BrainChip Holdings Ltd

Revolutionary AI technology

Share Price: A$0.20

ASX: BRN

Sector: Technology

25 June 2024

BrainChip (ASX:BRN), listed on ASX, is commercialising a revolutionary neuromorphic technology. This processor, called Akida, is proprietary intellectual property (IP) added to a computer chip, enabling the chip to function in a similar way to the biological brain for AI inference. Akida is unique in its ability to efficiently process data by taking advantage of sparsity, which eliminates unnecessary computation, improving power efficiency and performance. As a result, Akida doesn't need a continuous internet connection, which dramatically reduces latency. In other words, decisions can be made much faster compared to today's, software-based, energy-guzzling Artificial Intelligence (AI) solutions that run in the Cloud. And because Akida processes data in the same way the human brain does, the chip can learn autonomously, even when it's already deployed "in the field", resulting in ongoing improvements in outcomes.

Market cap. (A$ m)	371.2
# shares outstanding (m)	1,856
# shares fully diluted (m)	2,025
Market cap ful. dil. (A$ m)	405.0
Free float	100.0%
52-week high/low (A$)	0.49 / 0.15
Avg. 12M daily volume ('1000)	10,357
Website	www.brainchip.com

Source: Company, Pitt Street Research

Akida has immense potential

Because of its specific advantages, i.e. near-zero latency, extremely low energy consumption and autonomous on-chip learning, Akida has very substantial potential across a range of applications, especially in Edge AI. These include autonomous vehicles, drones, robotics, medical diagnostics, i.e. sensors that reside at the Edge of the Internet of Things.

Share price (A$) and avg. daily volume (k, r.h.s.)

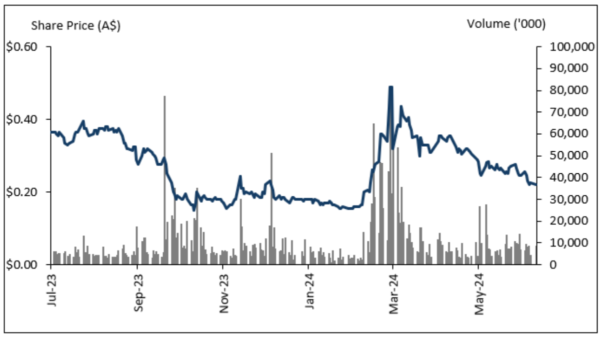

Source: Refinitiv Eikon, Pitt Street Research

The world is Akida's oyster

Although investors have become increasingly aware of the massive opportunity that AI presents, e.g. through the rapid rise of ChatGPT and NVIDIA's stellar growth, AI is still very much in its early stages. We believe the commercial opportunity for BrainChip is very substantial indeed, specifically because the technology underlying Akida is radically different compared to today's AI solutions and addresses the very large Edge AI market. We expect Akida to be able to provide AI capabilities to countless types of devices where previously this wasn't possible due to the restrictions of Cloud-based AI (cost, latency, energy consumption etc).

Valuation metrics	
Fair valuation range (per share)	A$1.59
Methodology	M&A metrics

Source: Pitt Street Research

Valuation of A$1.59 per share

We have valued BRN at A$1.59 per fully diluted share, based on industry M&A activity (please see page 19 for more detail). Investors have been awaiting additional commercial deals in the last few years and, to be fair, their patience has been tested. But we believe these types of deals will be the future catalysts for BrainChip's share price. And because of the very broad spectrum of potential applications for Akida and BrainChip's many ongoing commercial discussions with prospects, we are confident investors' patience will be rewarded. Please see page 20 for key investment risks.

Analysts: Marc Kennis, Nick Sundich

Tel: +61 (0)4 3483 8134

Marc.Kennis@pittstreetresearch.com

Nick.Sundich@pittstreetresearch.com

Disclosure: Pitt Street Research directors own shares in BrainChip.

Table of Contents

The BrainChip investment case

BrainChip is developing and commercialising its Akida Intellectual Property.

BrainChip (ASX:BRN) is developing and commercialising its Akida technology, IP (intellectual Property) for neuromorphic Systems-on-a-Chip (SOC). Akida is a technology added to a computer chip, enabling it to process data similar to the biological brain rather than how a conventional computer chip processes data.

Akida is a radically different approach to Artificial Intelligence

Rather than following a set program, Akida can work autonomously, like a human brain.

Neuromorphic processing emulates the structure and functions of neurons in nature, such as the human brain. Rather than relying on externally-set, software-based algorithms, Akida can convert traditional networks, like CNNs into Spiking Neural Networks (SNN). This allows customers to run their current networks, while taking advantage of our architecture. Spiking Neural Networks learn from events (the inputs, or spikes) it gets by attributing higher or lower importance (synaptic weights) to different spikes, similar to how the biological brain does that. BrainChip's Akida technology supports a variety of industry standard models, like CNN's.

Akida keeps getting better over time.

Akida provides the ability to add another element to existing trained data sets. This is a very powerful feature because it means that processing outcomes can continue to improve even after the chips have already been integrated into their respective applications.

TENNs are 50x smaller than CNN's

BrainChip has also developed a new type of Neural Network Architecture, called a Temporal Event Neural Network or TENNs. While not strictly an SNN, it captures event data in a similar way.

The predominant architecture for image processing are CNNs. In a recent paper BrainChip showed they can build networks of similar accuracy to CNNs, but 50 times smaller.

More efficient than ChatGPT

Additionally, the company recently presented at the Edge AI Vision Summit where it talked about a small LLM (Large Language Model) it built. When compared to a model built with the same basic architecture as ChatGPT (the Transformer), the BrainChip model was 5 times smaller and used 5,000 times less energy.

In another paper recently published on Eye Tracking, BrainChip showed state of the art results with 90% sparsity. Sparsity is a key component to how the human brain efficiently processes data. 90% sparsity means that 90% of the computation isn't required, further improving power efficiency and performance.

The combination of Akida's ability to process event data, while taking advantage of sparsity with the TENNs architecture could allow them to build solutions that are orders of magnitude more efficient than what's currently available today.

Those readers who which to learn more about the intricacies of neuromorphic computing, the benefits of SNN's over other types of neural networks and specifically how BrainChip's technology works, we'd like to refer to Appendices II, III and IV.

Key take aways about Akida

Extremely low latency and much lower power consumption.

1. Akida processes events on-chip, i.e. at the Edge of the Internet. So, it doesn't have to communicate with some far-away data center to access an algorithm to process the inputs the chip just received. Therefore, Akida doesn't necessarily need an internet connection to function. This aspect of Akida reduces latency[1] very substantially further aided by the fact that Akida processes data in parallel rather than sequentially.

   In other words, the chip's response time improves dramatically … very useful, for instance, in a self-driving vehicle that needs to avoid a pedestrian. But this feature is extremely useful for basically any Edge device.
2. Processing data at the Edge of the internet also preserves privacy by not having to send sensitive data to the Cloud.
3. Akida is also much less power-hungry compared to other, Cloud-based, neural networks. It requires only fractions of a watt to operate, which is extremely beneficial in Edge applications where battery life is a major issue.

Akida ticks all the boxes when it comes to requirements for successful Edge AI applications: More secure, real-time analytics at higher speeds and lower power consumption.

Where Akida shines

A wide range of application areas.

Akida has many potential applications and end markets. As we will outline in greater detail in later sections, these include aerospace and defence, industrial manufacturing, healthcare, smart homes and smart cities, automotive applications as well as consumer devices where TENNs architecture is showing exciting results.

Current solutions are either less efficient or more power-hungry, or both.

Many of these market opportunities for Akida are either not possible with today's computer chips or are only possible using substantial processing power and data resources. Most of the chips used in Edge applications today are power-hungry, general-purpose chips - such as Graphics Processing Units (GPUs) that consume orders of magnitude more power than Akida to solve the same problem.

BrainChip is ready to meet the emerging Edge AI opportunity

BrainChip has significant opportunity in Edge AI.

We believe BrainChip has created a significant market opportunity for itself with the development of Akida. Despite the promise of AI, the release of prominent products, such as ChatGPT, and continued development of solutions, such as autonomous cars that have given a glimpse of AI's capabilities, AI is really only at its early stages of development. Edge AI in particular is where we see very substantial, long-term growth and, in our view, Akida is ideally suited to meet the specific requirements of Edge AI.

We value BrainChip at A$1.59 per share

Using the Habana acquisition by Intel in 2019 as a guide to value BrainChip, we arrive at a value of A$1.59 per fully diluted BrainChip share (AUD/USD = 0.66). Please see page 17 for more detail around our valuation methodology.

[1] The delay that happens between when an action on a network or web application is taken and when it reaches its destination (and back).

The share price can potentially respond very strongly to positive news.

Even though this is substantially higher than today's share price, BrainChip shares have been trading as high as A$2.34 in January 2022. Back then, the share price was driven by news around a major car manfacturing using Akida in one of its concept cars. We believe this illustrates how well the share price can potentially respond to good news.

Multiple potential catalysts for BrainChip's share price

Investors are looking for new deals and revenues from existing ones.

We see several catalysts that can create shareholder value for BrainChip shareholders in the near future. The one thing many investors are waiting for is the company's first recurring revenues. This will require transitioning current commercial deals, e.g. with MegaChips, into the royalty generating phase, which requires Akida being designed into commercially available products.

But we believe there are other, nearer-term share price catalysts, including commercial deals with new customers and converting existing collaborations into formal, revenue generating agreements.

BrainChip's commercial progress to date

BrainChip (ASX:BRN) is an ASX-listed semiconductor company, developing and commercialising a revolutionary neuromorphic processor, called Akida. The company joined the ASX in March 2015 through the RTO of Aziana Ltd.

BrainChip was founded in 2004 by Peter Van Der Made after a stint at IBM where he was Chief Scientist for behaviour analysis security systems. He initially set out to create a fast parallel processor before realising it would be too slow to perform the required functions. By 2007, he had worked out a radically new chip design based on the biological brain and had been granted a patent for the technology by 2008 in Australia and the US.

First license deal with Renesas in 2020

BrainChip's first commercially available neuromorphic chip was the AKD1000 in April 2020. Later that year, the company secured its first IP license deal with Japanese chip manufacturer Renesas Electronics (TYO:6723). This was a game-changing event for the company and one that brought it to the forefront of many investors' minds.

Working with ARM, MegaChips and others

Since then, BrainChip has continued to develop Akida, releasing a new generation in early 2023. Big news came in early 2022 when a major car manufacturer announced it was using Akida hardware and software in a new concept electric vehicle.

Although BrainChip has been working with several signed and prospective clients, only a handful have been disclosed so far. Some of these include the aforementioned Renesas, Megachips, a Japanese fabless chip company, and the Space Machines Company.

Awaiting additional commercial deals

While the companies that BrainChip has been working with is impressive for a relatively small player in the semiconductor industry, investors have had to be patient in the last two years when it comes to the company signing additional license agreements that would generate revenues. Following the Renesas and MegaChips deals, it has been silent as far as actual signed contracts go.

Things take time in the semiconductor industry.

However, we believe that Akida's technical and commercial potential, as elaborated on below, is so substantial that BrainChip should be able to sign many additional deals going forward. But as always in the semiconductor industry, things take time. The industry may be extremely advanced as far as technology goes, when it comes to adopting new technologies, like Akida, it moves at a snail's pace.

One of the reasons for this is that a well-oiled chip production line is basically a license to print money during the good times. So, disrupting something that works well and makes them money is not in a chip manufacturers interest. Of course, they all need to innovate or risk going extinct, but many companies do so "reluctantly".

Prospective clients have to innovate, but they take their sweet time doing it.

In other words, they'll look at new technologies because they have to, but will only free up valuable engineering and production time when their production schedules allow them. Against such a backdrop, it may be easier to see why semiconductor IP companies on ASX, not just BrainChip, take longer to commercialise than some investors may want.

Akida 101: A revolutionary new chip design

Akida is a high-performance neuromorphic processor.

