

BrainChip Holdings Limited

Taking Deep Learning to higher levels

BrainChip Holdings Limited (ASX:BRN) is an ASX-listed semiconductor company that is currently in the early stages of commercializing Akida, its Neuromorphic System-on-a-Chip (NSOC). By integrating the Akida technology into their products, BRN's current and prospective customers can bring the benefits of one of today's most advanced technologies in Artificial Intelligence (AI) to their end-customers.

Addressing a market expected to grow to US\$ 66BN by 2025


The market for Deep Learning chipsets is expected to grow from around US\$ 4BN in 2018 to more than US\$ 66BN by 2025, according to Tractica, implying a CAGR of nearly 49%. Within this total market, BRN aims to sell dedicated Neuromorphic System-on-Chip (NSoC) devices and, more selectively, Intellectual Property (IP) blocks. Most of the chips used in AI applications being sold today are general-purpose chips, such as Graphics Processing Units (GPUs). Key target markets for BRN's NSoCs are primarily vision systems such as surveillance cameras, Advanced Driver Assistance Systems (ADAS) and Autonomous Vehicles (AV), vision guided robotics, drones, and Industrial Internet of Things (IIoT). More specialized edge applications such as cell phones are addressed with IP blocks.

Akida enables hardware-based Neuromorphic Computing

BRN has developed a hardware version of the biological neuron, like the ones found in the human brain, and has packed 1.2M of these artificial neurons and the accompanying 10BN artificial synapses on a neuromorphic computer chip called the Akida NSoC. The architecture and behavior of this chip are similar to that of a biological neuron, but implemented using a mainstream digital logic process.

The Akida technology is what is known as a Spiking Neural Network (SNN), i.e. similar to the human brain, Akida processes spikes or events instead of data. The advantage of processing spikes is that it can be done in an event-driven manner, compared to how "traditional" software-based neural networks, such as Convolutional Neural Networks (CNN), process data. The result is that the event driven network only processes data and consumes power when events are present as opposed to CNNs which consume power by processing all the input data, continuously.

This means SNNs are much faster and require only a fraction of the power consumed by CNNs. Additionally, BRN expects to commercialize the Akida NSoC at a substantially lower price point.



Readers should be aware that TMT Analytics has been engaged and paid by the company covered in this report for ongoing research coverage. Please refer to the final page of this report for the General Advice Warning, disclaimer and full disclosures.

Initiation of coverage

BrainChip Holdings Limited

(ASX: BRN)

Software & Services

Australia

Risk: High

BrainChip Holdings Limited (ASX:BRN) is an ASX-listed semiconductor company that is currently in the early stages of commercializing Akida, its Neuromorphic System-on-a-Chip (NSOC). By integrating the Akida technology into their products, BRN's customers and prospects can bring the benefits of one of today's most advanced technology in Artificial Intelligence to their respective end-customers.

**SUBSCRIBE TO OUR
RESEARCH HERE**

Please note TMT Analytics and/or its directors hold stock in BrainChip Limited as of the date of this report.

Speculative Buy

Current price: A\$ 0.041

Valuation range: A\$ 0.40-0.45

15 May 2019

Analyst: Marc Kennis

marc.kennis@tmt-analytics.com.au

+61 (0)4 3483 8134

Commercialization of Akida has already started

BRN has been working on the development of the Akida technology for ten years and anticipates it will have the technology available in test chips for prospective customers in the second half of 2019. However, commercialization of Akida has already started with BRN recently having released the Akida Development Environment (ADE), a software version of Akida that customers and prospects can 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.

Substantial pipeline of prospects

Per November 2018, BRN was involved in 21 active or committed pilot projects for Akida and its revenue-generating software-based SNN, BrainChip Studio. In addition, BRN has scored a large number of design wins and had more than 500 leads in its commercial pipeline. Upon successful completion of transferring the Akida IP from lab to fab, expected in the second half of 2019, we believe BRN should be in a position to secure a variety of different customer types for Akida, including cell phone manufacturers, semiconductor foundries, Automotive OEMs for Advanced Driver Assistance Systems (ADAS) and Autonomous Vehicles, third-party semiconductor IP providers, IDMs (Integrated Device Manufacturers) and companies in the Imaging space. Given that BRN will be commercializing the Akida NSoC through both device sales and an asset-light IP licensing model with recurring royalty revenues, we anticipate high gross margins.

Valuation range of A\$ 0.40-0.45 if Akida NSoC shows commercial viability

Looking at where Akida currently is in its development/commercialization process, we are of the opinion that BRN's share price does not accurately reflect the company's commercial potential described above. In our view, the company's valuation could move towards levels seen for comparable companies, such as Nervana and Movidius that were acquired by Intel Corp (NASDAQ:INTC), i.e. equivalent to A\$ 0.43 per fully diluted BRN share. Similarly, we believe AudioPixels' (ASX:AKP) A\$ 560M valuation, equivalent to A\$ 0.44 per fully diluted BRN share, also provides a gauge to BRN's potential valuation, given where AKP is positioned in its development process. In the near to medium term, though, we believe BRN will first need to demonstrate commercial viability of the Akida NSoC by signing one or more commercial agreements that involve integration of Akida and/or components of the Akida IP stack into customers' designs.

Given the strong upside we can potentially see for the BRN's share price if Akida is commercially validated, we start our coverage with a Speculative Buy recommendation.

Near-term share price catalysts

- Conversion of current prospect and discussion partners into paying customers and development partners, e.g. to develop and deploy Akida for specific applications such as for cell phones and ADAS.
- Updates on the Akida NSoC commercialization roadmap, specifically around first silicon test results.
- Additional customers for BrainChip Studio, especially in the Casino and Law Enforcement sectors.
- One of the Top-20 shareholders, Metals X Limited, seems to have largely sold its position on market recently. With selling pressure from Metals X now gone, we may see a bounce in BRN in the short term.

