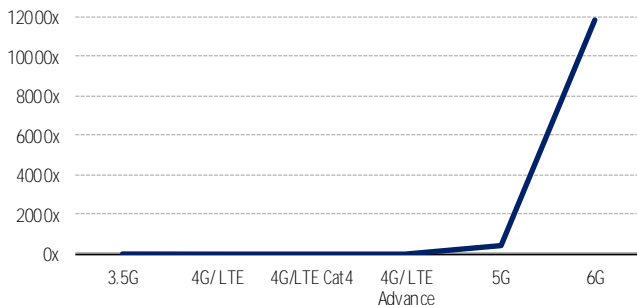
Did you know? At current average global mobile speeds, it will take someone c.181 million years to download all the data from the internet ...but it will take 838,000 years on a 5G network and “only” 17,000 years on 6G. Yet....data is doubling every 2-3 years....

Source: Statista, Unicorn Insight, IDC, BofA Global Research

**Challenges: complexity, prices, Moore’s Law, emissions...**

- So will Moore’s Law catch up with 6G and data? The total data we are creating is doubling every 2-3 years, or at a c50-60% CAGR, which is greater than the increases in CPU processing power. Now that we are entering the era of complex data, this problem could intensify. The 6G data transmission rate is 12,000x faster than 3.5G and 50-100x faster than 5G. This means that in a few years’ time CPU processing power will have to match the growth in the transmission of data. According to current projections, we are not there yet. We might actually reach the limit of Moore’s law this decade.

**Exhibit 142: Wireless generation download speeds compared to 3.5G**  
6G download speed will be c.12,000x higher, while global data growth continues to jump..

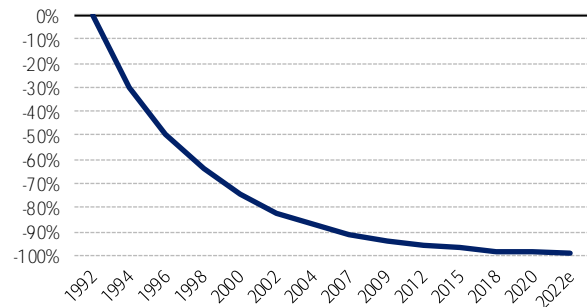


Source: BofA Global Research

BofA GLOBAL RESEARCH

**Exhibit 143: Change transistors gate nanometer length (1992 base year)**

...but at the same time we are reaching the limit of Moore’s law. And processing speed growth is moderating, reaching its limit in 2020s



Source: BofA Global Research

BofA GLOBAL RESEARCH

- More sites mean higher radiation, due to denser networks creating potential health risks. This issue has not been fully explored even in relation to 5G and will likely be in the spotlight the closer we get to 6G. 5G case studies will be crucial for the implementation of 6G.



- 6G will operate on a higher spectrum than 5G. The higher we go on the spectrum, the more data that can be carried. The downside is that the distance we can carry the data is shorter. Currently, in lab conditions, the Tera waves range is only about 10 metres. As with the deployment of 5G, the density of the network will continue to increase the higher we go up the generation. Bottom-line, we will need more sites. How many? It's not yet clear. Some estimates suggest 10x more sites, towers or micro-sites. Samsung is exploring the idea that personal mobile phones will work as transmitters that will amplify the signal. A Samsung whitepaper also discussed the idea of using satellite signals to cover more areas, splitting spectrum etc

Exhibit 144: Suggested solutions for 6G network architecture

Sites, microsites, towers, satellites are some of the solutions needed to increase network signal



Source: BofA Global Research, Samsung, Ericsson

BofA GLOBAL RESEARCH

## Healthcare: from GPT to GLP-1 transformation is here

Transforming World: 6-7% of global GDP (\$8.5-\$9 trillion) is spent on healthcare annually, but \$1tn is wasted in the US alone. Success in drug development is very low across all therapeutic categories worldwide. For example, 97% of cancer drugs fail during clinical trials. Increasing life expectancy, population growth, antibiotics resistance and the “explosion” of healthcare data all increase demand for new healthcare technologies, practices and approaches.

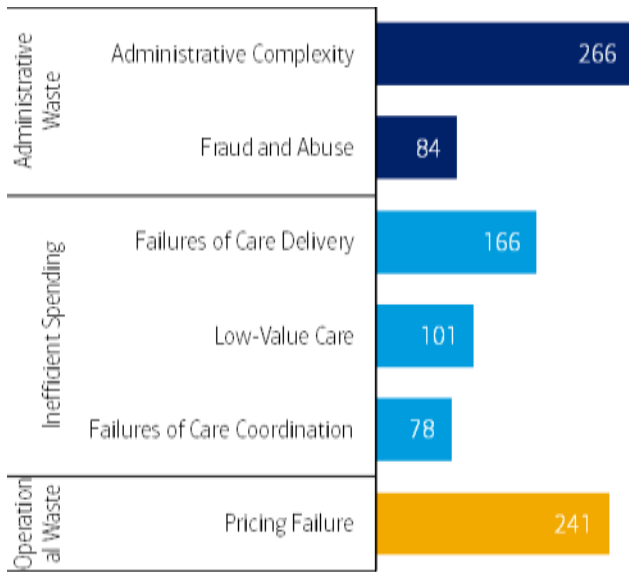
Did you know? Each person will generate enough health data in their lifetime to fill 300 million books, or the NY public library 6x...(source: KPCB)

Future of healthcare: AI healthcare, drug discovery, diagnostics, genomics, CRISPR, GLP-1



Nearly \$1 trillion of US healthcare spend is wasted each year. The US has one of the most expensive health systems in the world with spending totalling \$4.3 trillion in 2021 alone. One reason for the high cost of the US health system is waste. Approximately 25% of healthcare spending in the country is considered wasteful, and about 1/4 of that could be recovered through interventions. According to the Journal of the American Medical Association, the annual cost of wasteful spending in healthcare has ranged from \$760 billion to \$935 billion in recent years. The largest source of healthcare system waste, roughly \$266 billion, is administrative costs.

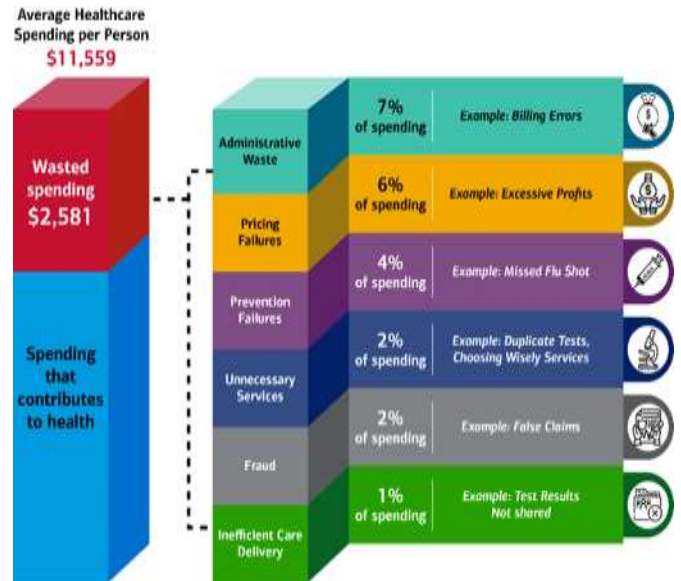
Exhibit 145: Types of Wasteful healthcare spending (billions of dollars)  
Wasteful healthcare spending can reach upto \$935 billion a year



Source: Peter G. Peterson Foundation (2023). Journal of the American Medical Association, Waste in the US Healthcare system: estimated costs and potential for savings (Oct'19)  
Notes: Data represent the upper threshold of estimates by Shrank and colleagues. Total sum may be different due to rounding.

BofA GLOBAL RESEARCH

Exhibit 146: 23% of healthcare spending is wasted. Administrative waste is the largest category and is accumulated through excessive administration costs and profits, including claims processing, ineffective use of IT, staffing turnover and paper prescriptions, costing approximately \$130 billion per year.



Source: Healthcare Value Hub

BofA GLOBAL RESEARCH



## AI revolutionizing Healthcare

What is it? Using AI to improve drug discovery, diagnose disease faster and reduce the costs of healthcare to improve patient outcome

Did you know? Generative AI could create \$60-110bn of economic value in healthcare (source: McKinsey)

AI could save up to \$360bn in US healthcare spend each year alone

The National Bureau of Economic Research (NBER) estimates that wider adoption of AI could lead to savings of 5-10% in US healthcare spending, roughly \$200-\$360 billion annually in 2019 dollars. These estimates are based on specific AI-enabled use cases that employ today's technologies, are attainable within the next five years, and would not sacrifice quality or access. These opportunities could also lead to nonfinancial benefits such as improved healthcare quality, increased access, better patient experience, and greater clinician satisfaction.

Exhibit 147: The AI healthcare trend in recent years

AI on the rise since 2020, expected to become even more relevant in 2026+



Source: BofA Global Research, CodeGlo Journal

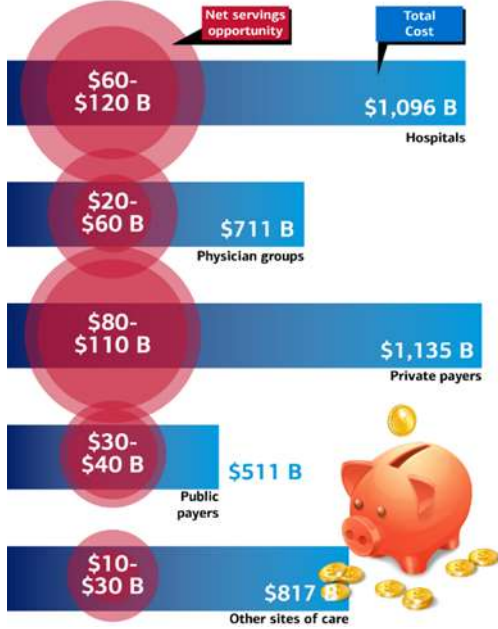
BofA GLOBAL RESEARCH

Based on the AI-driven use cases, private payers could save roughly 7% to 9% of their total costs, amounting to \$80 billion to \$110 billion in annual savings within the next five years. Physician groups could save 3% to 8% of costs, amounting to between \$20 billion and \$60 billion in savings. Meanwhile, hospitals could see savings of between 4% and 11%, or \$60 billion to \$120 billion each year, the report estimates.



Exhibit 148: AI could save the healthcare industry up to \$360 billion a year  
 Hospitals could see savings of between 4% and 11% each year while private payers could save roughly 7% to 9% of their total costs

Savings opportunity per stakeholder group, 2019, in USD



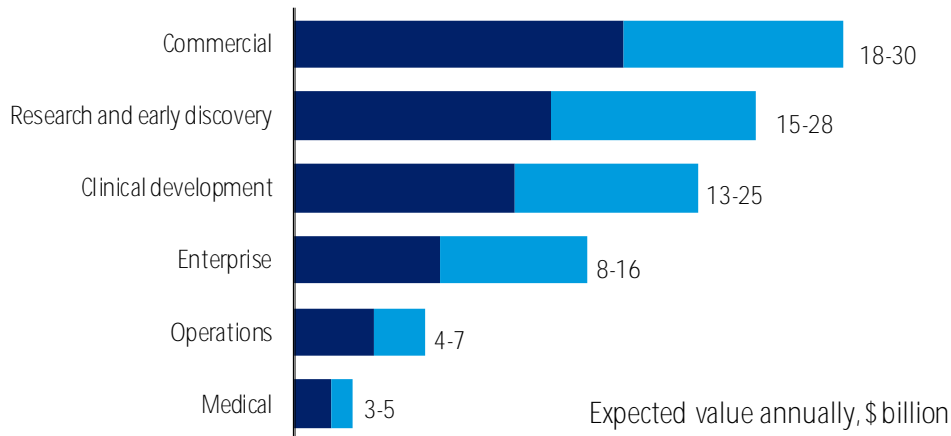
Source: National Health Expenditures, MedTech Pulse

BofA GLOBAL RESEARCH

Generative AI could create \$60-\$110bn of economic value in healthcare  
 At the same time, McKinsey estimates that Generative AI could create \$60 billion to \$110 billion a year in economic value for the pharma and medical product industries going forward. This is largely because it can boost productivity by accelerating the process of identifying compounds for possible new drugs, speeding their development and approval, and improving the way they are marketed. Pharmaceutical companies have long been in the vanguard of artificial intelligence. Even before last year's explosion of ChatGPT interest, researchers were applying complex AI models to unlock the mechanisms of disease. AlphaFold2, ESMFold, and MoLeR, for example, all use deep learning to help predict the structures of nearly all known proteins, transforming our understanding of their underlying diseases.



Exhibit 149: Generative AI is expected to produce \$60 billion to \$110 billion in annual value across the pharmaceutical industry value chain  
 Largely because it can boost productivity by accelerating the process of identifying compounds for possible new drugs, speeding their development and approval, and improving the way they are marketed



Source: McKinsey analysis

BofA GLOBAL RESEARCH

Generative AI and LLMs to transform healthcare

Exhibit 150: Health-LLM, a groundbreaking framework for evaluating LLM performance on a diverse set of health prediction tasks, training and prompting the models with multi-modal health data  
 The researchers meticulously selected thirteen health prediction tasks across five domains: mental health, activity tracking, metabolism, sleep, and cardiology, to cover a broad spectrum of health-related challenges and to test the models' capabilities in diverse scenarios

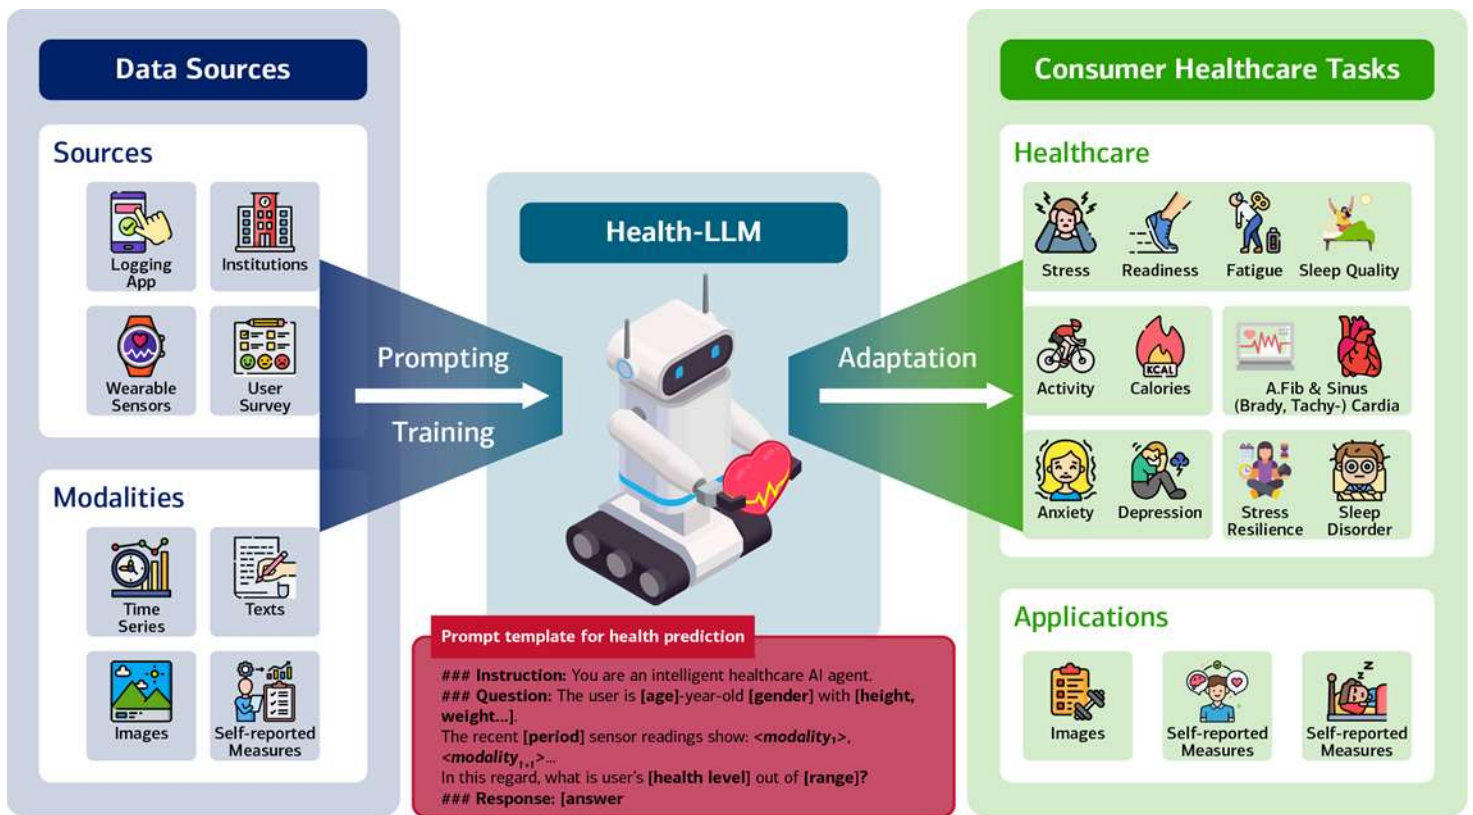
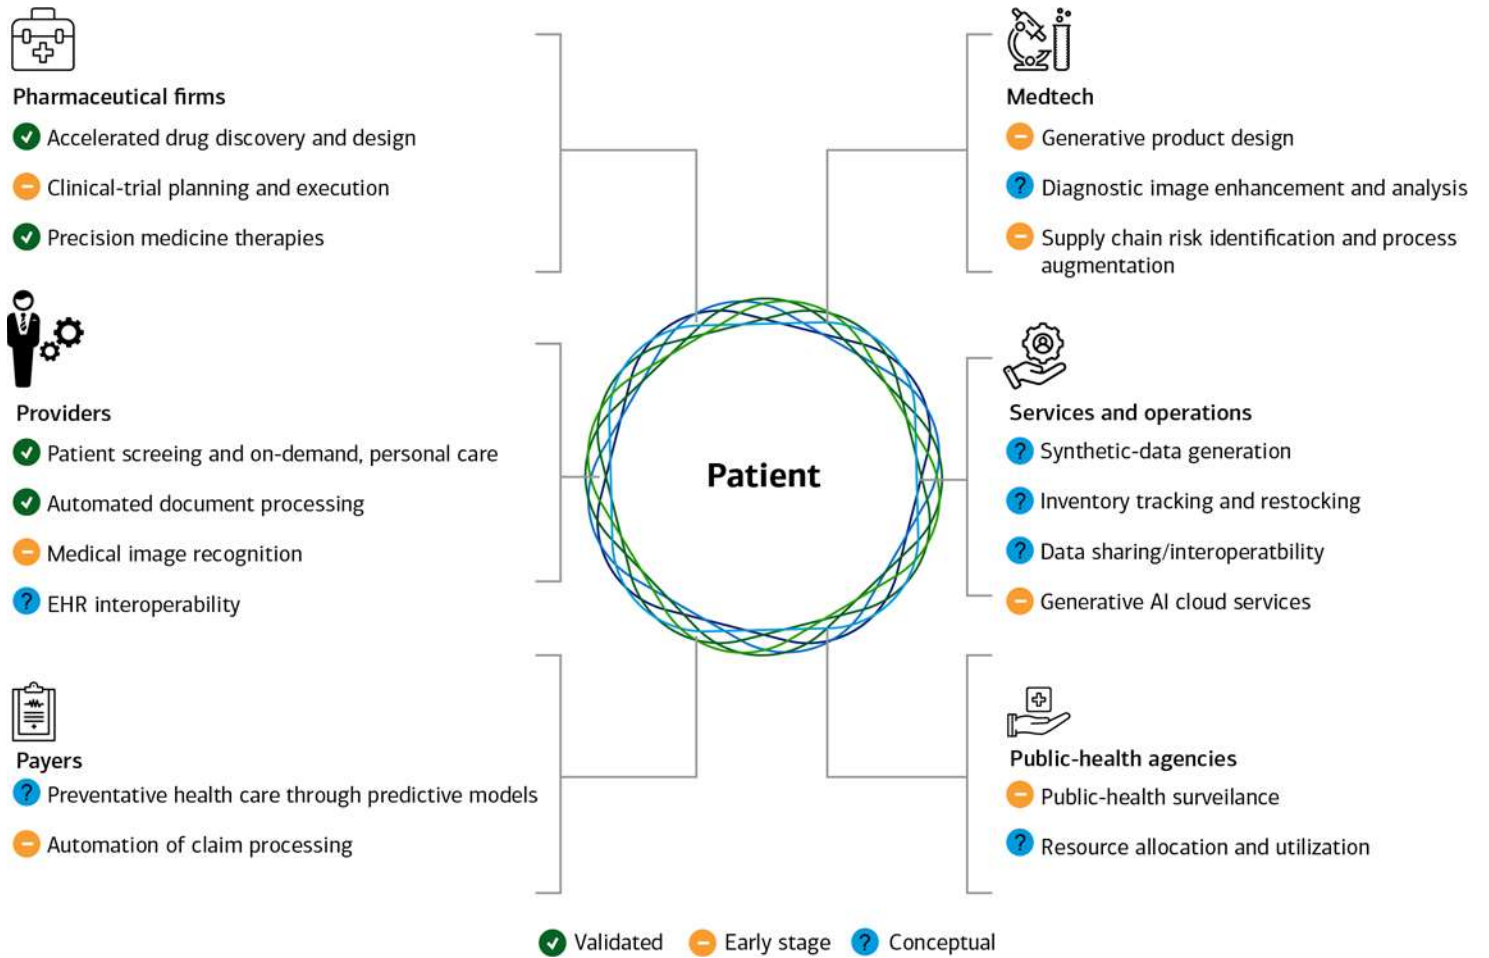


Exhibit 151: Generative AI has potential use cases across all health-care segments

Emerging generative AI use cases now exist in all health care segments, from providers and pharmaceutical firms to payers, medtech, service providers, and public health agencies



Source: BCG analysis

## AI Drug Discovery – on the cusp of disruption

What is it? Using AI to escape Eroom's Law (Moore's law in reverse) because Big Pharma invests \$50 billion annually in research for new medicines but the number of new drugs approved per billion US dollars spent has halved roughly every 9 years since 1950, falling around 80-fold in inflation-adjusted terms

Did you know? Scientists/humans can only hold 6-7 data points in their mind at one time whilst drug discovery's iterative process can be done millions of times within minutes by AI in the lab

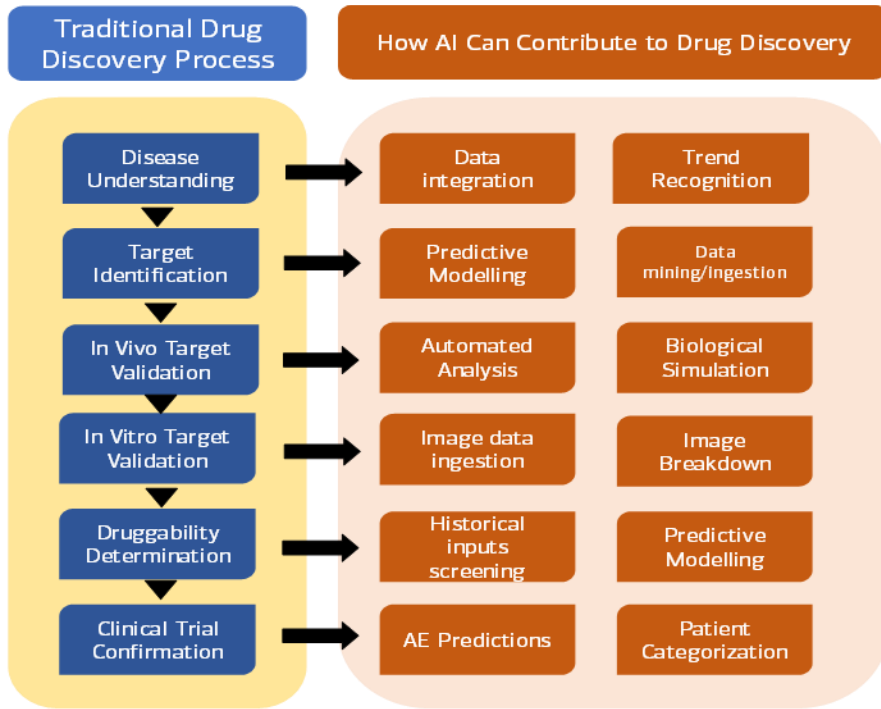
The drug development process has well-known inefficiencies that AI could disrupt. Interesting applications include: (1) AI-assisted therapeutic target selection; (2) personalized genomic screening for diagnosis and personalized therapy; (3) clinical trial optimization for streamlining enrollment, fragmenting patient populations, and precision targeting based on mutations to increase chance of response; and (4) optimizing healthcare systems resulting in efficiencies in time and cost. However given the complexity of biological systems, the data used to train AI models in healthcare will likely be much more densely packaged and more difficult to ingest compared to more generalist use cases of OpenAI's ChatGPT.

One of the first areas AI is likely to replace is the labor-intensive data mining work of drug discovery. Currently, the drug discovery process consists of several steps including: (1) data mining: a process in which the collection and compound profiling of data from various sources such as gene expression data, proteomics data, transgenic phenotyping can yield decisions for target identification/prioritization; and (2) genetic associations: a process that looks for connections between phenotypic evidence and genotypic markers (for example cross-linking the presence of certain mutative genes to the onset of certain cancers). That said, we think the major technological hurdle lies in the homogenization of multimodal data. Unlike LLM, which takes simple inputs (i.e., texts, images), AI models in healthcare deal with complex biological systems and the training difficulty is magnitudes higher than a LLM according to our KOL checks. We see major technical difficulty with homogenizing multi-modal data in a way suitable for ingestion (packaging genomics, clinical data, is dimensionally more difficult than packaging text or images). While LLM has proven to reach intelligence inflection points once the database reaches a certain scale, we see the same philosophy as not necessarily applicable for healthcare AI models with other moving parts also mattering.





Exhibit 152: Comparison of traditional drug discovery process and AI process  
 We expect future of drug discovery to leverage the use of AI in many areas

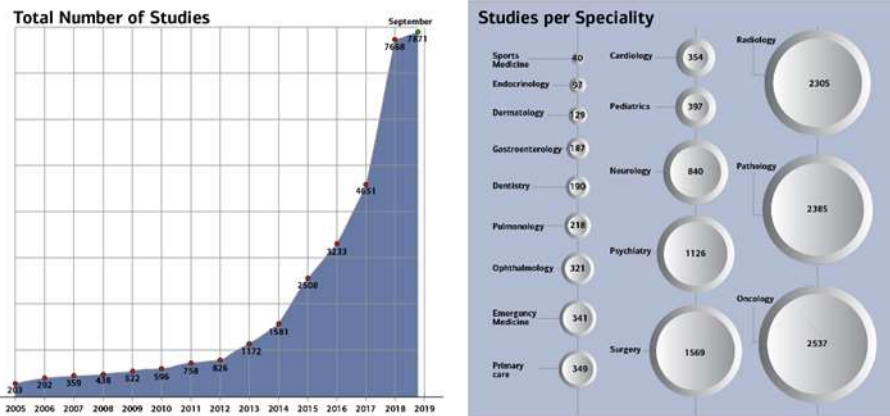


Source: BofA Global Research

BofA GLOBAL RESEARCH

Research into AI drug discovery is accelerating. In 2005 there were only 203 research papers on Pubmed.com in this medical field. By 2018 this number had risen to 7,668 (Source: MedicalFuturist)

Exhibit 153: Machine and Deep learning studies on Pubmed.com  
 Number of research papers on Pubmed.com rose from 203 in 2005 to 7,665 in 2018



Source: Medical Futurist

BofA GLOBAL RESEARCH

Simulation can already be used in pharma

Using life sciences as an example, AI and simulation technologies can take a molecular structure and simulate it billions of times, making small changes each time to see which structure is optimal. According to Jack Hidary, CEO of SandboxAQ, we can now do in a matter of weeks and months tasks that would take 10 years in the physical world, particularly in pharma. Drug discovery involves high costs and has a high failure rate (90%), which makes the business model unpredictable. The average investment is US\$1 - 4bn and it takes 10-15 years to develop the drugs. And revenues from the 10% of drugs that are successful have to pay for the 90% that failed.



**Exhibit 154: Challenges in the drug discovery market**

There is a high cost and failure rate in the drug discovery market e.g. US\$1-4bn of average investment, 10-15 year development timeline, 90% clinical trial failure rate



Source: SandboxAO

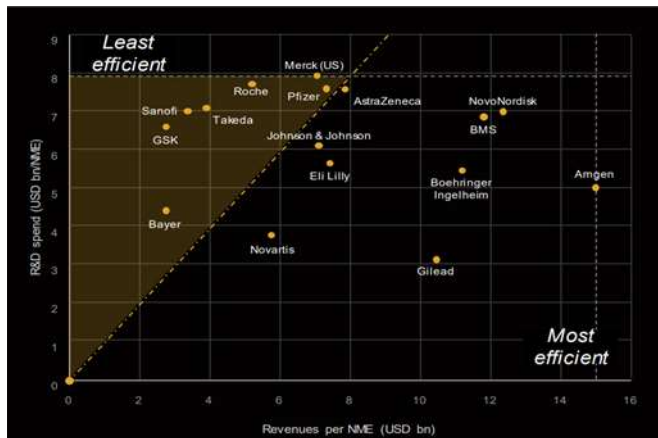
BofA GLOBAL RESEARCH

**Eroom's law to be reversed**

This helps explain Eroom's law (the observation that drug discovery becomes slower and more expensive over time), which entails a decrease in R&D efficiency. This is where companies spend more per new drug on R&D than they make in revenue. Per Hidary, new AI simulation work can change sectors such as life sciences from a lumpy business of mostly failure to one with predictable revenues. It allows companies to take all the data and molecular information and run billions of simulations de-risking the molecules, cutting down drug development time dramatically.

**Exhibit 155: R&D spend (US\$ bn/NME vs US\$ revenues per NME)**

Large pharma companies often spend more per new drug on R&D than they make in revenue

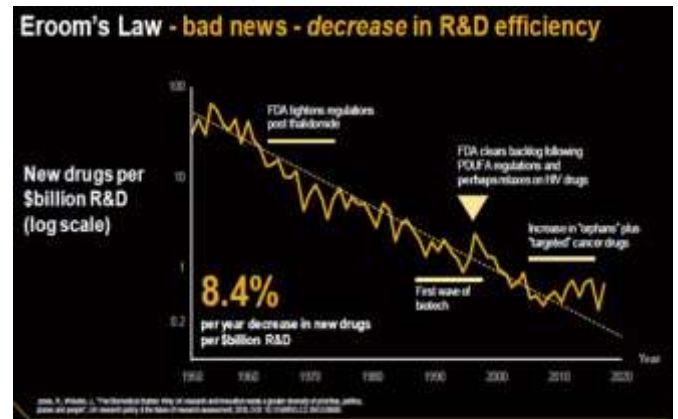


Source: SandboxAO; Drug Discovery Today

BofA GLOBAL RESEARCH

**Exhibit 156: Eroom's Law: new drugs per \$bn R&D continue to fall**

The pharma industry has seen a decrease in R&D efficiency over several decades despite the availability of new tech



Source: SandboxAO; Drug Discovery Today

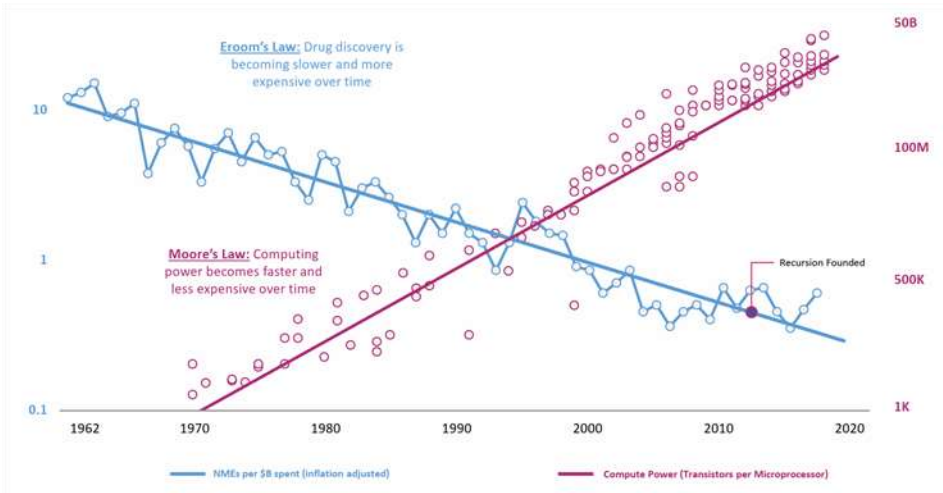
BofA GLOBAL RESEARCH

**Tech revolutionising healthcare**

Sandbox's molecular simulation division (called AQBioSim) helps biopharma and research institutions achieve breakthroughs in treatments for cancer, Alzheimer's, Parkinson's and other conditions per Hidary. He discussed a case study where a biotech company took 8 weeks to achieve results that would have taken 8 months and US\$2.5bn in expenses without its assistance. The drug is now going to clinic in the next 24 months; without molecular simulation, this wouldn't have happened until 2032. Sandbox did all of this using GPUs. Once quantum computers get better, it will use them as well.



Exhibit 157: Contrasting trends in computing power and drug discovery efficiency  
 With computing power increasing to an all-time high, drug discovery has become notably slower



Source: Company Filings, Recursion Pharmaceuticals Download Day 2023 BofA Global Research

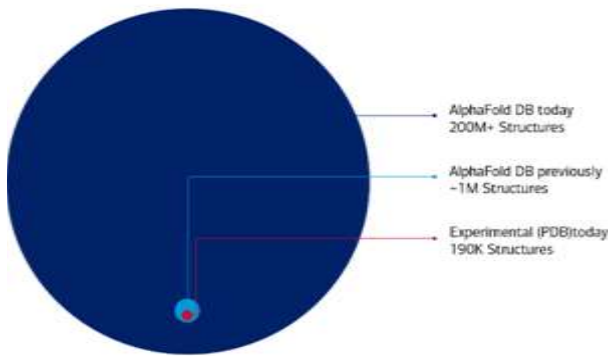
BofA GLOBAL RESEARCH

Google DeepMind: AlphaFold Isomorphic Labs

In 2022 the AlphaFold database expanded by over 200x from nearly 1 million structures to more than 200 million – with the potential to dramatically increase our understanding of biology.

Exhibit 158: Number of Protein Structures  
 AlphaFold database expanded by over 200x from nearly 1 million structures to more than 200 million in 2022

Number of Protein Structures



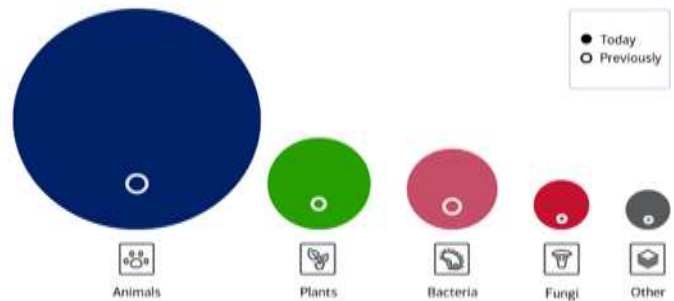
Source: Google DeepMind

BofA GLOBAL RESEARCH

Exhibit 159: Number of species represented in AlphaFold DB  
 Total increase from ~10k to ~1M

Number of species represented in AlphaFold DB

Total increase from ~10K to ~1M



Source: Google DeepMind

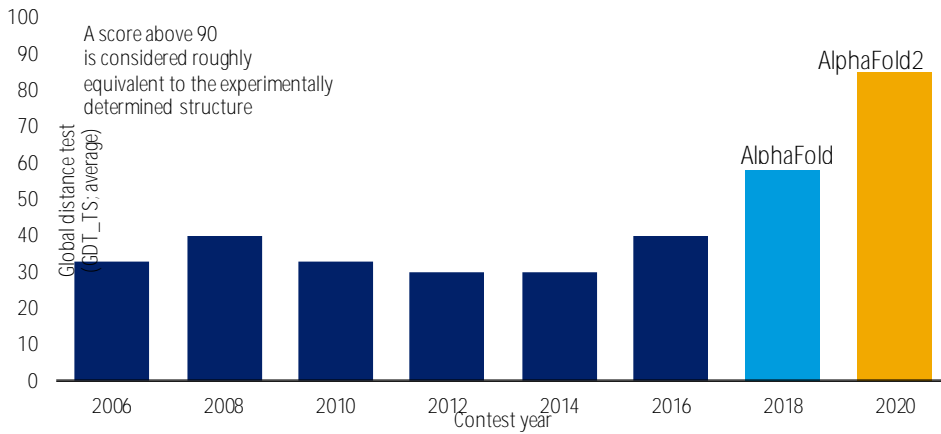
BofA GLOBAL RESEARCH

In November 2020, DeepMind's AlphaFold won the 14th Community Wide Experiment on Critical Assessment of Techniques for Protein Structure Prediction (CASP14), a biennial competition in which entrants receive amino acid sequences for about 100 proteins whose 3D structures are unknown. By comparing the computational predictions with the lab results, each CASP14 competitor received a global distance test (GDT) score. Scores above 90 out of 100 are considered equal to experimental methods. AlphaFold produced models for about two-thirds of the CASP14 target proteins with GDT scores above 90.