Akida (Greek for "spike") is BrainChip's flagship technology. It is a neuromorphic computer processor that provides high-performance, on-chip artificial intelligence capabilities at the edge of the internet (Edge AI) at extremely high energy efficiency and very low latency.

Energy efficiency is very important for devices that sit at the edge of the internet because battery size, and thus battery life, is usually a limiting factor for these types of devices, many of which don't have a fixed power connection.

Neuromorphic computing is different from "traditional" AI

Neuromorphic computing mimics the biological brain.

Neuromorphic computing emulates the structure and functions of neurons found in nature, for instance in the human brain. It does this very differently than today's algorithm-based AI, which is limited in terms of functionality to what its programmers intended. We have summarised the key benefits of neuromorphic processing in Figure 1.

Figure 1: Benefits of Neuromorphic processing

Key advantages	Details
Energy efficiency	• Neuromorphic processing can potentially reduce energy requirements for AI and machine learning (ML) tasks.
Cognitive abilities	• Neuromorphic computing aims to replicate cognitive abilities, such as pattern recognition, sensory processing and decision making.
Real-time processing	• Neuromorphic computing can perform complex computations in real time. This makes it suitable for applications that require quick decision making and responsiveness, like certain (autonomous) automotive applications.
Brain-inspired learning	• Neuromorphic hardware is designed to support synaptic plasticity, which, in turn, enables learning and adaptation (similar to how the human brain learns from experience).
Reduced latency	• Neuromorphic computing can reduce latency in processing tasks by performing computations close to data source (Edge AI) instead of in the Cloud.

While neuromorphic processing in itself is a big step forward for AI applications, BrainChip has taken things one step further. Its accelerator technology supports different neural networks in a very high-performance, low-power manner.

We have explained BrainChip's technology and how exactly neuromorphic processing works in great detail in appendices II through IV.

Akida's superior approach to data processing

Processing (learning) happens on-chip by attributing different weights to different signals.

In a nutshell, an Akida chip can process a range of different analogue inputs, such as sound and vision, in the same way the human brain does. Instead of having pre-programmed algorithms process incoming data according to how the programmers intended it, Akida processes digitised analogue inputs, so-called spikes, by having each individual artificial neuron in its network attribute more or less weight (so-called synaptic weight) to individual spikes depending on how other artificial neurons further up in the network respond to the signal that this particular neuron just propagated. This is the learning component of Akida.

And because this type of learning is much more autonomous than how pre-programmed algorithms work, the outcomes can potentially be much more radical or unexpected in a good way, i.e. Akida can potentially recognise patterns that its users weren't even looking for. In practice, though, a batch of Akida chips will have been trained for a specific application before it reaches the customer.

Parallel processing provides a tremendous speed advantage.

The processing of spikes is done in parallel, i.e. multiple spikes are processed at once rather than sequentially, which is how most of today's AI processing is done. This parallel processing delivers a tremendous speed advantage for Akida over today's algorithm-based AI.

Less need for additional memory on-chip or close to the processor.

And because the synaptic weights essentially form the aggregated memory of the Akida chip, there is much less need for additional system memory (on chip or very close to it), which speeds up processing even further.

Improved Cyber security.

This also means that there is less need to communicate over the internet, e.g. to a data center, which again increases processing speeds. But it also means improved Cyber security because there is less data in transit across the internet that could potentially be intercepted.

Summarising Akida's key advantages

In summary, we believe there are multiple reasons why the Akida technology is superior to CNN's, all of which are inherent to its neuromorphic nature and its ability to learn on-chip without the help of a data centre:

- **Very low power consumption**: Akida needs fractions of a watt to operate compared to CNN's. Consider for instance that it can efficiently analyse and categorise the 1.2 million images of the ImageNet dataset with just 300 milliwatts (0.3 watts), whilst a GPU, e.g. from NVIDIA, requires 300 watts![2]
- **Very low latency**: Because the artificial synapses in Akida's artificial neurons essentially remember the data that has been processed previously (synaptic weights), the chip can limit interaction with external memory during processing.

  On top of that, it can operate stand alone, i.e. without an internet connection to a data center. As a result, Akida has a very low latency, i.e. extremely high response time.
- **Ease of operation**: No internet connection is required for operation, so it can be seamlessly integrated into clients' Edge devices. It is available to customers as IP (Intellectual Property) through IP licenses.
- **Improved security**: Because Akida operates independently of the Cloud and stores data on-chip, it is more secure than competing technologies that upload data to the cloud and store it there for processing, sometimes keeping it for extended periods of time.

[2] https://brainchip.com/wp-content/uploads/2022/08/BrainChip-Learning-how-to-Learn.pdf

The market opportunity is set to quadruple by 2030

Very substantial growth ahead.

As we will cover in greater detail in the next chapter, the neuromorphic computing and edge AI markets were worth US$6.4bn and US$17.5bn respectively in 2023. Both markets are expected to roughly quadruple in the next six years, to US$24.5bn and US$62.9bn respectively by 2030.

We expect BrainChip will be able to benefit from this massive growth through collaborations with industry partners and direct deals with customers, because ultimately Akida can be useful in essentially any electronic device that would benefit from on-board AI.

A vast range of application areas.

And we believe this would not just be restricted to everyday electronic equipment. The recent launch of Akida into space as part of the ANT61 Brain computer in the Optimus-1 spacecraft shows Akida has utility in industries as significant and comprehensive as aerospace.

As more and more functionality is expected of electronic devices with each new generation, data processing and storage requirements seem to grow almost exponentially. The world cannot just ‘build more data centres’ because of their environmental footprint. The traditional approach to AI just uses more force (a next generation of an NVIDIA AI chip for instance) in the same old sequential computer architecture.

In order to address future AI requirements, we believe step changes are required, and Akida is the first of those.

A significant market opportunity awaits

BrainChip's market opportunities lie in neuromorphic computing and Edge AI, both of which we will discuss below.

The neuromorphic computing market is set to grow significantly

According to US research and consulting firm Verified Market Research, the global Neuromorphic Computing market was worth US$6.4bn in 2023 and is expected to reach US$24.5bn by 2030, representing a CAGR of ~21% during the forecast period (Figure 2).

Figure 2: The global Neuromorphic Computing market has huge potential (US$ BN)

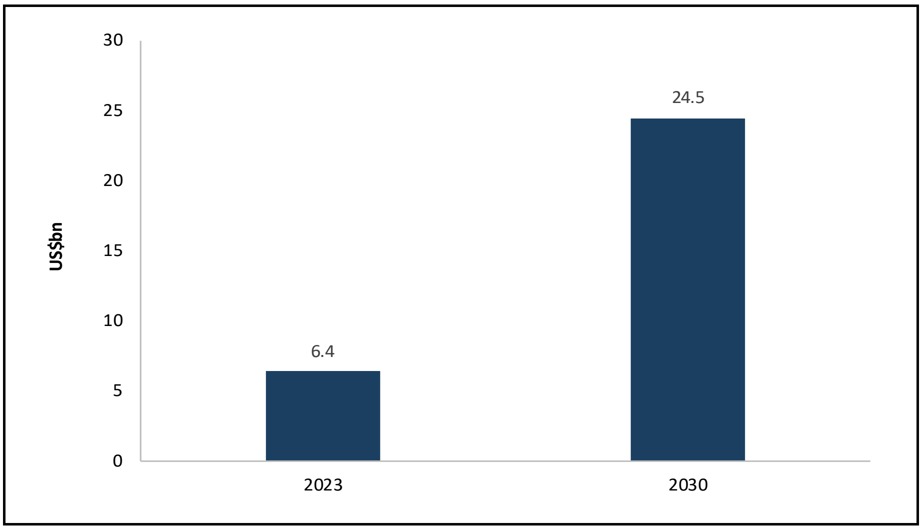

Source: Verified Market Research, Pitt Street Research

Rising demand for AI and ML technologies, and a shift away from traditional ICs towards neural architecture are driving growth in the Neuromorphic Computing market.

Several factors are driving this market growth, including the rise in demand for services based on AI and ML technologies, rising demand for high-performance integrated circuits (ICs), e.g. for High Performance Computing (HPC), and the shift away from traditional ICs towards neuromorphic architecture that is better suited to some of today's computing requirements.

The Consumer Electronics segment accounted for the largest market share, with a revenue of more than 57.9% in 2022. The increasing demand for several electronic devices, including laptops, PCs and tablets, is leading to an increase in demand for neuromorphic chips, which in turn, lead to a boost in the growth of the segment.

Many, high-value application areas

Additionally, the increased adoption of neuromorphic technology in deep learning applications, transistors, accelerators, next-generation semiconductors and autonomous systems, such as robotics, drones, self-driving cars will further drive market growth.

Neuromorphic computing has wide ranging applications

Emerging applications such as big data and IoT need substantial computational power in order to provide the services and information required by customers. So, why not just build more data centres? The answer is simple. Because of their substantial energy consumption and inefficiency. While data centres are set to consume over 20% of the world's energy by 2050 at current trajectory, decentralised AI processing would make a substantial difference to carbon emissions from Information Technology.

To fulfil these objectives, neuromorphic computing is a viable option for High Performance Computing and ultra-low power consumption. Neuromorphic computing has a host of applications areas (Figure 3).

Figure 3: Applications areas for Neuromorphic Computing

Application	Details
Artificial Intelligence (AI)	• Neuromorphic computing enhances the efficiency and performance of AI algorithms by simulating neural networks more closely with the human brain's architecture.
Brain-Machine Interfaces	• Neuromorphic computing can be applied in brain-machine interfaces (BMIs) and neural prosthetics, in order to restore or enhance nervous system functionality. • The use of neuromorphic instead of traditional devices is likely to create a more seamless experience for those with prosthetics.
Signal and image processing	• Neuromorphic circuits and platforms are becoming more widely available, which is advancing the field of signal processing. • Image processing is another common application of neuromorphic computing systems. These systems can be used in the analysis and extraction of information from images in real time and are particularly useful for tasks such as object recognition and facial recognition.
Advanced Robotics	• Neuromorphic computing contributes to advanced robotic systems by enabling more natural and adaptive behaviours, improved navigation, learning capabilities and effective human interactions.
Autonomous vehicles	• Neuromorphic computing enhances perception and decision-making capabilities in autonomous vehicles, making them better equipped to handle dynamic environments. • Driverless cars might not always be connected to the internet, especially when they are in the outskirts of a town during a commute. Neuromorphic computing could help them respond more effectively to their surroundings when not connected to a stable source.
Pattern recognition	• Neuromorphic computing can be utilised in fields such as image and speech recognition for efficient identification and adaptation to data patterns.

Source: Pitt Street Research

Neuromorphic computing serves a range of end markets

The combination of neuromorphic computing and AI is likely to result in a host of new products in the smart homes vertical

Neuromorphic processing is currently at an early stage, but could serve many end-user industries over time. These include the following:

- **Aerospace and defence -** Neuromorphic computing is expected to aid the military and defence industries in handling sensitive information related to battlefields due to the high processing speed of neuromorphic technologies. Additionally, neuromorphic computing can enhance the performance of drones and help with surveillance and security, tracking and targeting. Akida has recently taken its first steps into this market with the aforementioned launch of Optimus-1, which has Akida built into its onboard computer.
- **Industrial -** Neuromorphic computing holds the promise of improving automation and efficiencies in factories and large-scale operations. Processing of complex sensor data in parallel enables the real-time adaptation in smart manufacturing systems.

- **Healthcare -** Neuromorphic computing also has application potential in a range of use cases in the healthcare segment. Some of these include analysis of breath compounds, blood samples and patient statistics. Neuromorphic chips can also be used to improve drug delivery systems and in the analysis of Electroencephalogram (EEG) readings, which help in the diagnosis of neurological disorders by recording brainwaves and offering insights into cognition, emotions and disorders.
- **Smart Homes** - The installation of smart products in homes is likely to provide benefits to residents in terms of comfort, time, money and energy savings. Significant potential exists for neuromorphic computing in the IoT landscape for powering smart homes with intuitive systems.
- **Smart Cities** - The concept of a smart city is evolving rapidly and has transformed the way human beings and machines interact in urban settings. The use of neuromorphic technologies in use cases such as real-time monitoring for security, cyber-crime, environment, traffic and healthcare infrastructure is likely to improve the quality of life in these cities.