Table of Contents

Artificial Intelligence comes in different shapes and forms.....	3
Neuromorphic computing: hardware-based neural networks.....	9
BRN developed a hardware version of the biological neuron.....	11
Commercialization of Akida in various steps.....	16
BrainChip Studio: a software-based SNN.....	19
Deep Learning chipsets market to grow at ~60% CAGR.....	21
Very limited competition in Neuromorphic Computing chipsets.....	26
M&A provides guidance to BRN's potential valuation.....	29
Conclusion.....	31
SWOT Analysis.....	33
APPENDICES.....	34
GENERAL ADVICE WARNING, DISCLAIMER & DISCLOSURES.....	36

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 1.

FIGURE 1: ARTIFICIAL INTELLIGENCE CAPABILITIES AND 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: TMT Analytics

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 2) 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

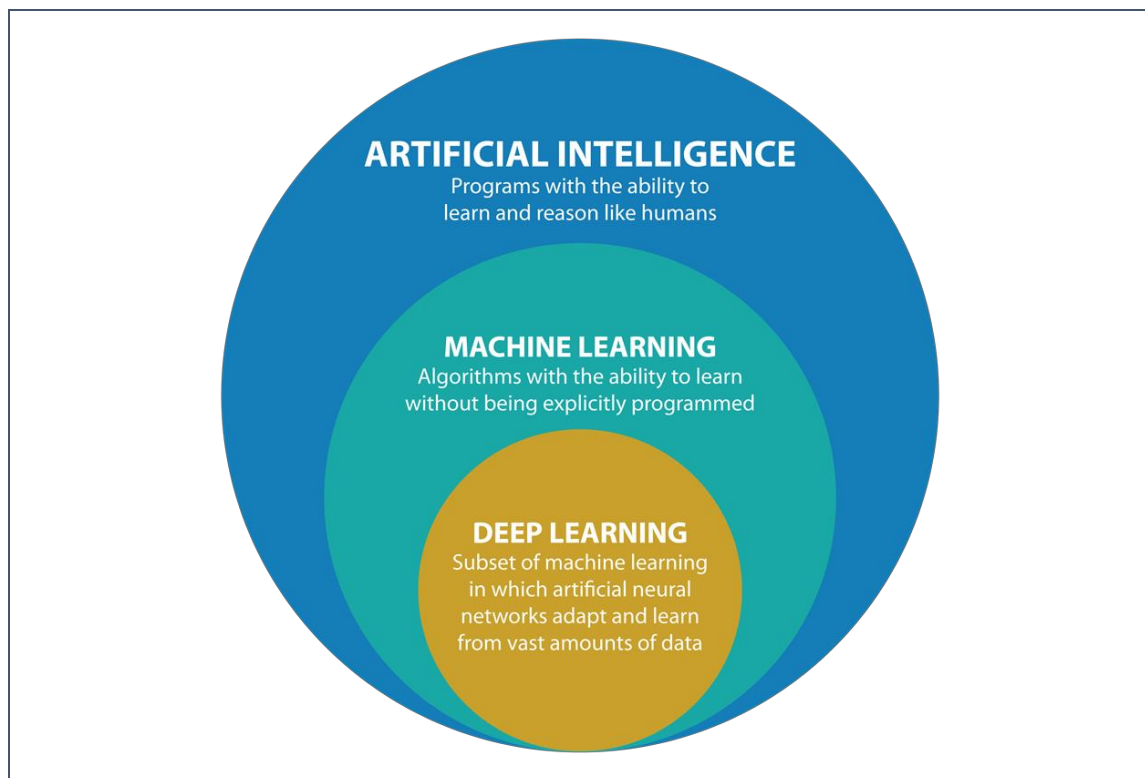
like performing accountancy tasks, in which the system runs the information through a set of static rules.

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.

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) kickstarted 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.

FIGURE 2: MACHINE LEARNING AND DEEP LEARNING SUBSETS OF AI



Source: Argility


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 3). 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 3: 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-weight 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.



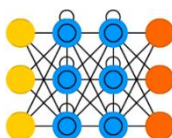
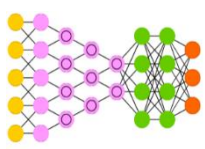
In the next chapter we will discuss this process in more detail.

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 4.

FIGURE 4: 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, TMT Analytics

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 5), though supervised learning is the most commonly used form of Machine Learning.

With supervised learning, each input fed to the system is labeled with a desired output value. A supervised learning algorithm analyzes 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.

In unsupervised learning, the training set submitted as input to the system is not labeled 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 5: 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: TMT Analytics

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

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. Home Land 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.

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 commercialization.

Neuromorphic computing: hardware-based neural networks

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 computer chips.

The human brain consists of roughly 86BN interconnected neurons (Figure 6) that send and receive electric impulses, or spikes, to and from neighboring neurons.

FIGURE 6: THE HUMAN BRAINS COMPRISES BILLIONS OF NEURONS




Source: Penn State, TMT Analytics

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 7), which can be seen as a gatekeeper. The human brain has approximately 150TR synapses (trillion = 1,000 billion).


FIGURE 7: SYNAPSES IN BETWEEN TWO NEURONS



Source: whfreeman.com, TMT Analytics

Synapses decide whether or not to pass on a signal to the next neuron depending on how strong, and thus relevant, the electrical impulse from the sending neuron is (Figure 8). In other words, 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 8: IMPULSES MUST EXCEED A CERTAIN THRESHOLD TO BE PROPAGATED



Source: BC Campus, TMT Analytics

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.