**Exhibit 160: Structure Solver**

DeepMind's AlphaFold2 algorithm significantly outperformed other teams at the CASP14 protein-folding contest, and its previous version's performance at the last CASP



Source: Nature

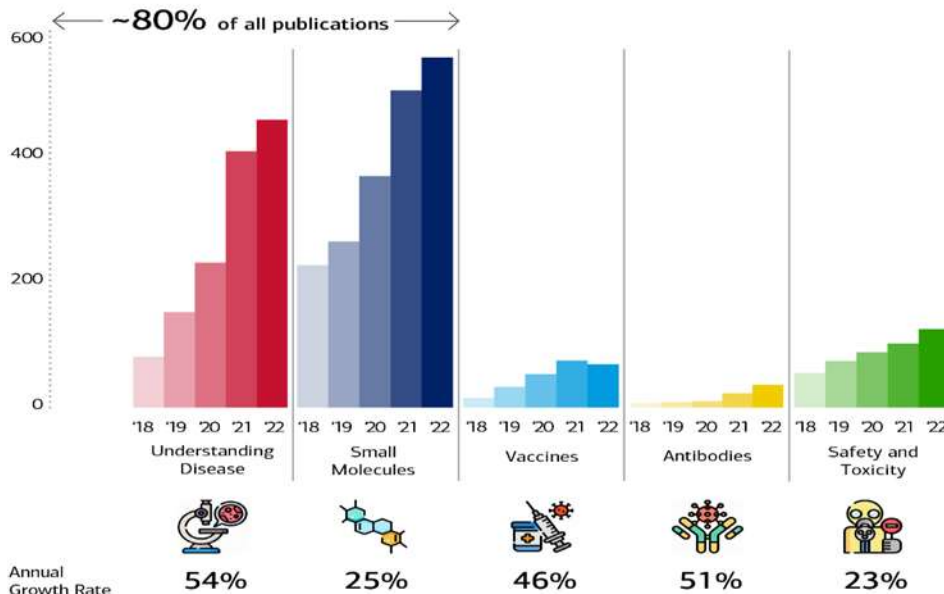
BofA GLOBAL RESEARCH

**AI drug discovery targeting multiple therapeutic areas**

**Exhibit 161: Publications on AI in drug discovery, by use case family and by year**

Nearly 80% of all publications are focused on understanding disease and small molecules

**Number of publications**



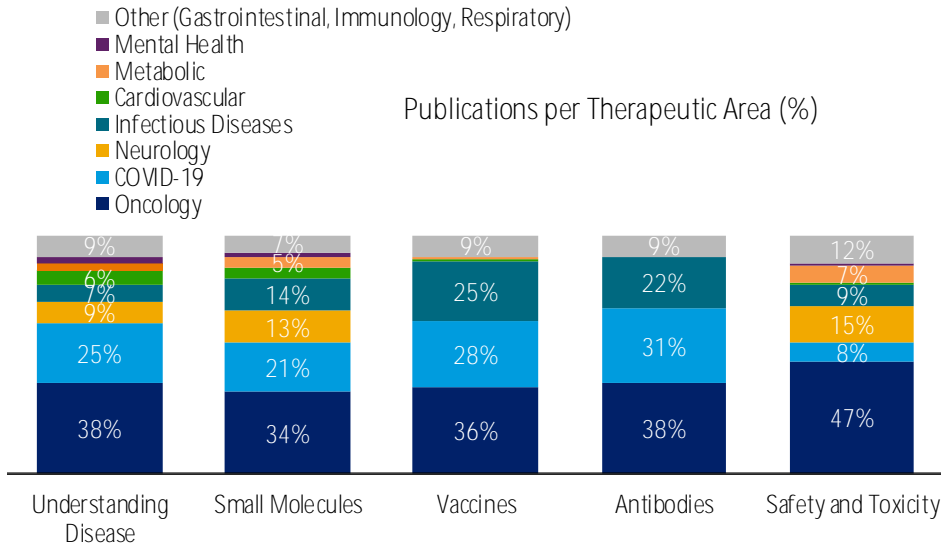
Source: BCG

BofA GLOBAL RESEARCH



Exhibit 162: Publications on AI in drug discovery, by therapeutic area

Most publications focusing on therapeutic areas that are both data-rich and commercially attractive, such as Oncology (37%) and Covid-19 (23%)

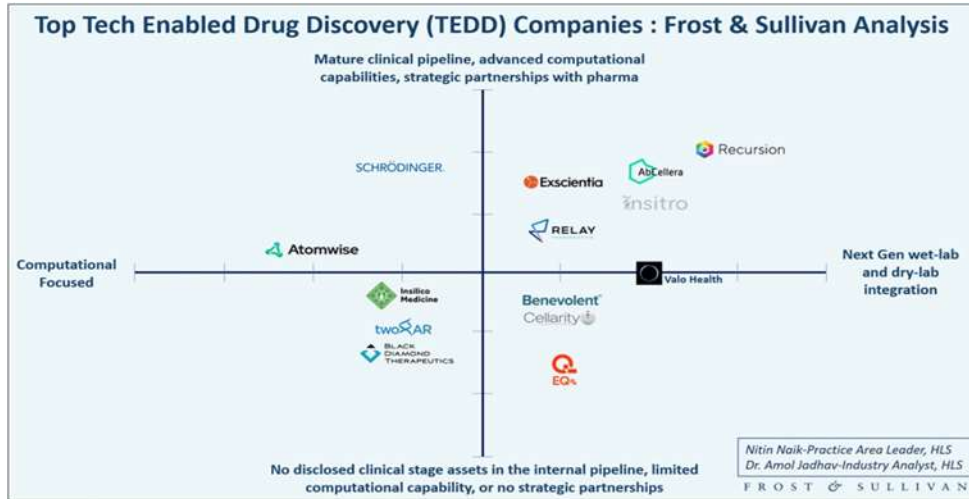


Source: BCG

BofA GLOBAL RESEARCH

Exhibit 163: There exists a diverse set of Tech-Enabled Drug Discovery companies

Despite what appears to be an already crowded field, we see little direct head-to-head competition in the near term



Source: Company filings, Frost & Sullivan

BofA GLOBAL RESEARCH



## AI Diagnostics – revolutionizing disease detection

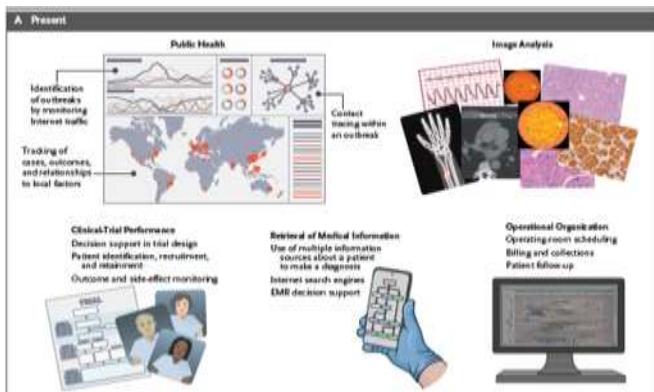
What is it? Leveraging AI to automatically diagnose diseases, making diagnostics cheaper and more accessible

Did you know? Google’s DeepMind correctly diagnosed eye diseases 94.5% of the time in a trial with Moorfields Eye Hospital

Correctly diagnosing diseases takes years of medical training. Even then, diagnostics is often an arduous, time-consuming process. However AI has recently made huge advances in automatically diagnosing diseases, making diagnostics cheaper and more accessible. AI is particularly helpful in areas where the diagnostic information a doctor examines is already digitized such as: detecting lung cancer or strokes based on CT scans, assessing the risk of sudden cardiac death or other heart diseases based on electrocardiograms and cardiac MRI images, classifying skin lesions in skin images, finding indicators of diabetic retinopathy in eye images.

The application AI in diagnostics is just beginning – more ambitious systems involve the combination of multiple data sources (CT, MRI, genomics and proteomics, patient data, and even handwritten files) in assessing a disease or its progression. AI can use image analysis of biological and medical images (microscopy, genomic heatmap, CT scans, MRI scans, etc.) to compute the effects of target modification in animal models for example. Furthermore AI compared to human scientists can process these images 10-100x faster.

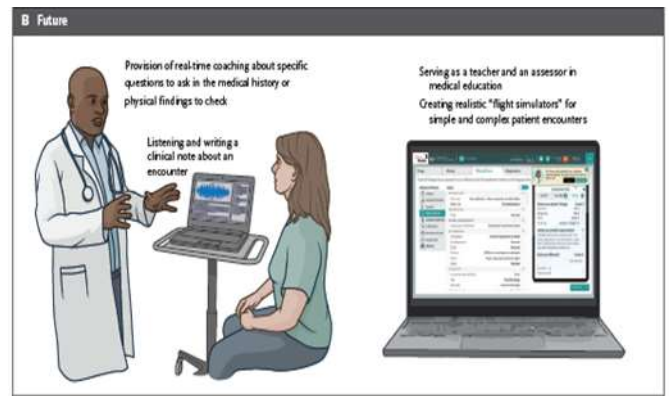
Exhibit 164: Present Healthcare  
From identification of outbreaks to patient follow-up



Source: New England Journal of Medicine

BofA GLOBAL RESEARCH

Exhibit 165: Future Healthcare  
Provision of real-time coaching, Listening and writing clinical notes, serving as teacher and as assessor, creating realistic 'flight simulators'



Source: New England Journal of Medicine

BofA GLOBAL RESEARCH





## Liquid Biopsy – disruptive diagnostics

What is it? Liquid biopsy is a “simple blood test” that can provide critical health information, especially for cancer

Did you know? Cancer accounts for nearly 10mn deaths annually and over 600k deaths in the US alone but detecting cancer earlier with liquid biopsy could drastically reduce fatalities<sup>1</sup>

Source: <sup>1</sup>WHO

Liquid biopsy (LB) is a disruptive technology in diagnostics – a “simple blood test” that can provide critical health information, especially for cancer. LB is emerging as a non-invasive alternative to tissue-based biopsies, aiding therapy selection in late-stage cancers and monitoring for cancer recurrence, and could potentially play a role in early-stage cancer screening. However, many technical and regulatory challenges remain but nonetheless is growing to become a \$45-70bn opportunity in the US alone

Exhibit 166: Liquid biopsy: a \$45-70bn market opportunity in the US  
Liquid biopsy addressable market is substantial and very underpenetrated

Application	US market size	Addressable US patient population	Other key assumptions
Therapy selection	\$2-3bn	700,000 late-stage cancer patients	85% molecular testing (MDx) penetration + 85% LB penetration among MDx; ASP \$3,500-6,000
Biopharma	\$1-2bn	360,000 patients	Over 300 targeted therapy assets x 3 trials / asset x 120 avg. patients / trial; 1,800 I-O trials x 100 avg. patients / trial; ASP \$4,000-6,000
Recurring monitoring	\$15-20bn	17mn cancer survivors	60-70% adoption rate for LB recurrence tests; ASP \$1,000-2,000 annually with various test frequency; not incl. RUO/biopharma opportunity
Early detection	\$25-50bn	43mn lung, liver, breast or colorectal cancer high-risk individuals; Or 107mn individuals age 50-79	Annual or biennial testing; ASP ~\$500; TAM would increase if targeting average-risk individuals
<b>Total est. market size:</b>	<b>\$45-70bn</b>		

Source: BofA Global Research, company reports

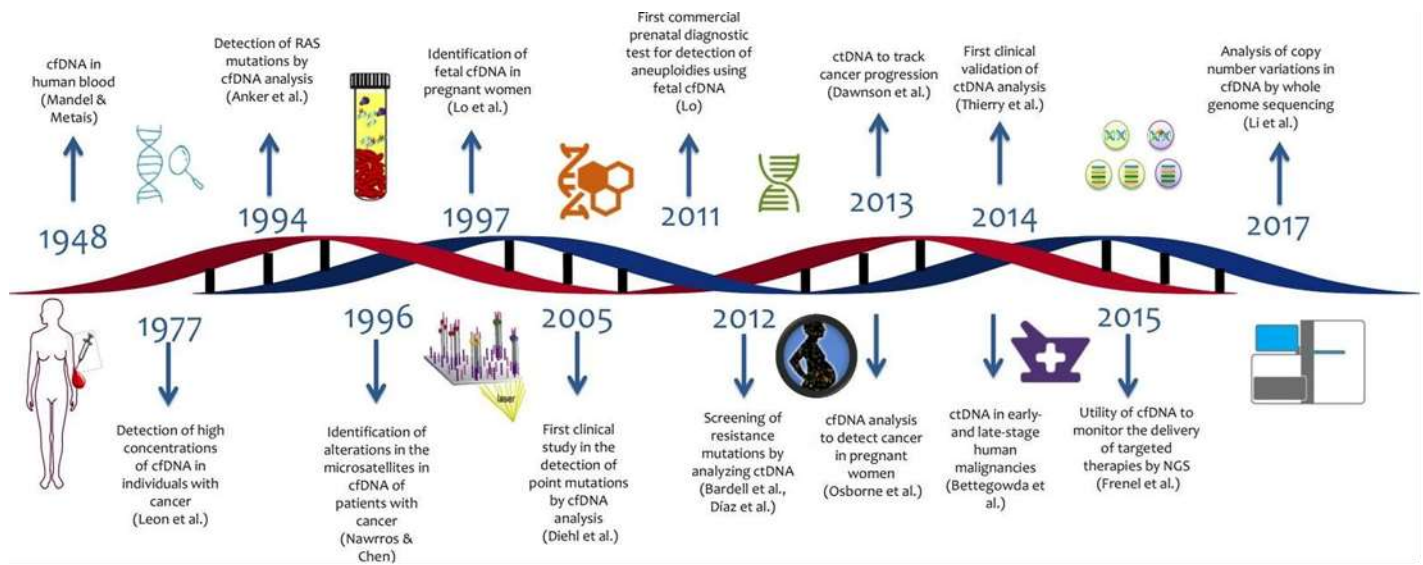
BofA GLOBAL RESEARCH

For starters the increasing accessibility and decreasing cost of next-generation DNA sequencing (NGS) throughout the 2010s catalyzed the exploration of liquid biopsy in cancer, as most of the LB technology is sequencing-based. With the promise of early disease detection and non-invasive disease monitoring moving from the lab to the clinic, the first clinical validation of ctDNA analysis occurred in 2014.



**Exhibit 167: Timeline of liquid biopsy development**

Despite over two decades of innovation in life science technologies, developing clinically useful liquid biopsies to detect cancer earlier and monitor cancer survivors continue to remain a scientific challenge

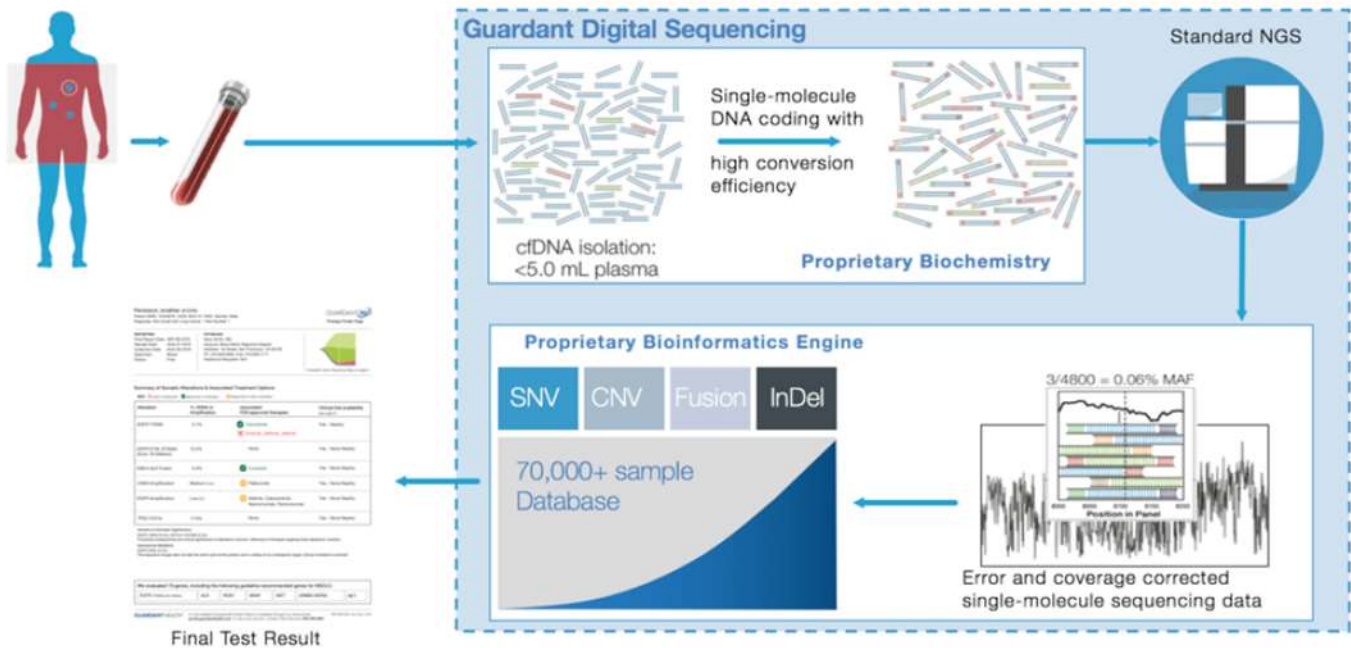


Source: Dominguez-Vigil et al.

Liquid biopsies start from collecting a patient sample, typically a blood draw where chemicals in the sample collection tube help stabilize any cfDNA. Several ‘sample preparation’ and analysis processes typically follow, such as plasma isolation and analyte extraction, target enrichment to select a subset of genes for analysis, NGS library preparation and DNA sequencing at depth (i.e., 500-1000x coverage), followed by data analysis, interpretation and reporting.

**Exhibit 168: Example workflow of liquid biopsy**

Process includes blood collection, shipping, laboratory processing, analysis and reporting



Source: Guardant Health



## Genomics – gene therapy, gene editing, CRISPR

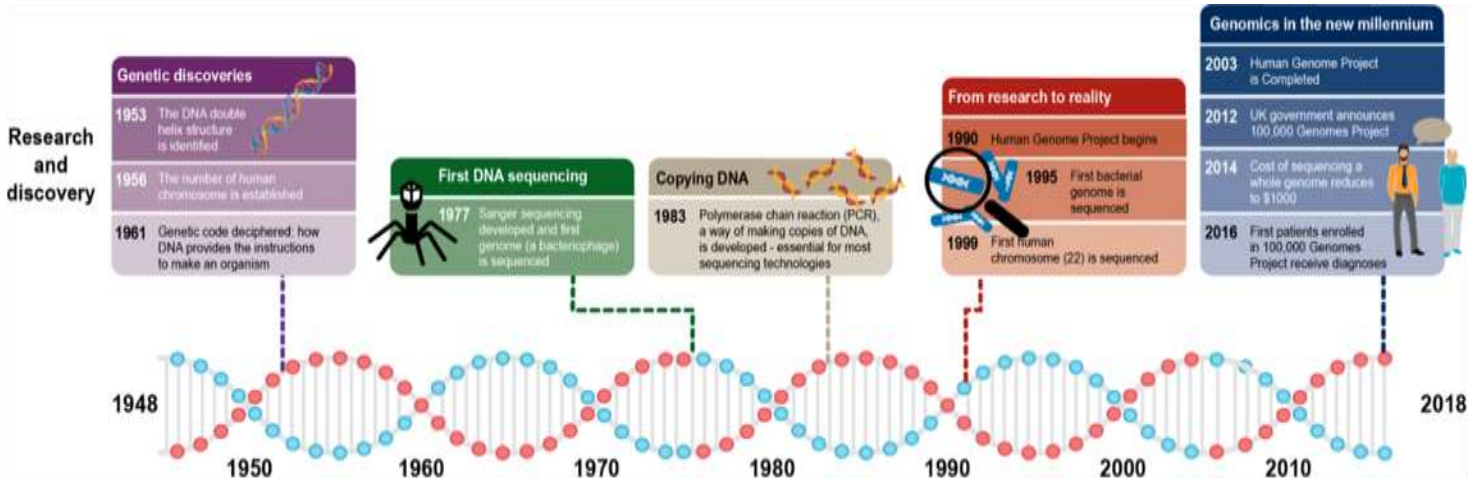
What is it? Genomics is the study of genes and their functions, and related techniques, whereas genetics is the study of heredity

Did you know? A coffee mug full of DNA could theoretically store all of the world's data (Mark Bathe, MIT)

Genomics is defined as the study of genes and their functions, and related techniques, whereas genetics is the study of heredity (source: WHO). Although DNA was first isolated as early as 1869, it took more than a century for the first genomes to be sequenced. The history of genomics dates back to the 1970s when the scientists determined the DNA sequence of simple organisms. However, the greatest breakthrough occurred in the mid-1990s when scientists sequenced the entire genome of a free-living organism (influenza). Since the turn of the millennium, the pace of innovation has accelerated further with the dawn of low-cost nextgen DNA sequencing techniques (source: Genomics.org).

Exhibit 169: Timeline of Genomics innovation

From the discovery of the DNA double helix to the Human Genome Project and more recently nextgen sequencing & CRISPR



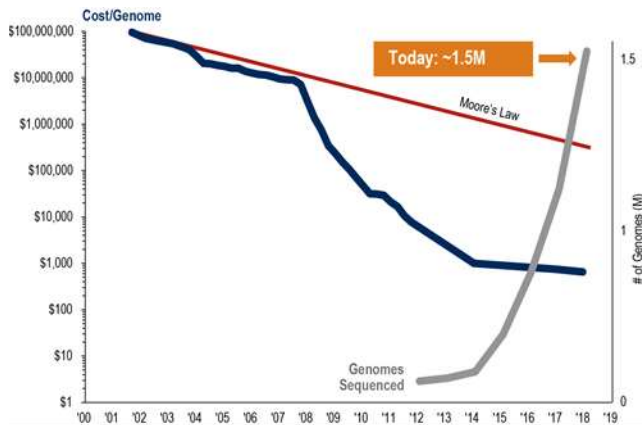
Source: Genomics.org

BofA GLOBAL RESEARCH

Recent advances in low-cost genome sequencing could take us to the next frontier in healthcare and medicine by enabling precision medicine. The combination of genetic information and proteomics, the study of proteins, means we can create treatments personalized for the individual. While the first sequencing of the human genome in 2003 cost around US\$2.7bn, this dropped to below US\$1,000 by 2014 with actors looking to reduce this to just US\$100 per genome and less than 1 hour in the coming years. Hence the ability to cheaply sequence has implications spanning drug discovery, medicine, and wellness, to food agricultural applications and synthetic biology.

**Exhibit 170: Cost declines of genome sequencing exceed Moore's Law**

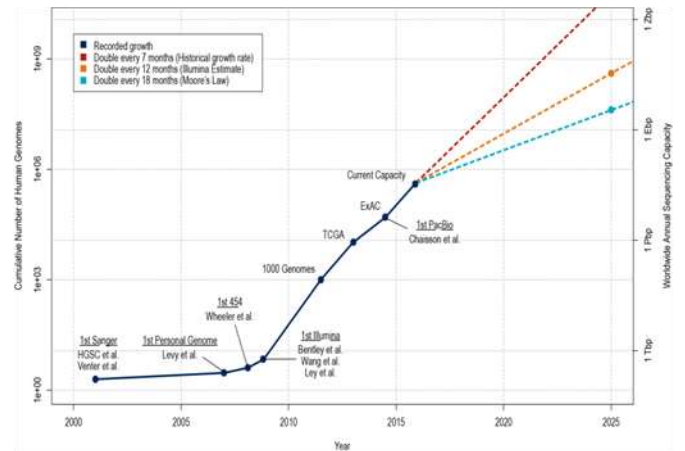
Cost dropped to below \$1,000 by 2014 with actors looking to reduce this to just US\$100 per genome and less than 1 hour in the coming years



Source: Illumina, NHGRI

BofA GLOBAL RESEARCH

**Exhibit 171: Growth of DNA sequencing**  
Growth rate has historically been doubling every 7 months



Source: PLoS Biology

BofA GLOBAL RESEARCH

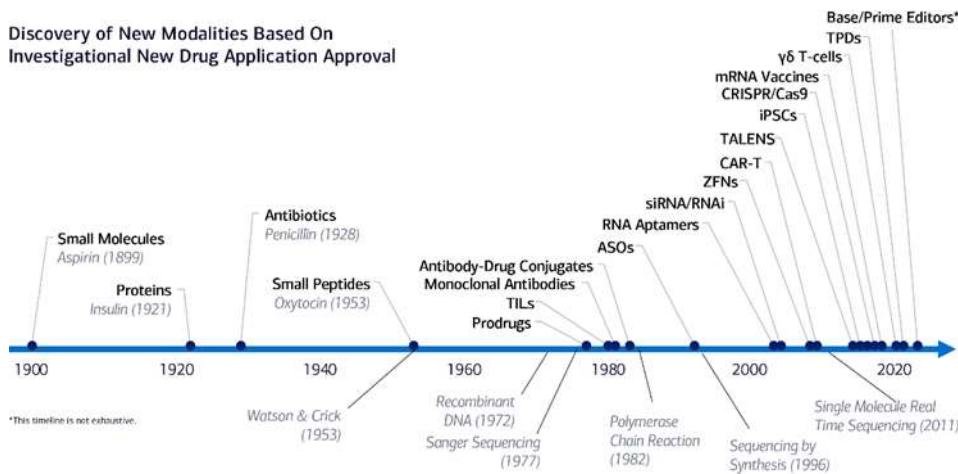
**Gene therapy: large opportunity with huge unmet need**

Gene therapy and gene editing are broadly speaking two approaches designed to correct the function of an abnormal gene. The key difference is that gene therapy aims to restore function by adding a new therapeutic genetic sequence, while gene editing focuses on directly altering specific genetic sequences in the genome. These approaches are best suited to monogenic diseases, which are caused by specific mutations in a single gene. There are an estimated 4,000 genetic-related disorders, meaning the gene therapy and gene editing markets represent a large opportunity with great unmet need. Although there are only 2 FDA-approved gene therapies in the US today, thousands of clinical trials for gene therapy and gene editing are currently being run worldwide.

**Exhibit 172: What is next in Genomics?**

New genomic tools are growing quickly

**Discovery of New Modalities Based On Investigational New Drug Application Approval**



Source: ARK

BofA GLOBAL RESEARCH

The two main approaches for gene delivery are viral and non-viral vectors. Viral vectors are currently a delivery vehicle used in FDA-approved gene therapies. Non-viral techniques are currently being studied as a safe and effective way to deliver genetic material to cells for therapeutic effect. Each approach has advantages and disadvantages in terms of efficacy, safety, manufacturing and packaging capacity. Viral vectors are the most used method for gene therapy and gene editing in the clinic. Nevertheless, several non-viral vector technologies are currently being developed to



overcome the current limitations. Significant advantages from using non-viral vectors and safety concerns surrounding viral vectors could pave the way for the former to become more widely used in gene therapy and gene editing.

**Exhibit 173: Key characteristics of viral vectors**  
 Different viral vectors have specific properties and applications

	AV	AAV	LV
Packaging capacity	8-36 kb	4.5kb	8kb
Gene expression	Transient	Transient	Stable
Cells infected	Dividing and non-dividing	Dividing and non-dividing	Dividing and non-dividing
Transduction efficiency	High	Moderate	Moderate
Integration	Non-integrating	Non-integrating	Integrating
Immune response	High	Low	Moderate
Approach	In vivo/ex vivo	In vivo/ ex vivo	Ex vivo

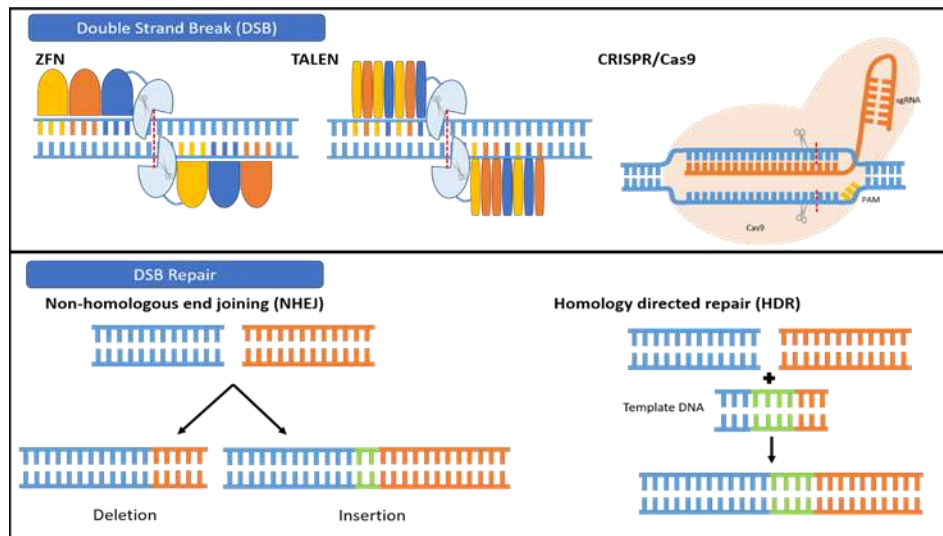
AV: adenovirus, AAV: Adeno-associated virus, LV: Lentivirus, kb: kilobase  
 Source: BofA Global Research

BofA GLOBAL RESEARCH

**Gene editing: directly modifying the genome**

In contrast to gene therapy, gene editing (or genome editing) directly modifies a patient’s existing genome to achieve a therapeutic effect. It has multiple applications to modify a patient’s genome including gene insertion, deletion (knockout), correction, modification, activation, or a combination of editing targets (multiplex editing). Gene editing can be accomplished through a variety of methods. Similar to gene therapy, it can be performed ‘in vivo’ by infusing the necessary biological machinery directly into a patient, or ‘ex vivo’, where cells are removed, modified, and transplanted back into the patient. Below we discuss the core technologies used in modern gene editing therapies, review the companies that have prominent gene editing programs, assess clinical data reported to date, and discuss current challenges in the field.

**Exhibit 174: Gene editing technologies**  
 Different technologies to generate DSB and repair mechanisms



Source: Beam Website

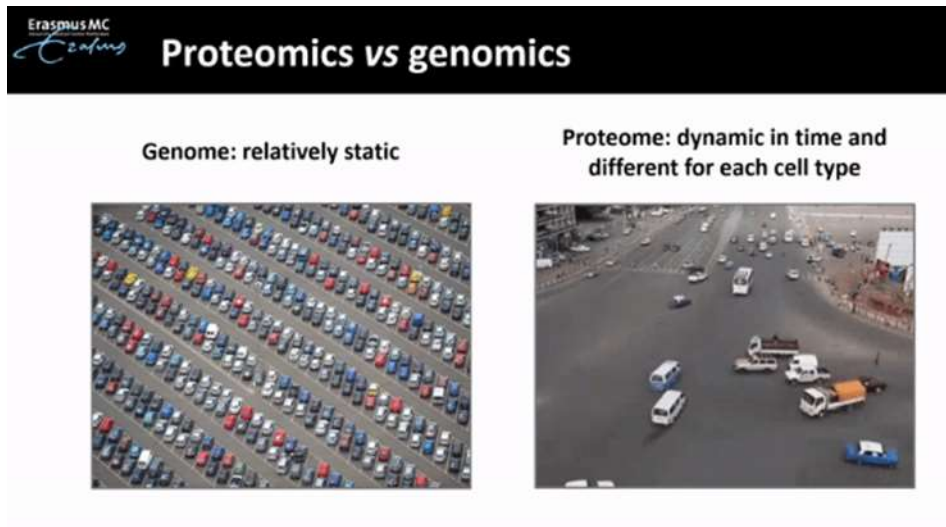
BofA GLOBAL RESEARCH



**Proteomics: next big thing after genomics**

Proteomics is roughly defined as large scale study of proteins. The proteome is defined as the entire set of proteins produced or modified by an organism or system. The key technology to study proteins at a large scale is mass spectrometry. The term proteomics was coined to make an analogy with genomics, the study of the genes. Proteomics, however, is much more complicated than genomics, mostly because while an organism’s genome is rather static/constant, the proteome differs from cell-to-cell and constantly changes dynamically through its biochemical interactions with the genome and the environment.

Exhibit 175: Genomics vs Proteomics  
Static vs Dynamic

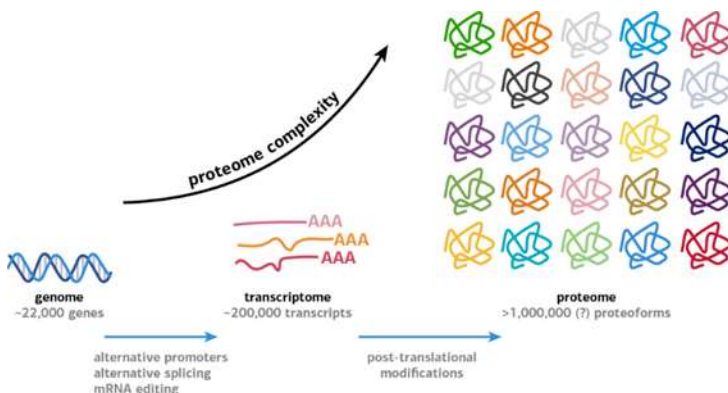


Source: Erasmus MC

Proteomics is an important new field because it can help identify and monitor biomarkers and can also facilitate drug development by providing a comprehensive map of protein interactions associated with disease pathways. For instance proteomics has previously been used to expose undesirable, off-target effects in the drug discovery process. An example of this was the use of mass spectrometry to uncover the molecular link between liver toxicity and Troglitazone, a since withdrawn type-II diabetes medication.

Exhibit 176: Basics of Proteomics

Proteomics is the next step in the study of biological systems and is loosely defined as the global analysis of proteins in a protein complex, organelle, cell, tissue or complete organism



Source: Erasmus MC





## CRISPR – DNA 2.0 and 3.0

What is it? CRISPR 1.0 is like a pair of genetic scissors whilst CRISPR 2.0 is more like a pencil and eraser and CRISPR 3.0 would allow scientists to insert entirely new genes into a person's genome.

Did you know? The human genome contains roughly 3-6 billion DNA base pairs letters for editing

“CRISPR” is an acronym for “Clustered Regularly Interspaced Short Palindromic Repeats,” which are unique DNA sequences found in some bacteria and other microorganisms. The most studied CRISPR system is associated with the “Cas9” protein and is known as CRISPR-Cas9. In 2012 and 2013, researchers modified CRISPR-Cas9 to serve as an effective and efficient technology for editing the genomes of plants, animals, and microorganisms. Since then, CRISPR-Cas9 has been used to modify the genomes of a variety of species, ranging from mice and fruit flies to corn and yeast.

Hence “CRISPR-Cas9” is a gene editing technology that offers the potential for substantial improvement over other similar technologies in ease of use, speed, efficacy, and cost. These characteristics led Science magazine to name CRISPR-Cas9 gene editing technology “Breakthrough of the Year” in 2015. It is also important to differentiate between the use cases of CRISPR: Somatic gene editing (occurs in single cell, offspring cannot inherit mutation) and Germline (occurs in gametes, offspring can inherit mutation).

Many in the scientific, engineering, and business communities believe that CRISPR-Cas9 could offer revolutionary advances in the investigation, prevention, and treatment of diseases; understanding of gene function; improvement of crop yields and development of new varieties; production of chemicals used in biofuels, adhesives, and fragrances; and control of invasive species. For instance, CRISPR-related approaches (gene drives) are being considered to reduce or eliminate the mosquito, which serves as the primary vector for the transmission of Zika or malaria, thereby improving public health (source: US Congressional Research Service).

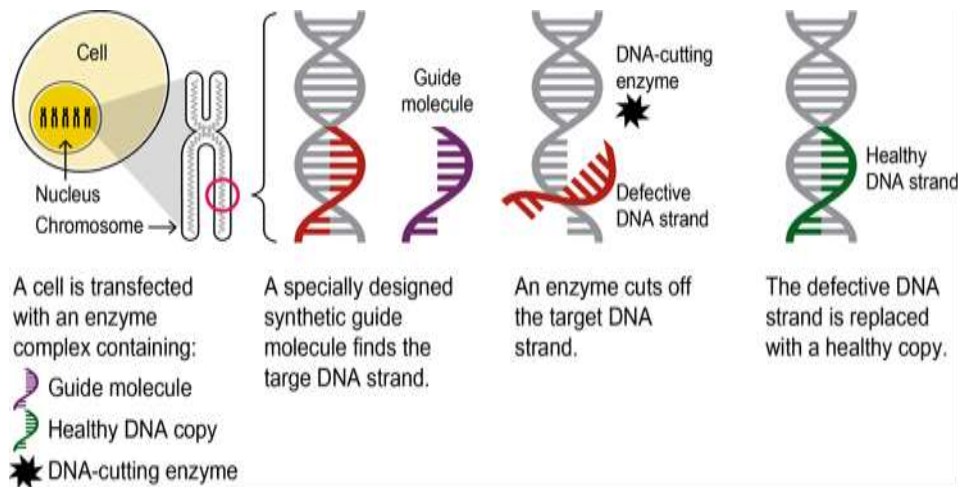
If CRISPR 1.0 is akin to a pair of scissors, then CRISPR 2.0 (base editing) is more like a pencil and eraser

CRISPR 1.0, the first generation of the technology, essentially makes cuts in the DNA. Cells repair these cuts, and this process usually stops a harmful genetic mutation from having an effect.



**Exhibit 177: How the CRISPR technique works**

CRISPR 1.0, the first generation of the technology, essentially makes cuts in the DNA. Cells repair these cuts, and this process usually stops a harmful genetic mutation from having an effect



Source: Genetic Literacy Project

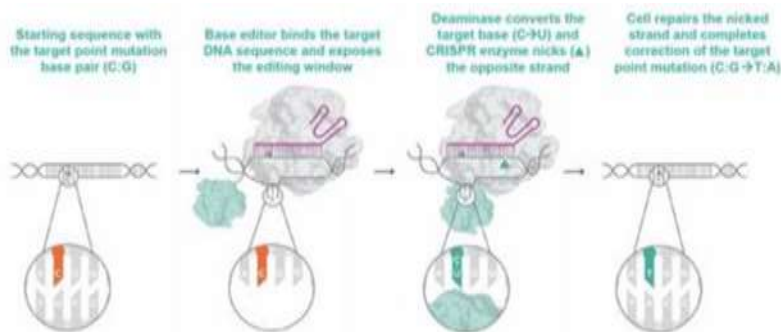
BofA GLOBAL RESEARCH

**CRISPR 2.0: base editing**

CRISPR 2.0 is the base editing technique that targets the core building blocks of DNA, which are called bases. The human genome contains roughly 3 billion base pairs (letters) or 6 billion DNA in a cell. Overall there are four DNA bases known as A, C, G and T. These letters pair off — A with T and C with G — to form DNA’s double helix. Instead of cutting the DNA, CRISPR 2.0 can convert one base letter into another. Base editing can swap a C for a T, or an A for a G. In theory, base editing should be safer than the original form of CRISPR gene editing. Because the DNA is not being cut, there’s less chance that you’ll accidentally excise an important gene, or that the DNA will come back together in the wrong way.

**Exhibit 178: Base editing mechanism**

At the molecular level base editing has a precise, reproducible editing mechanism



Source: Beam 10-K

BofA GLOBAL RESEARCH

**CRISPR 3.0: prime editing**

This technique allows scientists to replace bits of DNA or insert new chunks of genetic code. It has only been around for a few years and is still being explored in lab animals. But its potential is huge because prime editing vastly expands the options. “CRISPR 1.0” and base editing are somewhat limited – you can only use them in situations where cutting DNA or changing a single letter would be useful. Prime editing could allow scientists to insert entirely new genes into a person’s genome.

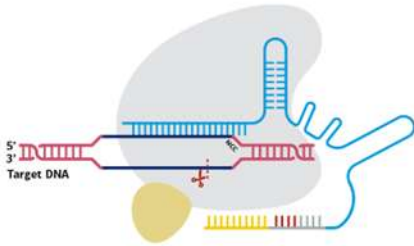


Prime editing is a 'search-and-replace' genome editing technology in molecular biology by which the genome of living organisms may be modified. The technology directly writes new genetic information into a targeted DNA site. It has received mainstream press attention due to its potential uses in medical genetics. It utilizes methodologies similar to precursor genome editing technologies, including CRISPR/Cas9 and base editors.

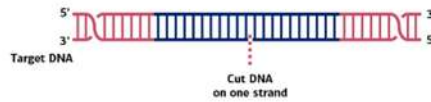
**Exhibit 179: How does Prime Editing work?**

In the newest PE system, PE3 and PE3b the unedited strand is nicked by a Cas9 nickase and the newly edited strand is used as a template to repair the nick, thus completing the edit

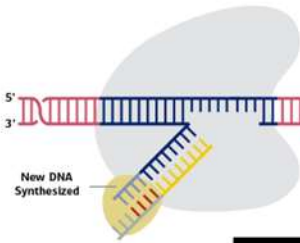
**1. PE:pegRNA Complex Binds to Target DNA**



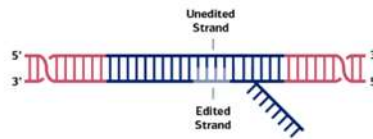
**2. Cas9 Nickase Cuts One Strand of DNA**



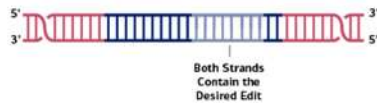
**3. Reverse Transcription of RNA Incorporates Desired Sequence Into Target DNA**



**4. Edited Strand is Incorporated. Original DNA is Cleaved by Cellular Endonuclease**



**5. Unedited Strand is Repaired to Match Newly Edited Sequence**



Source: Synthego

BofA GLOBAL RESEARCH



## GLP-1 – heading towards ‘GLP-2.0’

What is it? Glucagon-like peptide 1 (GLP-1) are hormones responsible for the “incretin effect” that makes you feel ‘more full’ and eat less

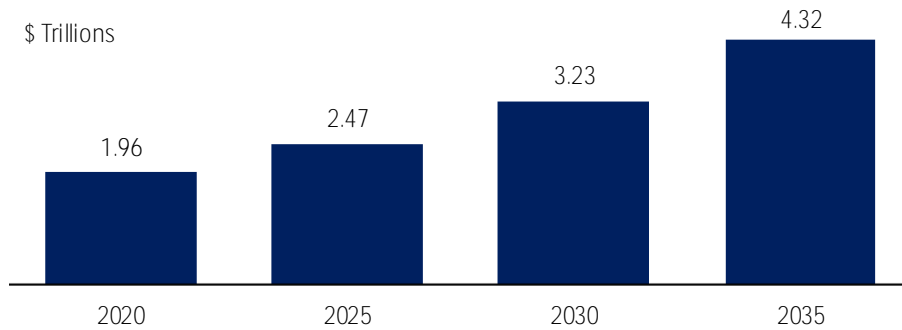
Did you know? The costs of obesity are greater than Germany’s GDP

Obesity could cost the global economy \$4.3 trillion by 2035

Obesity remains one of the most pressing public health challenges of our time. More than half of the world’s population is projected to be overweight or obese by 2035, with an estimated economic impact of \$4.3 trillion (around 3% of global GDP). For the US specifically, according to the CDC (Centers for Disease Control and Prevention), obesity costs the healthcare system nearly \$173 billion a year.