Edge AI has a plethora of applications

Edge AI's ability to produce real-time analytics more securely at higher speeds, lower costs and with lower power consumption has made it a viable alternative to Cloud-based AI. These advantages open up many additional application areas for Edge AI (Figure 4).

Figure 4: Applications of Edge AI

Application	Details
Self-driving cars	• Autonomous vehicles are using Edge AI learning at high speeds to continuously update and define safety parameters that make it easier for onboard systems to detect anomalies in the car and scan the car's surroundings.
Health Monitoring Devices	• Edge AI allows hospitals and other healthcare providers to reap the benefits of AI without having to transmit sensitive patient information across networks. • All data collected from health monitoring devices, such as cardiac trackers and blood pressure sensors, can be processed and analysed locally, thereby enabling real-time analytics and helping medical professionals provide better patient care.
Advanced medical robots	• Edge AI finds its applications in advanced medical robots, which use sophisticated sensors to see, hear, smell, touch and taste. • It further enables gesture control with faster response times, paving the way for doctors and therapists to enable interactions between people with disabilities and sophisticated robotic assistance devices.
Facial Recognition systems	• Facial recognition systems rely on computer vision algorithms as well as data analysis collected by a camera. And Facial recognition applications operating for security need to operate reliably even when not connected to the Cloud.
Autonomous drones	• Since autonomous drones are not piloted by human operators, they share similar requirements as autonomous cars. Drones are increasingly flying Beyond Visible Line of Sight (BVLOS) and out of range of an internet access point but still need AI capabilities. • Against this backdrop, Edge AI systems are likely to be indispensable for services such as Amazon Prime Air, which aims to deliver packages via drones.
Digital assistants	• Digital assistants, such as Google Assistant, Alexa and Siri, must be able to operate on smartphones and other digital devices even when not connected to the Internet.

Source: Pitt Street Research

The Edge AI market to grow to US$63BN in 2030

Edge AI market growing by 20% annually.

According to Maximize Market Research, an India-based market research and business consulting firm, the Edge AI market was valued at US$17.5bn in 2023 and is expected to grow at a CAGR of 20.1% from 2024 to 2030 to reach approximately US$62.9bn by 2030 (Figure 5).

North America dominated the market in 2023 with a revenue share of 39.9%. The growth is primarily due to the adoption of advanced technologies, such as AI, ML and deep learning in the region. The region is expected to continue to growth strongly with the advent of superior 5G network technology. The revenue in Asia Pacific is also anticipated to grow significantly as many Asian countries, including China, India and Japan, are using the AI and ML technologies in various industries.

Figure 5: The Edge AI market has robust growth prospects

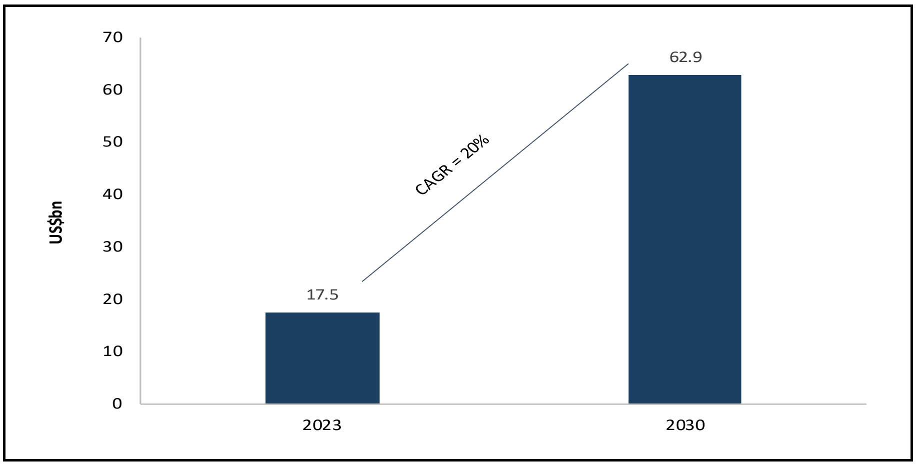

Source: Maximize Market Research, Pitt Street Research

The Edge AI market is driven by several key factors

Demand for real-time decision-making and reduced latency across industries along with privacy and security concerns are key drivers for the market growth. Cost optimisation is another factor, as Edge AI reduces the need for data transmission to the cloud, resulting in cost savings. These drivers, along with the growing use of edge AI applications contribute to the growth of the Edge AI market.

Edge AI serves many end markets

IT and Telecommunications sector occupied the highest revenue share of the Edge AI market in 2022.

The Edge AI market extends across various industries. In 2023, the IT and telecommunications segment held the largest revenue share of 21.4% of the global Edge AI market. New opportunities for telecommunication operators emerge with the proliferation of connected IoT devices and the transition of telecom networks to 5G, thereby fuelling the telecommunications sector.

The Consumer Electronics industry is likely to grow significantly over the rest of this decade due to the proliferation of smart wearables, smart speakers, and other digital devices, facilitated by increasing availability of advanced Edge AI hardware components. Moreover, data processing on Edge AI chips also improves privacy and security on consumer devices (Figure 6).

Figure 6: The Edge AI market has robust growth prospects

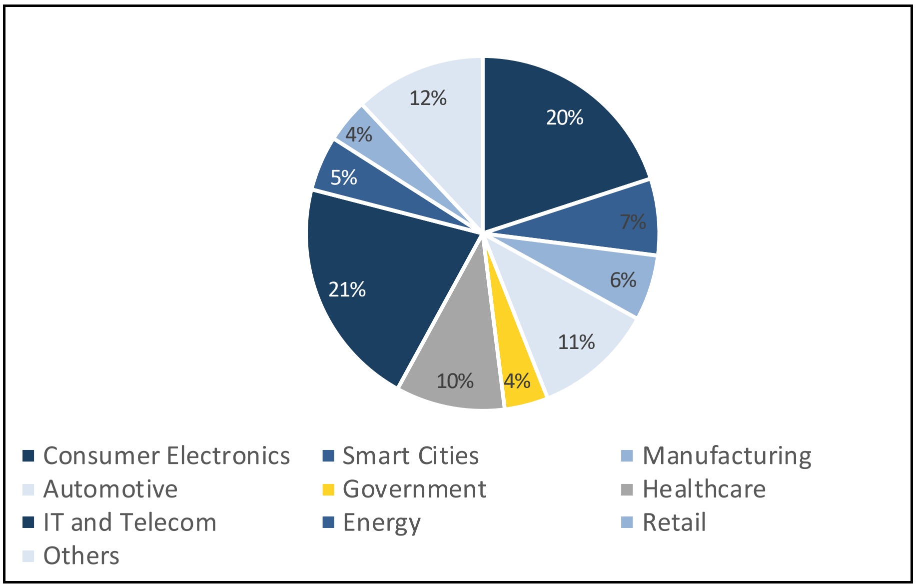

Source: Grandview Research, Pitt Street Research

BrainChip's Tech Journey

Akida 1.0 / AKD1000

BrainChip was the first ever company to file a digital neuromorphic chip patent back in 2008. When it reverse-listed on the ASX in 2015, its technology was known as SNAP (Spiking Neural Adaptive Processor), which is what Akida essentially is. The first commercial implementation occurred in late 2017 when the company shipped its Accelerator card to a major European automobile manufacturer.

The Akida name for the technology was adopted in 2018, derived from the Greek word for Spike. It was also in that year that the first Akida System-on-a-Chip (SoC) was launched – this could be used either as a stand-alone embedded accelerator or as a co-processor, whilst its predecessor was only the former.

At first, customers received a software version of Akida that customers and prospects could use to create, train and test neural networks destined for the Akida NSoC as well as run inference (outcome-based processing) to determine the performance and accuracy of the neural network. The first production chips were only received in 2021, although engineering samples were shown to prospective customers before that.

The announcement that got investors most excited was a collaboration with a major car manufacturer which announced the use of Akida hardware and software in a research prototype electric vehicle that has a >1,000km range. BrainChip's share price reached A$2.34 on the back of that announcement.

Akida 2.0 addresses many customer requirements

A 2nd generation of Akida was launched in March 2023. It allows designers and developers to do things that were not previously possible (i.e. with Akida 1.0), but were wanted by customers. The most notable included:

- 8-bit processing.
- Temporal Event Based Neural Net (TENN) support for vision (object detection, segmentation), audio, language models, speech and sensor data. This supercharges the processing of raw time-continuous streaming data[3]. It allows for radically simpler implementations by consuming raw data directly from sensors – drastically reducing model size and operations performed, while maintaining very high accuracy. These shrink design cycles and dramatically lower the cost of development.
- Support for accelerating vision transformers. This is primarily used for image classification, object detection and semantic segmentation.

How Akida stacks up against other technologies

The biggest way Akida stands out from the competition (Figure 7) is actually having a product on the market. Any products that purport to be similar actually aren't commercially available or are only available for general research purposes. This is in spite of the fact that several big names in technology have been developing neuromorphic chips for a number of years now.

[3] Examples include video analytics, target tracking, audio classification, analysis of health monitoring data such as heart rate and respiratory rate for vital signs prediction, and time series analytics used in forecasting and predictive production line maintenance.

Figure 7: Akida is well-positioned against competition

brainchip	Micro- to Mw Power use	Real-time on-chip learning & training	TensorFlow Compatible	Stand-alone possible (No CPU required)	On-chip Convolution	Available as IP	Green Technology
BrainChip Akida AKD1000	✓	✓	✓	✓	✓	✓	✓
IBM TrueNorth	✓	NONE	LEARN COREL	NO	NO	NO	✓
Intel Loihi	✓	PROGRAM	LEARN NEF	✓copack	NO	NO	✓
Google Coral TPU	2-5W	Math chip	✓	NO	NO	NO	NO
DLAs (Nvidia, others)	>5-10W	Math chip	✓	NO	NO	NO	NO

Source: Company, Pitt Street Research

IBM has had multiple neuromorphic chips. First there was TrueNorth that debuted in 2014, followed by NorthPole in 2023. TrueNorth was revolutionary at the time, boasting a power efficiency four orders of magnitude lower than conventional microprocessors. The trouble was the capital cost and this was resolved with NorthPole, which could operate 4,000 times faster. NorthPole is optimised for 2-bit, 4-bit and 8-bit low-precision operations. It operates at a frequency range of 25-425 megahertz and can perform 2,048 operations per core per cycle at 8-bit precision and 8,192 operations at 2-bit precision.

The trouble is that some applications demand neural networks too big to fit in a single NorthPole chip. As opposed to Akida, it has to pull data from memory it contains and cannot draw on outside sources.

Intel has a neuromorphic research chip called Loihi. The first generation was introduced in 2017, with the second generation released in 2021 (Loihi 2). Loihi 2 synthesizes 1 million neurons and 120m synapses. The former figure was a big improvement from the 128,000 neurons of the earlier generation, although the number of synapses went down from 128m. For now, it is not being commercialised, not even sold to external researchers yet. It will perform specific algorithms rather than being a general-purpose chip.

Google has the Coral TPU. It provides high performance machine learning inferencing for low-power devices. This is not for larger devices. You also need certain software to run inferences, dependant on the target platform and programming language.

NVIDIA has attracted the most attention from global investors, at least so far as its share price re-rating in the last year has demonstrated. It manufactures GPUs (Graphic Processing Units) and AI-optimised chips that are useful for a range of purposes, mostly in Cloud-based settings, including cryptocurrency, pharmaceuticals and other HPC applications.

Of course, Nvidia is a much larger company than BrainChip and it is not a competitor, because Nvidia is focused on data centres, rather than Edge AI applications.