BRN developed a hardware version of the biological neuron

BRN'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 behavior are similar to that of a biological neuron.

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 9), 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 9: NEURON STRUCTURE IN A SPIKING NEURAL NETWORK



Source: AGH University of Science and Technology, TMT Analytics

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 in chapter 1 (see Figure 3). 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 attributing 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 favor certain connections, while they may inhibit others over time. BrainChip has significantly expanded upon the base STDP learning rule for its technology.

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 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 10). Spikes are simply 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 10: AKIDA'S SEQUENTIAL PROCESSING OF SPIKES



Source: BrainChip, TMT Analytics

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 inside the synapses provide Akida with a tremendous speed advantage compared to traditional CNNs.

For instance, training a neural network for a specific task, such as image recognition, which might take an SNN several hours, could take a CNN several days or even longer.

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.

Due to less and faster processing and less data traffic to and from memory units, power consumption of SNNs 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.

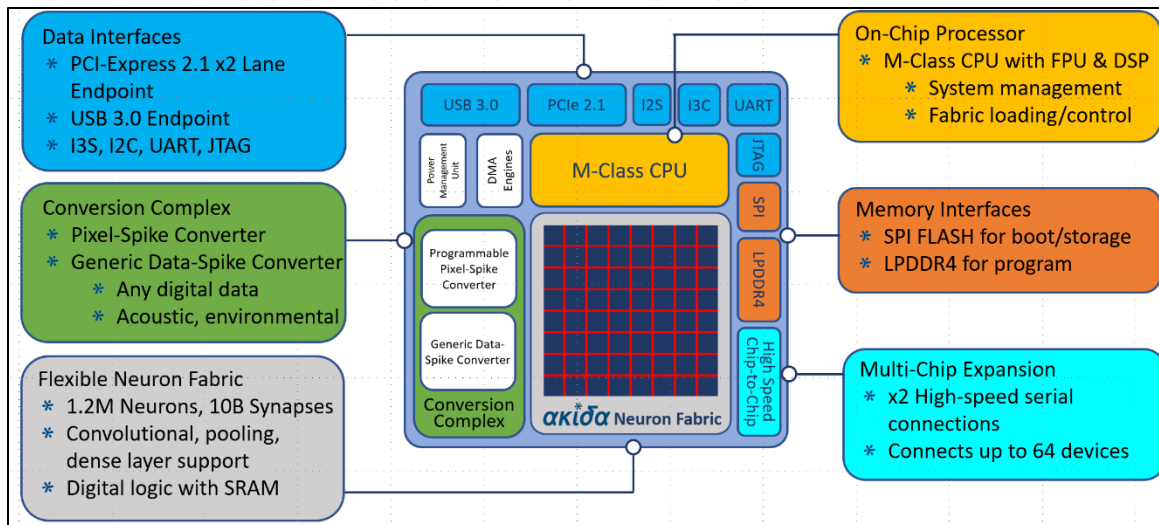
But energy consumption is also extremely important in large scale data centers. In other words, we believe hardware-only SNN's have several key competitive advantages compared to CNNs.

Akida is a Neuromorphic System-on-a-Chip

In semiconductor terminology, a System-on-a-Chip (SoC) is a complete chip solution that includes everything from a processing unit, different types of memory, I/O ports, controllers etc. In the case of Akida, with its neuromorphic computing capabilities, the complete system is dubbed Neuromorphic System-on-a-Chip (NSoC).

In its current specification, The Akida NSoC has up to 1.2M neurons and 10BN synapses embedded in the chip (Figure 11). Additionally, the chip has 7MB (Megabyte) of on-board RAM (Random Access Memory), controllers and multiple interfaces, for instance to facilitate co-processing and interconnectivity with other devices.

FIGURE 11: AKIDA NEUROMORPHIC SYSTEM-ON-A-CHIP ARCHITECTURE



Source: BrainChip, TMT Analytics

Akida also has embedded converters that convert binary data into spikes and spike trains. Using compilers and software, these converters take binary data, for instance video data, and convert it into spikes that Akida can process in its neural core. In its current iteration, Akida has

embedded data-to-spike converters for visual (pixel) data and generic data-to-spike conversions.

Learning before doing

In commercially rolling out the technology, BRN envisions that customers and customers' customers (in the case of semiconductor foundries), will specify and tune Akida chips for specific purposes, for instance image recognition in surveillance cameras or mobile phones.

Before being deployed "in the field", the chips will be operating in training mode. Once the learning phase is completed the chips can be switched to inference mode, i.e. the outcome-oriented production mode. This learning phase can also be performed off-chip.


By transferring the learning results from the trained chips to other Akida chips, these other chips will not all have to be trained individually before they can be assembled in finished end-products. This transfer of learnings is nothing more than copying the IOBN synaptic weights to an untrained chip. And even though chips in inference mode are essentially in production mode, they may still fine-tune their synaptic weights "on-the-job", i.e. individual chips will keep on learning and get better at their specific tasks over time.

Akida stacks up very favorably compared to other technologies

In addition to Akida, there are a number of other neuromorphic technologies being developed. The most comparable in terms of performance are Intel's Loihi and IBM's TrueNorth NSoCs. Intel's research test chip Loihi has 128 neuromorphic cores and has been manufactured at a resolution of 14 nanometers (nm). It has 131,072 neurons and 130 million synapses. IBM's TrueNorth NSoC has just over 1M neurons and 268M synapses and has been used in configurations of 64 chips resulting in a massive 16BN combined synapses.

As mentioned earlier, Akida has 1.2M neurons and IOBN synapses.