Exhibit 180: **Global economic impact of high BMI (BMI ≥ 25kg/M2), 2020-2035**

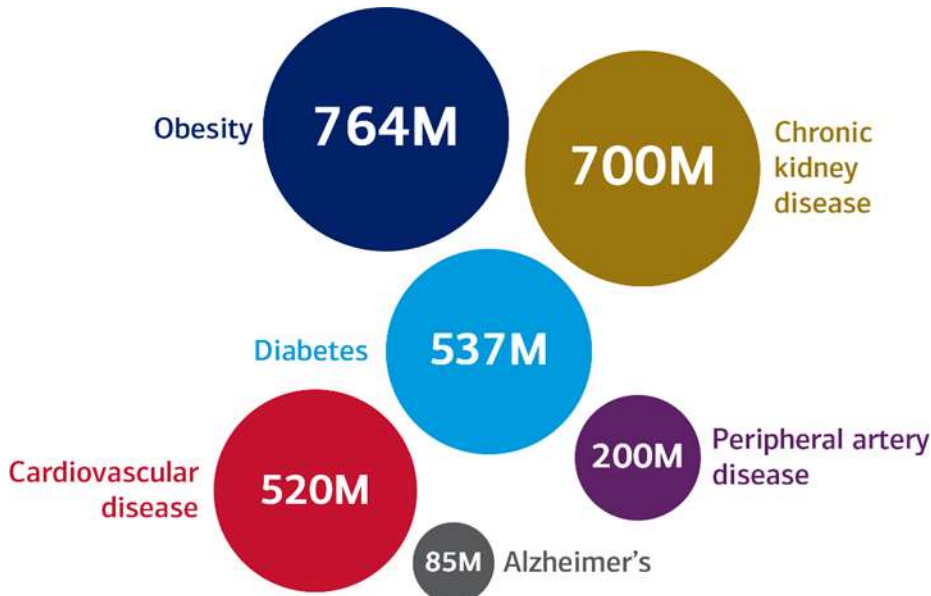
Obesity could have an estimated economic impact of \$4.3 trillion (around 3% of global GDP) by 2035



Source: World Obesity Federation, World Obesity Atlas 2023; Candesc Research and Analysis

BofA GLOBAL RESEARCH

Exhibit 186: Global patient populations across potential treatment areas for GLP-1 class drugs  
A new era of weight loss drugs may have far-reaching benefits



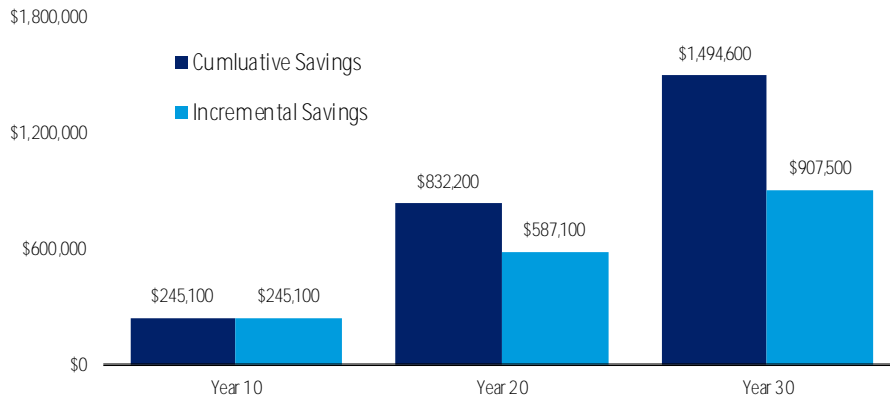
Source: Capital Group, Alzheimer's Disease International, American Heart Association, International Diabetes Federation, National Institutes of Health. Figures represent latest estimates available, as of November 16, 2023. Peripheral artery disease figure from NIH

BofA GLOBAL RESEARCH



GLP-1 reducing obesity could save US healthcare \$1.5 trillion over 30 years  
 A USC Schaeffer white paper found that coverage for new obesity treatments could generate \$245bn in cost offsets to Medicare in the first 10 years alone. By 30 years, these offsets would increase to \$1.5 trillion assuming that 100% of obese people takes GLP-1. As of mid-2023, about 4m people used one of the drugs, which represents only 1% of the US population.

Exhibit 181: Potential \$1.5 trillion in savings over 30 years from GLP-1  
 Estimated savings from GLP-1s in Medicare over time (millions)



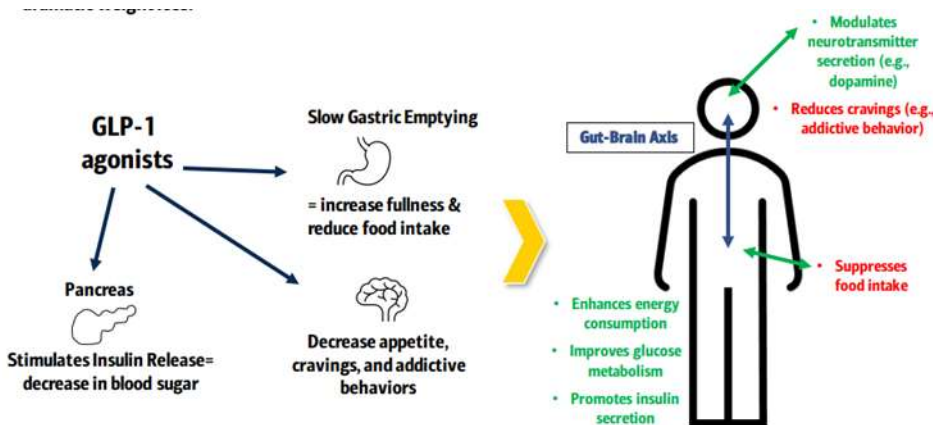
Source: USC Schaeffer Center White Paper Series

BofA GLOBAL RESEARCH

How does GLP-1 work?

Glucagon-like peptide 1 (GLP-1) and glucose-dependent insulinotropic polypeptide (GIP) are hormones responsible for the “incretin effect” or the augmentation of insulin secretion. There are many downstream effects of GLP-1 drugs such as Novo Nordisk’s Wegovy (and Ozempic) and Eli Lilly’s Mounjaro, which can lead to dramatic weight loss.

Exhibit 182: How does GLP-1 work?  
 Schematics of GLP-1



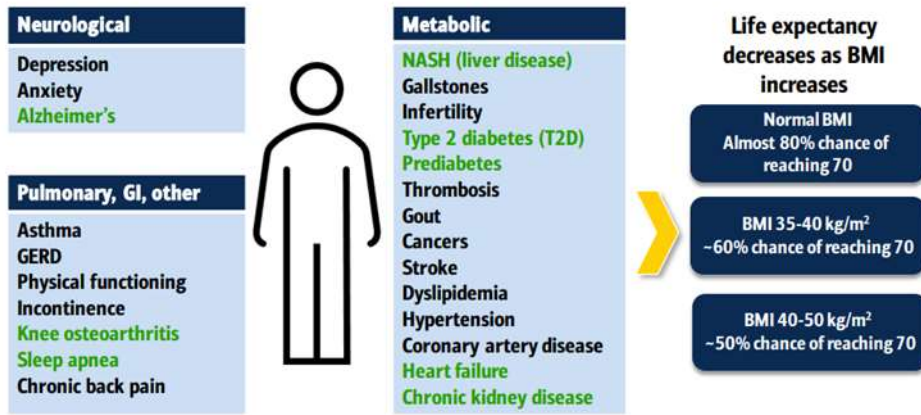
Source: BofA Global Research

BofA GLOBAL RESEARCH

An individual with obesity is at higher risk for a broad range of other diseases (e.g., comorbidities) and all-cause mortality.



Exhibit 183: Implications of obesity lower life expectancy with higher BMI

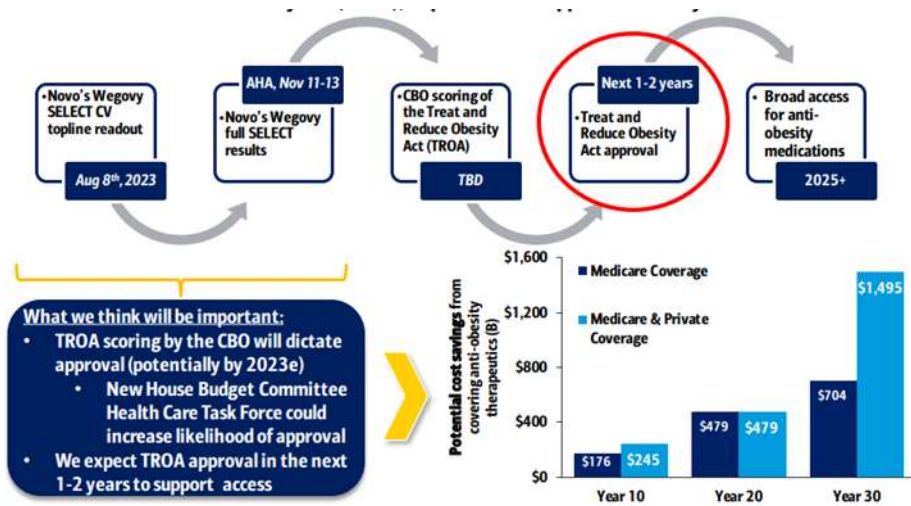


\*Green= indications beyond obesity that are in development. GLP-1s are approved already in T2D.

Source: BofA Global Research, BMI: body mass index, GERD: Gastroesophageal reflux disease, NASH - Nonalcoholic steatohepatitis  
BofA GLOBAL RESEARCH

According to the CDC, obesity costs the US healthcare system >\$173 billion annually. So, while treating obesity initially could weigh on healthcare spending, not treating it could bankrupt the system. Importantly, we now have a finish line in sight for Medicare / Medicaid coverage for anti-obesity medications with the Treat and Reduce Obesity Act (TROA), expected to be approved in 1-2 years.

Exhibit 184: Obesity Access  
What is the tipping point of demand?



Source: Source: BofA Global Research, CDC: Centers for Disease Control and Prevention, CBO: Congressional Budget Office, pt: patient, CV: cardiovascular, AHA: American Heart Association

BofA GLOBAL RESEARCH

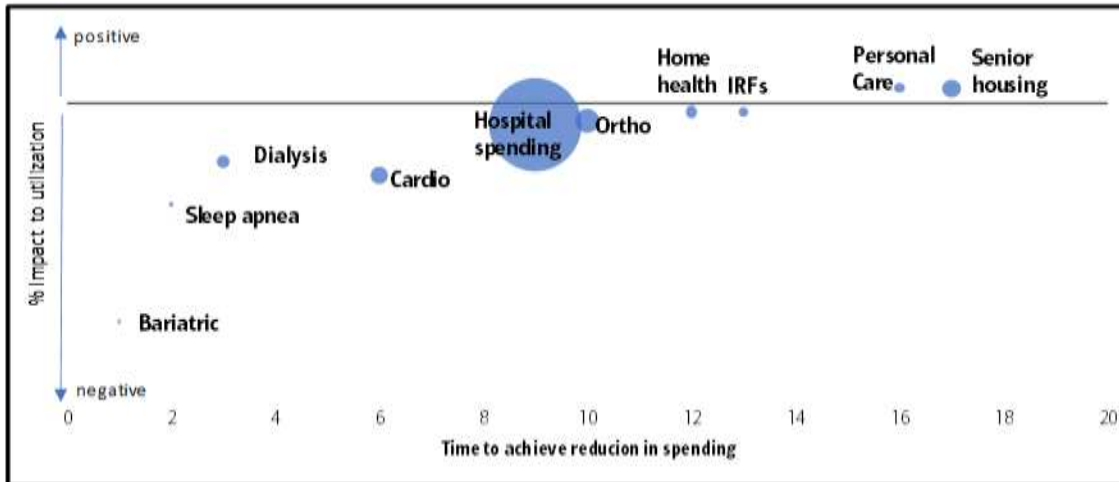
Pharma Impact: dialysis, sleep apnea

Among healthcare providers, many view cardio, dialysis, and sleep apnea as the areas that would see the most pressure on utilization of the healthcare / hospital / patient bed system. In the mid-term if there were broad and consistent use of GLP-1 drugs. We see GLP-1s impacting cardio, sleep apnea, and dialysis ahead of ortho volumes/spending. Hospitals would likely see pressure on utilization, but this would take time to build. As noted above, improvements in longevity could mitigate volume pressure and some subsectors could even benefit if people live longer (e.g., personal care, senior living), while other subsectors, that take risk, such as capitated physicians, would benefit in a way similar to MA companies.





Exhibit 185: GLP-1s are likely a headwind to sleep apnea, dialysis, some surgeries and post-acute. Could be slightly positive for personal care, senior housing  
 Estimated impact of GLP-1s by provider type. The size of the bubble represents relative category spending



Category:	Sleep apnea:	Dialysis:	Cardiology:	Hospitals:	Ortho surgeries:	Home Health IRFs:	Personal care:	Senior housing:
Companies impacted (tickers)	AHCO	DVA, FMS	CYH, HCA, THC, UHS, SGRY	CYH, HCA, THC, UHS	CYH, HCA, THC, UHS, SGRY	EHAB, AMED	EHC, SEM	ADUS, BKD

Source: BofA Global Research estimates

BofA GLOBAL RESEARCH

Non-Pharma Impact: consumer food staples, apparel, gambling, longevity  
 GLP-1s also have downstream effects such as lower appetite and a reduction in cravings/addictive behaviors. The impact of these effects is unclear but could change consumer behavior over the longer term, with some industries benefitting while others are at higher risk. To this end, we've evaluated the potential impact of broad adoption of GLP-1 drugs on consumer-focused sectors, including 1) snacks and beverages, 2) restaurants, 3) tobacco, 4) gaming, 5) apparel, and 6) food retail.

Consumer Staples: less snacking, less drinking, less unhealthy food  
 Americans consume, on average, ~1.3bn calories per day (3.8k calories/capita per the UN's FAO). Assuming uptake rates on anti-obesity medicines (AOMs) of 25-50mm people by 2030 at 15-20% caloric reduction, total calorie consumption could face a 1-3% headwind across food and beverage. 1) Intermittent eating occasions (e.g., snacking) could potentially face a more pronounced headwind vs meals (lunch/dinner), 2) In that context at-home consumption is likely more at risk than away-from-home (you'll still go out to eat). 3) Are diet/wellness products helped or hurt in this scenario? One could argue diet-friendly frozen meals/protein shakes could potentially be cut out altogether. 4) Beverages could be broadly impacted by wider consumer adoption of GLP-1. Alcohol consumption appears most at risk if consumers less and potentially experience less of an impulse to drink. The pressure on non-alcohol beverage consumption is more likely tied to the overall reduction in calories/consumption.

Exhibit 186: BofA Food/Bev team estimates a 1-3% total calorie reduction if uptake hits 25-50mm people  
 This assumes an avg per capita intake of ~3,864 calories per day and 15-20% reduction from GLP-1 usage

US Obese population uptake (# of persons)	Caloric intake reduction										
	12%	13%	14%	15%	16%	17%	18%	19%	20%	21%	22%
5,000,000	-0.2%	-0.2%	-0.2%	-0.2%	-0.2%	-0.3%	-0.3%	-0.3%	-0.3%	-0.3%	-0.3%
10,000,000	-0.4%	-0.4%	-0.4%	-0.4%	-0.5%	-0.5%	-0.5%	-0.6%	-0.6%	-0.6%	-0.7%
25,000,000	-0.9%	-1.0%	-1.0%	-1.1%	-1.2%	-1.3%	-1.3%	-1.4%	-1.5%	-1.6%	-1.6%
50,000,000	-1.8%	-1.9%	-2.1%	-2.2%	-2.4%	-2.5%	-2.7%	-2.8%	-3.0%	-3.1%	-3.3%
100,000,000	-3.6%	-3.9%	-4.2%	-4.5%	-4.8%	-5.1%	-5.3%	-5.6%	-5.9%	-6.2%	-6.5%

Source: UN FAO, Census Bureau, BofA Global Research estimates

BofA GLOBAL RESEARCH



Apparel Retail: weight loss = wardrobe replacement spending cycle

Eventual weight loss in the broader population could spur a wardrobe replacement cycle. An adoption rate of 38mn individuals (midpoint of BofA’s estimated 2030 TAM), combined with the assumption that a new wardrobe is needed every time someone goes up or down two sizes, results in \$50bn of new apparel spending. The annual US apparel market is worth \$323.6bn (2022A) according to Euromonitor International. Using this base, \$50bn would equate to an added 2.1% CAGR through 2030 from wardrobe replacements. Around 140m people in the US are classified as obese (110mn adults).

The average US womenswear size is 16-18 (XL-XXL), according to the US Department of Health and Human Services. This is up from a size 14 a decade earlier. Weight loss could aid demand for offerings from traditional retailers (who often carry sizes up to 14) and reduce demand for plus-size retailers. Conversely, athletic apparel brands could benefit given the healthier lifestyles shown to be supported by GLP-1 drugs.

A key assumption here is that the consumer can afford this expenditure. If individuals using GLP-1 save money through lifestyle changes, they could use this to fund spending on clothing. As adoption moves towards lower-income consumers, this added growth looks less likely, and we would expect the customer to trade down to lower-priced items to fund spending on new clothes.

Exhibit 187: Apparel spending necessary for wardrobe replacement for every two sizes of weight loss (in \$Billions)

A 38mn individual adoption (midpoint of BofA’s estimated 2030 TAM), assuming a new wardrobe is needed every time someone goes up or down 2 sizes, correlates to \$50bn of new apparel spending

		Weight reduction (lbs)						
		5	20	30	44	50	60	70
Population usage (in millions)	5	\$0	\$3	\$5	\$7	\$8	\$9	\$11
	15	\$0	\$9	\$14	\$20	\$23	\$27	\$32
	25	\$0	\$15	\$23	\$33	\$38	\$45	\$53
	38	\$0	\$23	\$34	\$50	\$57	\$68	\$80
	50	\$0	\$30	\$45	\$66	\$75	\$90	\$105
	75	\$0	\$45	\$68	\$99	\$113	\$135	\$158
	100	\$0	\$60	\$90	\$132	\$150	\$180	\$210
	110	\$0	\$66	\$99	\$145	\$165	\$198	\$231
	140	\$0	\$84	\$126	\$185	\$210	\$252	\$294

Source: BofA Global Research estimates, multiple as listed in table

Limitations: Our assumptions could prove conservative as we assume one wardrobe refresh per individual no matter the sizes changed, and realistically greater size changes could warrant multiple wardrobe refreshes. Our assumptions could also prove aggressive as individuals may shift regular apparel budget into new clothes instead of immediately purchase new ones, making the wardrobe refresh spend non-incremental.

BofA GLOBAL RESEARCH

Gambling: reduction in addiction to casinos

In 2013, the American Psychiatric Association recognized pathological gambling as an addiction instead of an impulse control disorder, which makes gambling problems (GP) relevant to GLP-1 and related addiction research. Gambling disorder is defined as persistent and recurrent problematic gambling behavior leading to clinically significant impairment or distress. Based on detailed research from the Australian Government, the prevalence of severe gambling problems is 0.7 per 100k, while an additional 1.7 per 100k have moderate gambling problems. Similar surveys in the UK and US have put problem gambling prevalence at 0.6%-3.6%.

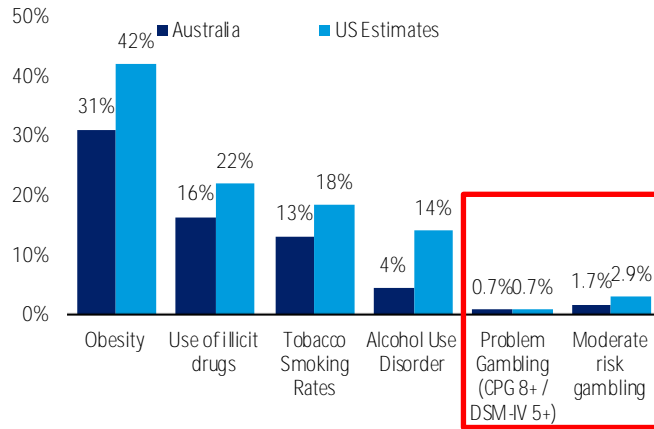
Slot machine players may account for around 75-80% of problem gamblers. Table games are viewed as less risky, and lottery-only players appear the least susceptible to problem gambling. Academic research has suggested problem gambling could account for as much as 30% of all gambling revenue and 40% of electronic gaming machines. Our own industry discussion suggests these estimates are rather high. We apply a more conservative range of 10-30%, which could lean towards a downside scenario.

Possible risks to our analysis and assumptions include 1) comorbidities may be overstated as gender and genetic proclivities for problem gambling may reduce overlap,



2) we are extrapolating data from Australia to the US, and Australia may have different population and gambling characteristics, 3) there is relatively limited academic research on problem gambling, 4) there is limited information on the true concentration of problem gambling within the top cohort of casino gambling.

Exhibit 188: Prevalence of health concerns in Australia and the US  
Gambling addictions have the lowest prevalence compared to other health concerns



Source: WHO, Australia Government, CDC, US Government, BofA Global Research  
Note: Adult (18 years old+) obesity estimates from 2017 and use of Illicit Drugs from 2019 (14 years old+) are sourced from the CDC and Australia government. Tobacco estimates from 2020 (15 years old+) and Alcohol Use Disorders (15 years old+) estimates from 2016 are both sourced from the World Health Organization. Problem Gambling and moderate risk gambling sourced from Government research in Australia and the US.

BofA GLOBAL RESEARCH

Exhibit 189: Gambling problem and GLP-1 Sensitivity analysis  
In our base case scenario, we expect a revenue headwind of -3% to the gaming industry from GLP-1s

	Low	Mid	High
US Adults Taking GLP-1s	5	15	38
US Obese Population	115	128	140
Penetration among obese population	4%	12%	27%
US Adults GP Prevalence	0.7%	1.6%	2.4%
Likelihood of obesity given problem gambling	49%	49%	49%
Prevalence of GP and Obesity	0.3%	0.8%	1.2%
US Adults with GP and on GLP-1s	0.04	0.43	0.82
Total US Problem Gamblers	1.8	4.0	6.3
% of problem gamblers on GLP-1s	2%	11%	13%
% of revenue from problem gamblers	10%	20%	30%
Revenue headwind	0%	-2%	-4%

Source: Australia Government, US Government, CDC, National Institute of Health  
Note: Low, mid, high estimates for adults taking GLP1s based on BofA estimates for current, 2025, and 2030. Prevalence of gambling problems is based on various studies from the Australia severe gambling problems at the low end to moderate gambling problems at the high end. We use Australia data because we found the Australia Government conducted the most comprehensive research related to the prevalence of gambling problems/addiction. Comorbidity rate is from research from the National Institute of Health. Percentage of revenue from gamblers estimates are based on academic research and government data from the US, Australia, and England.

BofA GLOBAL RESEARCH

### Senior Living, Home Care: longevity boost

In addition, improved longevity could drive higher utilization for post-acute care providers focused on elderly care. While the improved health of society could reduce the number of surgeries and cardio events, there would still be demand for senior care. Personal care might be among the sectors to see increased demand. These providers do not offer therapies but rather attendant care for activities of daily living (ADL) such as personal hygiene, dressing, meal preparation and medication adherence.

Similarly, senior housing facilities provide housing and support with activities of daily living to the elderly. Independent living (IL) facilities are hospitality-oriented and provide the lowest level of care. Services provided typically include meals, 24-hour emergency response, housekeeping, concierge services, transportation, and recreational activities. Most residents move in between the ages of 75 and 84. Move-ins are typically driven by: (1) a desire for a less demanding, more carefree lifestyle; or (2) a desire for the socialization afforded by a community of peers. Assisted living (AL) offers housing and the same hospitality services as IL but also helps with normal ADLs. In addition, AL provides residents with limited health care services. The latter part of the business could suffer as health improves. Nevertheless, AL specializing in Alzheimer's/ dementia would be least disrupted.

### Oral GLP-1 pill will be a gamechanger

Injectable formulations have dominated the GLP-1 market. The development of oral drugs could expand the type-2 diabetes (T2D) and underpenetrated obesity market. At the American Diabetes Association (ADA) meeting in June 2023, data were presented from multiple investigational oral daily non-peptide GLP-1 agonists. Most of these drugs, if approved, will be available towards the end of 2025 and beyond. They are small molecules that do not degrade in the GI tract and there are no stringent restrictions on administration. Oral GLP-1s could potentially broaden overall access to the market and not necessarily replace injectables. Below we discuss the oral GLP-1 landscape and recent data updates.



Exhibit 190 Overview of rough estimated weight loss profiles for key assets in development for obesity, through time  
Estimates are combination of numerical data and chart data, with estimations made therein

Drug	Mechanism	Stage of dev	Admin	Dosing	Rough estimated placebo adjusted weight loss, %				
					4wks	8wks	12wks	20wks	52wks
Zealand									
Survodutide	GLP1 + Glucagon	PIII	Subcut	Weekly	1	3	5	9	17 <sup>^</sup>
Novo									
Wegovy	GLP1	Approved	Subcut	Weekly	1	2	4	7	13
CagriSema	GLP1 + Amylin	PIII	Subcut	Weekly	2	4	6	7	
Rybelsus	GLP1	PIII	Oral	Daily	1	2	4	7	14
NN9487	Amycretin	PI	Oral	Daily	4	8	12		
NN9541	GLP1 + GIP	PI	Subcut	Daily			12 <sup>*</sup>		
LLY									
Zepbound	GLP1 + GIP	Approved	Subcut	Weekly	2	4	7	11	19
Retatrutide	GLP1 + GIP + Glucagon	PIII	Subcut	Weekly	4	7	11	13	22 <sup>*</sup>
Orforglipron	GLP1	PIII	Oral	Daily	3	5	7	9	13 <sup>**</sup>
Altimmune									
Pemvidutide	GLP1 + Glucagon	PII	Subcut	Weekly	2	4	6	8	14 <sup>*</sup>
Amgen									
MariTide (AMG133) 420mg	GLP1 + GIP	PII	Subcut	Every 4 weeks	7	12	16		
Structure									
GSBR-1290 120mg	GLP1	PI/II	Oral	Daily	3	5			
Roche									
CT388	GLP1 + GIP	PI/II	Subcut	Weekly	8				
Viking									
VK2735	GLP1 + GIP	PII	Subcut	Weekly	6		13 <sup>'</sup>		
VK2735	GLP1 + GIP	PI	Oral	Daily					
Innovent									
Mazdutide	GLP1 + Glucagon	PIII	Subcut	Weekly				15 <sup>^^</sup>	19 <sup>*</sup>
Pfizer									
Danuglipron	GLP1	PK data 1H24							

Source: Source: Journal info, Company info, BofA Global Research

Grey highlight: CagriSema placebo arm is on top of semaglutide

<sup>\*</sup>results at 48wks <sup>\*\*</sup>results at 36wks <sup>^</sup>results at 46wks <sup>'</sup>results at 24wks <sup>'</sup>results at 13wks <sup>'</sup> results at 15wks

Placebo arm contains semaglutide



# Energy and Materials meet DeepTech

A Transforming World: Installed renewable energy capacity rose from ~1.5TW in 2013 to 3.4TW in 2022. By 2030, this installed capacity could treble to reach 11TW. Our current infrastructure cannot cope and needs upgrading. Innovation can expand clean energy viability, from using AI and software to optimise the grid, to identifying new materials that could commercialise them faster.

Did you know? EV batteries alone could be used to satisfy short-term grid storage demand by 2030. One superconducting cable could transport 500x more electricity than copper wire, enough to power a large city.

Next Generation of Energy: Decarbonisation of energy supply (wind, solar, batteries), demand (EVs, heat pumps) and storage. Decentralised electricity grids and “Virtual Power Plants”. Breakthroughs in chemistry and materials (e.g. graphene, superconductors) accelerating energy breakthroughs such as nuclear fusion.

## Towards a decarbonised, decentralised, digital energy system

The global energy system is transforming to more diversified sources of power generation and storage on the supply side, and electrification of transport, buildings and industries on the demand side. This creates several difficulties in balancing the production, transmission, storage and usage of energy – particularly the electricity grid – not to mention the cost and minerals required to achieve it. Fortunately, technology, material science, chemistry and engineering are converging to mitigate these challenges and enable the next generation of energy.

## As capacity grows, clean energy prices continue to fall

Clean energy technologies continue to get cheaper as scale increases. Between 2010 and 2022 the price of wind, batteries, and solar fell by 55%/89%/90% respectively, per BNEF, with the “learning rate” (price reduction for each doubling in cumulative deployed capacity) 14%/18%/28%, respectively. Further decreases are expected – average capex of clean energy technologies could decline a further 15-20% by 2030 on average, per S&P Global.

Exhibit 191: Exponential energy: Rising scale and reducing cost of renewable energy and batteries  
Learning rates (price reductions for each doubling of installed capacity) were 14% for onshore wind, 18% for lithium-ion batteries, and 28% for solar PV between 2010-2022



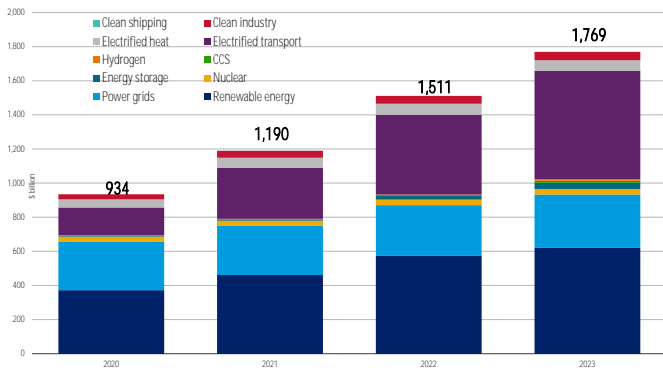
Source: BNEF

BofA GLOBAL RESEARCH



From Energy Transition to Energy Addition: clean energy capacity is multiplying  
 Global investment trends are beginning to reflect the shift to lower carbon energy, with close to \$2tn invested in a range of sectors, representing a larger share of investment than fossil fuels (BNEF, IEA). Power generation, EVs and grids made up ~90% of this investment in 2023, but the range of sectors is diversifying and seeing increasing investment in clean industry (steel, ammonia, bioplastics), storage and electrified heat, for instance. As such, total energy transition-related investment is now accelerating, almost doubling in 2020-23 alone.

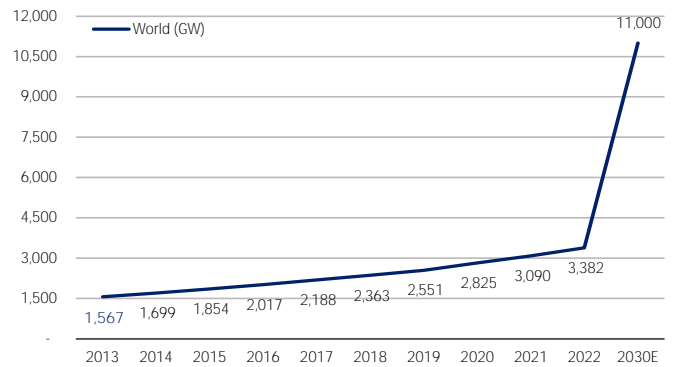
Exhibit 192: Energy Transition Investment doubled 2020-2023  
 Energy Transition investment accelerated from ~\$30bn annually in 2005 to almost ~\$1tn in 2020 and ~\$2tn in 2023 as the applicable sectors diversify



Source: BNEF Energy Transition Investment Trends 2024

BofA GLOBAL RESEARCH

Exhibit 193: Global Renewable Capacity more than doubled since 2013  
 Global renewable power capacity increased from ~1.6TW to 3.4 TW 2013-2022 per IRENA, and is projected to treble to 11TW by 2030 per IEA



Source: IRENA Renewable Energy Statistics 2023 (2013-2022), IEA WEO for 2030

BofA GLOBAL RESEARCH

Energy system will need a ~\$200tn overhaul

The total investment required to reach net zero emissions by 2050 could reach ~\$200tn between 2022 and 2050 per BNEF, roughly half of which from electric vehicle sales. However, significant investments will also be needed in the supply and transportation of energy – \$36tn in low carbon power generation, and \$21 tn on the grid. Even in a more conservative tech/economics led scenario, energy investment would have to double to \$4tn/yr on average, or \$120tn in total to 2050, per BNEF, owing to the increased electrification of the power sector. The outcome of whichever pathway is adopted will depend on policy and incentives, and the viability of earlier stage technologies, e.g. clean fuels such as hydrogen, carbon capture, and small nuclear reactors.

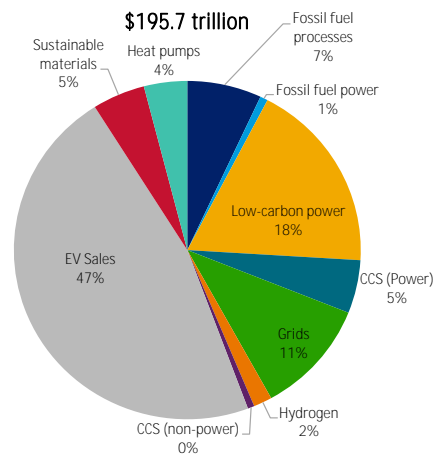
Exhibit 194: Global annual investment in energy to rise to up to \$7tn/yr  
 Global investment in energy supply & demand in range of \$120-196tn cumulatively in 2022-50 per BNEF



Source: BNEF NEO 2022

BofA GLOBAL RESEARCH

Exhibit 195: EVs would require almost half of the investment  
 Almost half of the investment would be required in EVs (\$92tn) in 2022-50



Source: BNEF NEO 2022 Net Zero Scenario

BofA GLOBAL RESEARCH





## Smart Grid Infrastructure

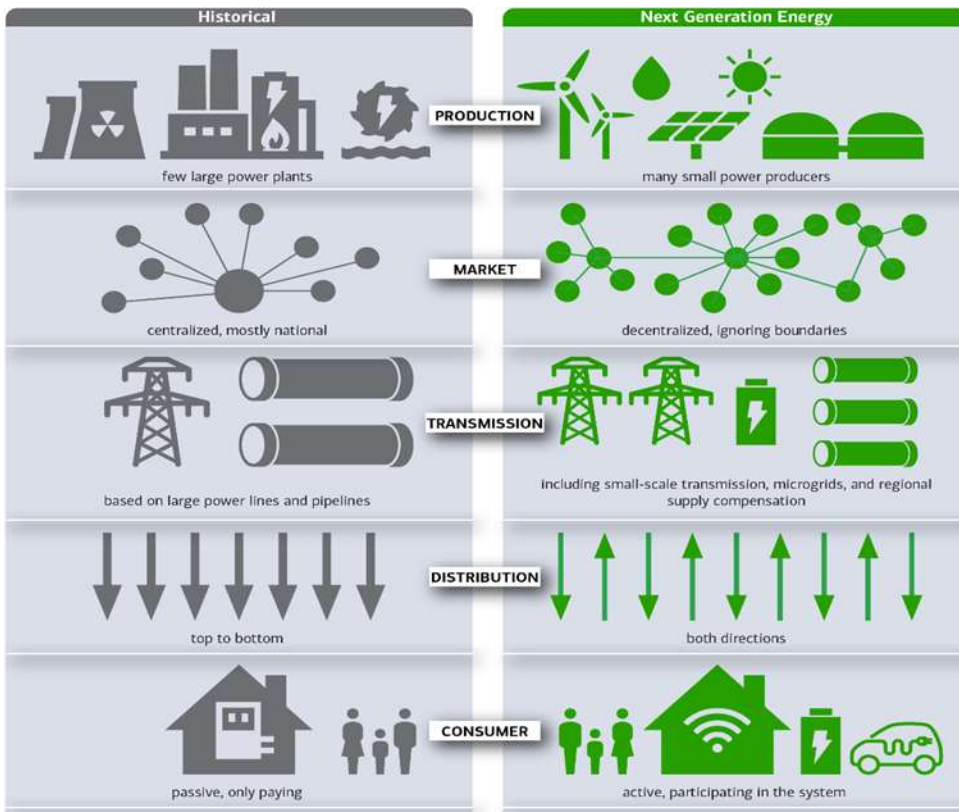
What is it? Expanding, upgrading and digitalising electricity grids to cope with the shifts in power supply (to more intermittent renewables), and demand (transport, buildings, industry) increasing the share of global energy consumption to electricity.

Did you know? Household electricity usage trebles once an electric vehicle and heat pump are added to existing consumption. The average age of grid cables in Europe is 45-50 years old. Electricity’s “market share” in final energy consumption is set to rise from 20% in 2022 to >50% in 2050 (source: Nexans, IEA)

Can the grid cope? The need for speed, security, and stability

The rising share of electricity in final energy demand will place increasing pressure on the grid to transport the volumes of electricity needed. This requires investment, but also changes to how grids are created and connected. Expanding the energy production, storage and interconnections are set to become essential to balance energy supply/demand. This also provides an opportunity, to shift from the current peak demand-led provision of energy to a more balanced supply, using a combination of price incentives, behaviour change and technology to combine the increasingly distributed energy resources.

Exhibit 196: Clean Technologies shifting from central to decentralized energy opportunity  
Structural changes in the energy system are being enabled by diversified energy sources and technology. Key trends: more power generation sources, localized grids, consumer participation in energy management



Source: Energy Atlas

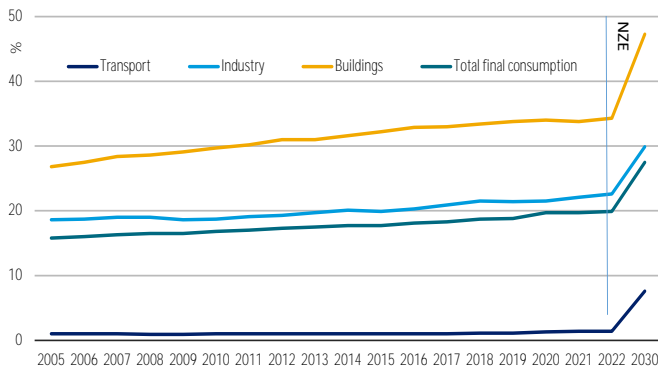
BofA GLOBAL RESEARCH



The problems: demand, diversification, delay, decay  
 increasing use of renewable sources combined with electrification of infrastructure & transport introduce far greater complexity than the installed grid can deal with as it stands. This presents several grid infrastructure-related challenges:

- The power we're supplying...it's electrifying. As more end markets electrify, we're going to need (a lot) more electricity supply. For example, electricity usage of a typical household triples once an electric vehicle (~3MWH/pa) and heat pump (~4MWH/pa) are added to existing consumption. Electrifying industry (e.g. steel, cement, ammonia and chemical production) will increase consumption even more. Indeed, electricity share in final energy consumption rise from 20% today to >50% in 2050 per the IEA. Volumes of electricity flowing through global electricity grids could rise by 2-3x from 23k TWH in 2020 to 39-74k TWH in 2050 per BNEF.