BrainChip's commercialisation journey

3 main configurations for the Akida IP.

BrainChip will license its Akida IP in 3 product classes:

- Akida-E, a highly energy efficiency version with low-node count optimised for operation very close to the sensor. This is designed for energy-sipping intelligent sensors that are even capable of performing without a supporting CPU.

- **Akida-S** which balances the power, performance and footprint. It is ideal for integration into MCUs (Micro Controller Units) or smaller ASICs (Application Specific Integrated Circuits) that are designed for detection and classification workloads.

- **Akida-P** a medium to high-performance version ideal for segmentation, prediction and complex model operations that can be accelerated in hardware. It can be optionally enhanced with vision transformers, enabling ASICs custom-built for specialised AI applications that deliver high-end performance cost-effectively and efficiently.

It is worth noting that Akida 1.0 was best suited for use cases of the E version only. Akida 2.0 is suitable for all 3.

Akida Edge Box recently launched.

BrainChip also recently started selling the Akida Edge AI Box at a list price of US$799, a product powered by the Akida NSoC chip and was optimised for running AI algorithms at the edge without the need for Cloud connectivity. The company expects to see the Edge Box used specifically in Retail and Security, Smart City, Automotive, Transport and Industrial applications.

BrainChip's commercial model

BrainChip is seeking to make revenues through IP licensing agreements.

BrainChip is seeking to make revenues through IP licensing agreements. These would involve customers integrating Akida chip architecture into their own infrastructure, paying both one-off licensing payments, engineering fees related to specific integration requirements. and ongoing royalty payments when the end products are sold.

Example

- Assume BRN entered into a licencing agreement with a customer to licence its technology with specific integration requirements.
- Terms could include US$500k of NRE fees, US$1m in one-time licencing fees and 5% royalty on sales from each of the customer's chips that incorporates Akida.
- Further assume that this customer will sell 1m, 5m and 10m of these chips in years 1 through 3, respectively, at US$25 each.

In this example BrainChip would recognise the following three revenue streams from this customer:

1) NRE fees of US$0.5m.

2) One-time licencing fees of US$1m before the start of production.

3) Royalties of US$1.25m (1m x US$25 x 5%), US$6.25m and US$12.5m, respectively, in the first three years of production.

3 burning questions that investors have right now

So, how long will it take for Akida to be commercialised?

Once a deal is signed, it typically takes 6-18 months of design, integration, manufacturing and chip testing/qualification work to create a complete SoC that is suited to the customer's own requirements and ready for mass production.

Weebit Nano's (ASX:WBT) deal with SkyWater Technologies is a good example of the fact that this takes time. A deal was signed in September 2021, although only in March 2023 did it become available to SkyWater customers, the first of which still has to sign a commercial agreement with WBT.

From signing a deal to generating royalties can take a number of years.

So, once a deal is signed, it will take multiple years before royalty revenues start coming in from a particular customer.

How are individual collaborations going?

This is arguably the biggest drag on the company's share price as well as the fact that BRN generally does not name clients or provides a running commentary on them. Even if it wanted to, it could not – many customers do not want to disclose that they are using Akida because of competitive issues. BrainChip would risk losing its customer.

As much as it would like to, BrainChip can't talk about individual customers.

Investors should be encouraged by the few collaborations or customers that have been named. It illustrates that there is indeed a market for Akida and that BrainChip has a solution that the market wants and needs.

However, as we noted above, it will take time for these collaborations to start generating recurring revenues.

What is next?

Investors should expect BrainChip's revenues to grow as it continues to expand its arrangements with customers, initially receiving license and NRE fees, followed by royalties later on.

BRN will also continue to invest in R&D work to ensure Akida remains competitive and suitable for the needs of existing and would-be clients. This would result in newer generations of Akida as well as individual products that could include Akida, such as the newly released Edge Box.

Balance sheet structure and funding requirements

As of 31 December 2023, the company had US$13m in cash and reported net operating cash outflow of US$4.4m. At current burn rates, this is enough to last for 3 quarters.

However, this does not account for the funding agreement the company has with LDA Capital that allows for LDA to subscribe for up to $12m in shares, subject to capacity limitations[4]. The issue price will be 91.5% of the higher of the average daily VWAP of shares over the pricing period and the minimum price notified to LDA Capital by the Company.

In other words, BrainChip is sufficiently funded in the short to medium term.

[4] Specifically ASX Listing Rule 7.1 whereby a listed entity can only place more than 15% of issued capital during a 12 month period if it can obtain shareholder approval.

Valuation of A$1.59 per share

Although BrainChip is technically not a pre-revenue company anymore following the licence fees it received from its initial customers, these revenues are non-recurring and royalty payments from these customers have not started yet. It is, therefore, not possible to capture the company's long-term value solely by looking at its current revenues and cashflow, i.e. without visibility of future revenues.

Using M&A activity as a valuation guide.

Rather, we prefer to look at what the semiconductor industry would be willing to pay for BrainChip's unique technology as a guide to how much its shares could be worth.

Prospects of an acquisition

The semiconductor industry has a long history of acquisitions where the key focus has been on the IP of the acquired companies, not their revenues at the time of acquisition. In our view, BrainChip's IP is so unique that there is a high likelihood of strategic interest in the company from other, large industry players. Whether or not that translates into an actual acquisition is impossible to predict, but we think there's a definite possibility of this happening in the medium term.

Industry deals as a guide

In our report on BrainChip from August 2021, we valued the company using the Habana acquisition by Intel in 2019. At the time, Intel paid US$2BN for Habana, a pre-revenue Israel-based AI startup. This translated into A$1.50 per fully diluted BrainChip share.

BrainChip is further advanced than Habana at the time of its acquisition by Intel.

At the time, Habana was slightly ahead of BrainChip in terms of product launches with its Goya and Gaudi chip launches in 2018 and 2019. Today, BrainChip also has products on the market and has received license fees from multiple customers. In other words, we believe that US$2BN valuation for Habana is still a pretty good guide for BRN potential valuation at the moment.

Value of A$1.59 per share

Using the Habana valuation to value BrainChip, we arrive at a value of A$1.59 per fully diluted share (AUD/USD = 0.66).

The share price can potentially respond very strongly to positive news.

While this is substantially higher than today's share price, BrainChip shares have been as high as A$2.34 in January 2022. Back then, the share price was driven by news around a major car manufacturer using Akida in one of its concept cars. We believe this illustrates how well the share price can potentially respond to good news.

Share price catalysts

We foresee BrainChip shares being re-rated towards our valuation by the following factors:

- Converting current license deals into royalty generating collaborations,
- New license deals,
- Converting collaborations into formal agreements,
- New IP patents being granted, and
- Ongoing positive sentiment around the rise of AI technologies in general.

Risks

We see the following key risks to our investment thesis:

- **Execution risk**: Any delay in Akida's commercialisation plan will likely negatively impact BrainChip's valuation.
- **Lower adoption rate of Akida technology**: While Akida is expected to be widely used across various end markets, lower than anticipated adoption rates by customers may hamper BrainChip's future growth.
- **Funding risk**: BRN will likely need to raise additional capital to support its development and commercialisation activities until it reaches cash flow break even, diluting current shareholders.
- **Technology risk**: We believe BrainChip is well-ahead of its competitors, with on-chip learning and on-chip convolution, which no other company has been able to do. The company also owns the patents on those technologies. As with all types of technology, though, there is a risk that the rapid pace of development in the Artificial Intelligence space may lead to BrainChip's technology becoming obsolete, or at least less unique.
- **Share overhang**: LDA Capital can sell the shares it owns in BrainChip (other than collateral shares) on market at any time, depending on certain daily volume restrictions, which may result in sudden and adverse share price movements.

Appendix I: BrainChip leadership and management

The company's current board and leadership composition is as follows:

Name	Designation
Board of Directors	
Sean Hehir	Chief Executive Officer
Peter Van Der Made	Founder
Antonio J. Viana	Chairman of the Board
Geoffrey Carrick	Non-Executive Director, Chair of the Audit & Governance Committee
Pia Turcinov	Non-Executive Director, Chair of the Remuneration & Nomination Committee
Duy-Loan Le	Non-executive Director
Management Team	
Anil Mankar	Co-Founder and Chief Development Officer
Ken Scarince	Chief Financial Officer
Tony Lewis	Chief Technology Officer
Steve Thorne	Vice President of Sales

Board of Directors

Sean Hehir has managed large global teams and been responsible for explosive revenue growth for global enterprise organizations such as Compaq and HP, as well as smaller, fast-growing companies like Fusion-io. Hehir is industry-recognized as a builder of trusted Strategic Alliances across diverse partners such as Systems Integrators, ISVs, and OEMs. He received his BS from the University of Massachusetts and MBA from Georgia State University.

Peter van der Made has been at the forefront of computer innovation for 40 years. Mr. van der Made is the inventor of a computer immune system and founded vCIS Technology where he served as CTO and later Chief Scientist when it was acquired by Internet Security Systems and subsequently IBM.

Previously, he founded PolyGraphic Systems, and designed a high resolution, high-speed colour Graphics Accelerator board and subsequent chip for IBM PC graphics. Mr. van der Made published a book, Higher Intelligence, which describes the architecture of the brain from a computer science perspective. Mr. van der Made designed the first generations of digital neuromorphic devices on which the Akida chip is based between 2004 and 2008, when he applied for a patent on this technology which was subsequently granted. He is actively involved in the design of the next generation of Akida chips and continues his research in advanced neuromorphic architectures.

Antonio J. Viana serves as a non-executive Chair and as a member of the Audit & Risk and Remuneration & Nominations committees. He is also non-executive director of Arteris Inc., a leading provider of network-on-chip (NoC) interconnect and other SOC-Fabric intellectual property. In 1999, Mr. Viana joined ARM Holdings, the global leader in semiconductor IP, serving in a number of leadership positions, most notably as the Global Director of the ARM Foundry Program and President of Commercial and Global Development. He was appointed to the ARM executive team as Executive VP of worldwide sales in 2008. At the beginning of 2013, his executive duties were expanded to include all of commercial and global development. Mr. Viana has also worked with Hughes Aircraft, Silicon Graphics, Encore Industries and was Senior VP of worldwide sales at Tensilica Inc. Most recently, Mr. Viana served as the Executive Chairman of QuantalRF AG, an emerging Swiss RF semiconductor company developing transformative wireless communication solutions.

Geoffrey Carrick held the positions of Head of Corporate Finance at Shaw and Partners Limited from March 2016 – July 2019, and Head of Equity Capital Markets at Commonwealth Bank from 2012-2015. From 1999 through 2011, Mr. Carrick was Division Director of Equity Capital Markets at Macquarie Capital. Mr. Carrick currently serves as Director of VCF Capital Partners Pty Limited and Non-Executive Director of Global Study Partners Holdings Pty Limited. Carrick is a graduate of the University of Sydney B.Ec, LLB. Mr. Carrick is the Chair of the Audit & Risk Committee and a member of the Remuneration & Nominations Committee at BrainChip.

Pia Turcinov has 30 years experience across private, not-for-profit and government sectors. Commencing her career as a corporate lawyer, Ms. Turcinov, later practiced commercial law in Western Australia. Her career progression evolved from legal advisor to commercial director, holding both senior-level positions and providing leadership and governance as a company director across a wide-range of businesses. She has also facilitated innovation through a number of State and Federal Government initiatives. Ms. Turcinov shares her expertise as an advisor, mentor and facilitator, as well as a keynote speaker on topics relating to future skills, innovation, diversity and STEM. As the mother of three daughters, she remains an ever-enthusiastic champion for diversity, equality and enabling female entrepreneurship. Ms. Turcinov was awarded the 2018 Excellence in Gender Equity Promotion Award by the United Nations Association of Australia (WA Division) Inc. In 2023, Pia was recognized for her significant service to technology and innovation, and to women in STEM, by the award of Member (AM) in the General Division of the Order of Australia.