FIGURE 12: AKIDA COMPARED TO OTHER TECHNOLOGIES



Source: BrainChip, TMT Analytics

In comparing the performance of neural networks, standardized datasets such as CIFAR-10 are typically used. CIFAR-10 is a collection of 60,000 images in 10 different classes, including cars,

dogs, ships, birds, airplanes etc., and is specifically used to train neural networks in object recognition.

In processing CIFAR-10 datasets, TrueNorth has shown to be less accurate than Akida (89% versus 82%) in recognizing objects in pictures while being able to process 6,000 frames per second per watt. This compares to 1,000 frames per second per watt for Akida (Figure 12). No data for Intel's Loihi is available.

However, at a price of US\$ 1,000 per chip, TrueNorth is 100x more expensive than the targeted commercial price point for the Akida NSoC of just ~US\$ 10. We believe this low price point will be very beneficial to BRN when it comes to driving adoption of the Akida technology.

Horses for courses

As is common in the semiconductor industry, customers may not always need a chip designer's full IP (Intellectual Property) stack. For instance, certain customers may only require certain IP blocks to complement their existing IP. This is often how cell phone manufacturers work, i.e. they embed third-party IP into their proprietary SoCs. Other future customers may want the full Akida NSoC, for instance to embed Akida in ADAS applications. Naturally, BRN's commercial pricing of Akida IP will depend on how much IP and which elements of its IP customers require.

So when it comes to the Akida specifications, a distinction needs to be made by Akida end-application. IoT-edge applications are devices at the edge of the Internet of Things (IoT), such as smart meters, soil moisture sensors, mobile phones, desktop computers, connected doorbells, robots in factories and smart TVs. The core of the IoT is made up of routers, switches as well as on- and off-premise data centers, collectively known as the Cloud.

From a specifications point of view, one of the key differences between edge devices and the Cloud is the scale of processing required. Data centers are Enterprise level computer farms in which individual chips in servers typically operate at more than 100W. This is why data centers consume vast amounts of energy resulting in high heat generation inside individual chips.

Chips in edge devices, on the other hand, are quite often very low power devices operating at 1W or less. Therefore, they operate at much lower temperatures.

Consequently, to cope with these distinct categories of end applications, the Akida chipset will need to have different specifications, i.e. to handle different transfers of energy and heat dissipation. More specifically, the circuitry for Akida in server farms will likely need to be different from the current low energy version of Akida. Additionally, in order to manage the different heat profile in server farms, Akida for data centers will likely require different packaging. Low energy chips are typically packaged using plastics or resins. High energy chips are typically packed using ceramics.

Strong IP protection through 10 granted and pending patents

As is illustrated in the appendix, BRN currently owns 10 patents, 3 of which granted, around neuromorphic computing, spiking neural networks, unsupervised learning, visual pattern recognition etc.

Commercialization of Akida in various steps

While the core development of the Akida NSoC is largely completed, a finished sample product in silicon, i.e. a chip, is not expected to be available before the second half of 2019. However, in July 2018, BRN launched the Akida Development Environment (ADE), which can be used for creation, training and testing of the Akida NSoC as well as to run inference (outcome-based processing) on Akida.

The ADE is a software simulation of the Akida chipset. Users can run various training methodologies, including unsupervised learning, using different datasets, such as CIFAR-10, MNIST and GoogLeNet. Data-to spike-converters are also included so users can feed pixel, audio and Big Data into the development environment.

The purpose of ADE is to provide prospective Akida customers and current BrainChip Studio customers (see below) the opportunity to familiarize themselves with the Akida architecture and its specifications as well as to illustrate the possibilities and opportunities that Akida presents.

By the time Akida becomes available in sample chips, in the second half of 2019, customers and prospects should already have a solid understanding of how they may implement Akida in their respective end-applications. It is expected that this will speed up the subsequent step in commercializing Akida, i.e. designing Akida into customers' system and chip architectures.

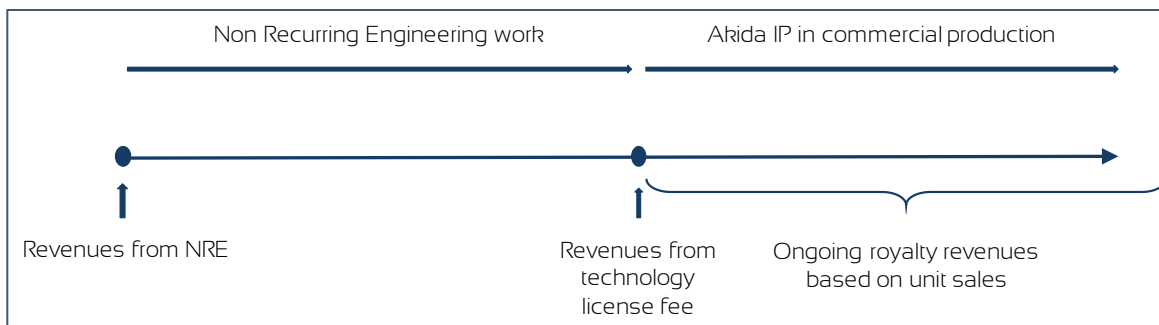
Once this design-in phase and subsequent testing are completed, sales of actual Akida NSoCs and third-party chips with embedded Akida technology can commence.

Akida to generate three separate revenue streams

Given the enormous costs to build computer chip factories, BRN itself will not be manufacturing Akida chips. Rather, the company will be using a partner, SocioNext, to manufacture the devices for BRN.

In addition, BrainChip will selectively license its Akida IP to other providers in the semiconductor industry. These IP licensing revenues from Akida are likely to take the form of Non-Recurring Engineering (NRE) work, license income and royalties from each chip sold by BRN's customers (Figure 13).