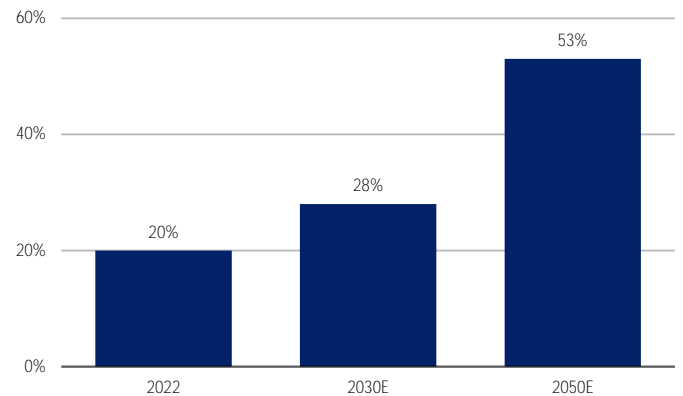
Exhibit 197: Share of energy consumed as electricity is set to accelerate  
 The IEA's net zero scenario projects a rise in the share of electricity in final energy consumption from 20% in 2022 to 28% in 2030



Source: IEA WEO 2022 Net Zero Scenario

BofA GLOBAL RESEARCH

Exhibit 198: More than half of energy could be electrified by 2050  
 Electricity expected to more than double its share of energy consumption to over 50% under IEA's net zero scenario



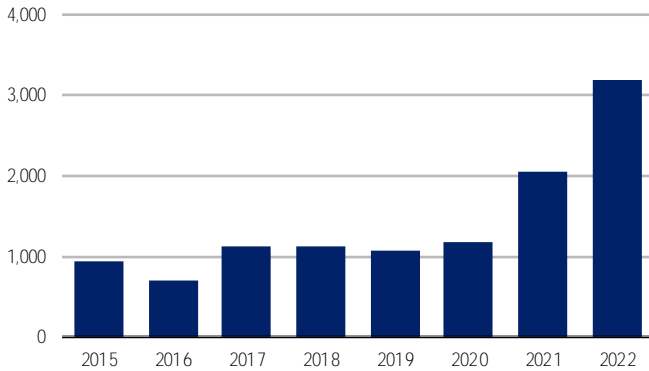
Source: IEA WEO 2022 Net Zero Scenario

BofA GLOBAL RESEARCH

- Shift in supply *and* consumption = complexity, new peaks: Shifting both the supply of energy (to a more diversified set of power generation sources vs centralised fossil fuel power generation) and demand patterns (when to charge EVs, heat water/homes, run heavy industrial equipment) brings complexity, particularly in projecting peak demand and dispatching the clean energy to accommodate it.
- Delayed connections & “congestion”: What we have cannot cope. There are significant bottlenecks related to:
  1. Solar and wind projects being connected to the grid – Europe had 600GW and US 900GW awaiting connection at end-2021, and new renewable plants can take 5-10 years to connect and integrate to the grid per BNEF.
  2. Increasing curtailment costs owing to a lack of grid connections to regions that could utilise the power: Germany’s congestion costs rose from €1bn in 2020 to €3bn in 2022 (per BNEF). In the UK, costs from wind curtailment alone were €1bn in 2022 and could rise to €4bn by 2030. Enough wind energy to power 1m homes was “wasted” in 2023, owing to insufficient grid capacity connections between Scotland and England (per Carbon Tracker).



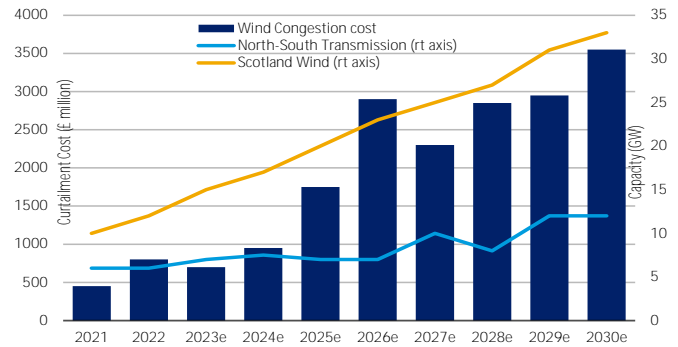
Exhibit 199: Congestion management costs in Germany (EURm)  
Getting power from source to consumption is increasingly expensive



Source: BNEF

BofA GLOBAL RESEARCH

Exhibit 200: Lack of grid infrastructure could increase UK costs 4x  
As wind deployment outpaces grid expansion, UK curtailment costs could increase from ~£0.9bn in 2022 to £3.5bn pa in 2030



Source: Carbon Tracker

BofA GLOBAL RESEARCH

- The grid is really old: The average age of cable in Europe was 45-50 years in 2020 and 35-40 years in the US (per Nexans). Close to 30% of Europe’s power lines are >40 years old. Much of the infrastructure needs to be replaced or upgraded to more efficient equipment with digital integration.
- Scarcity: Critical minerals volume is set to rise significantly for the energy transition. However, despite concerns over the impact from EV batteries, metal volumes for the grid would be far higher. Copper required for electricity networks could double in 2020-40, per the IEA.

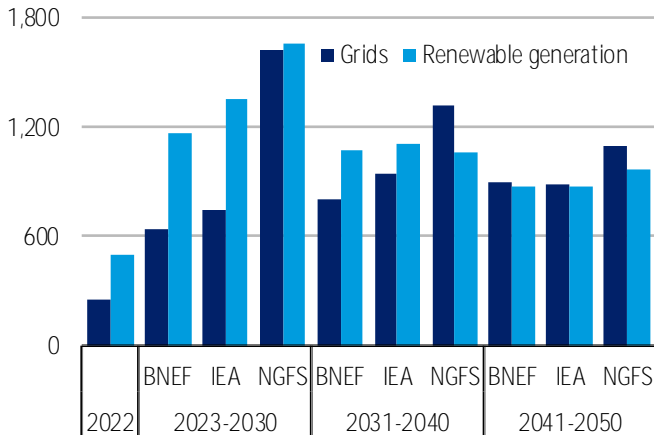
The good news: Investment & tech can solve this

The infrastructural challenge of meeting rising electrification is solvable with a combination of investment and technologies to adapt to the changing supply/demand.

- Annual grid investment to treble by 2030, quadruple by 2050: A range of external projections from BNEF, IEA and NGFS suggest the \$250bn spent annually on the grid in 2022 could more than treble between 2023 and 2030, and quadruple on average between 2041 and 2050 reaching ~\$1tn annually. This would exceed the annual investment into power generation. Cumulatively, between 2022 and 2050, \$21.4tn investment in the grid is required to reach net zero emissions per BNEF.
- Wired for growth: cable length to double and cables to become more efficient: To electrify the economy at this pace and scale requires a supersized grid. Total grid cable length is projected to double from 82m km in 2022 to between 132-174m in 2050 per BNEF – if laid in a single line that would be long enough to reach the sun. This will consist of scaling up overhead and underground cables, plus submarine cables (projected to rise 15x over that period) and interconnection lines (set to treble) to facilitate improved transportation of electricity. Furthermore, the next generation of superconducting materials could reduce resistance and losses in transmission dramatically, increasing energy throughput.



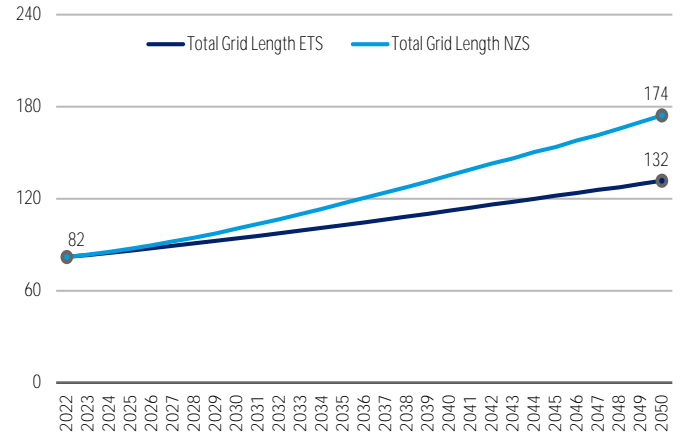
**Exhibit 201: Average annual grid spending needs to treble by 2041-50**  
 Whichever scenario, significant growth in grid capex is required, increasing from half that of power generation (2022) to exceed it by 2040-50



Source: BNEF (\$USD BN)

BofA GLOBAL RESEARCH

**Exhibit 202: Total grid length could more than double by 2050**  
 Total grid length was 82m km in 2022; 132-174m could be required by 2050 depending on the pace of transition to electrification



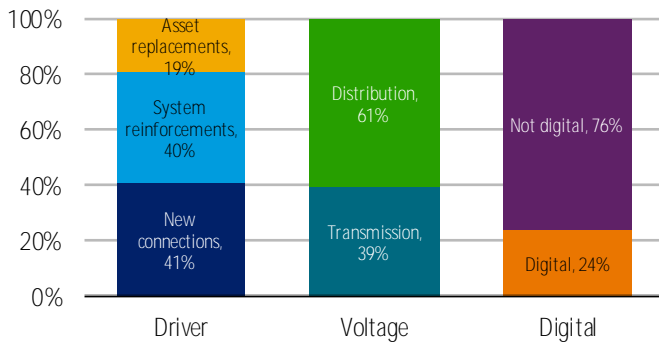
Source: BNEF, ETS = Economic Transition Scenario, NZS = Net Zero scenario

BofA GLOBAL RESEARCH

**Grids get digital**

The changing energy generation/consumption patterns will require better monitoring, optimisation and control, all of which can be enabled through digitalisation and automation. A third of grid investment to 2050 is projected to be in digital technologies per BNEF, shifting from the largely monitoring spend today to technologies required to enable more distributed energy resources. Automation & control would make up the majority of these investments (38% by 2050 per BNEF).

**Exhibit 203: Global grid investment by type in the NZS (2022-50)**  
 >60% the grid investments over 2022-50 will be directed to Power Distribution in BNEF Net Zero Scenario

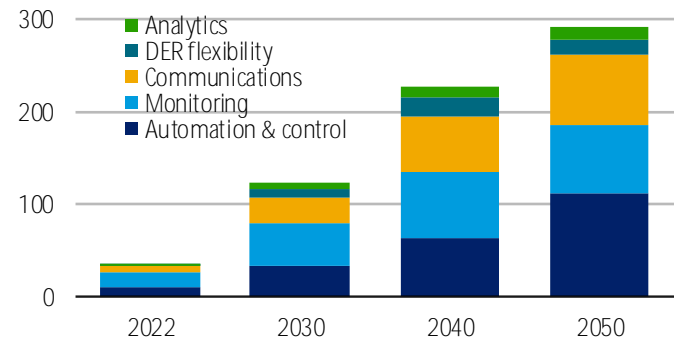


Source: BNEF New Energy Outlook Grids (2023)

Note: Data includes China ETS: Economic Transition Scenario; NZS: Net Zero Scenario

BofA GLOBAL RESEARCH

**Exhibit 204: Annual capex on grid digitalization (\$bn)**  
 38% of the digitalisation costs could be directed to automation & control of the power system in 2050



Source: BNEF New Energy Outlook Grids (2023) DER = Distributed Energy Resources

BofA GLOBAL RESEARCH

The shift to a digital grid presents several opportunities:

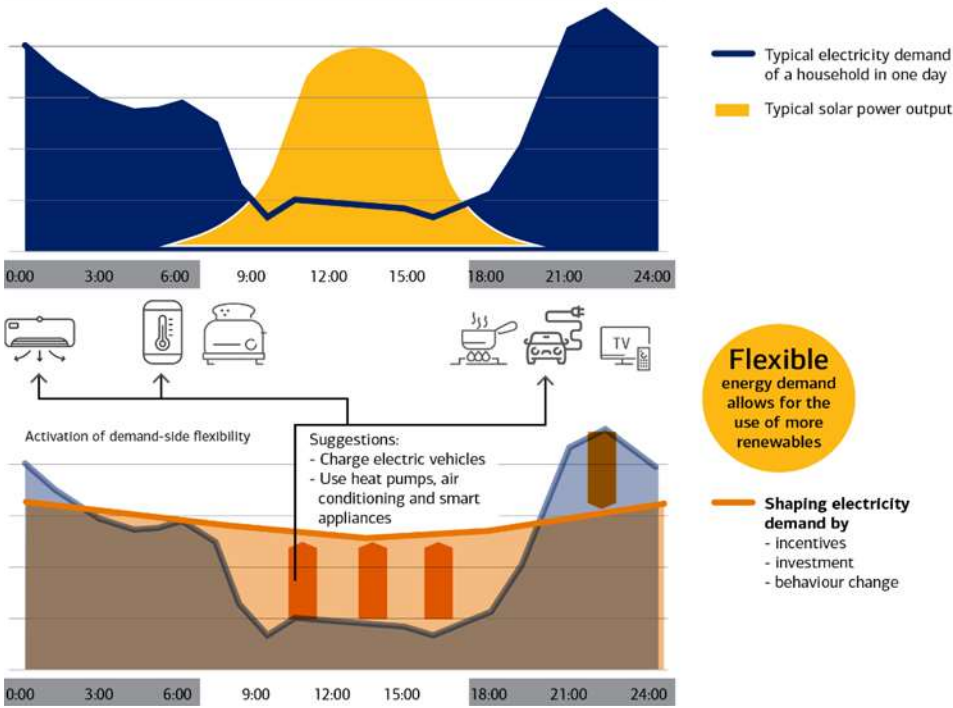
- **Lights stay on:** Digital solutions could help reduce risk of power outage: default detection, cybersecurity, response to adverse weather:
- **Reduction in congestion/saturation:** A digital grid improves stability, resilience and data management, incentivises more or less energy use to manage peak demand and reduces costs both to the operator and customer.
- **Peak/load shifting:** technologies, incentives, behaviour change or enforced restrictions can impact typical daily power consumption trends. Adding more direct electrification demands, such as EVs/heat pumps can help if they are programmed or controlled to consume energy at off-peak times. In the case of



EVs and home energy storage systems, they can also provide power back to balance the grid if they have the enabling bi-directional chargers, reducing the overall grid and power-generating capacities required at peak times.

Exhibit 205: Future grids require both supply and demand side flexibility  
 A combination of technology, behaviour change, incentives & enforcement could change energy peaks

Electricity demand and solar power output (one day)



Source: REN21 Renewables Global Status Report 2022

BofA GLOBAL RESEARCH

## Virtual Power Plants

What are they? A VPP aggregates or controls a network of power generating or consuming assets to be used collectively to balance an electricity grid.

Did you know? The US already has up to 60GW aggregated power capacity available through VPPs. By 2030, they could serve 10-20% of the peak electricity load, saving \$10bn per year.

Electric + digital: the sharing economy for energy is in sight

The changing supply towards more decentralised power generation sources and rising electrified end demand creates both a problem and opportunity for electricity grids. Technology offers a perfect opportunity to combine resources and manage the electricity network more effectively. Virtual Power Plants (VPPs, also referred to as Distributed Energy Resources (DERs)) offer huge potential to achieve this, a \$21bn annual investment opportunity by 2040 per BNEF in the technology required to enable them. They also bring the potential to save billions more in avoiding curtailment, congestion, and deferring transmission upgrades by managing peak demand and increasing utilisation of the existing grid assets.

VPP – aggregating resources of homes & businesses to support the grid  
A VPP is an aggregated network of decentralised power generating units, comprising thousands of households and businesses that offer the potential of their collective resources to support the grid (e.g. EVs, thermostats, heat pumps, solar production or battery storage). These devices can be flexibly charged, discharged, or remotely managed/scheduled to meet grid needs. When aggregated they can offer similar services to traditional power plants.

VPPs offer real cost and infrastructure savings

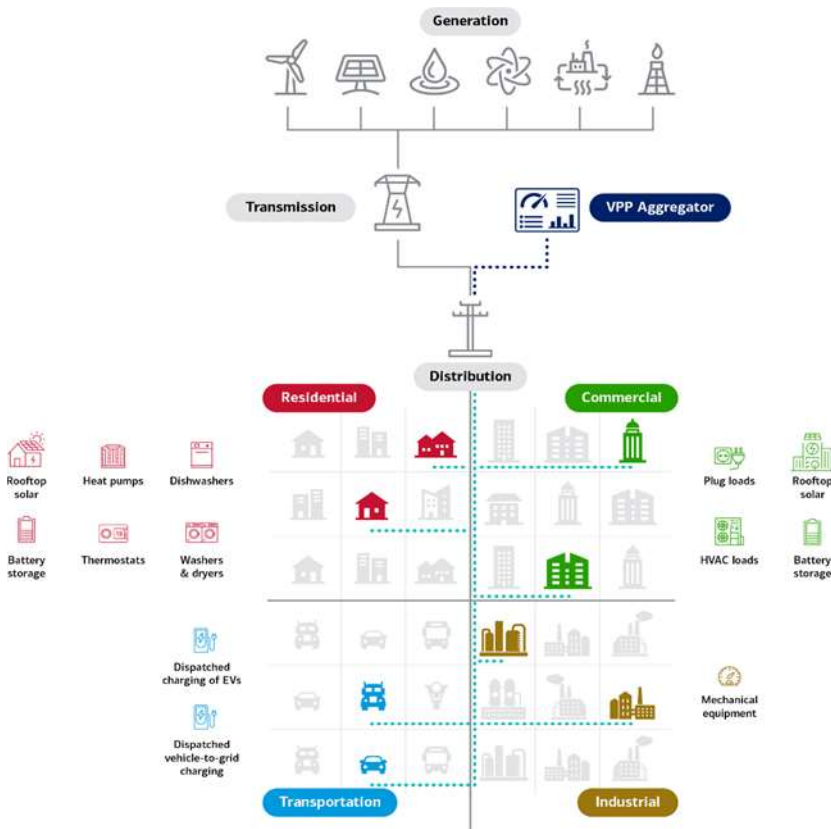
Per the US DOE, there are several business models emerging for the use of VPPs including a combination of: 1) access to residential customers' solar and battery storage; 2) access to battery storage only; 3) managed EV charging (scheduling, shifting, or directly accessing EV battery charge via vehicle to grid programmes); 4) flexible industrial demand; and 5) smart thermostats. Despite the relatively early shift to electrification, VPPs already have 30-60GW aggregated capacity in the US, per the US DOE, evenly split between utilities and wholesale markets. A trebling of VPP capacity could serve 10-20% of peak load nationally, saving \$10bn per year in the US and \$17bn globally per RMI. How? Whilst 200GW of new peak demand is expected by 2030, hundreds more GWh of battery capacity could be made available from EVs alone, if they can be accessed. Such is the potential, the US Federal Energy Regulatory Commission (FERC) has set out regulation to allow such aggregation, setting requirements to allow VPPs to earn the same compensation as other types of resource like power plants.

There were 93 VPP companies in 2022 per BNEF, specialising in various parts of the value chain across equipment, software or market trading.





Exhibit 206: Virtual Power Plants (VPPs): aggregating distributed, grid-interactive electric devices  
 Proliferation of connected devices and technologies can be aggregated by VPPs to balance power supply/demand, e.g. by accessing distributed power sources or incentivising reduced customer demand



Source: RMI

BofA GLOBAL RESEARCH

**‘She’s electric, can I be electric too?’ A consumer electric revolution is underway**

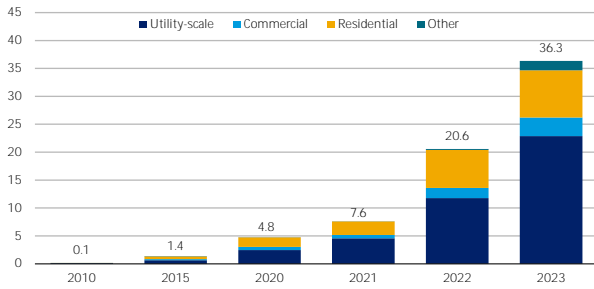
Key to the adoption of VPPs is interest and behaviour change from customers. Either by shifting times of their electricity consumption or sharing energy production and storage resources with the grid, and by a combination of price incentives paid to customers to make those resources available and lower energy prices. Already VPPs are paying energy customers in mostly pilots, such as Tesla in the US (paying customers \$2/kWh of battery storage), and Octopus in the UK (~\$1k per year savings on tariffs in a vehicle to grid trial, making EV batteries accessible to the grid). While price may be competed away as more customers make power resources available, the considerable infrastructure savings made in shaving the peak grid requirements could reduce average energy prices longer term.

**EV batteries alone could balance the grid as soon as 2030**

As the volume of EVs increases alongside the trend of increasing battery sizes, the volumes of battery capacity deployed cumulatively are set to grow exponentially from ~0.4TWh EV battery capacity deployed in EVs globally in 2020 to over 41TWh in 2040 (BNEF). Even accessing a small percentage of that energy could be useful in virtual power plants and grid balancing. Indeed, per NREL and Leiden University, *EV batteries alone could be used to satisfy short-term grid storage demand across most regions by 2030* with as little as 12% of participation rates in vehicle-to-grid programmes. Total capacity made available from EVs could be 32-62TWh by 2050 in their estimates, much higher than the 3-19TWh required for short-term grid storage capacity.

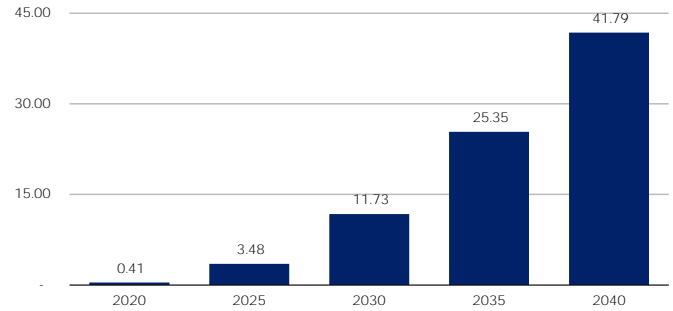


Exhibit 207: Energy storage investments are accelerating  
Investment in energy storage has increased from ~\$5bn in 2020 to >\$36bn in 2023, 1/3 of which is residential or commercial resources



Source: BNEF, Global new investment in energy storage by sector (\$bn)  
BofA GLOBAL RESEARCH

Exhibit 208: Cumulative EV battery capacity: >11TWh by 2030  
Increasing EV volumes could become a valuable energy storage resource to balance the grid, with expected capacity >11TWh by 203, >41 TWH 2040

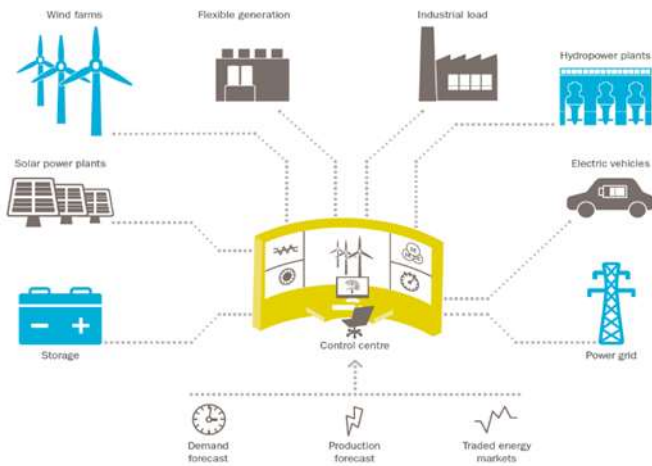


Source: BNEF  
BofA GLOBAL RESEARCH

### The role of AI in grid management

Shifting power supply (to more renewables and decentralised sources) and new electricity demands (EVs, heat pumps, industry) make more accurate demand, production and price forecasts more valuable to balance the grid. AI could play a critical role in enabling this, through real-time modelling and automation in particular.

Exhibit 209: Virtual Power Plants need data, forecasting and AI to enable them  
Shifting power supply (to more renewables and decentralised sources) and new electricity demands (EVs, heat pumps, industry) make more accurate demand/production/price forecasts more valuable to balance the grid – AI could play a critical role in enabling this, through real-time modelling and automation in particular.



Source: Statkraft  
BofA GLOBAL RESEARCH

Per MIT, AI is already being used or considered in grid operations for faster and better decision-making, by helping grid operators understand conditions and improve predictions. Every day, grid system operators run complex calculations to predict electricity requirements for the following day to recommend the most cost-effective way to supply it. In 2019, Argonne National Laboratory partnered with US grid operator MISO to test a machine learning model to optimise this daily planning. The model showed *this can be done 12 times faster than without AI*, reducing time required to 60 seconds. Given these calculations are done several times a day, time savings could be significant. The model is also being adapted to forecast power outages, by incorporating other factors such as weather. Other practical applications of AI already in place between grid operators and companies include aggregating smart meter/EV charging/IOT data for more accurate home energy predictions, and creating automated responses to vary power supply, and dynamic pricing incentives.



## The next generation of solar technology

What are they? Next-gen materials with strong light absorption and electric charge properties to increase the conversion efficiency of solar cells.

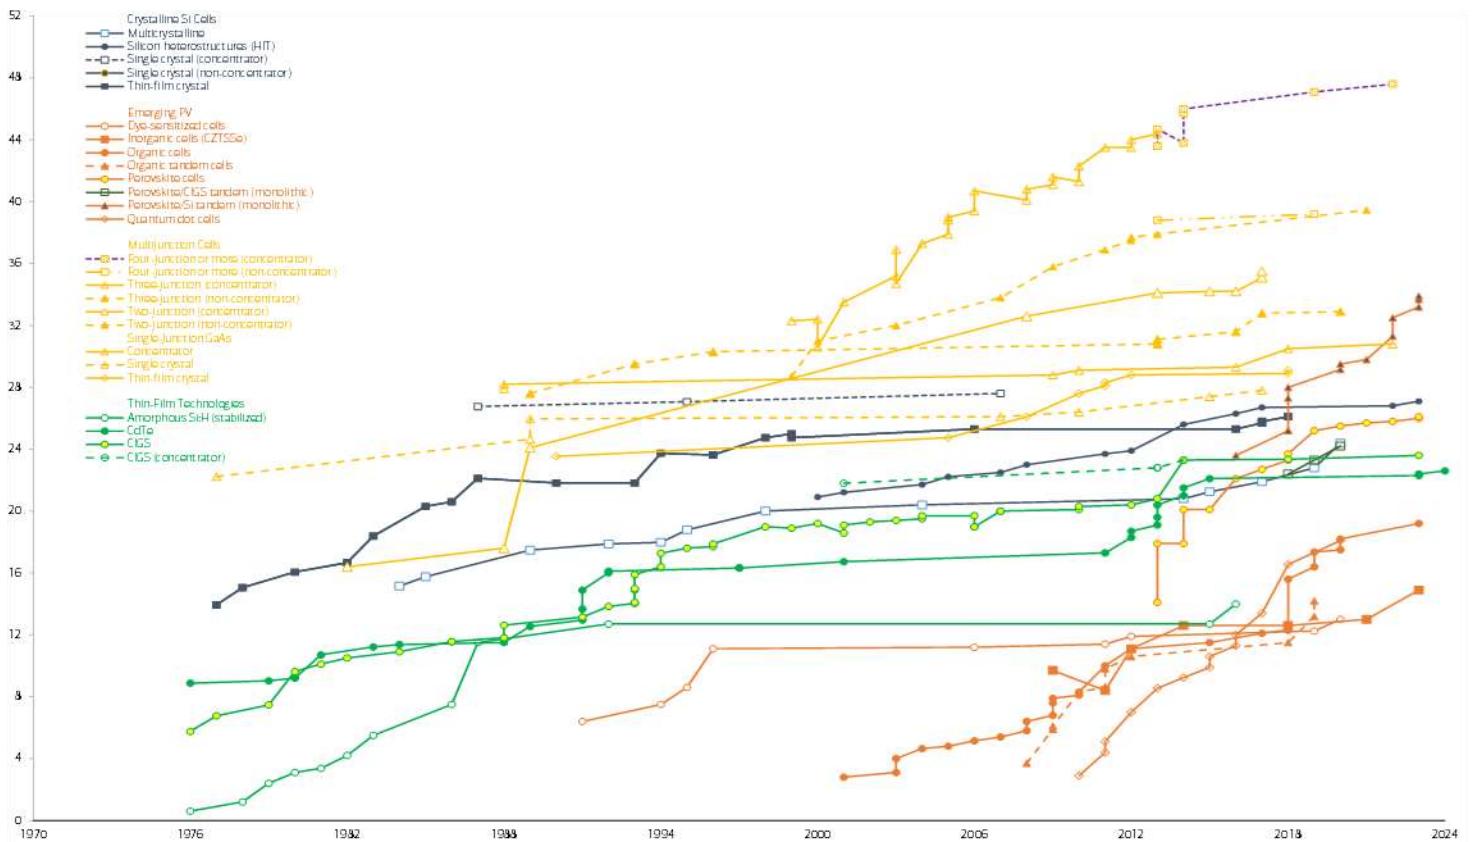
Did you know? Perovskite solar cells have increased in efficiency from 3% in 2009 to 26% in 2024.

### What's next in solar: efficiency, efficiency, efficiency

While the cost of solar PV (photovoltaic) modules has dropped tenfold in the past decade, the rate of efficiency improvements in the most commonly used silicon wafer technology has flatlined since the 1990s, potentially approaching its theoretical limit. In contrast, emerging PV technologies are showing near linear expansion, particularly in perovskite cells that are improving faster than any other PV material. Efficiencies are already in the low 30% range and showing continued improvement in both lab-stage and small-scale manufacturing.

Exhibit 210: Solar PV Techceleration – some solar cells now approaching 50% efficiency

The highest confirmed conversion efficiencies for research cells continues to increase across a diversifying range of photovoltaic technologies, particularly emerging PV perovskite cells



Source: NREL (as of February 2024)

### Perovskites 101

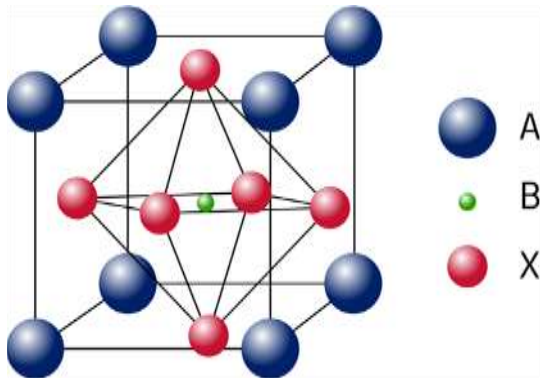
Perovskites are a family of materials with a crystal structure, with strong light absorption and electric charge properties. They can be tuned to respond to different colours in the solar spectrum by changing material composition. Continued research and innovation is showing increasing performance when used in solar panels, thus they are



at the forefront of research seeking to replace or complement silicon. Combining perovskites with other materials, such as silicon or other perovskite materials can deliver more power from the same device (referred to as tandem device architectures). While perovskite on silicon offers the most optimised route to commercialisation today (given silicon’s commercial presence in the solar supply chain and additive stability), combining with other perovskite materials is likely where the industry will trend longer term to increase efficiency further.

The key potential benefits they offer over current silicon based panels would be 1) already higher observed efficiency of converting solar rays into usable energy and rapid progress being made, 2) lighter materials, 3) ability to deposit onto most surfaces, expanding the potential to incorporate solar panels into e.g. windows, buildings and other surfaces (not just rooftops).

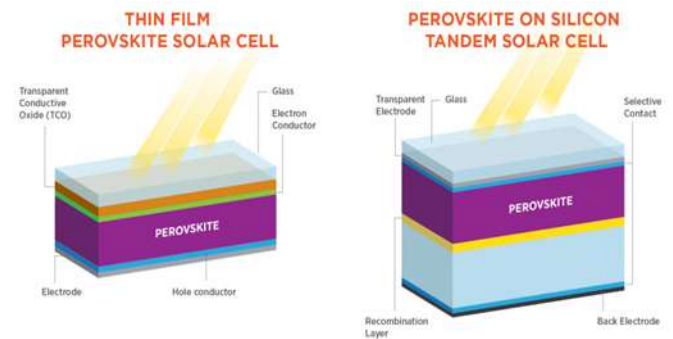
Exhibit 211: Illustration of Perovskite  
Perovskite is a calcium titanium oxide mineral



Source: BofA Global Research

BofA GLOBAL RESEARCH

Exhibit 212: Perovskite solar cells: thin film or tandem configurations  
A variety of Perovskite formulations can be used in solar cells as thin film, or combined in tandem with other materials to deliver more power



Source: US Department of Energy

BofA GLOBAL RESEARCH

“In 2004 it took a year to install 1GW of solar PV. By 2010 it took a month, in 2016 it took a week. 2023 saw single days on which 1GW was installed. Wind capacity has doubled 6x from 8GW in 2024 to 110GW in 2023. In 2023 enough wind and solar capacity was installed to deliver an expected 800TWh each year, 2.8% of global power demand. The growth of wind and solar is likely to look exponential for a long time to come.” *Michael Liebreich, via BloombergNEF*

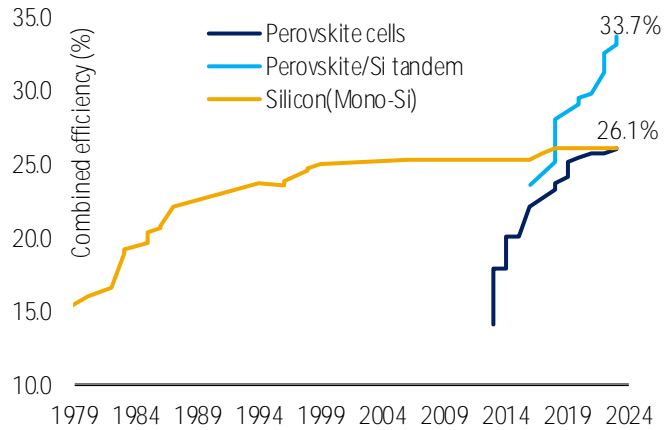
Perovskite breakthroughs: from research lab to near commercialisation  
Scientists and corporations continue to demonstrate this potential for solar PV cells, with efficiencies approaching 34% when deployed in tandem device architectures. Theoretically, the highest conversion efficiency of a perovskite cell can be 45%. Versatility sets the perovskite cells apart because they are light and flexible, and hence can be installed on walls or curved surfaces.

There are 22 companies attempting to commercialise perovskite solar cells (per Tracxn). While the majority of these developments have been lab-based thus far, some commercial breakthroughs are emerging. In January 2024, Fraunhofer & Oxford PV claimed a world record for the efficiency of a silicon perovskite tandem solar module in industrial format at 25% (claiming a 43% efficiency at the cell level). The researchers used manufacturing equipment already used in mass production, and aim to begin commercial production in 2024.



### Exhibit 213: Perovskite solar research breakthroughs are accelerating

Researchers & commercial developers of perovskite solar PV are announcing rapid efficiency improvements in single layer and tandem concepts



Source: NREL.

BofA GLOBAL RESEARCH

### Exhibit 214: Fraunhofer/Oxford PV perovskite tandem breakthrough

In January 2024 a world record for the efficiency of a commercial sized perovskite PV module was set – 43% at the cell and 25% at module level



Source: Fraunhofer / Oxford PV

BofA GLOBAL RESEARCH

### Challenges to commercialisation of perovskite

Whilst they've become highly efficient in a relatively short time, several challenges remain to commercialise the technology per the US DOE:

1. **Manufacturing scale:** Whilst lab tests of perovskite solar cells are showing rapidly increasing energy conversion efficiency rates, this is yet to be proven at commercial manufacturing scale. Fabrication efforts to mitigate this are ongoing but at early stage.
2. **Stability and durability:** Perovskites can decompose as they react to moisture and oxygen or when they are exposed to light/heat/applied voltage for an extended time. Manufacturers are looking at alternative materials and encapsulation methods to mitigate this and increase the lifespan.
3. **Materials:** Whilst perovskites could use significantly fewer materials than silicon cells, some formations include lead as the inorganic compound – a potential toxicity risk, given the solubility of the commonly used perovskite materials in solar.

## Democratising wind energy

What is it? Increasing power generation capacity per turbine and the geographical area wind energy can be deployed economically.

Did you know? One rotation of the highest-rated capacity wind turbine generates enough electricity to charge an electric vehicle.

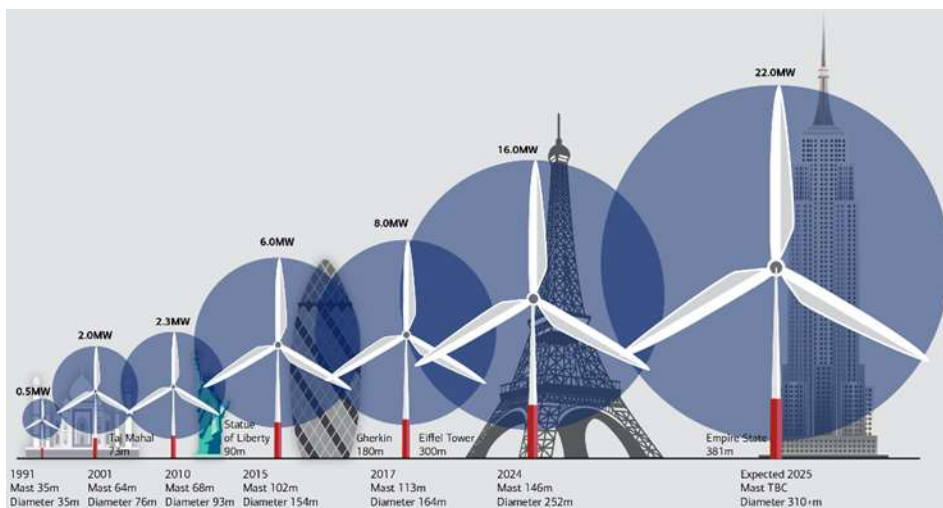
### Any way the wind blows...size does really matter

As the race to decarbonise accelerates, so has the race to deploy the biggest wind turbines. Since the first commercial wind farms of the 1990s, turbine capacities are more than 10x higher in offshore wind. The largest deployed commercially in 2024 are rated at 16 megawatts. To put this in perspective, the rotor diameters of these turbines now near the height of the Eiffel Tower. China's CTG-operated Fujian offshore wind farm claimed to have installed the world's first such turbine in 2023. One turbine can meet the annual electricity needs of 36,000 households, generating more than 66GWh.

In 2023, China's Mingyang announced the development of 18MW and 22MW turbines reaching 310m in diameter with the first deployment expected in 2025. This upscaling is important as it can achieve:

1. More power per unit: whilst output varies per manufacture and region/conditions, the increasing size of turbines can generate more energy per unit. Mingyang claims annual output can increase to 74GWh and 80GWh using 18MW and 22MW turbines, respectively (vs. 67GWh using its 16MW turbine) One rotation of the (18MW) blades would create almost 45kWh, enough to fully charge a small EV battery (vs 34kWh with the 16MW turbines).
2. Reduce cost/units: alternatively fewer turbines can be used to generate the same capacity. China's CSSC estimates that using its 18MW turbines would require 13% fewer units than deploying 16MW turbines, for example.