Duy-Loan Le has an extensive professional history, both technologically and in executive management, having retired from Texas Instruments (TI) as a Senior Fellow after 35 years. While at TI, she led the global R&D, manufacturing operation and high-volume production of TI's multi-billion-dollar memory, DSP, and base station product lines. Ms. Le holds 24 patents and serves on the board of two universities. In addition to BrainChip, she currently serves on the boards of Wolfspeed, National Instruments, Ballard Power Systems and Atomera. She was inducted into Women in Technology Hall of Fame and became the first engineer to be inducted into Asian Hall of Fame. She received numerous recognitions for her philanthropic contributions worldwide, including Congressional Special Recognition. Ms. Le serves on the BrainChip board as a non-executive director and is a member of the Audit & Risk and Remuneration & Nominations Committees.

Management Team

Anil Manker has spent 40 years developing products in the semiconductor industry. At Western Digital, Mr. Mankar developed PC core Logic chipsets. During his years at Conexant Systems Inc. in the position of VP of Engineering, he developed multiple products across industry segments and later became the company's Chief Development Officer overseeing all product development for V92 Modem, DSL, Set-top boxes, PC audio, and video 'System on a Chip' products. Mr. Mankar was SVP of VLSI Engineering at Mindspeed Technologies, responsible for Wireless and VOIP infrastructure product development.

Ken Scarince has extensive experience as a financial executive in a variety of industries, including high-technology. He most recently served as a consultant at 8020 Consulting, working on all aspects of finance at a variety of companies globally. Previously Mr. Scarince served as Controller at Virgin Galactic and Vice President of Finance and Chief Accounting Officer at Virgin America. He began his career as a senior auditor at Deloitte and Touche. Mr. Scarince received his Bachelor's Degree in Finance from Marquette University and his Masters Degree in accounting from the University of Wisconsin-Milwaukee.

Dr Tony Lewis is an executive, scientist and entrepreneur specializing in brain-inspired AI and robotics. Tony was Global Head of the Artificial Intelligence and Emerging Compute Lab at HP, Inc. At Qualcomm, Inc Tony headed Neuromorphic, Deep Learning and Robotics efforts. Tony has held faculty positions or held leadership positions at UCLA, UIUC, and the Univ. of Arizona. Tony holds a Doctor of Philosophy degree in Electrical Engineering from University of Southern California and a Bachelor of Science in Cybernetics/Applied Math from UCLA.

Steve Thorne is a senior sales and marketing executive with over 28 years' experience in AI and Data Center solutions within Habana Labs and Intel Corporation. In his last role as Head of Global Sales for Habana Labs he spearheaded the successful launch and ramp of three innovative Habana Gaudi AI accelerator systems into the market. At Intel he held multiple leadership roles for the Intel Xeon processor family and led teams to deliver on more than 25 product introductions in top Hyperscale CSP, OEM and Enterprise accounts. Steve has a Bachelor of Science in Computer Science from Georgia Tech

Appendix II – Introduction to AI, machine learning and deep learning

Artificial Intelligence comes in different shapes and forms

Artificial Intelligence (AI) is the science of training systems to perform human tasks through learning and automation. AI makes it possible for machines to learn to apply logic, adjust to new inputs and reason to gain an understanding from complex data. In simple words, AI provides machines with the ability to learn from data it receives by processing and recognizing patterns in that data.

AI is an overarching term and essentially consists of foundational building blocks and key elements, namely machine learning and deep learning, computer vision, natural language processing (NLP), forecasting and optimization, and machine reasoning. These building blocks or elements can be used independently or combined to build AI capability. Several AI capabilities and their use cases in a business context are illustrated in Figure 8.

Figure 8: AI use cases

AI capability	Use case
Pattern Recognition	For fraud detection by understanding typical trends/behaviors within customer financial transactions data or spot anomalies in an account's spending data
Prediction	For improving energy consumption forecasts by capturing short- and long-term variability in data
Classification	For supporting wildlife conservation efforts by examining animal track changes and grouping them by species type
Image recognition	For determining if nodes on a raw CT scan are malignant or benign; other applications include predictive diagnostics and biomedical imaging
Speech-to-Text	For transcribing voice messages to text for sentiment detection and further analysis in call centers
Cognitive search	For offering personalized recommendations to online shoppers by matching their interests with other customers who purchased similar items
Natural language interaction (NLI)	For generating automated financial reports on sales revenue predictions, which otherwise would be generated by the user itself
Natural language generation (NLG)	For automated generation of summaries after analyzing large sets of documents

Source: Pitt Street Research

These AI capabilities can be used either independently or combined with each other, depending on users' objectives and underlying data. For instance, in the banking industry, combinations of these capabilities are used for credit and risk analysis and to provide market recommendations by creating automated financial advisors. In the healthcare industry, such combinations are used for processing data from past case notes, biomedical imaging, health monitoring, etc. Other industries such as manufacturing and retail are also utilizing AI capabilities to optimize supply chains or to offer personalized shopping experiences and customized recommendations. In addition, governments across the globe are focusing on building smart cities and utilizing capabilities such as facial recognition for use in law enforcement.

Machine Learning and Deep Learning

While AI comprises all techniques that make machines perform tasks that require intelligence, Machine Learning specifically imitates how humans learn. Basically, Machine Learning is a subset of AI (Figure 9) and consists of the techniques that enable machines to learn from the data without being explicitly programmed to do so. Conversely, other AI techniques could be classified as rules-based or expert systems, which work on a pre-defined algorithm or logic – like performing accountancy tasks, in which the system runs the information through a set of static rules.

Figure 9: Machine Learning and Deep Learning – Subsets of AI

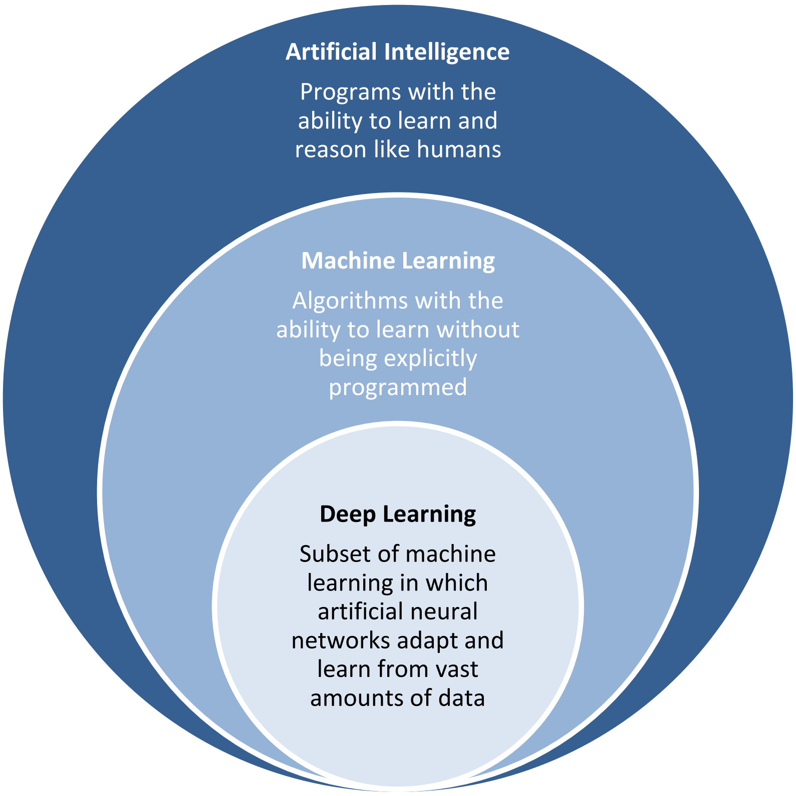

Source: Argility, Pitt Street Research

One aspect that separates Machine Learning from rules-based expert systems is the ability to modify itself when exposed to more data, i.e. machine learning is dynamic and does not require human intervention to make certain changes.

Machine Learning is dynamic in nature and has greatly evolved over the years

Though Machine Learning has evolved a lot over the years and is used to tackle many problems, for a long time it was still difficult for machines to perform many tasks such as speech, handwriting and image recognition, and more mundane tasks such as counting the number of items in a picture. The concept of Artificial Neural Networks (ANN) kick-started the development of Deep Learning, which provides machines the capability to perform tasks such as image recognition, sound recognition and recommender systems with much greater accuracy and speed.

Deep Learning itself is essentially a subset of Machine Learning and is all about using neural networks comprising artificial neurons, neuron layers and interconnectivity. Instead of organizing data to run through predefined equations, Deep Learning sets up basic parameters around the data and trains the computer to learn on its own by recognizing patterns using many layers of computer processing.

Artificial Neural Networks learn like the human brain does

Artificial Neural Networks (ANNs) are computing systems with a large number of interconnected nodes that work almost like neurons in the human brain. They use algorithms to recognize hidden patterns and correlations in raw data and then cluster and classify that data to solve specific problems. Over time, neural networks continuously learn from new data and apply those learnings to make future decisions.

A simple neural network includes an input layer, an output (or target) layer, and a hidden layer in between. The artificial neurons (or nodes) in these layers are interconnected and form a network termed as a neural network of interconnected nodes (Figure 10). As the number of hidden layers within a neural network increases, deep neural networks are formed. A simple ANN might contain two or three hidden layers, while deep neural networks can contain as many as 100 hidden layers.

Figure 10: Neural Network basic diagram

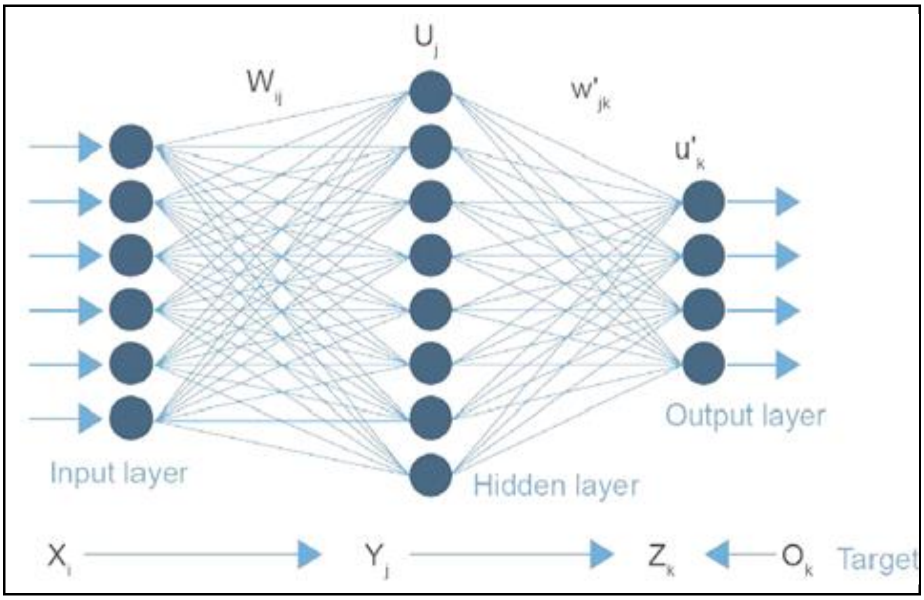

Source: ExtremeTech

In a typical neural network, a node is patterned after a neuron in a human brain. These nodes get activated when there are sufficient stimuli or inputs (just like neurons in a human brain). This activation spreads throughout the network, creating a response to the stimuli (output). The connections between these artificial neurons act as simple synapses, enabling signals to be transmitted from one to another. Signals across layers travel from the first (input layer) to the last (output layer) and get processed along the way.

While solving a problem or addressing a request, data such as text, images, audio and video, is fed into the network via the input layer, which communicates to one or more hidden layers. Each neuron receives inputs from the neuron to its left, and the inputs get multiplied by the weights of the connections they travel along. These input-weights are then summed up. If the sum is higher than a certain threshold value, the neuron fires and triggers the neurons it is connected to on the right. In this way, the sum of the input-weight product determines the extent to which a signal must progress further through the network to affect the final output.