FIGURE 13: ANTICIPATED REVENUE PROFILE FOR AKIDA COMMERCIALIZATION



Source: BrainChip, TMT Analytics

Revenues from Non-Recurring Engineering work

NRE work conducted by BRN on behalf of the customer will typically be focused on integrating the Akida chip architecture into the customer's proprietary chipset (design-in) to create a complete NSoC. Revenues from NRE are one-off while the NRE work itself can take anywhere from 6 to 18 months. The NRE fees vary widely across the industry, from several hundreds of thousands of dollars to several millions of dollars.

One-off license fees

Once the NRE work is completed and customers want to move into commercial production, they typically pay a one-off license fee for use of the technology. License fees can range quite broadly and will be different from one customer to another, depending on the intended application areas, the amount of IP to be used, expected volumes to be manufactured etc.

Recurring royalties from every unit sold by a customer

We believe the most lucrative future revenue stream for BRN will be royalties paid by customers for each chip they sell that includes Akida IP. These royalties are usually a percentage of the customer's selling price and typically range from 2% to 15%, again depending on intended application areas, the amount of IP used and expected volumes to be manufactured.

However, royalty percentages also depend on the uniqueness of the IP that is being licensed. We would argue that Akida's specifications and features are quite unique when compared to other technologies, such as Intel's Loihi and IBM's True North, which may result in BRN being able to charge above average royalty percentages for Akida.

Other royalty revenue models simply use a fixed dollar amount per chip sold, for instance if the customer only uses a limited IP stack.

Example: Assume a BRN customer agrees to pay a 5% royalty on sales from each of its 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.

Royalty income for BRN from this customer would amount to US\$ 1.25M, US\$ 6.25M and US\$ 12.5M respectively in the first three years of production.

In our view, the example above clearly illustrates the scalability of BRN's revenue model as well as the revenue potential for Akida. We believe there are likely to be multiple prospective customers in each addressable industry vertical, such as mobile phones, security cameras, visual inspection, IoT devices, cloud/server usage, cyber security, financial markets etc.

Selling direct and through third-parties

BRN is in discussions with various Akida prospects in different markets (see below), which may result in commercial agreements in due time. However, in addition to these potential direct sales, we believe future sales through third parties may potentially become a lot bigger than BRN's direct sales.

Firstly, should semiconductor foundries license BRN's technology, Akida may find its way into products from many different end-customers. Foundries such as TSMC (TPE:2330) and Global Foundries manufacture computer chips for third parties, including cell phone manufacturers, electronics companies and other chip manufacturers. Over the last 15 years foundries have evolved from "simply" manufacturing chips according to customers' chip designs into value-

added chip manufacturers that can add their own IP to customers' designs. Foundries develop this additional IP themselves, but also source third-party IP to complement their own offerings.

Secondly, BRN could potentially license Akida IP to providers of IP building blocks, like Lattice Semiconductor (NSDQ:LSCC), Cadence (NSDQ:CDNS) and Mentor (NSDQ:MENT). These companies provide core IP as well as separate IP blocks to their customers and we believe Akida IP could be a welcome addition to their offering.

In both channel partnership models, BRN would likely earn license fees and royalties, either a fixed amount or a percentage of revenues.

ADAS and Autonomous Vehicles among high-priority target markets

BRN is currently targeting specific customers and applications in Embedded Vision, Cyber Security and general IoT.

We believe Embedded Vision for the Automotive industry in particular should be considered of high interest when it comes to commercial opportunities for Akida. ADAS and Autonomous Vehicles require very substantial capabilities in the areas of object recognition and processing of radar and LiDAR (Light Detection and Ranging) data.

BRN is currently working with an unnamed European Automotive company in this area and we would expect BRN to be able to sign additional Automotive-related customers going forward, car manufacturers as well as Original Equipment Manufacturers (OEMs) that supply car manufacturers.

Mobile phone industry is increasingly incorporating AI

Another high-growth area for Embedded Vision is mobile phones. Manufacturers are increasingly incorporating AI into mobile phone camera software. However, we believe this could be complemented (or potentially be largely displaced) by hardware-based neuromorphic computing inside the mobile phone. Additionally, we see potential for software-based mobile applications, such as Google Assistant and Machine Learning applications, that currently require Cloud-access to function, to be replaced by on-board NSoCs, such as Akida.

We will elaborate on BRN's commercial opportunity in the chapter on the market for Deep Learning chipsets below.