Exhibit 215: The race to bigger wind turbines: +30x capacity per turbine since 1991  
Offshore wind turbines continue to increase in size and power per unit: capacities increased from 0.5MW in 1991 to the largest now rated 16MW in 2024, with their blade diameters now exceeding skyscrapers



Source: BofA Global Research, Visual Capitalist, FT, Bloomberg, Ineltec

BofA GLOBAL RESEARCH





Right-sizing supply chains to keep up may require \$100bn

There are challenges to the rapid iteration and deployment of larger turbines that is hindering efforts to establish supply chains to scale up the industry, rendering some elements of the supply chain obsolete before they're fully utilised. Each new turbine platform requires investment in new tools, moulds, trucks/cranes, and shipping vessels so large that don't fit into existing harbours, per European turbine manufacturers Vestas and ZF. Some manufacturers and industry groups in the US/EU propose a slowdown, standardisation, or even a temporary (regulated) cap on turbine sizes to mitigate this. The Dutch wind industry association (Sept 2023) proposed such a cap on turbine tip size for example to 1,000ft (~305m) for at least a decade to allow for such standardisation

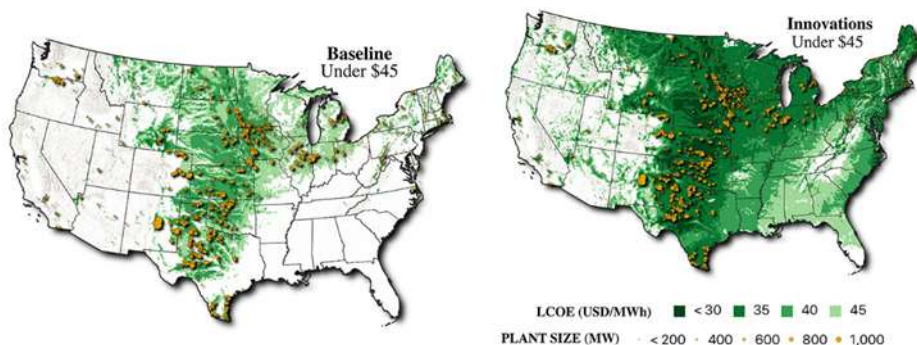
Per Wood Mackenzie, this could require \$27bn of investment to upgrade supply chains to meet a 5x growth in annual offshore wind installations of 30GW capacity additions by 2030 (their base case assumption), but as much as \$100bn if cumulative government targets of 80GW per year additions are to be achieved globally, mostly in installation (\$48bn) but also the foundation, blades, towers and nacelles.

Democratising wind energy: tech could unlock 80% more potential this decade  
 Whilst the increasing size of turbines and rated capacity is important for areas with high wind speeds, several other innovations could unlock greater potential to achieve cost effective wind power in a far wider geographic area. These innovations include:

- Low-specific power turbines: larger rotor size relative to generator size and at lower hub heights, this enables them to start generating electricity and capture more energy at lower wind speeds
- Advanced manufacturing: e.g. spiral welding and 3D printing enable on site creation to reduce cost/transportation constraints
- Wake steering: controls to tilt turbine direction and change generator speed to redirect turbines, can achieve 1-2% annual energy production gains

A recent (2023) study by NREL identified that applying these technologies in addition to longer blades, taller towers and improved cranes in the US could unlock an additional 80% economically viable wind energy capacity as soon as 2025. In addition to reducing power emissions, this could reduce the need for utilities to import as much energy by producing more locally, enabling a deferral of grid transmission upgrades for example.

Exhibit 216: Technology advancements could expand US wind energy potential 80% by 2025  
 Innovations such as low-specific power ratings, advanced manufacturing, taller towers, longer blades and wake steering could unlock an additional 80% economically viable wind energy capacity as soon as 2025



Source: NREL

BofA GLOBAL RESEARCH



# Graphene

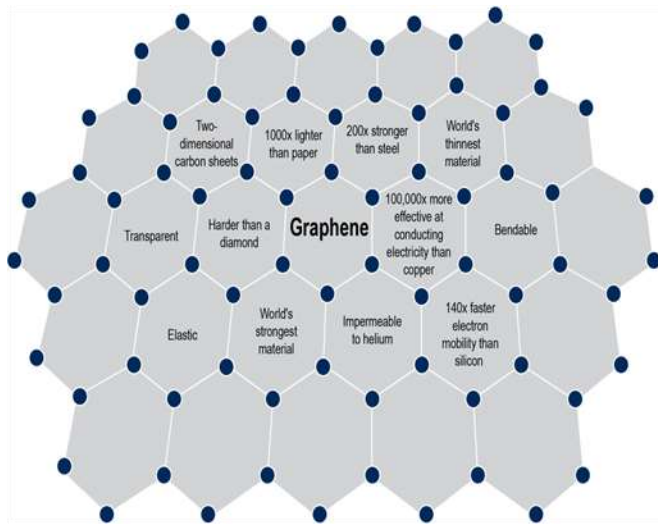
What is it? The thinnest and strongest material known to humans.

Did you know? Graphene was first isolated in 2004 using scotch tape. Graphene-coated materials can increase energy density of supercapacitors by 10x, and the efficiency of water filtration by 9x.

## Graphene: a record-breaking material

Graphene is the thinnest material in the world (only an atom wide), the strongest material (an elephant would need to stand on a pencil to break graphene that is one atom thick), and also a superconductor (capable of sustaining current densities of six orders of magnitude more than copper) (source: Columbia University, 2008, Chemical Society Reviews, 2015). This two-dimensional material made up of carbon molecules arranged in hexagonal shapes has the potential for multiple applications – from cheaper semiconductors and fast-charging batteries to lightweight cars, bulletproof vests, environmentally friendly concrete and medical applications.

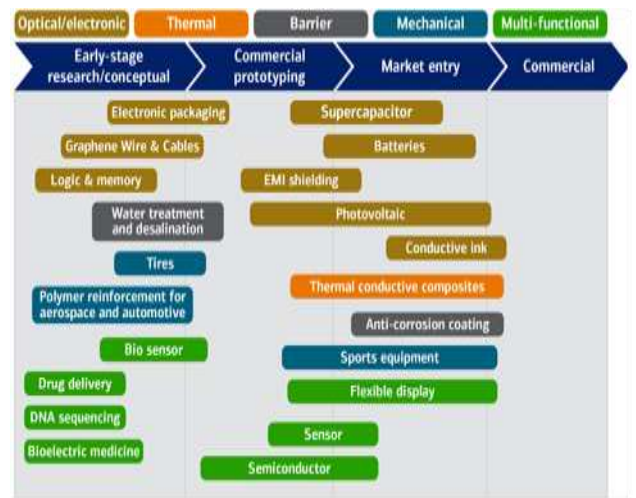
Exhibit 217: Properties of graphene – **the ‘wonder material’**  
Graphene is the world’s strongest material



Source: BofA Global Research

BofA GLOBAL RESEARCH

Exhibit 218: Several graphene applications beginning market entry  
Several Graphene applications at market entry stage nearing commercialisation



Source: IDTechEx Research

BofA GLOBAL RESEARCH

## So **it's strong, superconductive, thin ... but what can it be used for?**

Graphene and graphene derivatives are already used in functional ink, polymer additives, R&D, tyres, coatings and composites. The composites segment accounted for the largest share of the market in 2021 and they are used in industries such as manufacturing, aerospace, and autos (source: Spherical Insights). However, scientists have been producing proof-of-concept products, discovering various graphene applications despite difficulties in sourcing it, some of which are at or nearing commercialisation.

## EV batteries: improving several performance attributes with graphene

Graphene batteries could have higher energy densities, increased cycle life and faster charging. Several companies are already undertaking R&D towards this, for example in 2017, Samsung Advanced Institute of Technology researchers claimed to have developed a 'graphene ball' material that enabled a 45% increase in battery capacity and the ability to charge batteries 5x faster than lithium-ion batteries, taking only 12 minutes to fully charge. Graphene batteries could emerge in the early to mid-2030s as the price of graphene falls (per Focus, based on global patent data).



In the meantime, blending graphene materials into existing battery chemistries could improve performance across several attributes, and is already beginning to be commercialised:

- Electric two-wheeler company Yadea was the first to launch a commercial graphene battery in 2019 following a breakthrough deploying a graphene “super conducting paste” that enables the strong conductivity and faster heat transfer than their previous batteries. Their 2022 upgraded graphene battery has 30% higher energy density than their previous lead acid battery but with triple the cycle life durability (1,000 cycles), and 6x faster charging capability (achieving 80% charge under one hour).
- Nanoxplore is constructing a 2GWh-capacity factory in Canada (planned opening 2026) to scale up production of graphene powder, that can be blended alongside graphite/silicon in anodes to increase energy density and cycle life.

Exhibit 219: Advantages of graphene enhanced EV battery cells  
Adding graphene to battery anodes could offer several benefits, particularly around energy density, cycle life and weight

18650 Battery Cell KPIs	Regular Anode with NMC Cathode	Graphene Enhanced Anode with NMC Cathode	
Energy Density Increase	X	Up to 10%	Improved range performance for EVs / e-bikes
Cycle Life at Normal Charge Rate (80% capacity retention)	Up to 1500	Up to 2000	Energy stability and extended life for battery packs used at lower temperatures
Cycle Life at High Charge Rates (80% capacity retention)	Up to 300	Up to 450	Improved charging time without compromising battery performance
Internal Resistance (mΩ)	<30	<28	Fast charging improvement with reduced safety concerns
Maximum Weight of Battery Cell	48g	45g	Lower vehicle weight and improved energy efficiency
Temperature at High Charge Rate	Under 30°C	Under 26°C	Reduced load on BMS and improved cooling properties on pack level
Anode Electrode Expansion Rate (at 100% state of charge)	Up to 80%	Up to 40%	Reduced risk of thermal runaway with graphene / silicon technology

Source: Nanoxplore

BofA GLOBAL RESEARCH

Exhibit 220: Graphene could be commercialised in specialty compounds  
Graphene’s versatility could be deployed across concrete, fluids, foam and plastics to e.g. reduce cost, emissions, and increase durability of assets

	VALUE PROPOSITION	GRAPHENE ADDRESSABLE MARKET SIZE
Concrete	<ul style="list-style-type: none"> <li>CO2 EMISSIONS</li> <li>FAST HEALING</li> <li>IMPROVED DURABILITY</li> </ul>	6,000 ktpa <sup>1</sup>
Drilling Fluids	<ul style="list-style-type: none"> <li>FASTER DRILLING SPEED</li> <li>EXTENDED TOOL LIFE</li> <li>REDUCED FLUID LOSS</li> </ul>	256 ktpa <sup>2</sup>
PU Foam	<ul style="list-style-type: none"> <li>RECYCLED CONTENT</li> <li>FLAME RESISTANCE</li> <li>SUSTAINABILITY</li> </ul>	170 ktpa <sup>3</sup>
PE/PP Compound	<ul style="list-style-type: none"> <li>BIODEGRADABLE PLASTIC</li> <li>MECHANICAL PERFORMANCE</li> <li>FAST HEALING</li> </ul>	90 ktpa <sup>4</sup>

Source: Nanoxplore

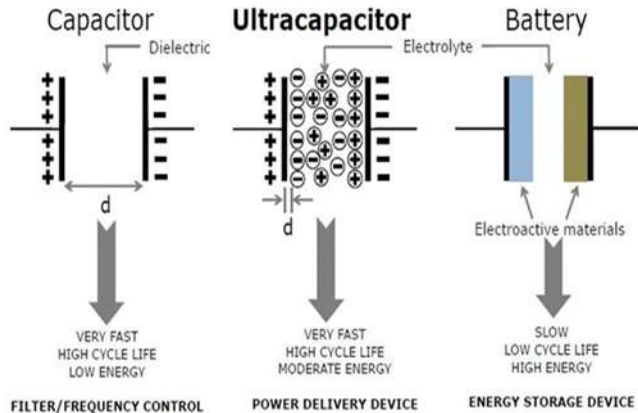
BofA GLOBAL RESEARCH

Graphene supercapacitor case study: future of hybrid energy systems?

Like batteries, supercapacitors store electricity but do so statically rather than chemically. While less energy-dense than batteries, they can charge/discharge far more quickly and last much longer, making them ideal for small bursts of power. Recent innovations in using graphene coated materials have demonstrated 10x improvement in their energy storage capacity (University College London & Chinese Academy of Sciences), claiming the highest reported energy density to date at 88Wh/L. Whilst this is still only ~10% of the energy storage capability of advanced lithium ion batteries, these breakthroughs could enable supercapacitors to be used in parallel to batteries, as 1) it could triple the life of batteries (given the supercapacitor could take care of acceleration and energy recovery, the stressful part of a battery’s life) 2) longer cycle life potential (up to 1m in supercapacitors vs 3,000-5,000 cycles in batteries per NaWa Technologies), 3) reduction in battery size/weight (by up to 1/3 without loss of range), and 4) the lack of rare minerals.



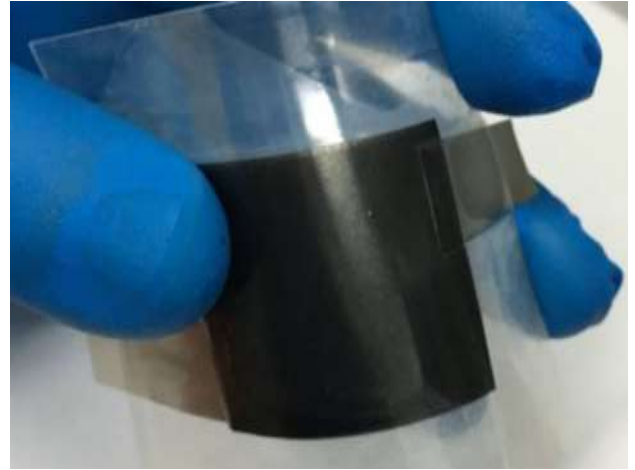
Exhibit 221: Differences between capacitors, supercapacitors, batteries  
While supercapacitors can store less energy than batteries, their ultra-fast charge/discharge could work in parallel with existing battery systems to elongate their usable life and reduce degradation



Source: Engineering.com

BofA GLOBAL RESEARCH

Exhibit 222: Graphene supercapacitors for increased energy storage  
Supercapacitor research has shown that using graphene-coated materials can increase energy density 10x vs current supercapacitors



Source: University College London

BofA GLOBAL RESEARCH

Can graphene supercharge the internet, energy, or industrials?

Graphene's many characteristics imply multiple other potential uses for the material:

- Telecoms – could speed up the transmission of data owing to graphene's high conductivity and ability to operate at lower power levels compared to silicon-based devices. Ericsson showcased graphene-based photonics switches in 2018, for example, where the first all-graphene optical communication link was also shown, and in 2021 the Graphene Flagship programme researchers demonstrated the world's first high-speed, graphene-based data communication, which could transmit data at 50Gb/s (source: Graphene Flagship, Orange).
- Semiconductors – graphene could replace silicon as processing power rates of development slow, improving the data processing, wireless communications, and consumer electronics markets.
- Construction – cement production causes 8% of the world's carbon emissions – 3x more than the aviation industry (source: International Energy Agency). If the world's cement were reinforced with graphene, the concrete savings would be equivalent to stopping all the annual emissions of Brazil (source: Grist).
- Image sensors – graphene can detect UV, visible, and infrared light in one sensor as it can absorb almost all wavelengths of light.
- Wearable tech – could monitor the heart rate, UV exposure, and blood oxygen through a UV patch with graphene inside it.
- Water tech: research is ongoing for the use of graphene in membranes for water filtration or desalination, owing to their high permeability (8-9x more efficient than current water filters per Nanografi) and high adsorption surface area that could enable it to adsorb a wider range of contaminants vs current filters.
- Solar cells: graphene could be incorporated into solar cells to enhance performance, efficiency and durability in a number of areas, e.g., the electrodes in existing thin-film cells, or integrated to next-gen perovskite solar cells to improve their stability (per graphene-info).



Challenges to commercialising graphene: separation, standards, safety  
Despite the vast potential, to commercialise graphene requires overcoming issues hindering more widespread development:

- Separation: producing high-quality graphene at scale is hard and not yet proven. The material was initially separated in 2004 by applying scotch tape to graphite repeatedly until left with a single layer of atoms. However, this “mechanical separation” is not replicable at scale. Much research is being done to find an alternative to current methods, which produce highly contaminated or oxidised graphene that erodes its characteristics. Several scientists claim to have found methods for large-scale graphene production, such as MIT using roll-to-roll chemical vapour deposition. Others include liquid phase or chemical exfoliation, and graphite intercalation and expansion, each with pros/cons.
- Lack of standardised material and processes, meaning the quality of graphene produced is varied. Without access to good-quality graphene, it is impossible to develop applications for it. For example, researchers from the National University of Singapore studied the graphene material from 60 suppliers and found that no company had more than 50% graphene content, while over 50% of them had less than 10% and instead contained graphite powder (source: Kauling et al, Advanced Materials, Nature, 2018). The National Physical Laboratory produced the first ISO graphene standard in 2017 aimed at mitigating this.
- Safety testing trials on the impact of graphene exposure to humans is in the early stages with the first controlled clinical trial only occurring in February 2024, by the Universities of Edinburgh and Manchester. This showed that inhaling the substance did not have any short-term adverse effect on the lungs or cardiovascular function. The trial used ultra-pure graphene oxide and involved 14 volunteers breathing in the material through a face mask for two hours while cycling. Further research would still need to be done ahead of widespread commercialisation.



## Superconductors: the grid “holy grail?”

What are they? Materials that carry electrical currents with zero resistance

Did you know? Superconducting wires are able to carry 500x more electricity than copper wires. One 17cm diameter cable can transport 3.2GW of power, enough to power a large city.

Zero resistance means there are energy-saving properties

Superconductors carry electrical currents with zero resistance. They become superconductive at the transition temperature. This saves a lot of energy being lost compared to traditional materials like copper, which loses part of the charge in the form of heat. The transition temperature varies depending on the material but the temperature at which this happens is typically below 20 Kelvin (-253 degrees Celsius). The material expels magnetic fields as they transition to the superconductive state.

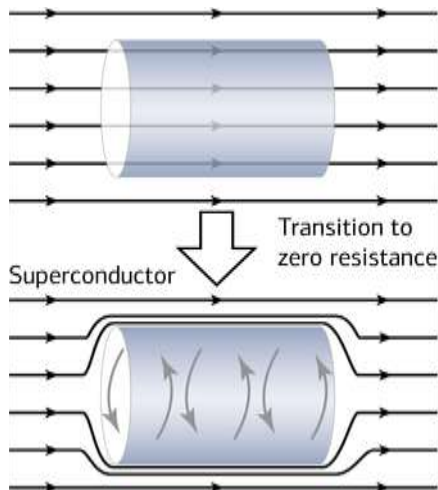
How is this possible? Some electrons are not bound to atoms but they can freely move throughout the material. This motion leads to an electrical current. In normal conductors, this does not happen because the electrons are scattered due to impurities, for example. But, in superconductors, there is ordering that prevents this scattering of electrons.

### Electromagnetism

Superconductors are diamagnetic, which means that they are highly repelled by and tend to expel a magnetic field. This is known as the Meissner effect. Superconductors refuse to let magnetism penetrate it. When a superconductor is put into a magnetic field, it makes an electric current flow through its surface. The currents create a magnetic field and cancel out the original field trying to get inside the superconductor and this repels the magnetic field outside. This explains how a magnet floats above a superconductor.

#### Exhibit 223: Illustrating the Meissner effect

The Meissner effect occurs when a material transitions from its normal state to a superconducting state and it excludes magnetic fields from its interior creating a magnetic field around the material



Source: Hyperphysics

BofA GLOBAL RESEARCH

### Superconductors vs conductors?

Superconductors only achieve superconductivity only when the temperature is brought down below the critical level. Superconductors have zero electrical resistance, while conductors have some finite resistance. Most common electrical conductors are made from metals, e.g., copper, silver and gold etc. Examples of superconductors (and their





transition temperatures) include mercury (4.1K or -269°C), niobium (9.2K or -264°C) and magnesium diboride (39K or -234°C).

Superconductors first discovered in 1911

Superconductors were discovered in 1911 by the Dutch physicist, Heike Kamerlingh Onnes. He found that the electrical resistivity of a mercury wire disappears when it is cooled below 4K (-269°C). He won the Noble Prize in Physics in 1913.

#### Limitations

- Too strong a magnetic field can cause loss of superconductivity: Strong magnetic fields above a certain value, cause a superconductor to revert to its normal/non-superconducting state, even if the material is kept well below the transition temperature.
- Cold temperature & high pressure required: Requires extremely cold temperatures or very high pressure for superconducting properties to occur which could be expensive to maintain. This makes it difficult to deploy more widely in ways that could change daily life. The highest recorded temperature was -23°C (-10°F or 250K), held by a lanthanum superhydride, LaH<sub>10</sub>, and was achieved in 2019. The challenge is to find materials that superconduct at everyday temperatures and pressures. A room-temperature operating superconductor would allow for all kinds of possibilities, such as improving existing technologies, e.g., medical imaging devices to moonshot technologies like nuclear fusion reactors.

The LK-99 hype – what happened?

On 23 July 2023, a team of Korean scientists claimed that copper-doped lead apatite, named 'LK-99', was superconducting up to temperatures of 370K (c.97 degrees Celsius). It made headlines after a research team claimed it was the first room-temperature superconductor, which could have drastic positive implications for our energy systems. Speculation was that it could create the perfect power grid or easily make trains levitate. However, there were vast amounts of scepticism from the academic world, due the lack of clear indicators of superconductivity and inconsistencies in the data.

The results from other academics so far show that LK-99 is not a superconductor, at room temperature or otherwise. With regard to levitation, the researchers shared a video of a fragment of the material appearing to partially levitate over a magnet supposedly due to the Meissner effect. But the reason it was levitating was because the material is ferromagnetic – meaning that it can be magnetised and then attracted or repelled by other magnetic materials – not because of the Meissner effect (source: Guo et al).

Exhibit 224: Image of LK-99  
LK-99 levitating over a magnet



Source: Hyun-Tak Kim

BofA GLOBAL RESEARCH

Relevance to energy? Superconducting cables can remove energy losses and reduce size/cost

Replacing the electrical grid with superconducting transmission lines could carry large currents over long distances with no losses. Superconducting materials can also act as fail-safe devices (known as fault current limiters) to guard a portion of the grid or substation against a sudden electrical outage. They can protect the grid against extreme weather and other catastrophic events. Also, as substation interconnects, they can prevent high fault currents in urban areas by sharing the load.

Superconducting cable benefits: up to 500x more electricity than copper

The reason for their importance in an electricity networks perspective is the combination of reduced energy losses and far higher current possible, thus higher efficiency, reduced size and cost required to deliver the same or higher electricity throughput.

#### Did you know?

Superconducting wires are able to carry 500x more electricity than copper wires. One 17cm diameter cable can transport 3.2GW of power, enough to power a large city.

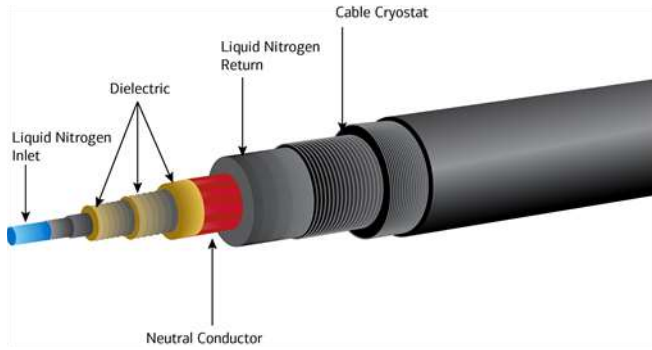
Source: Nexans

- Reduce losses: electricity grids “lose” around 10% of the energy transmitted owing to heat & resistance. Superconducting cables can almost eradicate this, leading to more energy throughput.
- Higher current: more importantly, superconducting cables can carry far higher currents than conventional copper/aluminium cables, making it possible to transport electricity at lower voltages, reducing the need for substations in city centres for example..
- No interference: they don't produce heat and are shielded electromagnetically (the Meissner effect noted above), meaning no interference with other cable networks, avoiding cable routing conflicts and making them appealing to deliver high volumes of power in urban areas in particular.
- Reduced size: cables can be smaller than existing ones, while carrying the same amount or more power, owing to higher current carrying capacity allowing for a smaller cross-section area of the cables, as well as less insulation required due to their operating temperature maintained at cryogenic levels (cooled with liquid nitrogen).
- Reduce cost up to 90%: per the German grid technology association currENT, these combined benefits of using superconducting cables can reduce energy redispatch volumes and costs by 60%, and up to 90% when combined with other technologies such as dynamic line rating (varying thermal capacity of power cables in real time to maximise load and better match energy supply/demand).



### Exhibit 225: Diagram of superconductor cable

To reach its ideal conducting temperature, the superconducting cable is cooled with liquid nitrogen



Source: Nexans

BofA GLOBAL RESEARCH

### Challenges to commercialisation: cooling & durability

Research and development of superconducting cables is ongoing, given these potential benefits, with government-funded projects such as the EU “SCARLET (superconducting cables for sustainable energy transition) project that combines 15 partners to design manufacture and test the liquid nitrogen cooled cables. However, the technology maturity is yet to be proven at grid scale. The key challenges that need to be overcome will be: 1) cryogenic cooling requirements: achieving and maintaining the low temperatures required for superconductivity, leading to 2) reliability and durability uncertainty over time in real-world environments, and 3) cost: ability to manufacture and maintain the materials required at grid scale projects price competitively.

Proving this technology maturity in overcoming those challenges is key to commercialisation to bolster the grid. However, commercial projects have begun in the rail network. Nexans installed two superconducting cables at French railway station Montparnasse in 2022 to secure the growing energy requirements – the first to be integrated into a rail network globally.

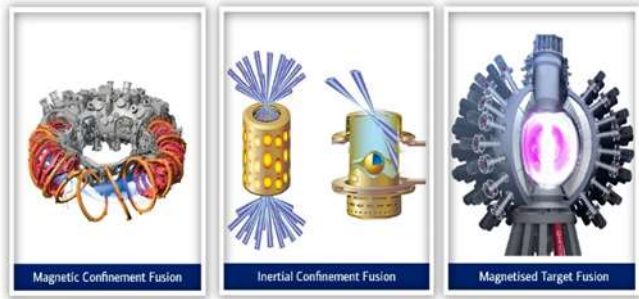
### Superconductors are enabling nuclear fusion breakthroughs

Nuclear fusion is the process joining atoms together in doing so releasing energy; it powers the sun and recreating it on earth for abundant clean cheap energy has been a goal of many since the 1950s. A combination of breakthroughs in materials alongside improving AI, simulation, quantum computing, supportive regulation and funding could change this.

- Superconducting material breakthroughs: Breakthroughs in both the magnetic (superconducting magnets to control a plasma in which fusion can take place) and inertial (using lasers to hit a fuel target producing a burst of energy each time) approaches to fusion are rising in frequency, enabled by the improving technology and research focus. In March 2024 superconducting magnets built by MIT and Commonwealth Fusion Systems confirmed they met requirements for an economic compact fusion plant. owing to the use of a new material named REBCO (rare earth barium copper oxide). Whilst superconducting magnets have been used in fusion tests before, they’ve all required operating temperatures of 4 kelvins (-270 degrees Celsius), but REBCO allows them to operate at 20 kelvins, bringing advantages in the material properties and practical engineering by eliminating the need for insulation, and greatly reducing the size. This breakthrough alone would reduce the cost per watt of a fusion reactor 40 times (source: MIT)
- AI & High-powered compute is improving simulation to iterate/test reactions and performance prior to expensive real-world experiments, and dynamic controls to suppress vibrations that cause energy to escape from the fusion to

maintain it without damaging the reactor. Princeton University developed an AI model (Feb' 2024) that could they claim could solve that – it predicts and figures out to avoid plasma becoming unstable and escaping the magnetic fields that hold it inside torus (donut shaped) reactors, avoiding plasma disruptions and expensive repairs/damage to reactors

Exhibit 226: Breakthroughs in Magnetic & Inertial fusion are rising  
Nuclear fusion methods include magnetic confinement, inertial confinement and magnetized target fusion; material and superconducting breakthroughs are rising, which could increase the chances of commercial success



Source: BofA Global Research, adapted from Engie

BofA GLOBAL RESEARCH

Exhibit 227: Investment in Nuclear Fusion companies is ramping up  
There were 43 companies pursuing nuclear fusion in 2023; 3x more capital was raised between 2021-23 than 2001-20 combined



Source: Fusion Industry Association

BofA GLOBAL RESEARCH

Industrial breakthroughs in healthcare, transport and computing

Beyond energy, superconductors could also unlock several industrial breakthroughs the imaging, computing and transportation sectors.

- Healthcare – magnetic resonance imaging equipment: Superconducting electrons can act as a superconducting quantum interference device (SQUID), a very sensitive detector of magnetic fields 100bn times smaller than those generated by an ordinary magnet. These can be used to map magnetic fields in the human brain created by currents within neurons.
- Ultra-high speed computer chips, memory chips & quantum computing: Current semis memory and logic chips are based on half-conductive (semis) silicon with various litho/deposition/etching materials. Room-temperature superconductors may lead to significant chip design and manufacturing opportunities, or could be used in high-end quantum computing. Superconducting qubits are quantum bits that use superconductors to store, process and transmit data. Quantum processors that use these qubits are developed from superconducting circuits cooled to low temperatures. But performance and working level temperature are not yet proven.
- High-speed ‘Maglev’ trains: These use superconducting magnets to replace typical steel tracks, the reduced friction increasing speeds to hundreds of miles per hour, achieved through the electromagnetic and electrodynamic suspension.



# Future Mobility: Are we nearly there yet?

A Transforming World: The future of mobility is going electric and autonomous, enabled by better batteries and rapidly advancing AI & computational capabilities. The convergence of technologies will also enable the first electric aviation, initially in drones, and soon for passengers too: 7 eVTOL aircraft are vying for certification in 2024/2025.

Did you know? Lithium-ion battery prices fell 99% between 1991 and 2023.

Next Generation of Mobility: Next-gen batteries, Autonomous Vehicle, Future Air Mobility

Converging technologies enabling electric, autonomous, and flying mobility Transportation remains the poster child for an industry that will be transformed with disruptive technologies, owing to the relative inefficiencies of the current transport network (congestion costs the US economy \$95bn in 2022), and huge addressable market. US households spend an average of 15% of disposable income on transportation (US Bureau of Transportation Statistics). Several converging technologies can unlock >\$10tn opportunity by 2035, with commercialisation beginning much sooner.

- Next Gen Batteries: enabled by improving chemistry and material sciences unlocking more use cases within all forms of transportation
- Autonomous Vehicles: Technology that can assist or fully replace human driving, enabled by a combination of sensors and software. The falling cost of compute and accelerating AI capabilities are enabling the first autonomous driving software systems.
- Future Air Mobility: Light aircraft powered by electric propulsion that can take off and land vertically, bringing lower-cost/higher-frequency urban air travel for passenger and cargo transit.



## Next-gen batteries

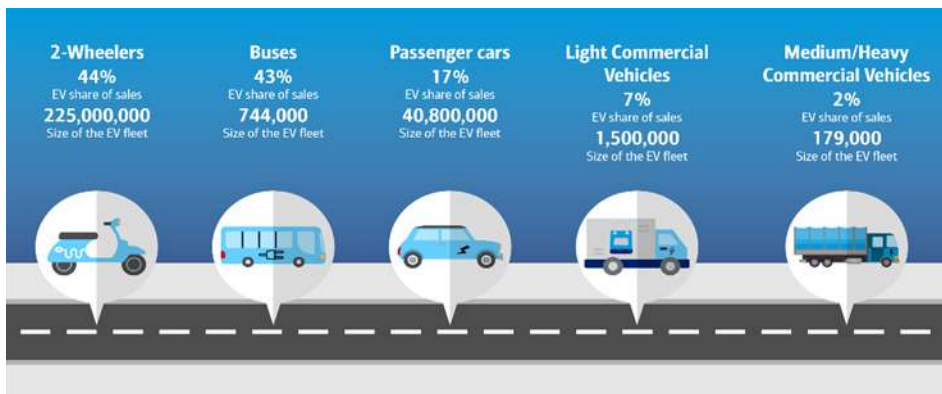
What are they? Batteries with increasing capacity to store energy and recharge faster and cheaper, enabled by improving chemistry and material sciences.

Did you know? Lithium-ion battery prices fell 99% by 1991-23. Battery demand has grown by 24x in the past decade (2013-23) but capacity increased 42x over the same period (RMI). Investment in battery factories is now larger than solar panel and wind turbine factories combined (\$45bn vs \$33bn in 2022).

### To Electric Vehicles and beyond: The battery domino effect

The transition to Electric Vehicles (EVs) is reliant on improving battery technology in particular. Battery energy density continues to increase, costs continue to fall, driving demand growth in several applications. EV adoption has spread to all areas of road transport, from two wheelers to cars to more recently commercial vehicles.

Exhibit 228: Next-gen batteries are at the heart of the EV transition, expanding addressable markets. Improving battery tech is enabling their deployment in larger vehicles. Over 40% of sales of 2 wheelers and buses are electric in 2023, 17% of cars, with vans (7%) and trucks (2%) also beginning to go electric.



Source: BofA Global Research adapted from BNEF data; cars includes battery electric plus plug in hybrid vehicles

BofA GLOBAL RESEARCH

However limitations remain in their viability owing to a number of trade offs - range, safety, durability and charging rates - and fierce competition as the global auto industry races to electrify. We highlight several innovations enabling the next-gen of batteries:

- Chemistry: increasing variants of nickel based batteries (e.g. high-voltage and manganese rich), material switching to more abundant alternatives (e.g. sodium ion), the rise of lithium ion phosphate, and solid state batteries nearing commercialisation.
- Structural & Material breakthroughs: Integrating batteries into structural objects (battery packs, vehicle chassis, or body structures) using new materials and manufacturing techniques
- AI: accelerating battery development with the use of Artificial Intelligence

### Tech & scale drives exponential cost reductions in lithium-ion batteries

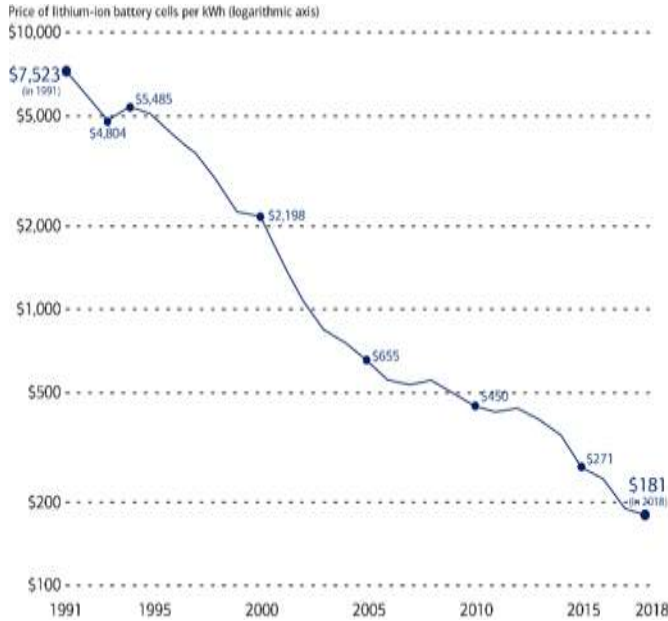
Whilst the first rechargeable lithium-ion batteries were created and manufactured in the 1970s, they were launched commercially for the first time by Sony in 1991, initially in camcorders followed by cell phones. These first iterations cost \$7,523 per kWh in 2018





prices, but prices fell 41x from 1991-2018, as a result of continued innovation and the scale of deployment. Indeed, prices dropped 18.9% for every doubling of installed capacity (the ‘learning curve’) between 1992 and 2016, similar to exponential reductions in the price of solar panel modules, for example (20.2% 1976-2019).

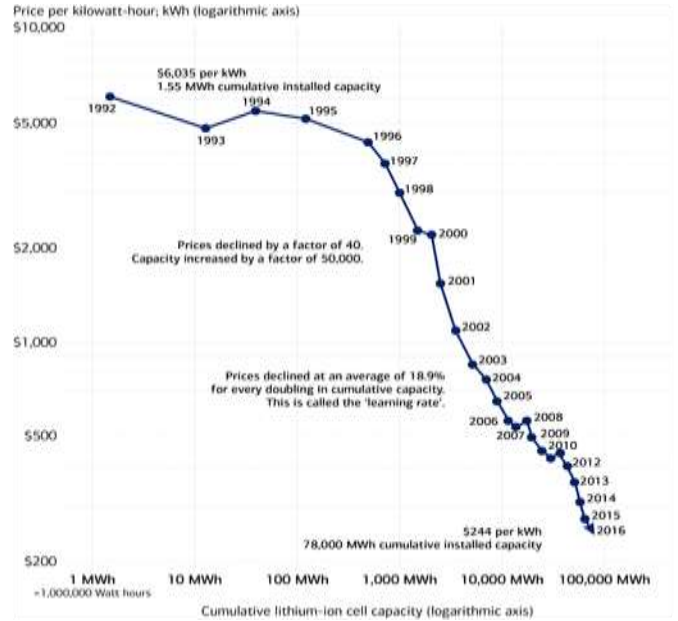
Exhibit 229: Lithium-ion battery cell prices fell 97% from 1991-2018. The price of the first commercially available lithium-ion battery cells per kilowatt hour (kWh) was >\$7,500 in 1991, falling 41x to \$181 by 2018



Source: Our World in Data, Micah Ziegler & Jessica Trancik (2021); prices adjusted for inflation given in 2018 US\$/kWh

BofA GLOBAL RESEARCH

Exhibit 230: Lithium-ion Learning Curve: 18.9% between 1992-2016. Comparing Price per kilowatt-hour (kWh) vs cumulative lithium-ion cell capacity shows prices fell 18.9% for every doubling of installed capacity



Source: Our World in Data, Micah Ziegler & Jessica Trancik (2021); prices adjusted for inflation given in 2018 US\$/kWh

BofA GLOBAL RESEARCH

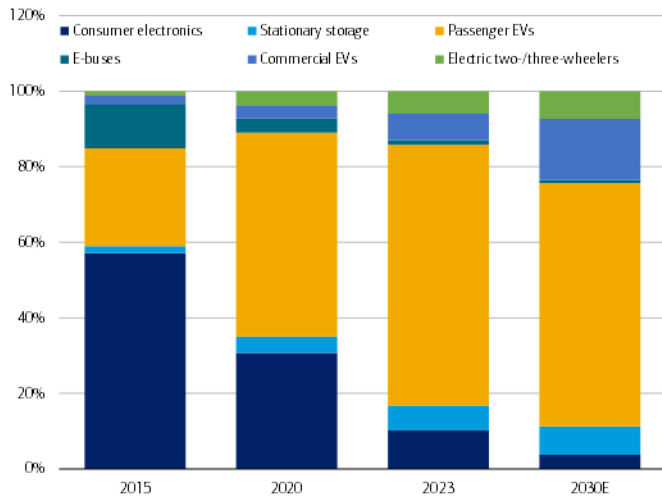
Goodbye gigawatt, hello terawatt: battery demand to rise 6x in 2023-35

This has had a demonstrable impact on the scale of battery demand, which rose to over 1 terawatt hour in 2023, up 12x from 83 gigawatt hours (or 0.08 terawatt hours) in 2015, and set to grow a further 6x to 2035 (to 5.7TWh) per BNEF. The key reason is the scale of capacity required in a lithium ion battery for transport applications. For example, on a per unit basis, the average pack size for an electric car battery grew from 40kWh in 2018 to 53kWh in 2022, and is projected to increase to 80kWh by 2030 per BofA APAC EV battery research team. By comparison, the capacity of an iPhone 15 smartphone is 0.013kWh – 4,083x smaller. The rising scale of EV demand for cars in particular is set to fuel this considerable demand growth for lithium-ion batteries.