Many types of neural networks

Over the past several years, many neural networks with different architectures and specifications have emerged. Feedforward Neural Networks (FNNs) are the simplest form of ANNs. For specific tasks, more complex ANNs have been invented, including the Convolutional Neural Networks (CNNs), which aim to mimic the human visual system, as well as the Recurrent Neural Networks (RNNs), which are used to interpret sequential data such as text and video. These major types of ANNs are described in Figure 11.

Figure 11: Types of Neural Networks Other Than Spiking Neural Networks

Neural Network	Description	Applications	Network Image
FNNs	Each perceptron (simplest and oldest form of neurons) in one layer is connected to every perceptron from the next layer. Information is fed forward from one layer to the next in the forward direction only. There are no feedback loops. Thus, the data is processed, and the results are calculated on every input sequence. This network may or may not have hidden layers.	Primarily used for animal recognition, digit recognition, cheque recognition, medical diagnosis, etc.	
RNNs	Use sequential information such as time-stamped data from a sensor device or a spoken sentence, composed of a sequence of terms. Unlike FNNs, inputs to RNNs are not independent of each other, and the output for each element depends on the computations of the preceding elements.	Primarily used in forecasting and time series applications, sentiment analysis and other text applications.	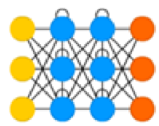
Long Short-Term Memory (LSTM)	A type of RNN that is explicitly designed to hold information for long periods of time and process the incoming data, along with the previously calculated results. LSTMs contain their information in a memory and can read, write and delete information from its memory.	Primarily used for text classification, machine translation, dialog systems, speech recognition, translating videos and images to natural languages, etc.	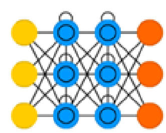
CNNs	Typically contain five types of layers: input, convolution, pooling, fully connected and output (more recent versions tend to be deep with more than seven or nine layers). Each layer has a specific purpose, like summarizing, connecting or activating.	Primarily used for image classification and object detection. Other applications include language processing, computer vision and video analytics.	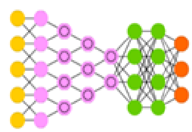

Source: Medium, SAS, Pitt Street Research

Supervised Learning versus Unsupervised Learning

Since the advent of Machine Learning, different algorithms or methods have been developed to process both structured and unstructured data. However, all Machine Learning methods can be broadly classified into either supervised learning or unsupervised learning (Figure 12), though supervised learning is the most commonly used form of Machine Learning.

With supervised learning, each input fed to the system is labelled with a desired output value. A supervised learning algorithm analyses the data and compares its actual output with desired output to find errors and modify the model accordingly. Supervised learning is commonly used in applications where future events are predicted based on historical data, e.g. determining fraudulent credit card transactions and predicting insurance customers likely to file claims.

Input dataset is labelled in supervised learning while in unsupervised learning the system structures the data

In unsupervised learning, the training set submitted as input to the system is not labelled with the historical data or a desired outcome. In simple words, unsupervised learning is used against data that has no historical labels. Therefore, the system itself develops and structures the data, identifies common characteristics, and modifies it based on knowledge gained during the process.

This form of Machine Learning is commonly used to segment customers with similar attributes who can then be treated similarly in marketing campaigns. It can also identify the main attributes that separate customer segments from each other. Other applications include segmentation of text topics, image recognition, pattern recognition in financial markets data, identification of data outliers, sound analysis, e.g. to detect anomalies and potential problems in jet engines etc.

Figure 12: Supervised Learning Vs. Unsupervised Learning

Parameter	Supervised Learning	Unsupervised Learning
Type of Input Data	Labeled	Unlabeled
Degree of Computational Complexity	High	Low
Accuracy of Results	High	Low to moderate
Timeliness of Analysis	Off-line	Real time
Commonly Used Algorithms	Random Forests, Linear Regression, Decision Trees, Naïve Bayes, Support Vector Machines, Neural Networks	Clustering (K-means, SVD, PCA, etc.), Association Analysis (Apriori, FP-Growth), Hidden Markov Model
Key Use Cases	Prediction and classification	Grouping and data interpretation

Source: Pitt Street Research

Convolutional Neural Networks are widely used today

CNNs are among the most widely used ANNs today given that they can learn unsupervised and require relatively little pre-processing. CNNs are used in a range of areas, including statistics, natural language processing as well as in signal and image processing, e.g. for medical image analysis.

However, CNNs are rather impractical for visual imagery classification given the large data sets that need to be processed, which consumes enormous amounts of energy while CNNs are relatively slow. With the advent of autonomous vehicles and the stringent requirements on image recognition capabilities by Advanced Driver Assistance Systems (ADAS) in cars, today's CNNs may not be the best solution.

Pros and Cons of Machine Learning and Deep Learning

Machine Learning and Deep Learning have various applications such as in defence and law enforcement

In summary, Machine Learning and Deep Learning have many applications, and organizations use these applications to drive automation for specific tasks and processes, e.g. to save cost, bring products to market faster, improve operational efficiencies, prevent fraud, gain new insights into data and enable new technologies to be deployed faster. Homeland Security (HLS) and law enforcement are other application areas for AI.

While Machine Learning supplements data mining, assists decision making and enables development of autonomous computers and software programs, Deep Learning, on the other hand, performs complex computations and is widely used for difficult problems that require real-time analysis, such as speech and object recognition, language translation and fraud detection.

However, these AI technologies do have their own limitations. Both Machine Learning and Deep Learning are susceptible to errors and whenever they make errors, diagnosing and correcting them can be difficult. In addition, it is impossible to make immediate accurate predictions with these technologies as they require substantial computational power and can be difficult to deploy, especially in real time.

Cons of Machine Learning and Deep Learning include susceptibility to errors and unintentional biases

Furthermore, the outcomes generated by these technologies are prone to hidden and unintentional biases, including racial biases, depending on the data provided to train them. Also, these technologies cannot always provide rational reasons for a prediction or decision.

Nevertheless, the utilization of Machine Learning and Deep Learning is anticipated to rise substantially as the potential of neural networks to solve problems, make predictions and improve decision-making are unparalleled.

The next iteration in neural networks is the emergence of Spiking Neural Networks that have multiple advantages over CNN's, including speed and power consumption, as we will elaborate on below. BRN is at the very forefront of Spiking Neural Network development and future commercialisation.

Appendix III – Neuromorphic computing 101

In our discussion of CNN's and ANN's so far we haven't specifically mentioned that these AI models are purely software-based. Given that these algorithms are so large and generally can't be executed locally, most queries such as Google Assistant and Siri queries on a mobile phone, need to be sent to the Cloud to be processed. The results then need to be sent back to the device. This takes time and processing such queries in the Cloud consumes tremendous amounts of energy.

However, in our view the most significant drawback of software-only neural networks is that the algorithms are designed by humans, i.e. software engineers, and hence the scope of the neural network is limited to the imagination of whoever designed the particular algorithm. We touched on this briefly when we discussed the differences between supervised and unsupervised learning, i.e. unsupervised learning allows the network to learn without any restrictions in what to look for.

Therefore, we believe the logical evolution of neural networks is a hardware-only solution that allows for unsupervised learning. Enter Neuromorphic Computing.

Neuromorphic chips mimic the biological brain in a hardware implementation

In simple terms, a neuromorphic chip tries to emulate the structure and functions of neurons found in nature, for instance in the human brain. But rather than using software, the artificial neurons are hardwired in digital chips. The human brain consists of roughly 86 billion interconnected neurons (Figure 13) that send and receive electric impulses, or spikes, to and from neighbouring neurons.

Figure 13: The human brain comprises billions of neurons

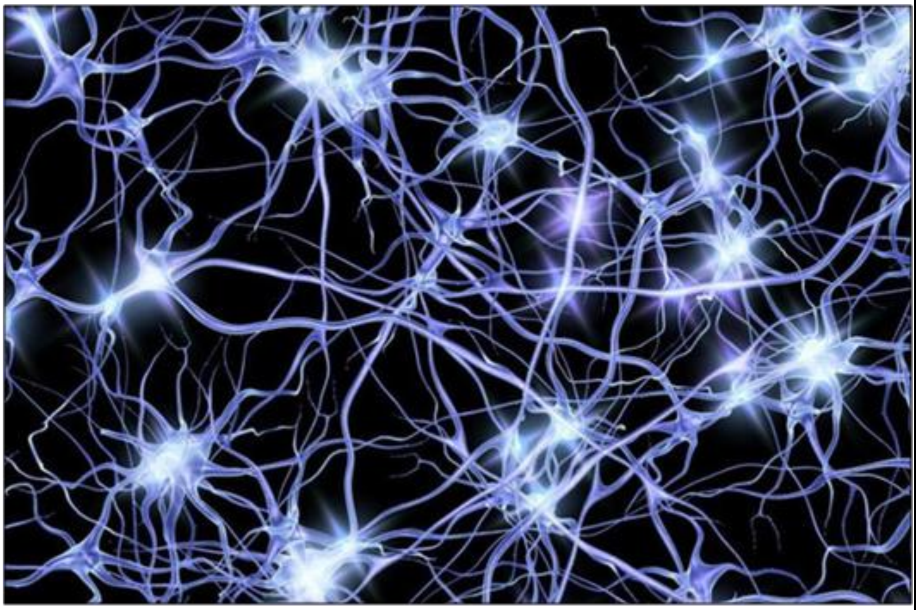

Source: Penn State, Pitt Street Research

A single neuron system comprises of the actual neuron, axons, dendrites and synapses. As we alluded to earlier, electric impulses are sent from one neuron to the next through connections known as axons. Before reaching the next neuron, though, each impulse arrives at the next neuron's synapse first (Figure 14), which can be seen as a gatekeeper. The human brain has approximately 150TR synapses (trillion = 1,000 billion).

Figure 14: Synapses in between two neurons

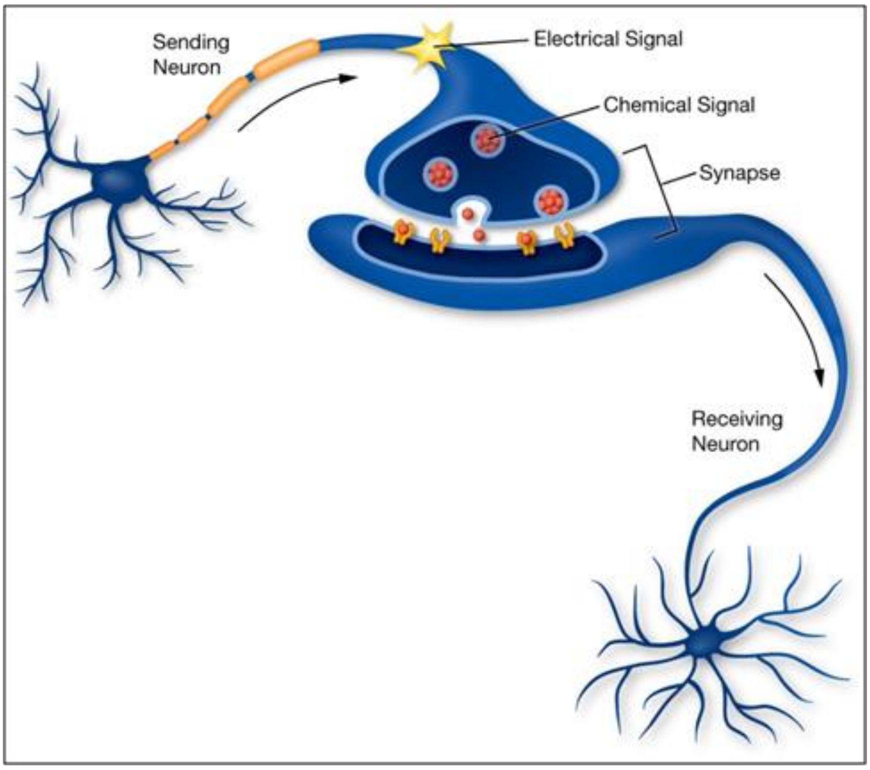

Source: whfreeman.com, Pitt Street Research

Synapses decide whether or not to pass on a signal to the next neuron depending on the weight stored in the synapses (Figure 15). The electrical impulse must reach a certain intensity threshold in order to be considered relevant and be allowed to pass on to the next neuron. As a particular input passes through the network of synapses, this network acts as a filter, activating some neurons while others remain inactive.