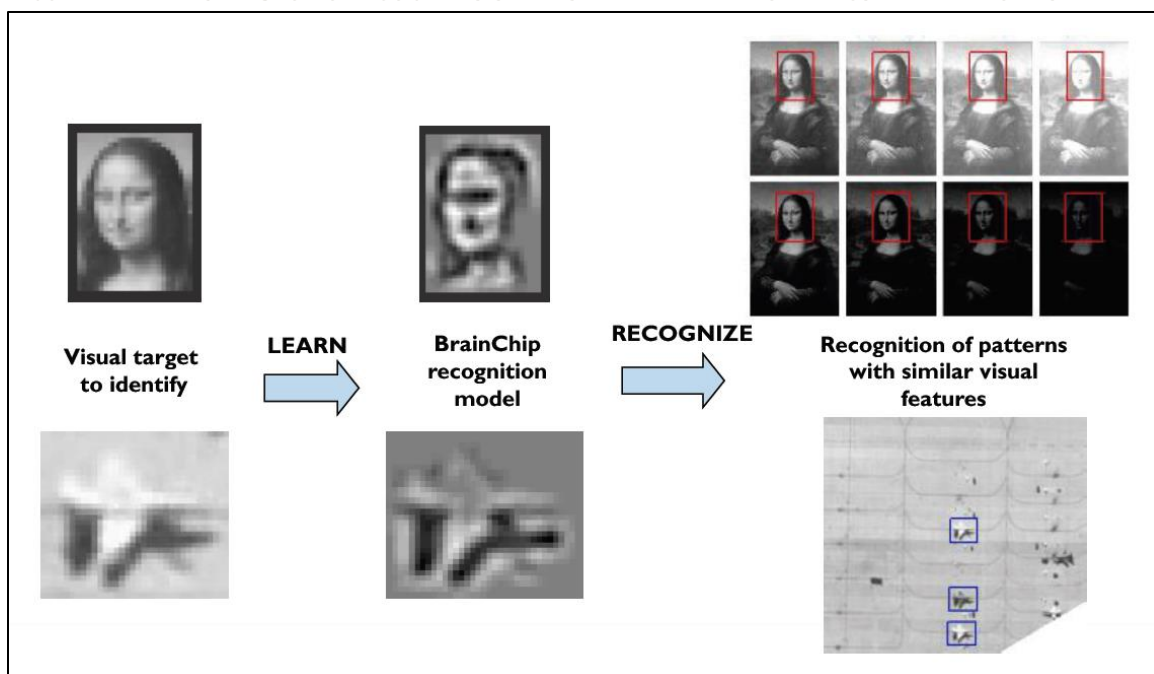
BrainChip Studio: a software-based SNN

In addition to Akida, BRN has another product line, BrainChip Studio (BCS), which the company acquired in 2016 (Spikenet acquisition). BCS is an AI-powered video analysis software suite that enables high-speed object search and provides facial recognition and classification capabilities. As opposed to Akida, BCS is software-based, i.e. video streams can be fed into BCS where they are processed using a software-based SNNs.

Contrary to most facial and object recognition systems on the market today, BCS works through shape recognition rather than through biometrics and measuring key facial points (Figure 14). BCS can operate using low resolution video and small image sizes (minimum of 24x24 pixels). Additionally, BCS can be trained in milliseconds using a single image, i.e. one-shot training.

Because of this unique training methodology, BCS is able to process vast amounts of data to extract faces and images in very short time frames.

FIGURE 14: BRAINCHIP STUDIO RECOGNIZES SHAPES RATHER THAN BIOMETRICS AND KEY POINTS



Source: BrainChip, TMT Analytics

Challenges with today's video analysis tools

The key challenges with this sort of image recognition and classification is the tremendous amount of video data available today. In addition to the massive video libraries that law enforcement and intelligence agencies already have, substantial amounts of video data are created every day.

Furthermore, many of the Deep Learning technologies in use today are CNNs that require substantially more compute time for the same tasks compared to the SNN used in BCS. Consequently, in addition to being slower, these algorithms consume lots of energy and are simply more expensive to operate than BCS for the same tasks.

Another challenge for object search, facial recognition and classification is video quality, which is often quite poor. This can lead to higher levels of inaccuracy in recognition and classification using today's CNNs.

Accelerator card provides up to 7x performance boost versus GPU acceleration

In order to further speed up the performance of BCS, BRN introduced the BrainChip Accelerator (Figure 15) in September 2017. This accelerator enables processing of 16 video channels simultaneously, which results in higher overall speed of the system, i.e. BrainChip Accelerator increases the number of processed frames per second (fps) by 5x to 6x.

FIGURE 15: BRAINCHIP ACCELERATOR CARD FOR BRAINCHIP STUDIO



Source: BrainChip, TMT Analytics

Additionally, compared to CNNs that use GPUs to accelerate image and video processing, BrainChip Accelerator delivers a 7x processing improvement of the number of frames per second per watt.

Many application areas for BrainChip Studio

The primary functionality of BCS is facial and image recognition and classification as well as forensic object search. There is a wide application area for this sort of functionality, including;

- Gaming, e.g. monitoring of casino gaming tables,
- Civil Surveillance & Intelligence, e.g. to monitor crowds and high-value areas (military bases, hospitals etc) by law enforcement agencies, and

BRN has been actively commercializing BCS for the last two years and has a number of customers and channel partners for BCS in the Gaming industry (Gaming Partners International) and law enforcement (multiple French police departments).

Through a global licensing, development and revenue sharing deal, BRN and its partner have developed a video analytics products for use in casino currency security (to prevent the use of counterfeit casino chips), game table operations and player behaviour analytics. BRN has received US\$ 500k in license fees and US\$ 100k for Non-Recurring Engineering work. The terms of the revenue sharing component have not been disclosed.

BRN is also working with DGSI, the French intelligence agency, and the French National Police on the implementation into the security applications these organizations operate. In our view, such collaborations provide substantial credence for BCS and should help BRN in its further commercialization of BCS to similar organizations outside of France.

Going forward, we see tremendous opportunity for BRN in the Automotive industry, both with car manufacturers and OEMs given that ADAS and Autonomous Vehicles require a lot of image and video processing.

In April 2018, BRN announced an agreement with Veritone (NASDAQ:VERI) to integrate BCS into Veritone's aiWare, a platform for cognitive AI engines. Veritone's customers can use aiWare for their specific purposes requiring AI capabilities. Through the Veritone Developer Program, developers can build custom applications on top of Veritone's platform and can now also incorporate BCS in their applications, which would generate revenues for BRN.


Deep Learning chipsets market to grow at ~60% CAGR

Deep learning, like other AI elements, has been undergoing a rapid evolution, and the semiconductor industry has followed suit with a wide array of hardware and software solutions. These offerings can be grouped under the deep learning chipset market or AI acceleration chipset market (as termed by BRN).