**Exhibit 231: Transportation now leads demand for lithium-ion batteries**

Demand for batteries used in transport overtook consumer electronics in 2020 at 83% of battery demand in 2023, projected to rise to 89% in 2030

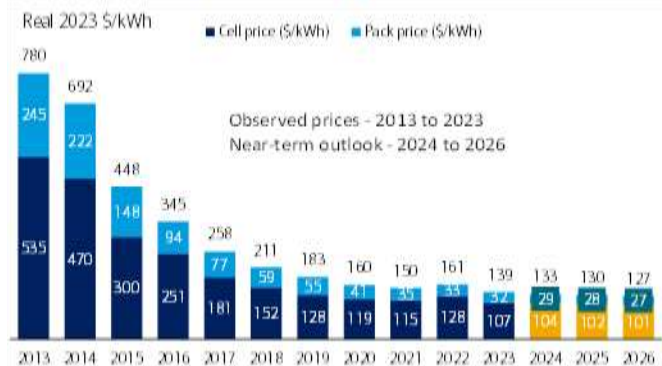


Source: BloombergNEF % share of demand for lithium-ion batteries  
BofA GLOBAL RESEARCH

**Energy density continues to rise, but at a slowing pace**

The average energy density of EV battery packs more than doubled from 88Wh/kg in 2010 to 188Wh/kg in 2023, increasing at a 6% CAGR. This has been achieved by a combination of advances in chemistry, improved cell/pack designs, and economies of scale in manufacturing. Energy density is projected to continue to increase, but at a slowing rate of 3% CAGR to 2035 (281Wh/kg), based on the current battery technology trajectory, per BNEF. Battery prices increased temporarily in 2022, due to rising material prices and component shortages. Whilst prices continued to decline in 2023 to \$107/kWh per BNEF, they will fall at a slower rate than the past five years to \$101 by 2026. *To achieve further more notable cost reductions would require disrupting the current technology at a faster pace.*

**Exhibit 233: Battery prices increased for the first time in 2022**  
Battery prices increased in 2022, due to rising material prices and component shortages but are now projected to fall again but at a lower rate

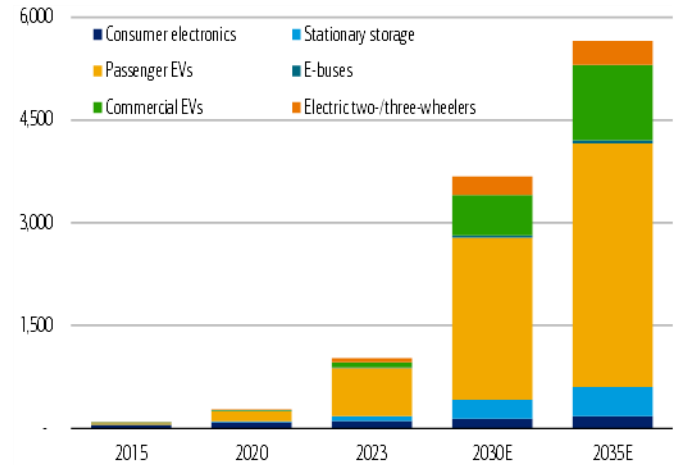


Source: BloombergNEF Volume weighted average lithium-ion battery cell and pack price  
BofA GLOBAL RESEARCH

**Advancing chemistry, engineering, materials for future batteries**

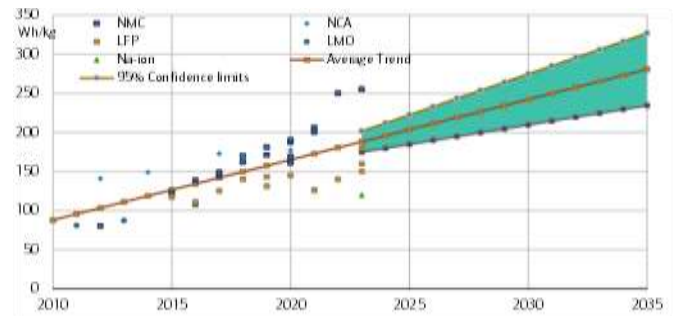
Several converging trends could enable improvements in the functionality and form factor of the next generation of batteries, with a convergence of improved chemistry, engineering and material science research and techniques. These are underpinned by increased funding and urgency, with batteries key to mitigating climate change as a

**Exhibit 232: 12x battery demand growth 2015-23, further 6x to 2035?**  
Rising demand for lithium-ion batteries largely driven by increasing battery size & volumes required in transportation



Source: BloombergNEF, Avicenne; Gigawatt hours lithium-ion battery demand outlook  
BofA GLOBAL RESEARCH

**Exhibit 234: Average EV battery pack energy density doubled 2010-23,**  
but growing at a slowing pace  
A further 50% increase in energy density is expected to 2035 based on the current technology trajectory

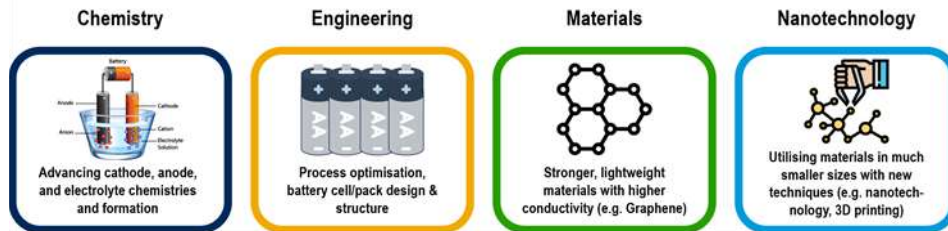


Source: BloombergNEF NMC = nickel manganese cobalt oxide; NCA = nickel cobalt aluminium oxide; LFP = lithium iron phosphate; LMO = lithium manganese oxide; Na-ion = sodium ion; Wh/kg = watt-hour per kilogram  
BofA GLOBAL RESEARCH



cleaner power source than fossil fuels. Furthermore, advances in AI and machine learning technologies could accelerate this, from identifying new materials or chemistry iterations at a faster pace, and simulation of their performance to speed up the development time.

Exhibit 235: Battery Breakthroughs enabled by chemistry, engineering, material science  
Advancing chemistry mix of battery components, engineering of cells/packs, and utilizing materials in smaller sizes and greater consistency are key enablers of next generation batteries



Source: BofA Global Research

BofA GLOBAL RESEARCH

Leading the charge with chemistry diversification; no longer a two-horse race  
The initial battery tech roadmap was categorised by nickel-based chemistries (particularly NCM: nickel, cobalt and manganese based) which is preferred in the West vs lithium iron phosphate (LFP) based chemistries launched in China. However, there are now multiple derivatives of each battery chemistry (defined by the concentration of materials used), each having unique benefits and drawbacks to their commercialisation.

Rather than a single chemistry emerging as the future “winner”, it is likely that each can serve different use cases and customers as the addressable market increases.

We highlight four emerging chemistry trends in the evolution of EV batteries:

- **Advanced Lithium-ion:** increasing nickel content in cathodes, replacing graphite with silicon in anodes, can enable a combination of increased energy density, fast charging and longer durability of batteries. Nanotechnology and manufacturing advances could enable this.
- **Rise of LFP/hybrid batteries:** the attributes of iron phosphate-based chemistries have improved considerably. Whilst they remain lower energy density than nickel-based chemistries, they're cheaper with longer cycle lives. Hybrid/derivatives of this chemistry are being launched ex-China for entry level EVs in particular (China accounted for 99% of LFP capacity in 2023 per the FT).
- **Material switching/rise of sodium ion:** replacing rarer/costlier materials such as lithium with relatively more abundant materials such as sodium, aluminium, or sulphur all offer potential for reducing cost. Challenges remain on conductivity and density of these alternatives, but sodium-ion chemistries in particular are beginning commercialisation, and could grow to 8% of battery capacity by 2030, per Roland Berger.
- **Solid state batteries:** replacing liquid electrolyte with a solid alternative, and lithium metal anodes (rather than graphite) could increase energy density and safety considerably, but challenges remain on achieving suitable cycle life and obtaining the materials at scale.

Exhibit 4: Summary of pros and cons of diverse battery chemistries

Major global battery manufacturers continue to diversify battery chemistries to differentiate rising demand in volume/entry EV models from premium models

High-Ni	High-voltage single-crystal mid-Ni	Mn-rich	LMFP	LFP	Na-ion	SSB
<b>Pros</b> <ul style="list-style-type: none"> <li>High energy density</li> </ul> <b>Cons</b> <ul style="list-style-type: none"> <li>High price</li> <li>High risk of TP due to structural instability (microcrack, gas generation)</li> </ul>	<b>Pros</b> <ul style="list-style-type: none"> <li>Decent energy density</li> <li>Lower price than high-Ni</li> <li>Lower risk of TP</li> </ul> <b>Cons</b> <ul style="list-style-type: none"> <li>Tech hurdle for mass production (single-crystal required for high voltage)</li> </ul>	<b>Pros</b> <ul style="list-style-type: none"> <li>Low price</li> <li>Higher thermal stability than high-Ni</li> </ul> <b>Cons</b> <ul style="list-style-type: none"> <li>Tech hurdle including lower stability, gas generation</li> </ul>	<b>Pros</b> <ul style="list-style-type: none"> <li>Low price as LFP</li> <li>High thermal stability as LFP</li> <li>Higher energy density vs LFP</li> </ul> <b>Cons</b> <ul style="list-style-type: none"> <li>Low cycle life vs LFP</li> <li>Lower energy density vs ternary battery</li> </ul>	<b>Pros</b> <ul style="list-style-type: none"> <li>Low price</li> <li>High thermal stability</li> <li>Long cycle life owing to less degradation</li> </ul> <b>Cons</b> <ul style="list-style-type: none"> <li>Lower energy density vs ternary battery</li> <li>Low space efficiency</li> </ul>	<b>Pros</b> <ul style="list-style-type: none"> <li>Lower costs and abundance of raw materials vs Li-ion</li> <li>Higher safety vs Li-ion</li> <li>Environmentally friendly</li> </ul> <b>Cons</b> <ul style="list-style-type: none"> <li>Low energy density</li> <li>Short cycle life</li> </ul>	<b>Pros</b> <ul style="list-style-type: none"> <li>High energy density</li> <li>High stability</li> <li>Long cycle life</li> <li>Fast charging</li> </ul> <b>Cons</b> <ul style="list-style-type: none"> <li>High costs</li> <li>Tech hurdle including dendrites</li> <li>Higher dependency on lithium</li> </ul>

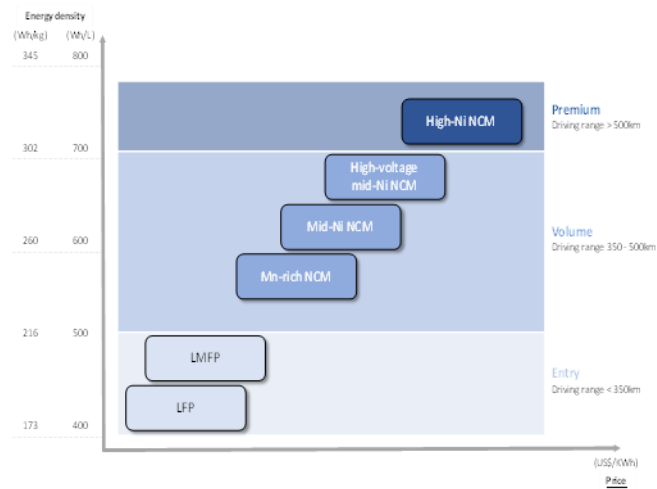
Source: BofA Global Research; Ni = Nickel, Mn = Manganese, LMFP = lithium manganese iron phosphate, LFP = lithium iron phosphate, Na-ion = sodium ion, SSB = solid state battery

BofA GLOBAL RESEARCH

Battery chemistry mix is diverging between premium/entry level

Based on the current state-of-the-art and next-generation battery chemistries, a divergence is expected across regions and vehicle segments, to differentiate products between entry level and premium vehicles, for example. An increasing share of high-nickel chemistries is expected for premium vehicles, with the share of LFP and sodium ion increasing particularly in entry level vehicles.

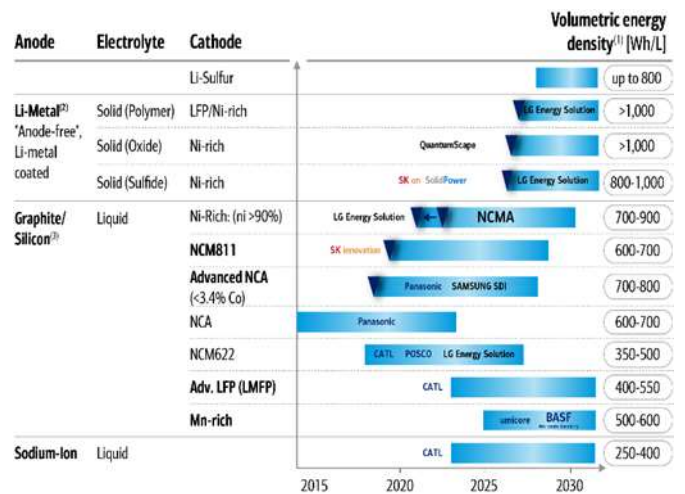
Exhibit 236: EV Battery energy density by various chemistries  
High-Ni batteries usually have the highest energy density of around 700-750Wh/L energy density, vs. LMFP/LFP's around 400-500Wh/L



Source: BofA Global Research

BofA GLOBAL RESEARCH

Exhibit 237: Li-B tech roadmap: EV battery chemistries diversifying  
Next gen battery tech: targeting cost reduction, increase energy density and charging rates.



1) Stacked electrodes (pouch or prismatic cell); 2) Foil or deposited; 3) Various ratios of Si/graphite and different Si morphologies possible

Source: Roland Berger

BofA GLOBAL RESEARCH

So what are the next-gen EV battery chemistries? No longer just LFP & NCM

High-nickel batteries for higher-energy density and longer driving range  
High-nickel NCM (nickel, cobalt, manganese) batteries use a high content of nickel – usually over 80%, but more recently many battery and cathode makers have commercialised nickel content over 90%, with cobalt and manganese below 10%, to maximise the energy density of a cell and thus driving range. However, as materials (cobalt and manganese) that support structural stability in a cell are reduced, it increases the chance of thermal propagation (TP), increasing the risk of explosion that needs to be mitigated with advanced cooling/battery management systems, adding cost and complexity.



### High-voltage single crystal mid-nickel: cheaper, stable, longer lifespan

Mid-nickel cathodes have 40-60% nickel content. Although they have lower energy density than high-nickel cathodes, they are cheaper and more stable owing to lower contents of nickel and cobalt. Battery manufacturers are increasingly attempting to apply high voltage to mid-nickel cathodes to raise the energy density close to that of high-nickel cathodes. Furthermore, using single-crystal cathode materials can extend battery life. These single-crystal cathodes are made of metals, such as nickel, cobalt and manganese, placed into a one-body structure through the heating process. Compared with conventional cathodes in a polycrystal structure of metal particles (which create fissures impacting a battery's lifespan), single-crystal cathodes have an enhanced thermal stability and lifespan, as structural damages during charging cycles are reduced.

### Manganese rich: cheaper/safer vs high nickel, higher energy density than LFP

Lithium-ion batteries with high content of manganese but lower contents of high-price metals, including nickel and cobalt are called 'high-Mn', 'manganese rich', or several other names such as 'lithium manganese oxide (LMO)', or 'cobalt-free' depending on the component makers and the composition of constituent metals. Manganese is the fifth most abundant metal in the earth's crust, which makes it relatively cheap and easily accessible, and it is often blended with other metals to improve stability and performance in batteries.

### LMFP: higher energy density than LFP

A lithium manganese iron phosphate (LMFP,  $\text{LiMn}_x\text{Fe}_{1-x}\text{PO}_4$ ) battery uses a mixture of lithium iron phosphate and lithium manganese phosphate in the same structure as in lithium iron phosphate. Adding manganese into LFP will lift energy density, by increasing the voltage within the same specific capacity, and will improve the low-temperature performance, the two main pain points of traditional LFP. Meanwhile, the incremental cost is minimal given the inexpensive metal used, thus LMFP is a pioneer technology and is considered an upgraded version of LFP.

### Rise of LFP batteries; stable, cheap, but lower energy density

A lithium ferrophosphate (LFP) battery is a type of lithium ion battery that uses lithium iron phosphate ( $\text{LiFePO}_4$ ) as the cathode material. LFP batteries operate in a similar way to other lithium ion batteries, moving between positive and negative electrodes to charge and discharge. However, as the name suggests, LFP requires three elements: lithium (Li), iron (Fe), and phosphorus (P). It doesn't contain nickel (Ni) or cobalt (Co), which leads to reduced raw materials costs and fewer environmental concerns. Moreover, it has lower energy density compared with ternary (NCM) batteries.

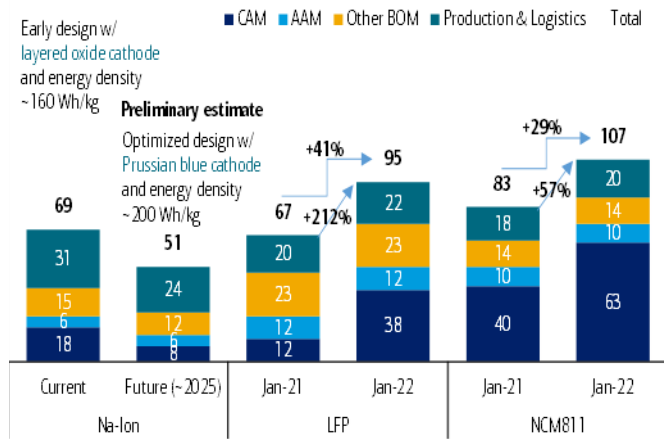
### Sodium-ion batteries: less prone to commodity volatility

Sodium-ion batteries use sodium-ions ( $\text{Na}^+$ ) as the charge carriers rather than lithium-ions ( $\text{Li}^+$ ), while its operational principles and structure of the cell are identical to those of lithium-ion batteries. Na-ion batteries are gaining prominence as a potential substitute for existing lithium-ion battery technology, mainly because they are cost-effective, using raw materials that are more abundant, and have lesser environmental impact given the ease of obtaining sodium relative to lithium. As the cost of manufacturing the cathode is relatively similar for both sodium-ion and lithium-ion battery technologies, the primary cost savings with sodium-ion batteries stem from the affordability of their raw materials: sodium and aluminium – sodium carbonate costs c.US\$0.3/kg vs. lithium carbonate c.US\$20/kg as of 2022. Given sodium-ion batteries use aluminium foil as a collector for anode materials instead of copper foil, which is c.3.2x as expensive as aluminium, the overall manufacturing costs for SiB are much more affordable vs. LiB, and could fall much further as more capacity is added and optimised.

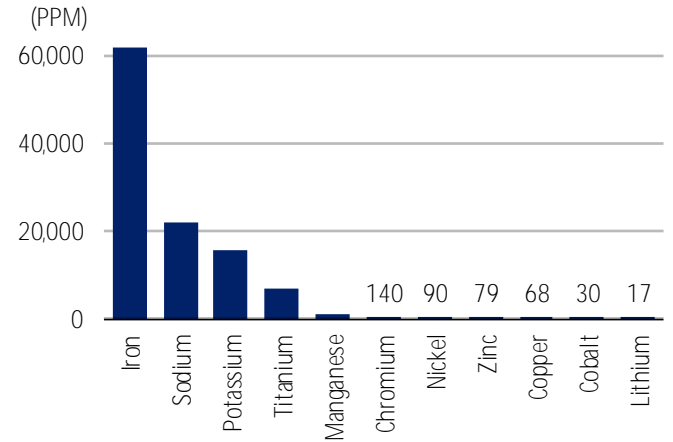


**Exhibit 238: Sodium-ion could de-risk from volatile battery metal prices**

Increased maturity and scale of sodium-ion batteries could reduce cell cost to \$51/kWh which could insulate manufacturers from volatile commodity prices such as nickel/cobalt/lithium given these metals aren't required



**Exhibit 239: Abundance in crust by materials**  
Sodium is the sixth most abundant element globally



Source: Company data, Press reports, BofA Global Research

BofA GLOBAL RESEARCH

BofA GLOBAL RESEARCH

**Solid-state batteries: safer, higher energy density, long life & fast charging**

Solid-state batteries have long been considered a crucial innovative technology for the future of electric vehicles, thanks to advantages including:

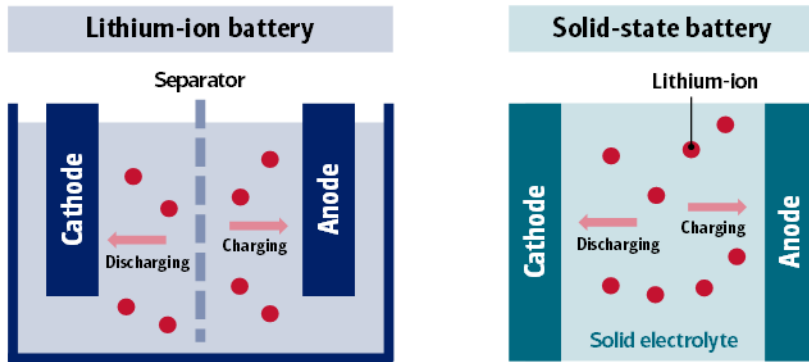
- **Higher energy density:** a solid-state battery uses an anode made of pure lithium metal. It can as much as double energy density compared with a traditional lithium-ion battery, resulting in a lighter and smaller battery.
- **Improved safety:** replacing the liquid electrolyte found in current lithium-ion batteries with solid components (oxide, sulphide or polymer-based), preventing safety issues related to electrolyte leaks or fires (given liquid electrolytes are more volatile and flammable).
- **Longevity:** a thicker separator allows the batteries to be more resistant to high temperatures, which makes the separation between the anode and cathode more reliable, with consistent performance over time and a longer service life.
- **Fast-charging:** the improved resistance to high temperatures also provides fast-charging features – a solid-state battery can be charged up to 6x faster compared with a traditional lithium-ion battery.

In conclusion, solid-state technology could make batteries lighter, faster to charge and eventually cheaper, as battery pack sizes could be reduced to achieve the same level of driving range. However, challenges remain to their commercialisation: 1) avoiding dendrite formation (tree-like structures of lithium metal that grow inside the battery and can reduce performance and lead to short-circuiting or failure); 2) achieving manufacturing cost parity to current industrial scale batteries; and 3) reducing lithium consumption – solid-state batteries could require 5-10x as much lithium (given their use of pure lithium metal anodes) as lithium-ion cells use currently, per consultancy firm Minviro.





Exhibit 240: Structure of solid-state battery vs traditional lithium-ion battery  
 Solid-state batteries use solid electrolyte instead of liquid electrolyte solution



Source: Company data, BofA Global Research

BofA GLOBAL RESEARCH

### Lithium sulphur batteries

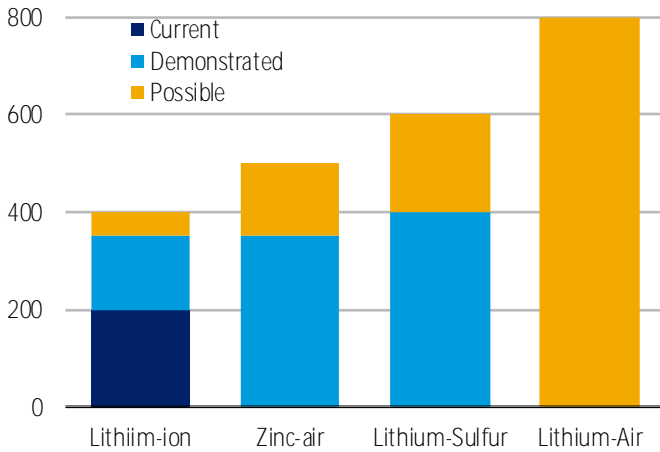
Lithium sulphur batteries have a lithium anode and a sulphur carbon cathode, and could provide high energy density at lower cost. The low atomic weight of lithium and moderate weight of sulphur make the Li-S batteries lighter too. Li-S batteries could offer an alternative to Li-ion batteries because of their high energy density and reduced cost due to the use of sulphur (instead of cobalt). Applications so far are mainly outside of transportation – for example, NASA has invested in solid-state Li-S batteries to power space exploration. On the downside, Li-S batteries have low electrical conductivity, thereby requiring extra conductive agents, which could increase the weight and stress on the battery pack. Current prototypes have lower cycle life as a result. So far, charging a Li-S battery causes a build-up of chemical deposits that degrade the cell and shorten the lifespan.

### Metal air batteries: getting batteries close to hydrocarbon energy density?

Metal air batteries have a pure metal anode (lithium, aluminium, zinc or sodium) and an ambient air cathode (mostly oxygen). They have a power density close to that of gasoline. It is a type of fuel cell battery that utilises oxidation of a metal at the anode and reduction of oxygen at the cathode to produce electricity. Pairing metal and oxygen (from the air) can theoretically lead to electrochemical cells with the highest specific energy. Research projects, such as Argonne National Laboratory, suggest 5x more energy can be achieved than for a Li-ion battery, making it one of the most sought-after battery breakthroughs. The metal-air battery is expected to be lightweight as the cathode that weighs the most is made up of air.

However, there are many challenges such as capturing enough volume and purity of oxygen from air, battery recyclability and cycle life. Being able to create the complex packaging and air breathing interfaces in the cell to enable them to recharge is also a key challenge. Improvements in zinc-air and lithium plating could mitigate this, but commercialisation will also depend on significant cathode improvements.

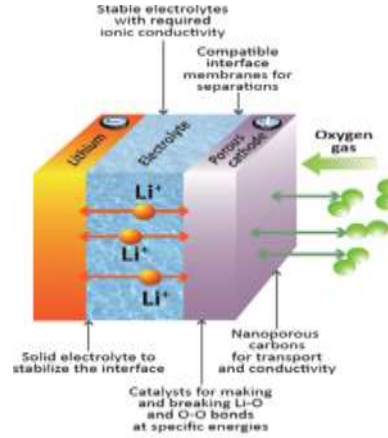
Exhibit 241: Next Gen battery energy density outlook  
Alternative battery chemistries to lithium-ion could double achievable gravimetric energy density with further breakthroughs



Source: Rocky Mountain Institute, denotes Wh/KG energy density

BofA GLOBAL RESEARCH

Exhibit 242: Lithium air batteries  
Replacing current lithium-ion cathodes with a porous alternative that can use oxygen as the reactant can theoretically achieve 4-5 the energy capacity at lower cost. Power output, safety and cycle life are key challenges



Source: Ecologic Institute

BofA GLOBAL RESEARCH

### Nanotechnology for extreme fast charging (XFC)

Varying anode chemistry can also enable extreme fast charging (XFC). Current (graphite) anodes are stable at normal charging rates but face high electrical resistance – fast charging can cause dendrites on the anode surface, damaging the cells. Replacement with materials such as silicon/tin can alleviate that, allowing faster conductivity. However, nanotechnology and new materials are required to enable the use of smaller-sized particles. These can give more area to penetrate the active material and lower resistance and heat.

### Battery Innovation: battery cells, packs, system & software

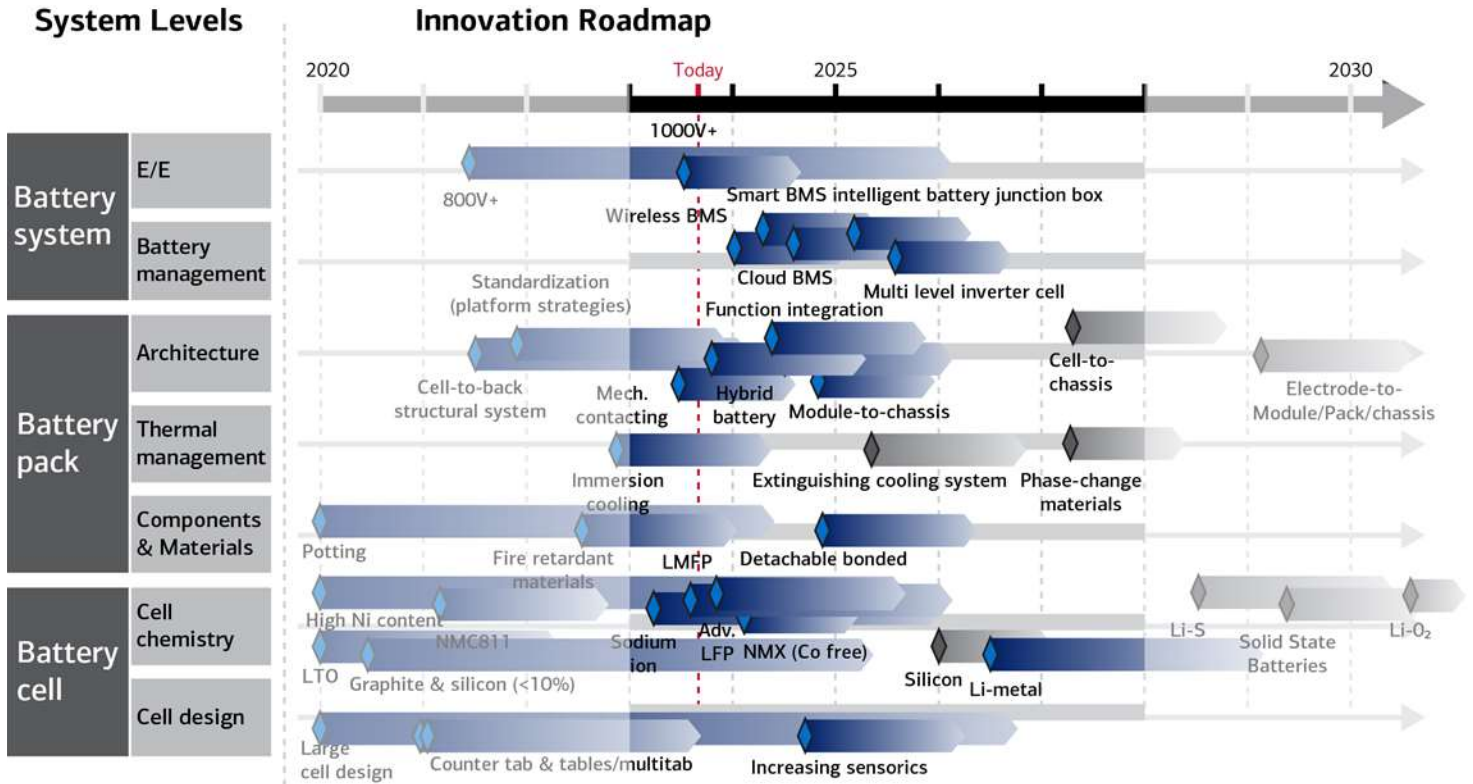
As the competitiveness and scale of deployment of batteries increases, manufacturers are increasingly looking beyond the battery cell to achieve improved efficiency and performance across the pack and entire system.

1. Battery pack: materials, thermal management, and architecture – e.g., embedding batteries into chassis or structure of cars
2. Battery system: software and electric architecture are growing in importance and sophistication as a means of differentiation.



Exhibit 243: Innovation roadmap in battery cells, packs and systems

Beyond cell chemistry, several innovations at the pack and system level are possible to improve battery performance

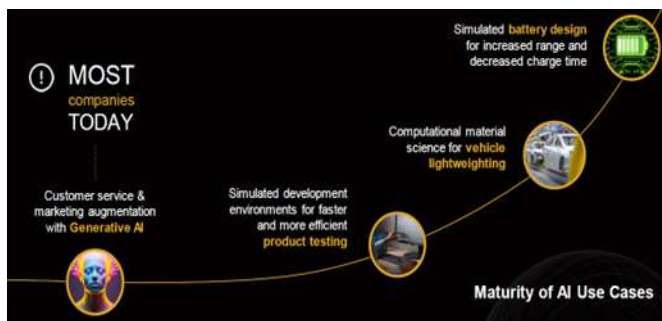


Source: Roland Berger, PEM RWTH Aachen University

AI could accelerate battery development

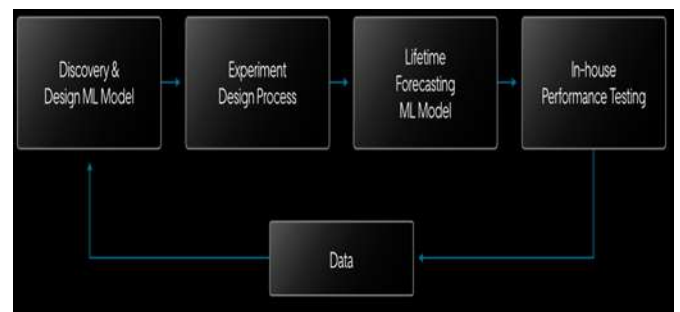
The current approach to battery design is slow, complicated and expensive for a number of reasons: 1) the infinite number of battery materials that could be combined; 2) long testing times that can take years to complete and require high compute resources; and 3) difficulty in predicting how change in battery design will impact performance. Combined, these challenges slow the path to next-generation batteries. However, recent breakthroughs in Artificial Intelligence (AI) technologies could accelerate battery development significantly, by combining data with machine learning and forecasting models to speed up development from years to months, or less.

Exhibit 244: AI use cases to extend beyond chatbots to chemistry  
Advancing AI with simulation and quantum computing could expand the maturity of use cases in battery development



Source: Sandbox AQ, presented at BofA Transforming World Conference 2023

Exhibit 245: AI and data could accelerate battery development  
Combining vast amounts of relevant data with machine learning and forecasting could accelerate feedback loops in battery development



Source: Chemix

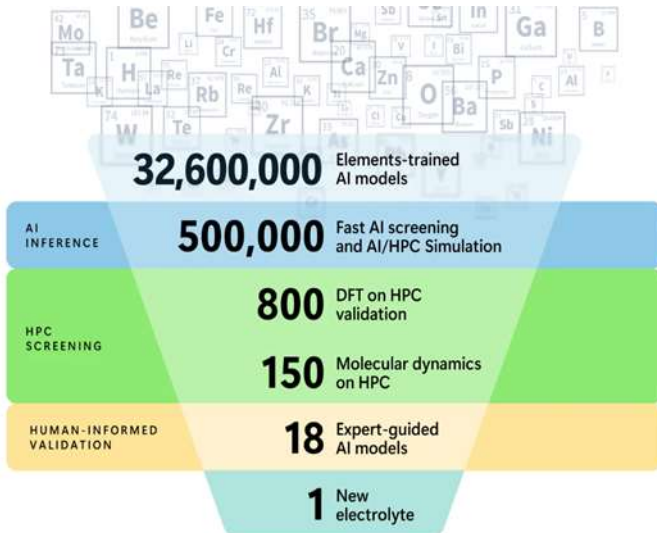


## When AI meets quantum tech: battery abundance

Furthermore, as computational capabilities continue to advance with high-powered compute getting cheaper, and eventually quantum computing able to be commercialised, the development of battery chemistries could accelerate manifold. The ability to parallelise multiple experiments is already being tested to analyse the structure of molecules and combination of materials and how they interact, with the aim to identify new battery materials and chemistries faster, and simulate potential outcomes and cycle life testing.

This work can begin to be tested prior to quantum computing becoming commercially available, by combining AI with cloud computing. For example, in January 2024, Microsoft announced a partnership with the Pacific Northwest National Laboratory (PNNL) in Washington, US to use advanced AI with high-performance computing. They used AI to winnow >32m potential materials to 18 promising candidates in just 80 hours, with PNNL scientists then testing the 18 with expert-guided AI models to filter down to one final new electrolyte candidate that uses 70% less lithium than lithium-ion batteries by replacing with the much more abundant sodium.

Exhibit 246: Combining AI/cloud to identify a new battery material  
In 2024 Microsoft and PNNL used AI to identify a new battery material, filtering from ~33m elements to 1 new electrolyte in weeks



Source: Microsoft Azure Quantum Elements; DFT = density functional theory, HPC = high-powered computer, AI = Artificial Intelligence  
BofA GLOBAL RESEARCH

While validation is ongoing, the research team was able to build a working solid-state battery tested at both room and high temperature, notable owing to 1) the speed of development and 2) dispelling prior scientific bias. It was thought sodium and lithium couldn't be used together in such a system owing to their different particle sizes. The scientists previously assumed the material couldn't support the movement of two different ions but found with testing they helped each other, as the molecular structure of the new material discovered has channels that enable both ions to pass through the electrolyte.



Structural & Material breakthroughs could bring next-gen batteries faster

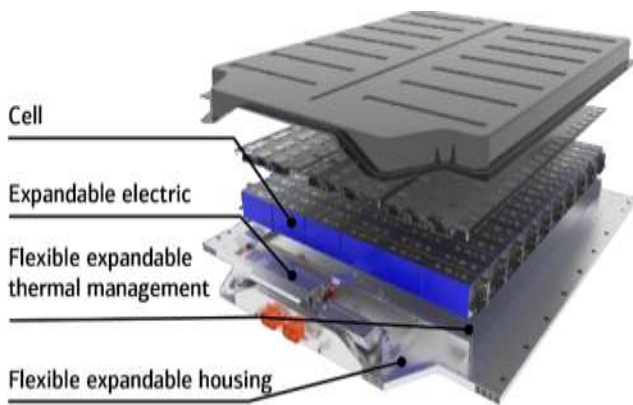
What are they? Integrating batteries into structural objects (battery packs, vehicle chassis, or body structures) using new materials and manufacturing techniques

Did you know? Integrating battery cells directly to packs can increase energy density by 20% and reduce number of parts by 40% for the same battery system.

Cell to pack and cell to chassis: the end of the battery pack?

Alongside battery cell chemistry improvements are methods to integrate the cells directly into battery packs to increase energy efficiency and lower costs. These “cell to pack” technologies can eliminate or simplify battery module structures. Usually battery cells are clustered into modules, which is then installed into a larger battery pack. Removing the modules can save space/weight and thus allow for more energy in the same space. Tesla and CATL announced plans to deploy this in 2021, and BYD is already doing so with Blade batteries. According to CATL this can increase gravimetric energy density by 10-15% and volumetric energy density by 15-20% for the same battery system. It also reduces the number of parts in a battery pack by 40% and increases production efficiency by 50%.

Exhibit 247: Cell-to-Pack improves volume utilization; less parts & cost  
Improving fabrication and battery management systems can enable cells to be integrated directly into battery packs, rather than housed in smaller modules first



Source: CATL

BofA GLOBAL RESEARCH

Exhibit 248: Increasing available energy in the same pack size  
Cell-to-pack can increase the energy density available in the same pack size, by removing the packaging/connectors between modules



**Low space utilization  
(~40% VCTP)  
Battery pack with modules**

Source: BYD; VCTP = volumetric cell-to-pack ratio



**High space utilization  
(~60% VCTP)  
Blade Battery without modules**

BofA GLOBAL RESEARCH

The key rationale behind this is improving energy density in the short term, but the trajectory of the technology is evolving towards cell-to-chassis/cell-to body techniques that also eliminate the need for dedicated battery packs, instead integrating cells within the structure of objects such as a vehicle chassis. Tesla demonstrated this at its 2020 battery day and CATL aims to commercialise this before 2030. If achieved, not only could this enable higher energy density, but it could begin to enable performance feedback loops within EVs by reducing weight, thereby the volume of energy required to travel the same distance, and therefore cost. This could accelerate the crossover point for cost parity of heavier vehicles in particular, by reducing the battery size/costs they would require to electrify.