Figure 15: Impulses must exceed a certain threshold to be propagated

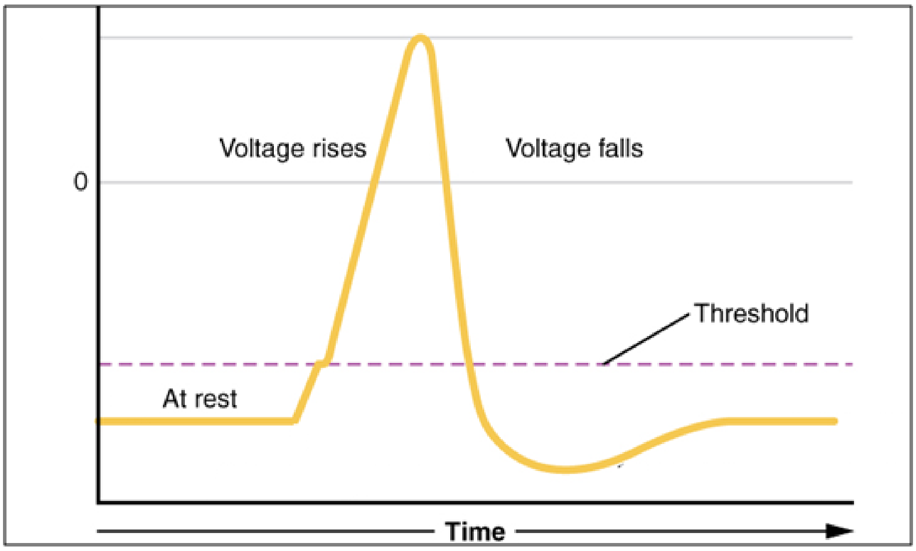

Source: BC Campus, Pitt Street Research

The threshold function of the synapse plays a key role in neuromorphic chips

The key function in the process above is performed by the synapse that decides which impulses to propagate and which impulses to terminate. That decision is based on the synaptic weight, which attributes a value to each incoming impulse upon which the decision whether or not to propagate the impulse is based.

This threshold function is equally critical in biological neurons as it is in hardwired artificial neurons as we will elaborate on below.

Appendix IV – BrainChip's solution for Neuromorphic Computing

BRN has developed a hardware version of the biological neuron

Company's development work in the last ten years has focused on creating a hardware implementation of the biological neuron described above, resulting in the design of a neuromorphic chip called Akida, which is Greek for spike. The cell architecture and cell behaviour are similar to that of a biological neuron.

Artificial synapses work like biological synapses in terms of threshold function

The architecture is such that incoming impulses, also known as spikes, are received by artificial synapses, which can autonomously decide whether or not to propagate the spike (Figure 16), based on how strong, or relevant, the spike is.

Similar to biological synapses, artificial synapses work with a threshold function, i.e. the sum of impulses received from multiple, connected, neurons must reach a certain threshold level in order to be considered relevant and to be fed forward to other neurons in the system.

Figure 16: Neuron structure in a Spiking Neural Network

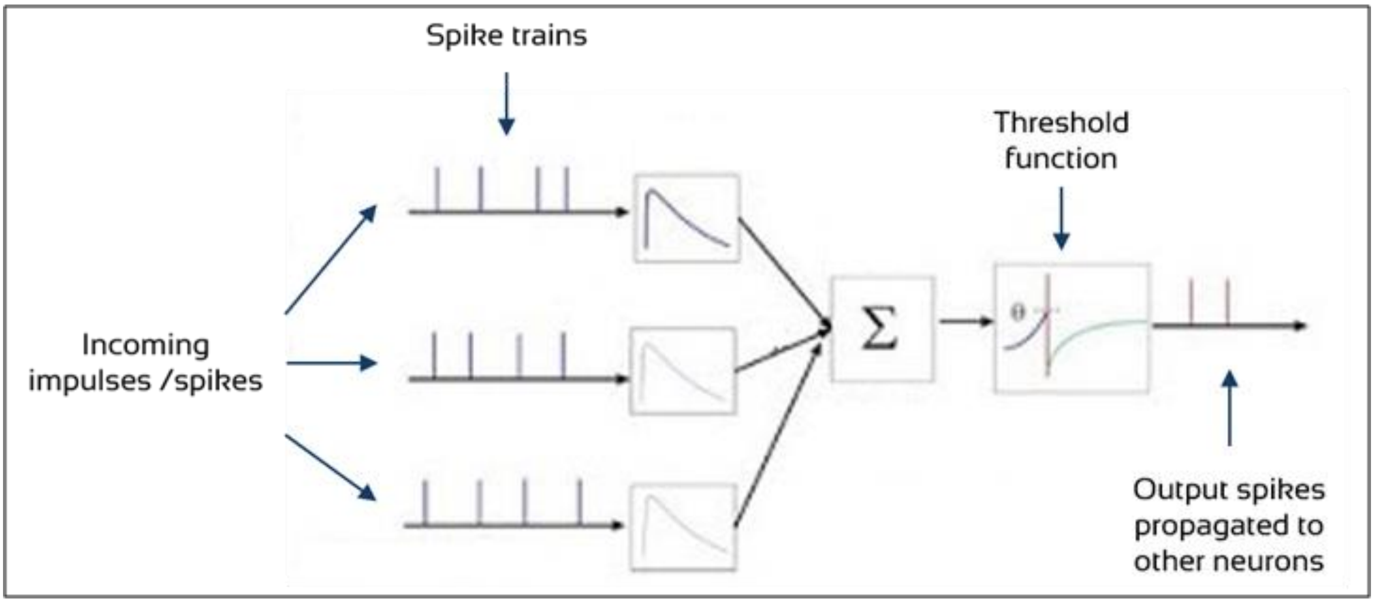

Source: AGH University of Science and Technology, Pitt Street Research

On the scale of the entire chip, BRN has a fabric of Neural Processing Cores (NPCs), where each NPC emulates 1,000s of the biologically-inspired neurons. These NPCs are connected in a network, so that they can emulate the multiple layers similar to the structure of the artificial neural networks we described previously (see Figure 10). This enables spikes to move through multiple layers of neurons before an output spike is generated. This type of neural network is known as a Spiking Neural Network (SNN).

Spike-Timing-Dependent-Plasticity: Learning by weights

The way in which an individual neuron in an SNN "learns" is by increasing or decreasing the weight attributed to each of its artificial synapses, depending on the relevance of the spikes the neuron received through each synapse. This relevance depends on how often and how many spikes are received. This process is known as Spike Timing Dependent Plasticity (STDP), i.e. in addition to a neuron actually firing, the frequency of spikes and the actual number of spikes fired is also relevant information. Based on this information, artificial synapses will start to favour certain connections, while they may inhibit others over time.

BRN has significantly expanded upon the base STDP learning rule for its technology. The use of STDP, the learning method of the brain, is unique in Artificial Neural Networks. Most ANNs use Deep Learning, which is a successive approximation method. STDP leads to much faster results and needs much less data than Deep Learning. The reason that CNNs use deep Learning is because they compute with numbers rather than spikes. BRN delivers a tool that can convert standard CNNs to fast-learning, low power spiking CNNs for execution on Akida, called CNN2SNN. This tool is free and delivered with MetaTF.

The memory of each neuron, i.e. what it has "learned", is embedded in the weight of its synapses, the so-called synaptic weights. Synaptic weights are dynamic and can change over time based on the relevance of the spikes it receives from preceding neurons in the network.

Spiking Neural Networks' advantages over CNNs

Spiking Neural Networks are much faster and use substantially less power

As opposed to traditional, software-based CNNs, that use integer data types (ones and zeros) as inputs, SNNs process spikes (Figure 17). Spikes are simply an encoding method of real-world sensory data points, such as sound and vision. A sequence of spikes is known as a spike train. Spike trains offer the ability to process this real-world sensory data in the same way the human brain does.

Figure 17: Akida's sequential processing of spikes

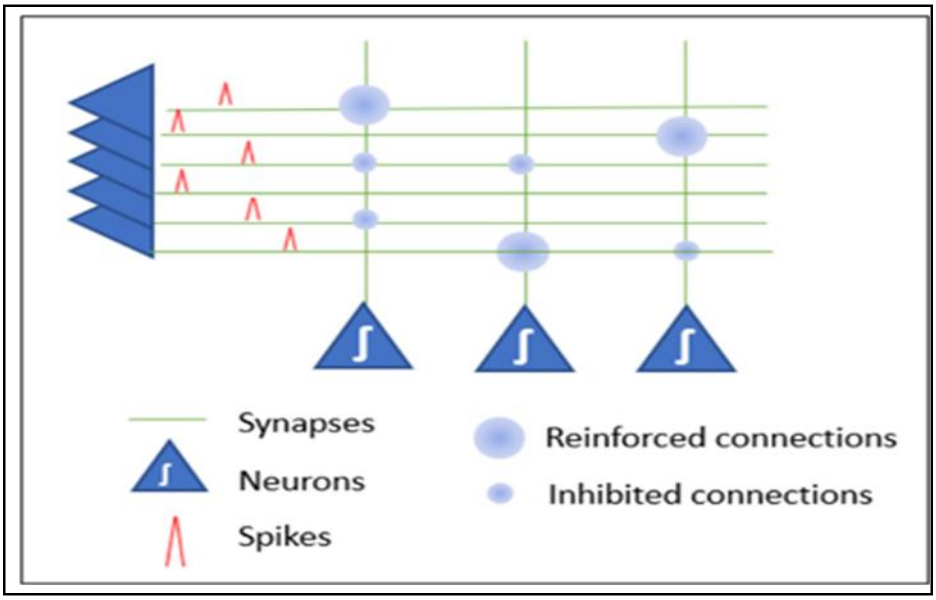

Source: Company, Pitt Street Research

Given that Akida's neurons receive incoming spikes from multiple other neurons at the same time, Akida can process incoming information in parallel, i.e. many spike trains are being processed by Akida simultaneously. This is a key distinction from software-based CNNs, which process data sequentially, i.e. one mathematical calculation after another.

Additionally, given that most of the memory of an SNN can be found in the synaptic weights, SNNs don't require much access to external memory, such as DRAM (Dynamic Random-Access Memory) to retrieve information on the weight of a particular synapse or to temporarily store the outcome of a calculation.

Parallel processing and "on-board" memory provide speed advantage over traditional CNNs

Parallel processing and "on-board" memory inside the synapses provide Akida with a tremendous speed advantage compared to traditional CNNs. In addition to speed advantages, the fact that SNNs don't need to move very much data from the processor to the memory and back with each calculation also results in SNNs using substantially less power compared to CNNs. The latter require very substantial traffic of data between processing units and memory units.

SNNs consume significantly less power as compared to CNNs because of sparcity

Due to sparsity of SNNs (many of the parameters, i.e. weights or connections in the network, are zero or close to zero), their power consumption can be up to 95% lower compared to CNNs. This makes SNNs ideally suited for IoT edge applications that are not permanently connected to a power source, such as mobile phones, Electric Vehicles and sensors. But energy consumption is also extremely important in large scale data centers. In other words, we believe sparsity in hardware-only SNN's provides several key competitive advantages compared to CNNs.

Akida is expected to take AI to the next level

Currently, AI processing is centralised, with connected devices relying on data center services, such as Amazon Web Services (AWS), to carry out calculations. However, Akida has taken a different approach by decentralising AI processing and carrying out computations at the edge of the connected devices (Edge AI).