Deep learning chipsets are broadly categorized into central processing units (CPUs), graphics processing units (GPUs), field programmable gate arrays (FPGAs), application-specific integrated circuits (ASICs), Systems-on-Chip (SoC) accelerators and other architectures such as neuromorphic and digital signal processors (DSPs). The choice of architecture while running a deep learning application depends on the level of performance and programming as well as power consumption requirements. In addition, the nature of problems (training or inference) also determines the suitability of an architecture. For instance, ASICs and neuromorphic processors provide maximum suitability for training purposes, followed by GPUs, although currently GPUs are more in use due to abundant availability and ease of programming.

According to Tractica, the deep learning chipset market was valued at US\$ 1.6BN in 2017 and is anticipated to grow at a CAGR of 59.3% during 2017-2025 to reach US\$ 66.3BN by 2025 (Figure 16). The growth is anticipated to be driven by increasing use cases for AI applications, growing AI training needs across enterprises, and rising adoption of inference use cases in smartphones, agriculture, manufacturing, cars and healthcare.

FIGURE 16: DEEP LEARNING CHIPSET MARKET (IN US\$BN)



Source: Tractica, TMT Analytics

Inference to experience the fastest growth

The deep learning chipset market for training was estimated at US\$ 0.6BN in 2017 and is poised to grow at a CAGR of 44% during 2017-2025 to reach US\$ 11.1BN by 2025. The inference market constituted ~50% of the overall market and generated US\$ 0.8BN in revenues in 2017. This

market is anticipated to grow at a CAGR of 68.7% during 2017-2025 to reach US\$ 52.4BN by 2025. The general-purpose category, which includes both inference and training, is projected to reach US\$ 2.8BN by 2025.


Edge computing to drive deep learning chipset sales

In terms of usage, the market is segmented between enterprise/data center and edge devices. The data center market consists of AI workstations and servers, while the edge devices segment includes Automotive, Drones & Robotics, Security cameras, mobile phones, tablets, and Gaming and Medical Devices. The data center market uses deep learning networks for both training and inference, while there is hardly any training being done on edge devices anymore as most edge devices are in "production mode", i.e. used for inference.

The market is currently dominated by data centers. However, Tractica predicts that the edge computing market, where AI processing is done on the device itself, will represent >75% of the total market by 2025, with cloud/data center environments expected to account for the rest of the market.

Tractica forecasts that AI-enabled edge device shipments will grow at a CAGR of 48.7% to reach 2.6BN units in 2025 from just 161.4M units in 2018 (Figure 17). Currently, edge computing largely takes place on mobile phones and PCs/tablets. However, during 2018-2025, several other categories are expected to emerge, including smart speakers, security cameras, head-mounted displays, drones, cars as well as consumer and enterprise robots.

FIGURE 17: AI-ENABLED EDGE DEVICE SHIPMENTS (IN M UNITS)



Source: Tractica, TMT Analytics

Data centers moving away from GPUs to FPGAs and ASICs

Deloitte estimates that deep learning chip unit sales for data centers will have almost quadrupled during 2016-2018 from 0.2M units in 2016 to 0.8M in 2018 with nearly 63% of sales estimated to comprise GPUs in 2018. The firm also predicts that by the end of 2018, >25% of all chips used to accelerate deep learning in the data center will be FPGAs and ASICs. These chips are anticipated to substantially increase the use of deep learning, enabling applications to use less power and at the same time become more responsive and capable. This is also likely to expand the total available market for both training and inference-based deep learning acceleration.

NSOCs to bring neuromorphic computing to the edge

Gartner anticipates that neuromorphic chips currently have <1% penetration of the total target market. However, it considers the technology transformational, due to the capability to

accelerate deep learning tasks, low power consumption and faster implementing of training compared to GPUs, which underlines the significant potential for growth of neuromorphic chip sales going forward.


In addition, Gartner expects neuromorphic chips to be initially used in IoT edge devices due to their ability to execute certain levels of neuromorphic computing at the edge, reducing bandwidth and central processing requirements.

We believe the increased penetration of neuromorphic computing architectures beyond 2019, owing to their inherent advantages over GPU-based systems, will drive BRN's share in the fast-growing deep learning chipsets market.

Neuromorphic computing market is still small but set to take off

According to Knowledge Sourcing Intelligence the neuromorphic computing market was valued at US\$ 28.3M in 2017 and is poised to grow at a CAGR of 51.2% during 2017-2023 to reach US\$ 338.1M by 2023 (Figure 18). The major factors driving the growth include increasing use of neuromorphic chips for developing brain-based robots and cognitive robots, and growing demand of AI for applications such as language and image processing, computer vision, non-linear controls & robotics and translation & chatbots.

FIGURE 18: GLOBAL NEUROMORPHIC COMPUTING MARKET (IN US\$M)



Source: Knowledge Sourcing Intelligence, TMT Analytics

In terms of applications, image recognition accounts for the largest (41%) share of the neuromorphic computing market. Image recognition applications are primarily used in the Aerospace & Defense and Surveillance & Law Enforcement industries where video monitoring and machine vision enable tracking and surveillance. The segment is expected to remain dominant during 2017-2023 owing to the rise in demand for digital cameras and other imaging systems. The signal processing segment accounts for ~30% of total market revenues, considerably boosted by rising demand for processing of audio and acoustics signals.

According to Grand View Research, North America accounts for the lion's share of the market, primarily due to the presence of key neuromorphic chip manufacturers in the region. During 2017-2023, North America is expected to continue to outperform other regional markets. The growth in this region can be attributed to increased penetration of devices with unique voice and image identification capabilities in defense, wearables, the IoT and in robotics technology.

In Europe, significant investments are being made in neuromorphic projects and related R&D activities by prominent universities, such as University of Manchester and Heidelberg University. These investments are expected to offer substantial growth opportunities in the region. In Asia Pacific and South America, the neuromorphic market is in a nascent stage, and primarily driven by growing automation demand in countries such as China, India and Brazil.