Massless energy storage: integrating batteries into structures

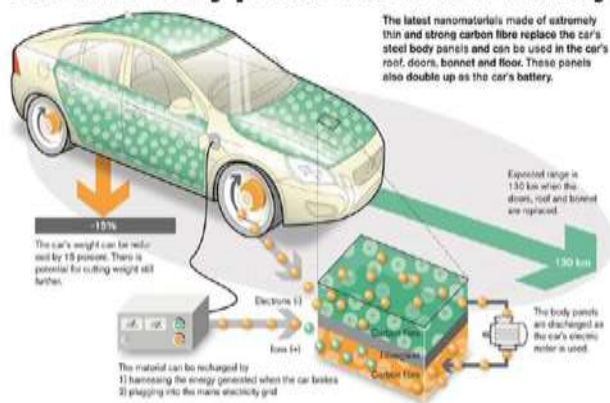
Further development of the technology to integrate batteries into the structure of objects is paving the way for “massless” energy storage in vehicles and other technologies. Current batteries account for a large proportion of vehicle weight without fulfilling any load-bearing function – the batteries in a Tesla Model S for example weigh >500kg, a quarter of the total vehicle weight. Integrating batteries into the structure of objects could remove that incremental weight of batteries (hence the term “massless” energy storage), having the effect of reducing: the volume of stored energy required to power the vehicle, and the raw material resource intensity of the vehicles.

To do so requires using materials that can offer both rigidity and conductivity such as carbon fibre, or graphene. In 2021 research by Chalmers University demonstrated a structural battery with 10x higher energy density than previous iterations. Whilst at 24 Wh/kg the density was still only 20% that of current lithium-ion battery technology, the battery as strong as aluminium but less than half the weight. The researchers claim up to 75Wh/kg would be achievable within a few years, making the technology viable for use in consumer electronics and electric bicycles. Longer term, electric cars, satellites and planes could adopt the technology as it matures.

**Exhibit 249: The car's body panels could serve as a battery**

Using thin lightweight materials safe enough to be incorporated into the structure of vehicles could enable the body of a car to double up as its battery

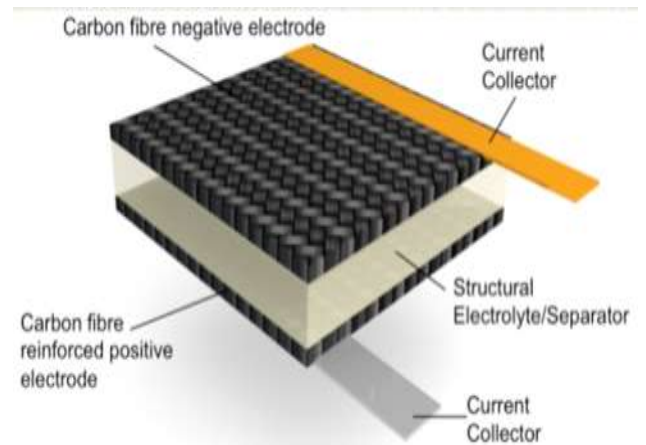
**The car's body panels serve as a battery**



Source: Rocky Mountain Institute

**Exhibit 250: “Massless” energy storage breakthrough**

Chalmers University of Technology demonstrated a structural battery in 2021 with an energy density of 24 Wh/kg using carbon fibre as a cathode, LFP coated aluminium foil anode and a fibreglass electrolyte separator



Source: Chalmers University of Technology





## Future Air Mobility: It's a bird? It's a drone? It's eVTOL?

What is it? Light aircraft powered by electric propulsion that can take off and land vertically. The answer to: increasing urbanisation and congestion by democratising lower-cost/higher-frequency urban air travel for passenger and cargo transit.

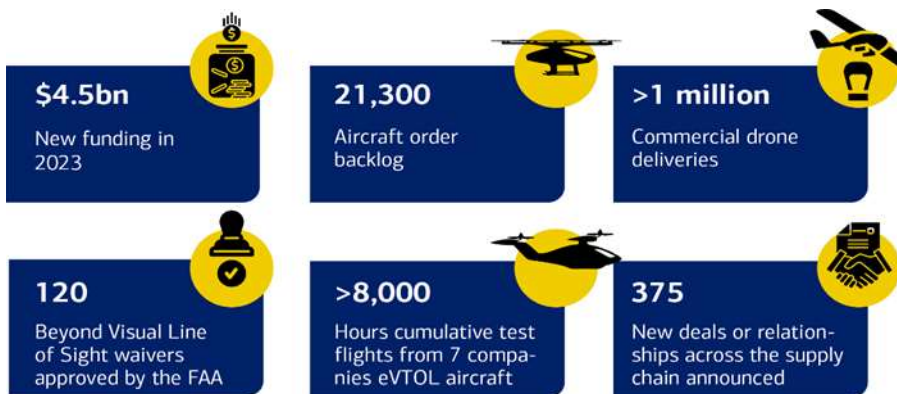
Did you know? The first passenger eVTOL aircraft are expected to be certified for commercial operations in 2024. Using drones can reduce per-package delivery emissions by up to 97%.

### Future Air Mobility: it's about to take off

A combination of improving tech (batteries, avionics, autonomous systems/sensors), regulatory frameworks and decarbonisation targets are accelerating the growth potential and commercialisation of "Future Air Mobility" industries. These include electric Vertical Take-Off and Landing aircraft (eVTOL), commercial drones, sustainable aviation and supersonic aircraft, per McKinsey. More than \$4.5bn in new funding was announced in 2023 across 151 deals, taking cumulative funding of these areas to \$22bn since 2004. We focus on drones and eVTOL given their near-term commercial trajectories.

- Drones: This is already a \$34bn market and projected to grow to \$55bn by 2030 as tech improves and use cases expand, particularly delivery services – >1m commercial drone deliveries were made in 2023, a market that has grown at over a 100% CAGR since 2018, per McKinsey.
- eVTOL: A new classification of light commercial aircraft for passengers and/or cargo is nearing commercialisation, following 8,000 hours of cumulative test flights from 7 aircraft that are targeting certification in 2024/2025. Some are based on drone platforms using multicopters, others larger aircraft using fixed-wing designs.

Exhibit 251: Future Air Mobility: >1m drone deliveries in 2023; >8,000 hours eVTOL test flights  
The Future Air Mobility industry achieved several milestones in 2023 around testing, funding and partnerships



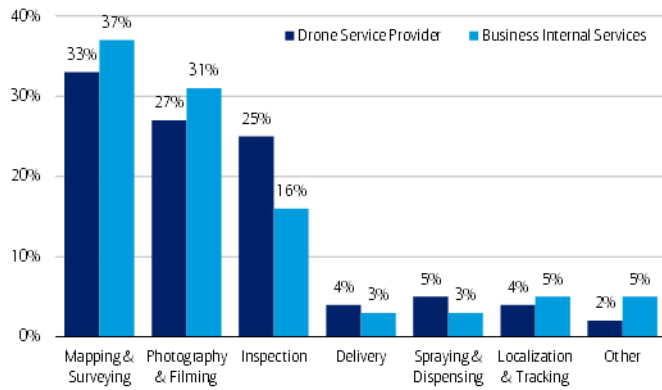
Source: McKinsey & Co: Future Air Mobility funding includes eVTOL, hydrogen, electric and hybrid powertrains

BofA GLOBAL RESEARCH

Commercial drone market: \$34bn in 2023, projected at \$55bn in 2030  
Commercial drone services (ex-military) are already a market worth \$34bn annually – 80% of which being services, 16% hardware, and 4% software revenues. It is projected to grow at a 7% CAGR to \$55bn by 2030 (source: Drone Industry Insights). The largest use case for commercial drones thus far is mapping & surveying (e.g. in the energy industry), worth \$10bn alone in 2023, followed by filming and inspection drones. The use cases are expanding, however, to areas such as food/goods delivery, agriculture and precision tracking in warehouse logistics.



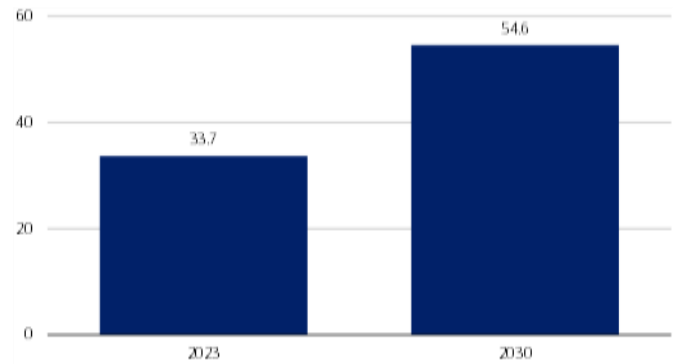
Exhibit 252: What are drones used for?  
Mapping, filming, and inspection were the key use cases for drones in 2023



Source: Drone Industry Insights

BofA GLOBAL RESEARCH

Exhibit 253: Global drone industry market value is \$34bn  
The value of the global drone industry is projected to increase from \$34bn in 2023 to \$55bn in 2030



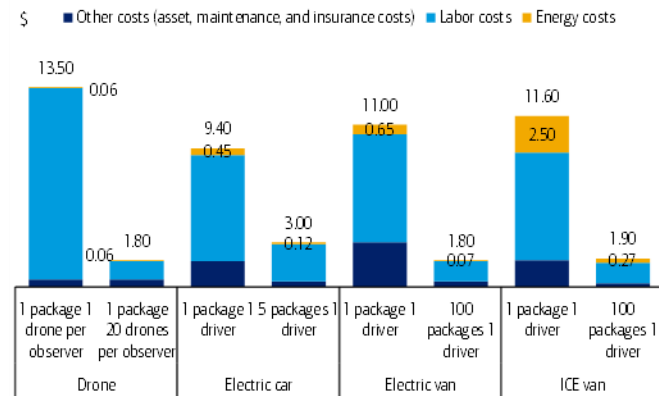
Source: Drone Industry Insights

BofA GLOBAL RESEARCH

Packages before People: Drone delivery market commercialising

Drone delivery is a rapidly growing commercial opportunity following regulatory approval and certification of several company drone designs, and a range of new operating platforms enabling the delivery of e.g. medicine, food and packages. The number of deliveries has swelled from 6k in 2018 to almost 900k in 2022, with >1m projected in 2023 (500k deliveries were made in the first half) per McKinsey, by companies including Zipline, Wing (Alphabet), Matternet, Swoop Aero, Flytrex and Manna. Growth is expected to accelerate as improving tech and favourable regulation improve economies of scale.

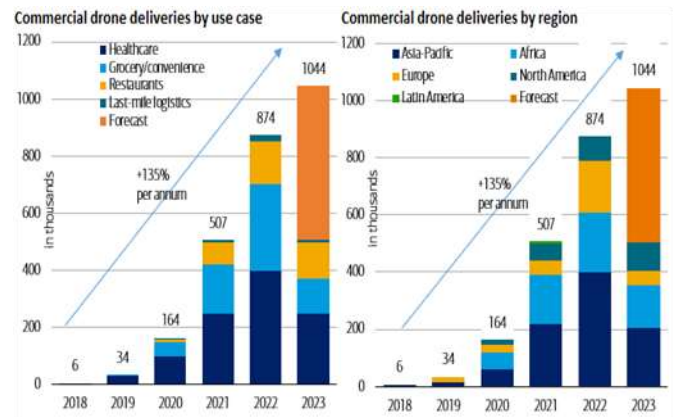
Exhibit 254: Drones could become cost-competitive with road transport  
Unit costs of drone delivery could become competitive with other ground transport modes per McKinsey as volumes increase and operational restrictions from regulators reduce



Source: McKinsey

BofA GLOBAL RESEARCH

Exhibit 255: Commercial drone delivery market growing rapidly  
There were ~900k commercial drone deliveries in 2022, expected to cross 1m in 2023 as use cases expand beyond healthcare to food delivery



Source: McKinsey

BofA GLOBAL RESEARCH

Why now? Convergence of improving tech, economics, and enabling regulation

A combination of improving technology and confidence in the capabilities of commercial drone platforms from regulators is leading to operational restrictions being removed. In turn, this is increasing the scale of operations permitted. The key enablers are:

- Technology: improvements and the falling cost of batteries, materials, and sensing for autonomous flight capabilities are enabling them to be commercialised, increasing the payloads that can be carried and extending flight time/range.



- Emissions: the per-package emissions for drone delivery are 85-90% lower than for a single package delivery via vehicles with combustion engines, per McKinsey, and as much as 97% lower per delivery according to Zipline.
- Regulation: regulators have placed several restrictions on commercial drone operation, but they are beginning to be loosened as more services are tested/proven. The US Federal Aviation Administration, for example, has a 5-stage certification process (Part 135) that Wing, Amazon, Flytrex, UPS and Zipline have achieved as of 2023. However, as well as limiting the number of drones they can fly simultaneously, operators are also required to deploy visual observers to oversee drone operations within the line of sight. These restrictions are also beginning to be loosened, though: in November the FAA approved beyond visual line of sight (BVLOS) operations for Zipline, with several other operators looking to achieve the same permissions that would enable them to scale their operations.
- Cost: drones are already cost-competitive with other forms of transport in areas with poor road infrastructure or where pooling deliveries doesn't make sense, but they could become more widely cost-competitive as volumes increase and the operational restrictions from regulators reduce. McKinsey estimates the cost of drone operations for a single package at \$13.50, 95% of which is labour. Increasing the volume of drones per observer would reduce this cost to \$1.80 (competitive with a van delivering 100 packages). Removing observers completely with more automation would reduce delivery costs even further.

#### Autonomous drones: coming to your neighbourhood soon

In order to fly, drones combine robotics and aeronautics technologies, using a mix of propellers and flight controllers to maintain stability. The improving capabilities of the onboard compute and sensing equipment enables them to perform certain functions or entire routes autonomously, using a combination of onboard GPS, cameras, and sensing equipment such as radar/LiDAR – similar to autonomous vehicles on the ground.

Drones are also categorized by the level of automation they can achieve without remote pilot control, from low to high. Most of the commercial drone operations have progressed to Level 3, equipped with sense and avoid technologies that enable them to navigate safely beyond an operator's line of sight, but still be piloted or controlled remotely if required. No full automation operations exist yet, but some highly automated (Level 4) systems are being commercialised for mapping/surveillance of challenging environments such as industrial inspections. Exyn drones were the first to achieve such capability in 2021, owing to their use of LiDAR sensors.

Exhibit 256: Autonomous drones: Increasing levels of drone autonomy driven by tech & regulation  
 Drones are categorized by the level of automation they can achieve without remote pilot control, from low to high. No full automation operations exist yet

Autonomy Level	Level 0	Level 1	Level 2	Level 3	Level 4	Level 5
Machine Involvement						
Human Involvement						
Degree of Automation	<b>No Automation</b>	<b>Low Automation</b>	<b>Partial Automation</b>	<b>Conditional Automation</b>	<b>High Automation</b>	<b>Full Automation</b>
Description	Drone control 100% manual	Pilot in control. Drone controls at least one vital function	Pilot responsible for safe operation. Drone can control heading, altitude in certain conditions	Pilot as a back-up. Drone performs all functions given certain conditions	Pilot out of the loop. Drone has backup systems & can operate if one fails	Drones able to use AI to plan their flights as autonomous learning systems
Obstacle Avoidance	NONE	SENSE & ALERT		SENSE & AVOID	SENSE & NAVIGATE	

Source: Drone Industry Insights

BofA GLOBAL RESEARCH

What do nextgen drone deliveries look like? From parachutes to wires

Per Tracxn, 143 companies globally already operate in the drone delivery space, using a combination of aircraft and operating models. Most initially deployed parachutes to deliver their cargo, or lowered the whole drone to the ground with its cargo. However this limits the range (as it uses more energy to land/take off than cruise) and is noisy, sparking complaints and restricting their application in urban areas. To mitigate this and expand coverage, many operators have iterated the technology to enable lowering devices by wire, or using a droid.

- Zipline has completed >960k commercial deliveries since its 2016 launch, covering 69m miles and delivering 10m items in the process, initially medical goods in several African countries and more recently goods/food in the US. Its new drone announced in 2023 hovers at 300ft while a droid delivers packages below.
- Wing (Google) has completed >350k deliveries with a fleet of 1k delivery drone aircraft across 3 continents since launching in 2019. Packages are attached to the bottom of the drone and lowered to customers via a wire. It announced a larger drone in January 2024 that should enable it to carry double the payload where required – up to 5 pounds (2.3kg) 2.5vs the current 2.5 pounds (1.1kg), given 30% of their deliveries thus far required two drones to fulfil.
- Flytrex: We met with drone food delivery company Flytrex on our 2022 Israel tech field trip. It offers delivery services in suburban areas that would otherwise not be economical using human drivers given the smaller scale than urban areas, dropping food directly into customers’ gardens in 5 US cities. The drones cost \$3k each to produce and make ~200 deliveries per week. They operate from hubs such as out-of-town retail parks with a launch location, carry out up to 5-mile round trips, travelling ~30mph for 5-minute deliveries. Food is lowered into customers’ gardens from 80 feet via a wire.



Exhibit 257: Zipline: drones meet droids  
 >900k commercial deliveries since 2016 initially by parachute; new platform uses a droid



Source: Zipline

BofA GLOBAL RESEARCH

Exhibit 258: Wing: expanding aircraft size  
 Wing (Google) >350k commercial deliveries since 2019, lowered to customers via wire



Source: Wing

BofA GLOBAL RESEARCH

Exhibit 259: Flytrex: first to wire-drop deliver  
 Flytrex was the first drone delivery system to operate in an urban environment, using a wire drop delivery



Source: Flytrex

BofA GLOBAL RESEARCH

**Is it a car? Is it a drone? Is it a plane? No, it's eVTOL.**

Electric Vertical Take-Off and Landing (eVTOL) is a new classification of light commercial aircraft for passengers and/or cargo. As the name suggests, they can take off and land vertically like helicopters, but use electric propulsion. Some are un-winged for short distances, like drones, while some are winged and fly forward like planes. The lighter weight, distributed propulsion and lower complexity are intended to make them more manoeuvrable, efficient, and lower cost vs traditional helicopters. They're designed to fly at lower altitudes than commercial aircraft, and to be piloted initially and eventually flown fully autonomously as regulation permits.

The key enablers thus far are converging technologies from the aerospace, autos and tech sectors, but challenges to commercialisation remain, relating to technical and regulatory issues along with public acceptance. Their key use cases are likely to be drone cargo (with increasing payloads) and passenger transit for short inner city and regional trips, hence the industry is often referred to as Urban Air Mobility (UAM).

Exhibit 260: Routes approved for eVTOL operation in Paris so far  
 eVTOLs need infrastructure ("Vertiports") and civil aviation authority approval for the routes: 5 have been approved for Paris launch phase in 2024



Source: Volocopter at BofA Transforming World Conference

BofA GLOBAL RESEARCH

Exhibit 261: Future electric autonomous aviation inspected by robots?  
 Supernal demonstrated S-A2 eVTOL concept aircraft at CES 2024 inspected by a Boston Dynamics robot for defects, due for launch in 2028



Source: BofA Global Research taken at CES 2024

BofA GLOBAL RESEARCH

Sooner than you think; commercial eVTOL launches start 2024/25  
 Certification and commercial launch of the first eVTOL services are expected in 2024 at the Paris Olympics, with testing underway already in several other cities. Up to seven



eVTOL aircraft are aiming for certification by 2025 (per McKinsey and press reports): Volocopter, Joby, eHang, Lilium, Archer, Autoflight and Vertical Aerospace. As such, funding, testing and partnerships are ramping up – there were >8,000 hours of test flights and >1m commercial drone deliveries in 2023 per McKinsey.

Why now? Converging technologies from Aerospace, Autos, AI  
Several technologies are converging that eVTOL companies are looking to exploit, particularly from the aerospace, automotive and technology industries.

Exhibit 262: EVTOL enabling technologies: combining Aerospace, AI, and Automotive  
Advancing capabilities in lightweight materials, compact flight control, AI-enabled systems, improving (and increasingly cheaper) Electric Propulsion, and on-demand mobility networks are key technologies enabling eVTOL UAM from the Aerospace, Tech and Auto/Mobility industries



Source: BofA Global Research

BofA GLOBAL RESEARCH

- **Aerospace:** Improving drone technology capable of increasing payloads, aircraft design (to accommodate multiple rotors and distributed propulsion), and materials (advanced lightweight composites) are all enablers of eVTOL aircraft from the aerospace industry. Improving software and digital capabilities that air-based vehicles can deploy already – such as avionics, navigation and flight control systems – are being advanced and redesigned for eVTOL.
- **Autonomy:** Combining fly-by-wire compute, navigation, and collision avoidance sensors such as radar could allow eVTOL vehicles eventually to fly autonomously. Advancing AI technologies, modelling and simulation can enable new Air Traffic Control systems specifically for Autonomous UAM that monitor and communicate with vehicles without human intervention. On the user experience side, platforms just as ridehail and on-demand mobility services are commonplace and can be leveraged to scale passenger demand.
- **Automotive:** Continually improving energy density and falling battery costs as Electric Vehicles scale are key enablers of the first eVTOL concepts to use lithium ion batteries for zero-emission travel cost efficiently. While current cell energy density/weight may limit the range of initial vehicles, improving and next-gen cell chemistries could allow longer UAM trips. eVTOL companies are already partnering with battery cell developers to design bespoke chemistries for UAM vehicles (e.g. Lilium with Customcells).
- **Manufacturing:** eVTOL companies envision higher scale (and lower cost) of manufacturing vs today's aircraft. The simpler design and fewer moving parts vs combustion-fuelled aircraft will be the key drivers of increased manufacturing. For context, Volocopter aims to be producing >1,000 aircraft by 2025, and external consultants project an addressable market as big as 250,000 vehicles by 2035 (Frost & Sullivan). This compares to ~32,000 helicopters in operation globally with annual production averaging ~1,000 per year (Statista). To enable this, many eVTOL companies have sought investments and partnerships with auto companies. For example, Joby with Toyota, Archer with



Stellantis, Volocopter with Daimler and Geely, and Hyundai in its own UAM division, to leverage strategic synergies including manufacturing. However, all manufacturing will need to be certified to aerospace standards.

Exhibit 263: Tech powering eVTOLs: lightweight airframe, electric propulsion, fly-by-wire autonomy  
UAM aircraft developers are competing for owned IP & partnerships to deliver a competitive market launch



Source: BofA Global Research

BofA GLOBAL RESEARCH

### >900 eVTOL concepts categorised by propulsion type

The number of eVTOL vehicle design concepts proliferated from six known designs in 2016 to >900 in 2023, according to the Vertical Flight Society. Spanning 413 companies, eVTOL aircraft are categorised by the type of design and propulsion they intend to deploy. For personal travel, there are hover bikes: single-person eVTOL aircraft piloted in a saddle or standing (all wingless configurations), and roadable aircraft (“flying cars”) – mostly winged vehicles requiring a traditional airstrip or similar to take off. Thus, the latter are not categorised as eVTOLs but rather eSTOLs (electric short take-off and landing).

For passenger transit-based UAM, in addition to electric rotorcraft (e.g. electrifying conventional helicopters), three new categories make up the majority of the active eVTOL developments: wingless multicopter, lift & cruise, and vectored thrust (open propeller or ducted fan-based) aircraft. These dominate current company efforts, accounting for 732 of the 129 eVTOL concepts in 2023, as per the Vertical Flight Society. The key difference between the three relates to the propulsion they use, impacting the range/speeds they can achieve, and the certification/technical requirements to operate:

1. Multicopters are wingless and more akin to larger commercial drones.
2. Lift & Cruise aircraft have separate propulsion units for take-off and cruise.
3. Tilting rotors or ducted fan-based vectored thrust aircraft use the same propulsion to take off and cruise.

Exhibit 264: Three emerging eVTOL categories  
 Multicopter, Lift & Cruise, and vectored thrust describe the propulsion used



Source: BofA Global Research, EASA

BofA GLOBAL RESEARCH

Exhibit 265: Diversifying attributes to increase addressable markets  
 The key difference between the three relates to the propulsion they use, impacting the range/speeds, and certification/technical requirements

Description	Wingless Multicopter	Lift + Cruise	Vectored Thrust
	Multiple rotors, Thrusters only for lift, cruise via rotor pitch	Independent thrusters for lift & cruise	Same thrusters used for lift & cruise that tilt between vertical & horizontal flight
Range (km)	35-65	50-150	200-300
Speed (km/h)	50-110	50-200	195-320
Benefits	High redundancy, simple controls, low noise, low maintenance, lightweight, low operating cost	Enhanced range and cruise efficiency vs multicopter	Longest range/speed, stabilisation by controlling thrust direction
Challenges	Limited range/low cruise efficiency, low passenger occupancy	Increased complexity, noise & certification requirements vs multicopter	High energy requirements at takeoff and landing, highest complexity for certification (all actuators need to be fully redundant & certified for multiple failures)
Companies (prototype name)	Volocopter (Volocity), Ehang (184)	Wisk (Cora), Embraer (Eve)	Joby (S4), Lilium (Jet), Archer (Maker)

Source: BofA Global Research, EASA, Altran, Volocopter

BofA GLOBAL RESEARCH

Regulating eVTOL: airline safety for urban environments

Certification is a lengthy and thus costly process as this is a new category of aircraft and will operate in higher-population areas, increasing collision risk. Europe's EASA guidance states that not only are eVTOL safety levels expected to match those of commercial airlines – a one in a billion chance of catastrophic failure – but the aircraft need to be able to undertake a controlled emergency landing in such an event. While regions vary slightly on their requirements, the certification to launch eVTOL passenger services commercially will require: 1) a type certificate (TC) indicating an aircraft design is safe and airworthy; and 2) a production certification authorising the manufacture of duplicate aircraft under an approved design.

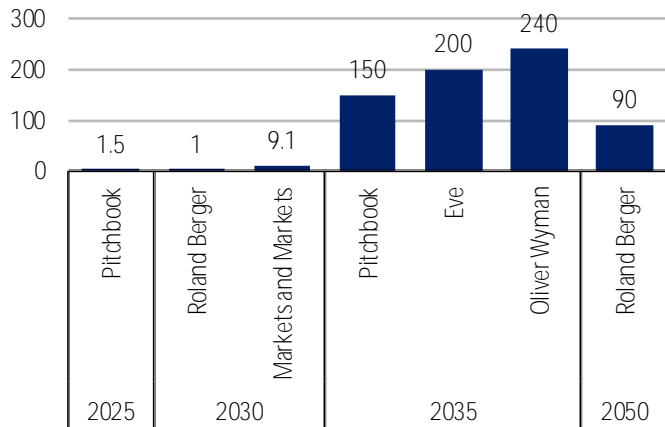
Taxis & airport shuttles obvious targets; >\$250bn by 2035

The addressable market for UAM eVTOL operations could be vast. How quickly that is realised will depend on the mitigation of regulatory, technical and public perception challenges. Thus external industry forecasts for eVTOL vary in their assumptions of revenues and number of aircraft required in the long term. 2025-30 revenue projections range from \$0.5-9bn, but rise rapidly post 2035, from \$90-150bn for passenger services (equivalent to 19% of current global airline revenue, as per Pitchbook), rising to \$250bn when including logistics drone delivery per Oliver Wyman. To achieve that, eVTOL aircraft would need to displace the equivalent of 25% of urban parcel deliveries, 10% of taxi & car rentals, 5% of city-to-city bus transit trips, and 1% of intercity train travel and private car use by 2035.

Per eVTOL manufacturer Volocopter, the key use cases would be 1) business travel (faster point-to-point transfers to airports/central urban locations), 2) tourism (for sightseeing or to travel to specific venues that have eVTOL landing capabilities that could become a differentiator), 3) commuting (from suburban locations), and 4) emergency services (an estimated potential 5,000 eVTOL aircraft market for this segment alone).



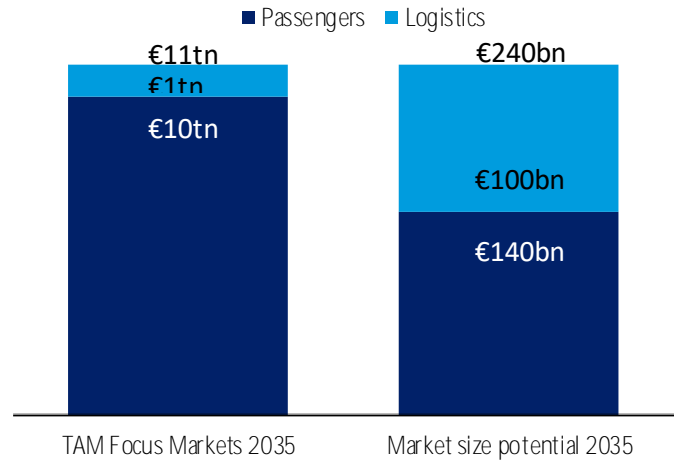
Exhibit 266: Long-term TAM projections range from \$90-240bn  
External forecasts for UAM market vary depending on assumptions on: 1) use cases addressed, 2) penetration/rate of adoption, and 3) utilisation and operating hours



Source: BofA Global Research \$bn UAM industry revenues review of cited sources: Pitchbook, Roland Berger, Markets & Markets, Oliver Wyman/Volocopter, EVE

BofA GLOBAL RESEARCH

Exhibit 267: Urban air mobility addresses a multi-trillion opportunity  
Oliver Wyman expects that UAM will reach market potential of €240bn by 2035



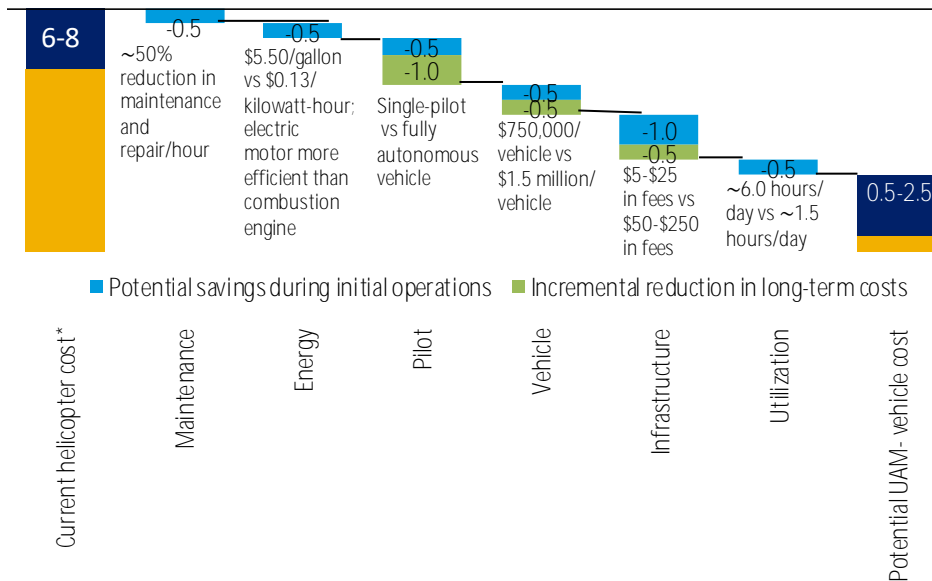
Source: Oliver Wyman

BofA GLOBAL RESEARCH

eVTOL costs could be half those of helicopters

To enable a higher addressable market, costs would need (and are expected) to fall. McKinsey projects UAM operating costs to decline towards \$0.5-2.5 per seat mile – an 80% reduction vs the cost of helicopters today – if further improvements in infrastructure and manufacturing are achieved with higher scale and pilot costs are removed as UAM vehicles become autonomous.

Exhibit 268: Operating costs could evolve for urban-air-mobility vehicles  
Potential evolution in operating cost per seat-mile for urban-air-mobility (UAM) vehicles, \$



Source: McKinsey

\*Current costs vary depending on various factors, including number of passengers and helicopter type

BofA GLOBAL RESEARCH

Physical infrastructure – vertiports the new airports?

To operate UAM vehicles commercially would require places to take off, land, recharge, and undergo any maintenance. These vertipads, bases and hubs could cost up to \$7m to



build and \$17m to operate per McKinsey. The number of locations required would vary by city, but naturally will define the UAM capacity in that location assuming the aircraft are regulated to take off and land only at these approved sites.

Operating ecosystem: to vertically integrate or not?

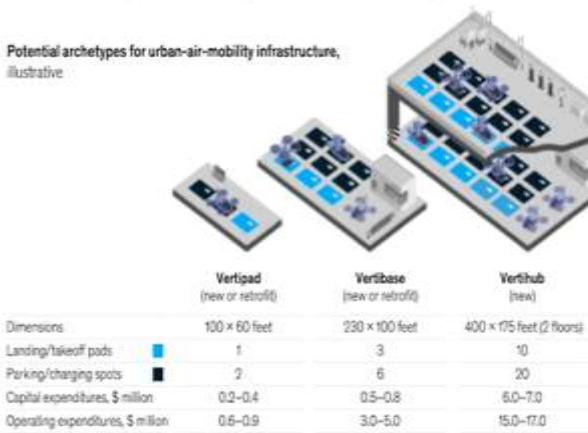
Two diverging business models are emerging: 1) a vertically integrated manufacturer/operator model; and 2) an ecosystem of manufacturers, operators and service providers more akin to current airline/transit offers.

- Vertical integration: Given the nascent stage of the eVTOL industry, some companies are targeting a “full stack” provision of aircraft development and manufacturing, flight operations, air traffic management, and managing customer user experience in the flight and during booking and payment. Partnerships are already being formed to try to exploit this. For example, Volocopter is creating “VololoQ” with Lufthansa and Microsoft to develop a cloud-based software platform to link flight and ground operations, battery management, and customer interface.
- Traditional manufacturer & operator model: As and when more eVTOLs are certified, a more traditional operating dynamic may be more viable to enable replicable and scalable operation in multiple cities. For example, United Airlines, Virgin Atlantic and Azul have all already announced partnerships or committed orders for eVTOL aircraft from Archer, Vertical Aerospace and Lilium, respectively, indicating the separate manufacturer & operator model may take precedence for regional eVTOL flights.

Exhibit 269: New Infrastructure for New Mobility: vertipads, bases, and hubs to enable UAM

Each city would require a different combination of single vehicle locations (Vertipads) with larger base and hub-style locations accommodating multiple vehicles taking off and landing (and waiting) simultaneously

There are three potential archetypes for urban-air-mobility infrastructure.

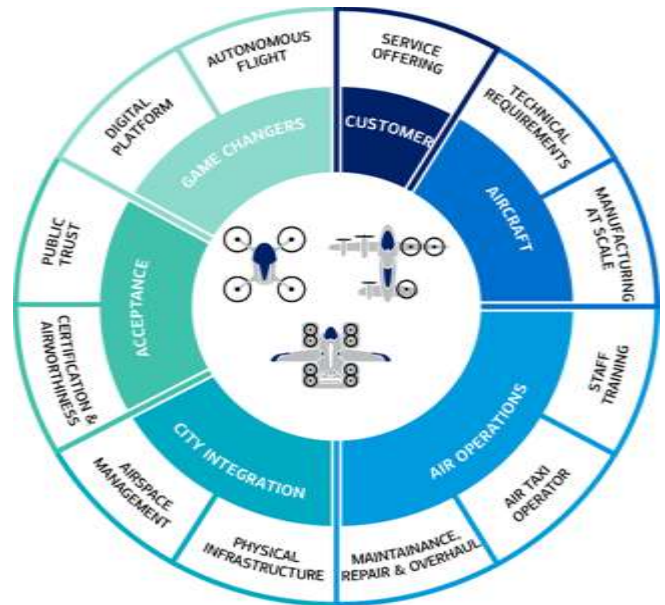


Source: McKinsey & Co

BofA GLOBAL RESEARCH

Exhibit 270: Enabling UAM operating system: aircraft, infrastructure, and user experience

Enabling commercial UAM services will require a combination of aircraft, customer service, city/air traffic infrastructure integration, and customer and regulator acceptance. Autonomy could be a game changer



Source: Volocopter, BofA Global Research

BofA GLOBAL RESEARCH

Challenges to commercialisation

Three key challenges remain to the near-term commercialisation of future air mobility services: 1) type certification: the official stamp that vehicles are safe, hence the heavy testing period required and the timing to obtain the necessary certification path, but also 2) infrastructure: ensuring cities are ready with landing infrastructure and



routes approved by the civil aviation authorities, and 3) public perception – testing in several locations globally, to raise awareness.



## Autonomous Vehicles

What are they? Technology that can assist or fully replace human driving, enabled by a combination of sensors and software.

Did you know? A Level 4 highly automated self-driving system uses up to 46% of an electric vehicle's power consumption, almost halving the usable driving range.

### Autonomous driving – **when, where and how? Depends who you ask...**

Autonomy remains the key tech goal for the mobility industry. Increased processing capability at lower cost and power consumption can enable this, accelerated by advances in machine learning and AI capabilities, albeit with longer timelines than initially hoped.

We highlight the key near-term trends in AVs:

1. ADAS first, full AV later: Increasing levels of autonomous features in private cars through advanced driver assistance systems (ADAS, e.g. self-parking and highway driving initially, with urban self-driving at a later stage) but with the driver either partially in control or as a backup. Achieving further levels of conditional automation remains dependent on regulation, technological approaches (which sensors to use and proven safety) and cost.
2. Sensor wars, LiDAR and/or Radar: an ongoing debate within the industry relates to whether light detection and ranging (LiDAR) sensors are needed to achieve autonomous driving and, if so, whether the cost could come down far enough to accommodate them profitably in privately owned vehicles. LiDAR unit costs are expected to halve in the next few years as volumes ramp up, but improving radar capabilities aided by software could suffice and be far cheaper. Related to sensor choice is the power required to run AVs: efficiency is becoming more important, especially with electric vehicles, to avoid utilising too much battery range.
3. AV1.0 is here: Fully autonomous driving, while taking longer to be commercialised than hoped, is beginning to be deployed for set routes/use cases (e.g. select ridehail cities, trucks), with a consolidating handful of providers piloting AV services. Proprietary data and computation are the key required competencies.
4. AV2.0 = embodied AI in AVs: generative AI advances can enable vehicles to be equipped with onboard intelligence to drive independently (vs the rules-based mapping approaches of most of the current AV programmes), specifically using the emergent AI tech of 1) reinforcement learning on-board cars, 2) synthetic training data to speed up the testing and safety validation work, and 3) multimodality to control and understand more vehicle functions.
5. “Driver as a service” or “AI Defined Vehicles”: Software subscription business models are emerging as several AV tech companies near commercialisation, either via a monthly subscription fee for private drivers, pay per mile for truck fleet operators, or licensing of software/sensor pods to vehicle manufacturers to enable them to create new experiences or “AI defined vehicles”. This extends beyond cars to industrial machinery, e.g. agriculture.

### Reducing the cost and hardware needed for autonomy

Two converging trends could accelerate the development of autonomous vehicle software: 1) utilising next gen AI and computational tools – these could make AVs 5-7x cheaper to launch and scale than using map-based approaches per Wayve, as well as faster, as it would negate the need to remap or recalibrate systems in every new





location or environment, 2) the reducing cost of sensors required to enable autonomous driving owing to rising volumes and using software to improve existing sensor capabilities (e.g. radar). Many connected high-end cars already in production now have the hardware/software integration potential to adopt self-driving technology with an over-the-air software update. While initial deployment is likely in driver assistance features, several companies are now developing technologies. They intend to grow into auto OEM manufacturing cycles towards Level 3-4 autonomous driving, leveraging the hardware already embedded into vehicles or with minimal additional equipment required.