This provides various advantages such as data privacy and reducing the dependency on internet bandwidth. Given that in current AI solutions, such as Siri and the Google Assistant, the data has to be first transferred to the cloud for processing at the data center, it makes the personal information in the dataset vulnerable to cybercrimes.

Furthermore, Akida is a fully-integrated, purpose-built neural processor (Figure 18). Traditional neural processing solutions require a CPU for running the neural network algorithm, a deep learning accelerator such as a GPU for carrying out multiplication and addition calculations (knows as MACs), and memory for storing network parameters.

However, by integrating these three functions into a single neural processor, can reduce energy consumption needed in traditional solutions to interact with the three components. Moreover, the consolidation also helps Akida reduce its size.

Figure 18: Traditional neural processing solution vs Akida

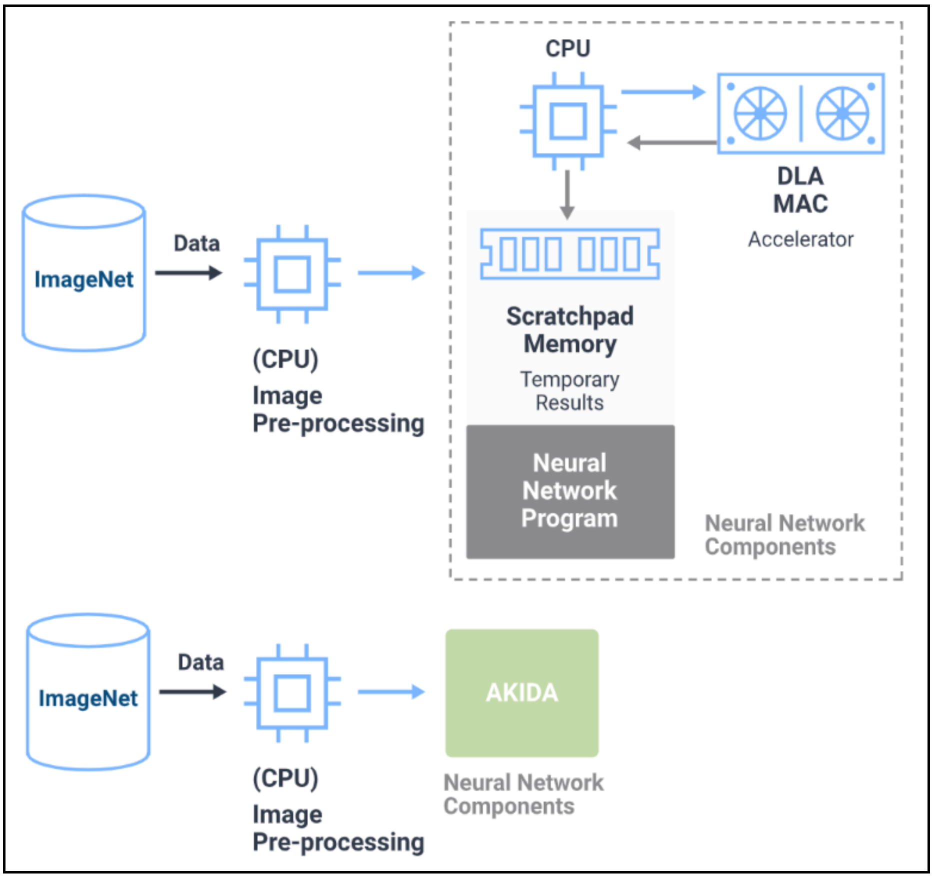

Source: Company

Appendix V – Capital structure

Class	In Millions
Fully paid shares	1,806
Options	63.4
Performance rights	4.8
Restricted stock	37.6
Diluted shares	**1,911.7**

Source: Company

Note: BRN excludes options, performance rights and RSUs from the calculation of diluted loss per share in its annual report because the company considers them anti-dilutive[5].

Appendix VI – Analysts' qualifications

Marc Kennis, lead analyst on this report, has been covering the semiconductor sector as an analyst since 1997.

- Marc obtained an MSc. in Economics from Tilburg University (The Netherlands) in 1996 and a post graduate degree in investment analysis in 2001.
- Since 1996, he has worked for a variety of brokers and banks in the Netherlands, including ING and Rabobank, where his main focus has been on the technology sector, including the semiconductor sector.
- After moving to Sydney in 2014, he worked for several Sydney-based brokers before setting up TMT Analytics Pty Ltd, an issuer-sponsored equities research firm, which merged into Pitt Street Research in 2018.

Nick Sundich is an equities research analyst at Pitt Street Research.

- Nick obtained a Bachelor of Commerce/Bachelor of Arts from the University of Sydney in 2018. He has also completed the CFA Investment Foundations program.
- He joined Pitt Street Research in January 2022. Previously he worked for over three years as a financial journalist at Stockhead.
- While at university, he worked for a handful of corporate advisory firms.

[5] Page 48 of BRN's Annual Report for FY23.

Appendix VII – Patent portfolio

Figure 19: BrainChip's patents

Patent Number	Country	Country
US 8,250,011	US	Autonomous learning dynamic artificial neural computing device and brain inspired system
US 11,157,800	US	Neural processor based accelerator system and method
US 10,157,629	US	Low power neuromorphic voice activation system and method
US 11,151,441	US	System and method for spontaneous machine learning and feature extraction
US 11,157,798	US	Intelligent Autonomous Feature Extraction System Using Two Hardware Spiking Neutral Networks with Spike Timing Dependent Plasticity
US 10,410,117	US	Method and a system for creating dynamic neural function libraries
US 11,238,342	US	Method and a system for creating dynamic neural function libraries
US 11,429,857	US	Secure Voice Communications System
US11,657,257	US	Spiking neural network
CN111417963	CN	Spiking neural network
AU2019372063-B2	AU	Spiking neural network
US 11,468,299-B2	US	Spiking neural network
AU2021254524-B2	AU	Spiking neural network
US 11,227,210-B2	US	Event-based classification of features in a reconfigurable and temporally coded convolutional spiking neural network
AU2020315983-B2	AU	Event-based classification of features in a reconfigurable and temporally coded convolutional spiking neural network
US 11,704,549	US	Event-based classification of features in a reconfigurable and temporally coded convolutional spiking neural network
AU2022203607	AU	Event -Based Extraction of Features in a Convolutional Spiking Neural Network
US2019/0286944	US	Digital electronic circuit and system implementing a spiking neural network for performing unsupervised detection of patterns
EP3324344	EP	Digital electronic circuit and system implementing a spiking neural network for performing unsupervised detection of patterns

Source: Company

Appendix VIII – Glossary

Artificial Intelligence (AI) - a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions, influencing real or virtual environment

Central Processing Unit (CPU) -

Convolutional Neural Networks (CNNs) – a type of neural network that is utilised for classification and computer vision tasks.

Deep Learning Networks (DNN) – Networks that use so-called deep learning to perform sophisticated computations on large amounts of data. It is called 'deep' because of its sophisticated method – by using the human brain.

Digital Signal Processors - Processors dedicated to number-crunching digital signals like audio. They capture signals such as voice, audio, or video as digital information, process them at very high speeds and then feed the information back for use

Edge Computing – Computing where data is processed at the point of, or closer, to where it is being generated. This enables processing at greater speeds and volumes and consequently leads to greater action-led results.

Intellectual Property (IP) – Intangible assets as a result of creativity. In the context of BRN, this alludes to the company's technology, patents and trademarks.

Internet of things (IoT) – A term for a network of interrelated sensing devices that connect and exchange data with other IoT devices and the cloud

Latency – The delay before a transfer of data begin following an instruction for its transfer.

Neural Networks - a machine learning program, or model, that makes decisions in a manner similar to the human brain, by using processes that mimic the way biological neurons work together to identify phenomena, weigh options and arrive at conclusions. They are sometimes called artificial neural networks (ANNs) or simulated neural networks (SNNs)

Neuromorphic – An adjective used to describe systems with neural principals (i.e. relying on the human brain).

Spiking Neural Networks (SNNs) - a two-layered feed-forward network with lateral connections in the second hidden layer that is heterogeneous in nature. are inspired by information processing in biology, where sparse and asynchronous binary signals are communicated and processed in a massively parallel fashion

Synapses - The places where neurons connect and communicate with each other. Each neuron has anywhere between a few to hundreds of thousands of synaptic connections, and these connections can be with itself, neighboring neurons, or neurons in other regions of the brain.

Systems on a Chip (SoC) - an integrated circuit that compresses all of a system's required components onto one piece of silicon, as opposed to systems that have multiple components.

Vision Transformer Networks (VTNs) – A neural network that specialises in image classification tasks. It does so by employing a Transformer-like architecture over patches of the image.

General advice warning, Disclaimer & Disclosures

Terms & Conditions

The information contained herein ("Content") has been prepared and issued by Pitt Street Research Pty Ltd ACN 626365615 ("Pitt Street Research"), an Authorised Representative (no: 1265112) of BR Securities Australia Pty Ltd. ABN 92 168 734 530, AFSL 456663. All intellectual property relating to the Content vests with Pitt Street Research unless otherwise noted.

Disclaimer

Pitt Street Research provides this financial advice as an honest and reasonable opinion held at a point in time about an investment's risk profile and merit and the information is provided by the Pitt Street Research in good faith. The views of the adviser(s) do not necessarily reflect the views of the AFS Licensee. Pitt Street Research has no obligation to update the opinion unless Pitt Street Research is currently contracted to provide such an updated opinion. Pitt Street Research does not warrant the accuracy of any information it sources from others. All statements as to future matters are not guaranteed to be accurate and any statements as to past performance do not represent future performance.

Assessment of risk can be subjective. Portfolios of equity investments need to be well diversified and the risk appropriate for the investor. Equity investments in a listed or unlisted company yet to achieve a profit or with an equity value less than $50 million should collectively be a small component of an individual investor's equity portfolio, with smaller individual investment sizes than otherwise. Investors are responsible for their own investment decisions, unless a contract stipulates otherwise.

Pitt Street Research does not stand behind the capital value or performance of any investment. Subject to any terms implied by law and which cannot be excluded, Pitt Street Research shall not be liable for any errors, omissions, defects or misrepresentations in the information (including by reasons of negligence, negligent misstatement or otherwise) or for any loss or damage (whether direct or indirect) suffered by persons who use or rely on the information. If any law prohibits the exclusion of such liability, Pitt Street Research limits its liability to the re-supply of the Information, provided that such limitation is permitted by law and is fair and reasonable.

General advice warning

The Content is General Financial Advice but has been prepared for general information purposes only and is not (and cannot be construed or relied upon as) Personal Financial Advice nor as an offer to buy/sell/subscribe to any of the financial products mentioned herein. No investment objectives, financial circumstances or needs of any individual have been taken into consideration in the preparation of the Content.

Financial products are complex, entail risk of loss, may rise and fall, and are impacted by a range of market and economic factors, and you should always obtain professional advice to ensure trading or investing in such products is suitable for your circumstances, and ensure you obtain, read and understand any applicable offer document.

Disclosures

Pitt Street Research has been commissioned to prepare the Content. From time to time, Pitt Street Research representatives or associates may hold interests, transact or hold directorships in, or perform paid services for, companies mentioned herein. Pitt Street Research and its associates, officers, directors and employees, may, from time to time hold securities in the companies referred to herein and may trade in those securities as principal, and in a manner which may be contrary to recommendations mentioned in this document.

Pitt Street Research receives fees from the company referred to in this document, for research services and other financial services or advice we may provide to that company. The analyst has received assistance from the company in preparing this document. The company has provided the analyst with communication with senior management and information on the company and industry. As part of due diligence, the analyst has independently and critically reviewed the assistance and information provided by the company to form the opinions expressed in the report. Diligent care has been taken by the analyst to maintain an honest and fair objectivity in writing this report and making the recommendation. Where Pitt Street Research has been commissioned to prepare Content and receives fees for its preparation, please note that NO part of the fee, compensation or employee remuneration paid will either directly or indirectly impact the Content provided.