“AI is rapidly surpassing humans at many tasks: driving is the next frontier... Gone are the days where you need a vehicle with five LiDAR, six radar, 30 cameras, and a 50 core CPU, and 4 GPUs. That kind of six figure vehicle cost is unsustainable and not commercially viable. We can operate with simply surround cameras, forward-facing radar, and a single GPU. Most high-end vehicle manufacturers are starting to put that in their vehicles on the road. So there's been a collision point now between the autonomous driving industry and the automotive industry with what is being built into vehicles. And that is enabling amazing commercialization opportunities”. Alex Kendall, CEO, Wayve at BofA Transforming World Conference

### Embodied AI: The ChatGPT moment is coming to driving

While generative AI models are at an early stage of commercialisation, their rapid progress is revealing the emergent capabilities that could accelerate autonomous driving. They could provide vehicles with embodied intelligence to understand what's ahead and act on it in real time, as opposed to the rules and mapping-based approaches adopted in the majority of self-driving vehicle programmes currently in operation. This technology is already being created and validated by a number of self-driving tech companies including Wayve, Ghost Autonomy, and Tesla. The premise is to switch from an engineered system interconnected by lines of code, to a purely data-driven approach with a single neural network (ChatGPT style). To commercialise this approach to self-driving thus requires a combination of algorithms, compute and data as the differentiators. Specifically:

- Generative AI could control the multiple software layers required for self-driving rather than each needing to be programmed individually.
- Synthetic data and simulation could speed up the testing, avoiding sole reliance on road/real-life testing of autonomous fleets to capture data (which is expensive).
- Multimodality – using a combination of text/image/video data inputs can enable operators to prompt vehicles to control/communicate with the car or understand the decisions being made in real-time.

### Autonomous cars need to go on a data diet

Whilst generative AI could accelerate the development of self-driving systems, risks include:

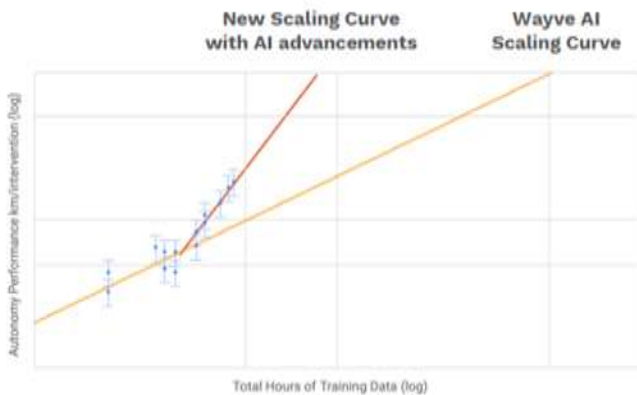
- Hallucination: the potential lack of transparency and controllability of such systems, meaning regulators may lack confidence that they won't 'hallucinate' (make misjudgements or catastrophic errors).
- Data volume/cost: Autonomous cars and control systems would also create vast amounts of data. Estimates vary depending on the sensors and compute deployed; Lucid estimate between 1.4TB/hour and 19TB/hour, Intel estimate 4TB per day, Tuxera estimate between 380-5,100 TB each year would be created. To put that in context, this volume of data at the higher range would fill the largest commercially available supercomputer in 5 hours, and cost up to \$350k per year to store at Amazon's 2023 storage prices.

To mitigate this, several companies are working to create the hardware and storage levels that would be required.



- Several tech hardware developers are creating purpose built chips to bring supercomputing powers to autonomous cars, such as Mobileye (e.g. EyeQ Ultra), Nvidia (DRIVE), Ambarella (CV3), aimed at achieving not only the processing power required but reducing power consumption and increasing efficiency.
- Tesla’s supercomputer (Dojo) was created in 2021 partially for this reason for example, at the time being the 5<sup>th</sup> largest supercomputer in the world per Forbes, with computing power equivalent to 100,000 Nvidia GPUs. This is set to treble by the end of 2024 owing to the increasing volume of data processing and calculations required to enable autonomous driving.

Exhibit 271: AV performance efficiently scales with data  
AI advancements are accelerating Autonomy performance in terms of kms between interventions



Source: Wayve

BofA GLOBAL RESEARCH

Exhibit 272: Multimodality & synthetic data to build trust/safety  
Generative AI models connecting vision, language and action for the first time, and simulating rare edge case scenarios



Source: Wayve

BofA GLOBAL RESEARCH

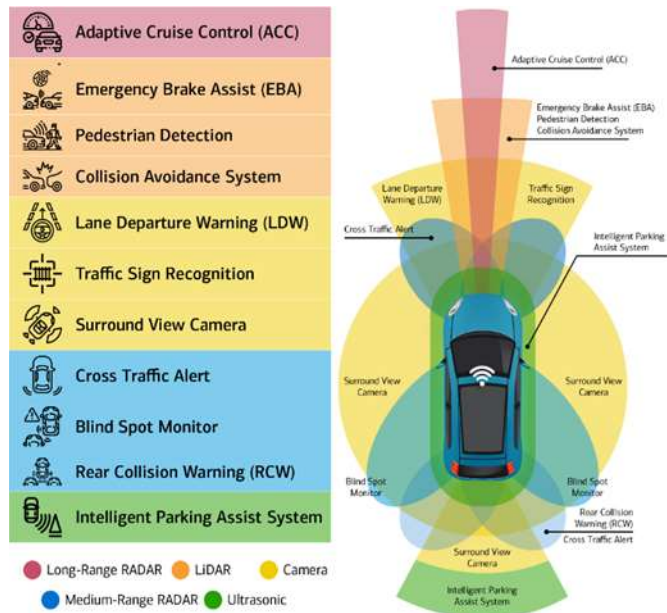
### Sensor wars: Radar vs LiDAR in the race to self-driving

The key hardware components of AV systems include radar, camera, ultrasound and lasers such as LiDAR (Light Detection And Ranging). The functions they enable range from driver alerts to physical intervention such as speed/lane changes and automated parking and self-driving. Market views on LiDAR have been polarised. Tesla, for instance, believes it is unnecessary to achieve full self-driving. German OEMs such as Mercedes/BMW and Auto suppliers such as Aptiv and Valeo suggest it is required in higher levels of automation owing to the limitations of camera/radar-only systems. Mobileye also proposes using both radar/LiDAR to ensure sensors serve as a backup for each other (“True Redundancy”).

Innovation is accelerating on both sides – to improve radar functionality using software and to reduce the cost of LiDAR. Ultimately each use case has different requirements that can be met by a combination of sensor configurations that need to be optimised with software. Emerging generative AI technologies could enable AV systems to operate with much less hardware – most of which is being equipped in the latest high-end cars.



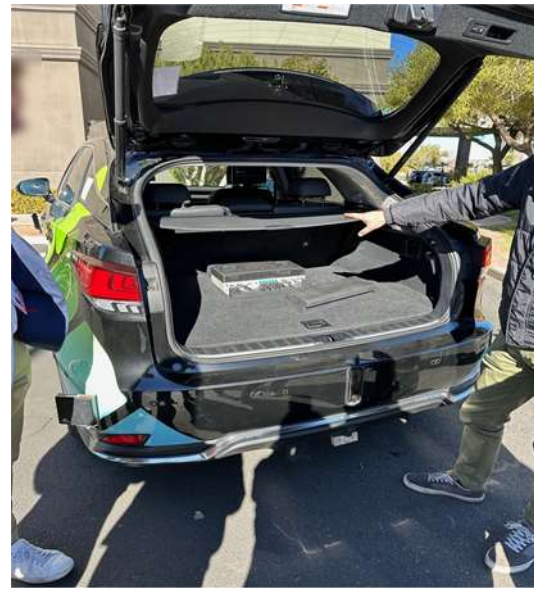
**Exhibit 273: Autonomous Vehicle Sensing Systems**  
 Fusion of Radar, LiDAR, Camera and Ultrasonic sensors targeting safer driving with higher levels of automation



Source: RGSBI

BofA GLOBAL RESEARCH

**Exhibit 274: AV Compute space requirements get smaller**  
 Advancing computational power means self-driving features can be achieved with more compact equipment, demonstrated by Ambarella at CES 2024



Source: BofA Global Research





BofA GLOBAL RESEARCH

More sensors, more safety, but cost and liability questions remain

The exact specifications of ADAS systems and the technological approach vary across providers, which include vehicle manufacturers and suppliers, as well as semiconductor, computing platform and specialised ADAS hardware suppliers. The capabilities are defined by the SAE Levels of Driving Automation – Level 0 being no driving automation through to Level 5 full automation. Levels 0-2 are classified as driver support features where the driver remains in control of a vehicle, with Levels 3-4 offering certain automated features that enable a vehicle to drive itself. The implied number of sensors/equipment – and thus the cost to achieve self driving in terms of content per vehicle – increases with these levels.

Exhibit 275: ADAS feature levels, core products and number of sensory devices required per ADAS level

Camera and radar systems are the most prevalent components at present, but increasing automated driving features require more sensors and software to integrate them

Components & applications	Level 1	Level 2	Level 2+	Level 3	Level 4
	Driver assistance	Partial Automation <small>Higher degree of sensor fusion required</small>		Conditional Automation	High Automation
<b>Camera</b> Sign recognition Lane departure warning Park assist Surround view Driver monitoring e-Mirrors 	2-3	>4	4-6	6-8	6-8
<b>Radar</b> Adaptive cruise control Cross traffic alert Rear collision warning Blind spot detection 	1-3	3-5	5-7	>5	6-10
<b>LIDAR</b> Emergency braking Pedestrian detection Collision avoidance 	-	0	Optional	1-3	1-3
<b>Ultrasonic</b> Park assist 	2-3	2-3	2-3	2-3	2-3
<b>Total sensors</b>	5-9	9-12	11-16	13-19	15-24
<b>Functions</b>	Automatic emergency braking (AEB) Adaptive cruise control (ACC) Blind-spot detection Lane departure warning Lane keep assist Lane change assist Distance control Automatic speed limiting Interaction assistance Collision avoidance	Automatic Emergency Braking (AEB) ACC + Lane keep assist Automatic park assist (APA) Automatic emergency steering (AES) Traffic Jam assist	Highway piloting Lane changes Manoeuvre emergency braking Remote parking Full environmental monitoring	Fully automated driving in defined situations i.e. Automated valet parking	

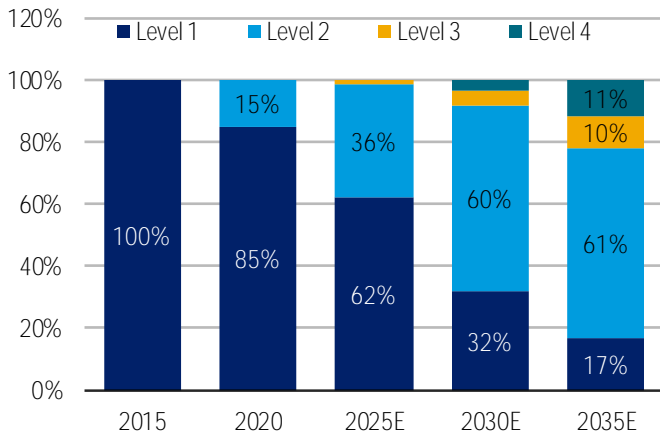
Source: BofA Global Research, Infineon, IHS, NXP

BofA GLOBAL RESEARCH

Self-driving in the long term; half of cars still have no ADAS in the short term  
 While full autonomy remains the long-term aim, implementing even the lowest levels of autonomy presents both a step change in safety potential and a lucrative revenue opportunity for the providers of such systems. In 2020, more than half of new cars produced globally still had zero ADAS content. Of those that did, 85% were still at Level 1. This gives suppliers an opportunity to double content per vehicle by 2025 relative to 2020, with the global ADAS market projected by BofA’s autos team to grow at a c.20% CAGR over 2020-25 to c.US\$32bn.



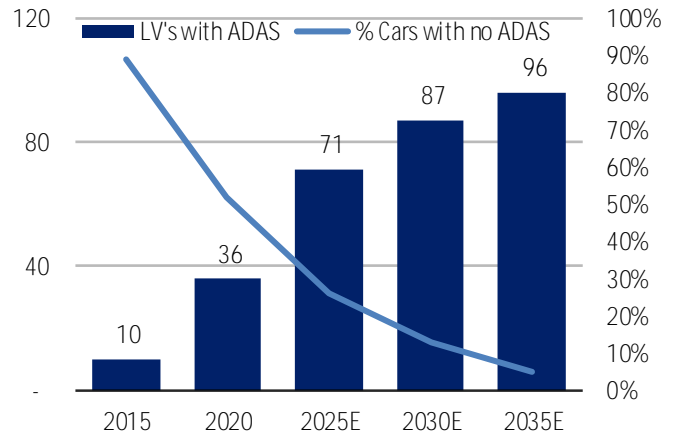
Exhibit 276: Segmentation of ADAS content\* in cars by level L1 and L2 (incl. L2+) to contribute majority (92%) of share until 2030



Source: Infineon

BofA GLOBAL RESEARCH

Exhibit 277: Cars produced with ADAS content (mn units) 52% of vehicles produced in 2020 still had zero ADAS content



Source: BofA Global Research estimates, Infineon, IHS

BofA GLOBAL RESEARCH

### Software Defined Vehicles for better features and faster upgrades

Several auto/tech suppliers have demonstrated visions of 'software defined vehicles' that are set to be a key enabler of self-driving systems, as well as improving the range of features cars can adopt, such as payments or better entertainment. To enable this requires a shift from the largely independent electronic control units (ECUs) that control each tech component in most vehicles today, towards a few high-powered computers controlling multiple functions within cars. This can enable companies to launch new features and content more quickly (requiring software changes rather than hardware), and use less wiring thus save weight/cost, and upgrade cars over time via the air updates.

This shift to 'software defined vehicles' is particularly important in enabling autonomous driving too, firstly the shift to centralised high powered compute can power all automated driving components centrally from one chip making them easier to integrate and programme. More importantly, they can be upgraded over the air remotely as self-driving software improves or bugs can be fixed without down-time impacting safety or commercial potential.

### Trucking may be first to market with commercial AVs

Tech breakthroughs in both the hardware and software needed (e.g. 4D radar, generative AI and edge computing) and a clear regulatory pathway may be required to achieve full self-driving at scale. However, some early applications are near commercial readiness such as trucks with closed loop applications and defined road transport corridors expected between 2025 and 2027 (Kodiak, Aurora).





Exhibit 278: Aurora & Continental targeting OEM integration  
Aurora is targeting a “driver as a service” pay per mile for its autonomous trucking business model, aiming for commercial production in 2027



Source: BofA Global Research

BofA GLOBAL RESEARCH

Exhibit 279: Kodiak self-driving trucks targeting fleets/truck customers  
Kodiak has 34 autonomous trucks driving pilot routes currently, aiming for commercial launch for select routes/customers in 2025



Source: BofA Global Research

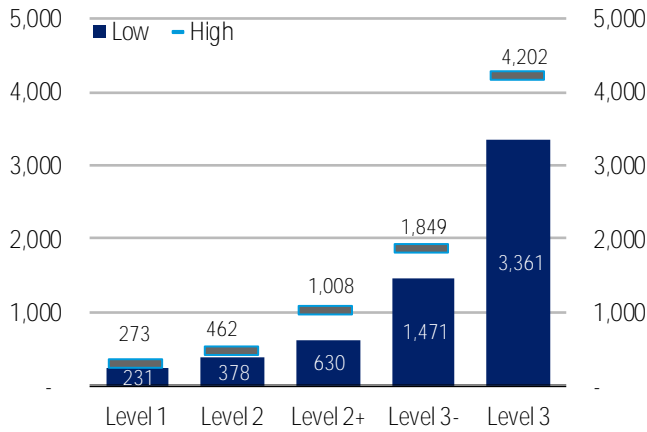
BofA GLOBAL RESEARCH

Challenges remain: regulation, cost, power consumption and cybersecurity  
Several challenges remain prior to full-scale commercialisation of autonomous driving, centred around regulation, reducing cost, power and cybersecurity threats.

- Regulation: Crossing to Level 3 automation is challenging owing to legislative and liability questions over such technologies. Finding a solution is complex and encompasses a wide range of factors including AI decision-making parameters, cybersecurity, insurance cover, and consent for monitoring. The ideologies around these factors vary by region and government. As such, current deployment of ADAS technologies has been limited to Levels 1-2 thus far, with higher penetration of Levels 3+ limited to premium vehicles barring meaningful price reductions.
- Cost: Far greater sensing perception is required for higher levels of autonomy. The number of sensors doubles at Level 2 and trebles at Level 2+, and is likely to rise much higher as partial/full autonomous features are added to cars. Level 1 ADAS requires c.US\$300 of sensing equipment and a Level 2-enabled vehicle c.US\$1k, but Level 3 jumps to >US\$4,000. This equates to 30% of the average bill of materials of an entire vehicle, severely impacting the affordability of higher levels of autonomy in the short term. This largely reflects the addition of LiDAR. Lower-cost/spec products are under development, but in the meantime AV tech developers are focusing on bringing cheaper self-driving software to market that can be commercialised in private cars initially.



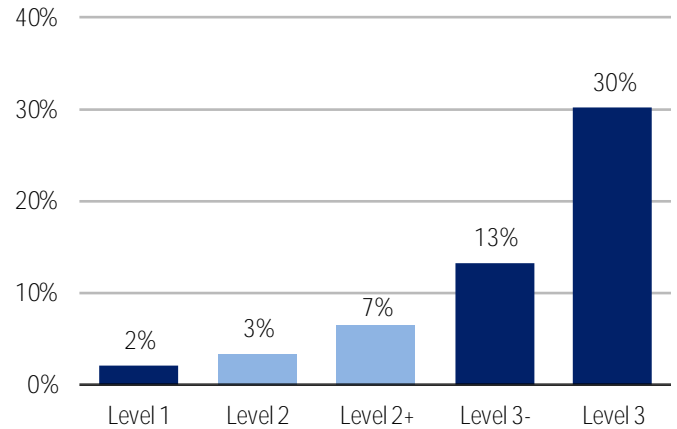
Exhibit 280: Content per vehicle range (EUR) by ADAS level  
Step-up change in cost from L3 upwards



Source: BofA Global Research estimates, Aptiv

BofA GLOBAL RESEARCH

Exhibit 281: ADAS content as a % of average vehicle bill of materials  
Level 2 and 2+ ADAS appears to be the sweet spot. At L3 affordability begins to be more severely impacted

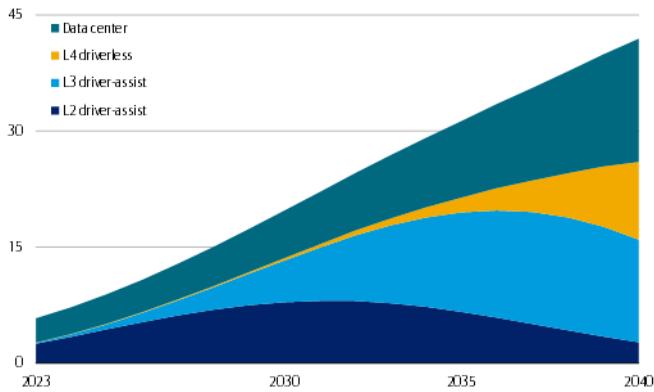


Source: BofA Global Research estimates

BofA GLOBAL RESEARCH

- Power consumption: limiting the battery drain from autonomous driving computation is becoming an increasing challenge for AV developers. As autonomous driving tasks and computation become more complex, the energy required to power them increases. BNEF estimates electricity used for AV capabilities could grow 7x by 2040. An EV equipped with an L4 autonomous driving system would use up to 46% of its power consumption, almost halving an EV range. A number of chipmakers are developing new hardware to simultaneously improve efficiency and compute.

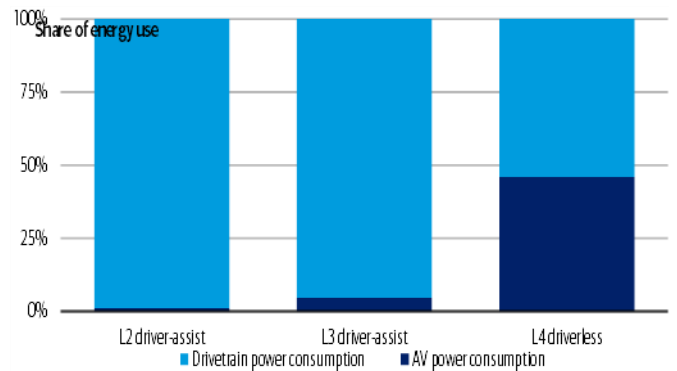
Exhibit 282: Electricity used for AV capabilities could grow 7x by 2040  
In-vehicle computers could consume 26TWh in 2040, with a further 16TWh in data centres



Source: BNEF, Terawatt hours (TWh) electricity consumption

BofA GLOBAL RESEARCH

Exhibit 283: **AV compute could consume nearly half of an EV's range**  
An EV equipped with an L4 autonomous driving system could "lose" up to 46% of its range due to the power required for the AV system



Source: BNEF

BofA GLOBAL RESEARCH

Cybersecurity: ADAS and software defined vehicle architectures could require 300-500m lines of software code per Roland Berger (vs cars running with an average of 10m in 2010). This dramatically raises the potential for cyberattacks – not only on the car itself but on all components of its ecosystem, e.g., back end, infrastructure. The automotive cybersecurity market is projected to double over 2020-30 from US\$5bn to US\$10bn amid efforts to mitigate against this increased threat, per McKinsey.



# Disclosures

## Important Disclosures

**FUNDAMENTAL EQUITY OPINION KEY:** Opinions include a Volatility Risk Rating, an Investment Rating and an Income Rating. **VOLATILITY RISK RATINGS**, indicators of potential price fluctuation, are: A - Low, B - Medium and C - High. **INVESTMENT RATINGS** reflect the analyst's assessment of both a stock's absolute total return potential as well as its attractiveness for investment relative to other stocks within its Coverage Cluster (defined below). Our investment ratings are: 1 - Buy stocks are expected to have a total return of at least 10% and are the most attractive stocks in the coverage cluster; 2 - Neutral stocks are expected to remain flat or increase in value and are less attractive than Buy rated stocks and 3 - Underperform stocks are the least attractive stocks in a coverage cluster. An investment rating of 6 (No Rating) indicates that a stock is no longer trading on the basis of fundamentals. Analysts assign investment ratings considering, among other things, the 0-12 month total return expectation for a stock and the firm's guidelines for ratings dispersions (shown in the table below). The current price objective for a stock should be referenced to better understand the total return expectation at any given time. The price objective reflects the analyst's view of the potential price appreciation (depreciation).

Investment rating	Total return expectation (within 12-month period of date of initial rating)	Ratings dispersion guidelines for coverage cluster <sup>R1</sup>
Buy	≥ 10%	≤ 70%
Neutral	≥ 0%	≤ 30%
Underperform	N/A	≥ 20%

<sup>R1</sup> Ratings dispersions may vary from time to time where BofA Global Research believes it better reflects the investment prospects of stocks in a Coverage Cluster.

**INCOME RATINGS**, indicators of potential cash dividends, are: 7 - same/higher (dividend considered to be secure), 8 - same/lower (dividend not considered to be secure) and 9 - pays no cash dividend. Coverage Cluster is comprised of stocks covered by a single analyst or two or more analysts sharing a common industry, sector, region or other classification(s). A stock's coverage cluster is included in the most recent BofA Global Research report referencing the stock.

Due to the nature of strategic analysis, the issuers or securities recommended or discussed in this report are not continuously followed. Accordingly, investors must regard this report as providing stand-alone analysis and should not expect continuing analysis or additional reports relating to such issuers and/or securities.

BofA Global Research personnel (including the analyst(s) responsible for this report) receive compensation based upon, among other factors, the overall profitability of Bank of America Corporation, including profits derived from investment banking. The analyst(s) responsible for this report may also receive compensation based upon, among other factors, the overall profitability of the Bank's sales and trading businesses relating to the class of securities or financial instruments for which such analyst is responsible.

## Other Important Disclosures

From time to time research analysts conduct site visits of covered issuers. BofA Global Research policies prohibit research analysts from accepting payment or reimbursement for travel expenses from the issuer for such visits.

Prices are indicative and for information purposes only. Except as otherwise stated in the report, for any recommendation in relation to an equity security, the price referenced is the publicly traded price of the security as of close of business on the day prior to the date of the report or, if the report is published during intraday trading, the price referenced is indicative of the traded price as of the date and time of the report and in relation to a debt security (including equity preferred and CDS), prices are indicative as of the date and time of the report and are from various sources including BofA Securities trading desks.

The date and time of completion of the production of any recommendation in this report shall be the date and time of dissemination of this report as recorded in the report timestamp.

<sup>R1</sup> Recipients who are not institutional investors or market professionals should seek the advice of their independent financial advisor before considering information in this report in connection with any investment decision, or for a necessary explanation of its contents.

Officers of BofAS or one or more of its affiliates (other than research analysts) may have a financial interest in securities of the issuer(s) or in related investments.

Refer to [BofA Global Research policies relating to conflicts of interest](#).

**"BofA Securities" includes BofA Securities, Inc. ("BofAS") and its affiliates. Investors should contact their BofA Securities representative or Merrill Global Wealth Management financial advisor if they have questions concerning this report or concerning the appropriateness of any investment idea described here in for such investor. "BofA Securities" is a global brand for BofA Global Research.**

**Information relating to Non-US affiliates of BofA Securities and Distribution of Affiliate Research Reports:**

BofAS and/or Merrill Lynch, Pierce, Fenner & Smith Incorporated ("MLPF&S") may in the future distribute, information of the following non-US affiliates in the US (short name: legal name, regulator): Merrill Lynch (South Africa): Merrill Lynch South Africa (Pty) Ltd., regulated by The Financial Service Board; MLI (UK): Merrill Lynch International, regulated by the Financial Conduct Authority (FCA) and the Prudential Regulation Authority (PRA); BofASE (France): BofA Securities Europe SA is authorized by the Autorité de Contrôle Prudentiel et de Résolution (ACPR) and regulated by the ACPR and the Autorité des Marchés Financiers (AMF). BofA Securities Europe SA ("BofASE") with registered address at 51, rue La Boétie, 75008 Paris is registered under no 842 602 690 RCS Paris. In accordance with the provisions of French Code Monétaire et Financier (Monetary and Financial Code), BofASE is an établissement de crédit et d'investissement (credit and investment institution) that is authorised and supervised by the European Central Bank and the Autorité de Contrôle Prudentiel et de Résolution (ACPR) and regulated by the ACPR and the Autorité des Marchés Financiers. BofASE's share capital can be found at [www.bofam.com/BofASEdisclaimer](http://www.bofam.com/BofASEdisclaimer); BofA Europe (Milan): Bank of America Europe Designated Activity Company, Milan Branch, regulated by the Bank of Italy, the European Central Bank (ECB) and the Central Bank of Ireland (CBI); BofA Europe (Frankfurt): Bank of America Europe Designated Activity Company, Frankfurt Branch regulated by BaFin, the ECB and the CBI; BofA Europe (Madrid): Bank of America Europe Designated Activity Company, Sucursal en España, regulated by the Bank of Spain, the ECB and the CBI; Merrill Lynch (Australia): Merrill Lynch Equities (Australia) Limited, regulated by the Australian Securities and Investments Commission; Merrill Lynch (Hong Kong): Merrill Lynch (Asia Pacific) Limited, regulated by the Hong Kong Securities and Futures Commission (HKSF); Merrill Lynch (Singapore): Merrill Lynch (Singapore) Pte Ltd, regulated by the Monetary Authority of Singapore (MAS); Merrill Lynch (Canada): Merrill Lynch Canada Inc, regulated by the Canadian Investment Regulatory Organization; Merrill Lynch (Mexico): Merrill Lynch Mexico, SA de CV, Casa de Bolsa, regulated by the Comisión Nacional Bancaria y de Valores; BofAS Japan: BofA Securities Japan Co., Ltd., regulated by the Financial Services Agency; Merrill Lynch (Seoul): Merrill Lynch International, LLC Seoul Branch, regulated by the Financial Supervisory Service; Merrill Lynch (Taiwan): Merrill Lynch Securities (Taiwan) Ltd., regulated by the Securities and Futures Bureau; BofAS India: BofA Securities India Limited, regulated by the Securities and Exchange Board of India (SEBI); Merrill Lynch (Israel): Merrill Lynch Israel Limited, regulated by Israel Securities Authority; Merrill Lynch (DIFC): Merrill Lynch International (DIFC Branch), regulated by the Dubai Financial Services Authority (DFSA); Merrill Lynch (Brazil): Merrill Lynch SA. Corretora de Títulos e Valores Mobiliários, regulated by Comissão de Valores Mobiliários; Merrill Lynch KSA Company: Merrill Lynch Kingdom of Saudi Arabia Company, regulated by the Capital Market Authority. This information: has been approved for publication and is distributed in the United Kingdom (UK) to professional clients and eligible counterparties (as each is defined in the rules of the FCA and the PRA) by MLI (UK), which is authorized by the PRA and regulated by the FCA and the PRA - details about the extent of our regulation by the FCA and PRA are available from us on request; has been approved for publication and is distributed in the European Economic Area (EEA) by BofASE (France), which is authorized by the ACPR and regulated by the ACPR and the AMF; has been considered and distributed in Japan by BofAS Japan, a registered securities dealer under the Financial Instruments and Exchange Act in Japan, or its permitted affiliates; is issued and distributed in Hong Kong by Merrill Lynch (Hong Kong) which is regulated by HKSF; is issued and distributed in Taiwan by Merrill Lynch (Taiwan); is issued and distributed in India by BofAS India; and is issued and distributed in Singapore to institutional investors and/or accredited investors (each as defined under the Financial Advisers Regulations) by Merrill Lynch (Singapore)



(Company Registration No 198602883D). Merrill Lynch (Singapore) is regulated by MAS. Merrill Lynch Equities (Australia) Limited (ABN 65 006 276 795), AFS License 235132 (MLEA) distributes this information in Australia only to 'Wholesale' clients as defined by s761G of the Corporations Act 2001. With the exception of Bank of America N.A., Australia Branch, neither MLEA nor any of its affiliates involved in preparing this information is an Authorised Deposit-Taking Institution under the Banking Act 1959 nor regulated by the Australian Prudential Regulation Authority. No approval is required for publication or distribution of this information in Brazil and its local distribution is by Merrill Lynch (Brazil) in accordance with applicable regulations. Merrill Lynch (DIFC) is authorized and regulated by the DFSA. Information prepared and issued by Merrill Lynch (DIFC) is done so in accordance with the requirements of the DFSA conduct of business rules. BofA Europe (Frankfurt) distributes this information in Germany and is regulated by BaFin, the ECB and the CBI. BofA Securities entities, including BofA Europe and BofASE (France), may outsource/delegate the marketing and/or provision of certain research services or aspects of research services to other branches or members of the BofA Securities group. You may be contacted by a different BofA Securities entity acting for and on behalf of your service provider where permitted by applicable law. This does not change your service provider. Please refer to the [Electronic Communications Disclaimers](#) for further information.

This information has been prepared and issued by BofAS and/or one or more of its non-US affiliates. The author(s) of this information may not be licensed to carry on regulated activities in your jurisdiction and, if not licensed, do not hold themselves out as being able to do so. BofAS and/or MLPF&S is the distributor of this information in the US and accepts full responsibility for information distributed to BofAS and/or MLPF&S clients in the US by its non-US affiliates. Any US person receiving this information and wishing to effect any transaction in any security discussed herein should do so through BofAS and/or MLPF&S and not such foreign affiliates. Hong Kong recipients of this information should contact Merrill Lynch (Asia Pacific) Limited in respect of any matters relating to dealing in securities or provision of specific advice on securities or any other matters arising from, or in connection with, this information. Singapore recipients of this information should contact Merrill Lynch (Singapore) Pte Ltd in respect of any matters arising from, or in connection with, this information. For clients that are not accredited investors, expert investors or institutional investors Merrill Lynch (Singapore) Pte Ltd accepts full responsibility for the contents of this information distributed to such clients in Singapore.

#### General Investment Related Disclosures:

Taiwan Readers: Neither the information nor any opinion expressed herein constitutes an offer or a solicitation of an offer to transact in any securities or other financial instrument. No part of this report may be used or reproduced or quoted in any manner whatsoever in Taiwan by the press or any other person without the express written consent of BofA Securities.

This document provides general information only, and has been prepared for, and is intended for general distribution to, BofA Securities clients. Neither the information nor any opinion expressed constitutes an offer or an invitation to make an offer, to buy or sell any securities or other financial instrument or any derivative related to such securities or instruments (e.g. options, futures, warrants, and contracts for differences). This document is not intended to provide personal investment advice and it does not take into account the specific investment objectives, financial situation and the particular needs of, and is not directed to, any specific person(s). This document and its content do not constitute, and should not be considered to constitute, investment advice for purposes of ERISA, the US tax code, the Investment Advisers Act or otherwise. Investors should seek financial advice regarding the appropriateness of investing in financial instruments and implementing investment strategies discussed or recommended in this document and should understand that statements regarding future prospects may not be realized. Any decision to purchase or subscribe for securities in any offering must be based solely on existing public information on such security or the information in the prospectus or other offering document issued in connection with such offering, and not on this document.

Securities and other financial instruments referred to herein, or recommended, offered or sold by BofA Securities, are not insured by the Federal Deposit Insurance Corporation and are not deposits or other obligations of any insured depository institution (including Bank of America, N.A.). Investments in general and, derivatives, in particular, involve numerous risks, including, among others, market risk, counterparty default risk and liquidity risk. No security, financial instrument or derivative is suitable for all investors. Digital assets are extremely speculative, volatile and are largely unregulated. In some cases, securities and other financial instruments may be difficult to value or sell and reliable information about the value or risks related to the security or financial instrument may be difficult to obtain. Investors should note that income from such securities and other financial instruments, if any, may fluctuate and that price or value of such securities and instruments may rise or fall and, in some cases, investors may lose their entire principal investment. Past performance is not necessarily a guide to future performance. Levels and basis for taxation may change.

This report may contain a short-term trading idea or recommendation, which highlights a specific near-term catalyst or event impacting the issuer or the market that is anticipated to have a short-term price impact on the equity securities of the issuer. Short-term trading ideas and recommendations are different from and do not affect a stock's fundamental equity rating, which reflects both a longer term total return expectation and attractiveness for investment relative to other stocks within its Coverage Cluster. Short-term trading ideas and recommendations may be more or less positive than a stock's fundamental equity rating.

BofA Securities is aware that the implementation of the ideas expressed in this report may depend upon an investor's ability to "short" securities or other financial instruments and that such action may be limited by regulations prohibiting or restricting "shortselling" in many jurisdictions. Investors are urged to seek advice regarding the applicability of such regulations prior to executing any short idea contained in this report.

Foreign currency rates of exchange may adversely affect the value, price or income of any security or financial instrument mentioned herein. Investors in such securities and instruments, including ADRs, effectively assume currency risk.

BofAS or one of its affiliates is a regular issuer of traded financial instruments linked to securities that may have been recommended in this report. BofAS or one of its affiliates may, at any time, hold a trading position (long or short) in the securities and financial instruments discussed in this report.

BofA Securities, through business units other than BofA Global Research, may have issued and may in the future issue trading ideas or recommendations that are inconsistent with, and reach different conclusions from, the information presented herein. Such ideas or recommendations may reflect different time frames, assumptions, views and analytical methods of the persons who prepared them, and BofA Securities is under no obligation to ensure that such other trading ideas or recommendations are brought to the attention of any recipient of this information.

In the event that the recipient received this information pursuant to a contract between the recipient and BofAS for the provision of research services for a separate fee, and in connection therewith BofAS may be deemed to be acting as an investment adviser, such status relates, if at all, solely to the person with whom BofAS has contracted directly and does not extend beyond the delivery of this report (unless otherwise agreed specifically in writing by BofAS). If such recipient uses the services of BofAS in connection with the sale or purchase of a security referred to herein, BofAS may act as principal for its own account or as agent for another person. BofAS is and continues to act solely as a broker-dealer in connection with the execution of any transactions including transactions in any securities referred to herein.

#### Copyright and General Information:

Copyright 2024 Bank of America Corporation. All rights reserved. iQDatabase® is a registered service mark of Bank of America Corporation. This information is prepared for the use of BofA Securities clients and may not be redistributed, retransmitted or disclosed, in whole or in part, or in any form or manner, without the express written consent of BofA Securities. BofA Global Research information is distributed simultaneously to internal and client websites and other portals by BofA Securities and is not publicly-available material. Any unauthorized use or disclosure is prohibited. Receipt and review of this information constitutes your agreement not to redistribute, retransmit, or disclose to others the contents, opinions, conclusion, or information contained herein (including any investment recommendations, estimates or price targets) without first obtaining express permission from an authorized officer of BofA Securities.

Materials prepared by BofA Global Research personnel are based on public information. Facts and views presented in this material have not been reviewed by, and may not reflect information known to, professionals in other business areas of BofA Securities, including investment banking personnel. BofA Securities has established information barriers between BofA Global Research and certain business groups. As a result, BofA Securities does not disclose certain client relationships with, or compensation received from, such issuers. To the extent this material discusses any legal proceeding or issues, it has not been prepared as nor is it intended to express any legal conclusion, opinion or advice. Investors should consult their own legal advisers as to issues of law relating to the subject matter of this material. BofA Global Research personnel's knowledge of legal proceedings in which any BofA Securities entity and/or its directors, officers and employees may be plaintiffs, defendants, co-defendants or co-plaintiffs with or involving issuers mentioned in this material is based on public information. Facts and views presented in this material that relate to any such proceedings have not been reviewed by, discussed with, and may not reflect information known to, professionals in other business areas of BofA Securities in connection with the legal proceedings or matters relevant to such proceedings.

This information has been prepared independently of any issuer of securities mentioned herein and not in connection with any proposed offering of securities or as agent of any issuer of any securities. None of BofAS or any of its affiliates or their research analysts has any authority whatsoever to make any representation or warranty on behalf of the issuer(s). BofA Global Research policy prohibits research personnel from disclosing a recommendation, investment rating, or investment thesis for review by an issuer prior to the publication of a research report containing such rating, recommendation or investment thesis.

Any information relating to the tax status of financial instruments discussed herein is not intended to provide tax advice or to be used by anyone to provide tax advice. Investors are urged to seek tax advice based on their particular circumstances from an independent tax professional.

The information herein (other than disclosure information relating to BofA Securities and its affiliates) was obtained from various sources and we do not guarantee its accuracy. This information may contain links to third-party websites. BofA Securities is not responsible for the content of any third-party website or any linked content contained in a third-party website. Content contained on such third-party websites is not part of this information and is not incorporated by reference. The inclusion of a link does not imply any endorsement by or any affiliation with BofA



Securities. Access to any third-party website is at your own risk, and you should always review the terms and privacy policies at third-party websites before submitting any personal information to them. BofA Securities is not responsible for such terms and privacy policies and expressly disclaims any liability for them.

All opinions, projections and estimates constitute the judgment of the author as of the date of publication and are subject to change without notice. Prices also are subject to change without notice. BofA Securities is under no obligation to update this information and BofA Securities ability to publish information on the subject issuer(s) in the future is subject to applicable quiet periods. You should therefore assume that BofA Securities will not update any fact, circumstance or opinion contained herein.

Certain outstanding reports or investment opinions relating to securities, financial instruments and/or issuers may no longer be current. Always refer to the most recent research report relating to an issuer prior to making an investment decision.

In some cases, an issuer may be classified as Restricted or may be Under Review or Extended Review. In each case, investors should consider any investment opinion relating to such issuer (or its security and/or financial instruments) to be suspended or withdrawn and should not rely on the analyses and investment opinion(s) pertaining to such issuer (or its securities and/or financial instruments) nor should the analyses or opinion(s) be considered a solicitation of any kind. Sales persons and financial advisors affiliated with BofAS or any of its affiliates may not solicit purchases of securities or financial instruments that are Restricted or Under Review and may only solicit securities under Extended Review in accordance with firm policies.

Neither BofA Securities nor any officer or employee of BofA Securities accepts any liability whatsoever for any direct, indirect or consequential damages or losses arising from any use of this information.