As far as the end-use segments are concerned, Consumer Electronics accounted for 60% of the total market in 2016, followed by the Automotive, Healthcare and Military & Defense segments. The Automotive end-use segment is expected to see the highest growth during 2017-2023 driven by the increasing roll-out of autonomous cars and smart vehicles. The Military & Defense segment is also forecast to grow considerably, owing to rising demand for neuromorphic chips in satellites for aerial imagery and surveillance.

BRN's key end-use markets to witness substantial growth

Major breakthroughs are being achieved in the fields of image and speech recognition, and signal processing through various competing AI technologies. Therefore, through its proprietary Akida NSoC offering, BRN plans to cater to emerging use cases and application areas within computer/embedded vision, cybersecurity and IoT. The market for neuromorphic processing in all three verticals is poised to grow at healthy double-digit rates, offering significant growth potential (Figure 19).

FIGURE 19: OPPORTUNITY ANALYSIS BRN'S KEY END-USE MARKETS

End-use Market	AI-based Market Size (in US\$BN)		CAGR (2017-2023)	BRN's Focus Areas
	2017	2023		
Computer Vision	2.4	25.3	48.1%	Object classification, advanced driver assistance systems (ADAS), surveillance, vision guided robotics
Cybersecurity	5.0	22.8	28.8%	Packet inspection, file property classification, anomaly detection
AI hardware at the Edge	610M units in 2019	1.6BN units by 2024	21% ('19-'24)	
AI at the Edge software	US\$ 356M in 2018	1.2	26.5% ('18-'23)	

Source: TMT Analytics, Technavio, MarketsandMarkets

ADAS and Autonomous Vehicles are driving the computer vision segment

AI-based computer vision is gaining momentum supported by rapid evolution of deep learning and CNNs. These technologies are empowering machines with technical capabilities such as identification, classification, positioning, detection and segmentation of objects and scenes within pictures and video. This will enable their wider use in video surveillance, automated driving, vehicle/face recognition, autonomous robotic navigation, medical image analysis, augmented reality/virtual reality development, localization and mapping, etc.

Within the AI-based computer vision market, the Automotive vertical is projected to witness the fastest growth boosted by the ongoing need to improve road safety. In addition, various

Automotive and IT giants are making huge investments to develop Autonomous Vehicles, which in-turn is propelling the demand for AI-based embedded vision solutions. Key examples include Amazon working with Toyota to develop a fleet of autonomous vehicles for last-mile delivery and the ongoing "autonomization" of high-end cars from manufacturers such as Tesla, Audi and BMW.

Cybersecurity market driven by Banking, Financial Services and Insurance

Simultaneously, the AI-driven cybersecurity market is projected to grow primarily due to digital transformation of companies and the rapid increase in cloud-based services driving the need for enhanced security against cyber-attacks.

Companies are expected to increasingly deploy AI-based cybersecurity solutions for real-time identification, response and defense against network attacks, prevention of malware, network traffic anomaly detection, application security detection and network risk assessment.

According to CB Insights, 11 of the top 100 AI startups are focused on different aspects of cybersecurity from anomaly detection for web security. Furthermore, the growing number of mobile phones and IoT edge devices is leading to a rise in demand for deep packet inspection (DPI) services, which primarily highlight network trends, help internet service providers (ISPs) optimize bandwidth and throughput, and can reveal user behavior.

In terms of end-use segments, the Banking, Financial Services and Insurance industries are the largest consumers of AI-based cybersecurity services owing to rising demand for financial fraud detection services. These industries accounted for ~30% of the market in 2017 and are expected to maintain their dominance in the 2018-2023 timeframe.

AI at the Edge to drive uptake of small AI-enabled devices

Artificial Intelligence at the Edge essentially means that IoT devices at the edge of the network, e.g. sensors, smart watches, security cameras etc, are able to locally process the data generated by these devices themselves, rather than having to send all data to the Cloud for processing. Only in circumstances where a Cloud connection is required, for instance in the case of an alarm, alert and reporting, will the Edge device connect to the Cloud.





We anticipate strong growth for AI at the Edge driven by newly emerging applications in areas such as Smart Cities, Smart Homes and Industrial IoT. MarketandMarkets expects AI at the Edge software to grow from US\$ 356M in 2018 to ~US\$ 1.2BN in 2023.

AI hardware at the Edge is expected to grow from 610M units in 2019 to an estimated 1.6BN units by 2024, according to Research and Markets. In our view, BRN should be able to play a significant role in this market segment given the Akida attributes compared to other AI hardware solutions out there.

Very limited competition in Neuromorphic Computing chipsets

The global deep learning market is highly fragmented with many large and small players. There are approximately 80 companies competing in the market including the world's leading semiconductor companies, tech giants and start-ups. In the next few years, the competition is expected to intensify as new vendors with differentiated products and capabilities continue to foray into this market. The competitors in this market can be broadly categorized into four categories: IC vendors, tech giants & high-performance computing (HPC) vendors, IP vendors and start-ups (Figure 20).

FIGURE 20: DEEP LEARNING VENDOR LANDSCAPE

Vendor Type	Description	Number of Vendors	Sample Vendors
IC Vendors	<ul style="list-style-type: none"> Traditional semiconductor manufacturers Currently, leading the hardware acceleration in the deep learning market, which was dominated earlier by software 	13	
Tech Giants & HPC Vendors	<ul style="list-style-type: none"> Leaders in deep learning software Focus is on building own deep learning hardware to reduce dependencies on IC vendors and capitalize on anticipated hardware acceleration in the market 	12	
IP Vendors	<ul style="list-style-type: none"> Traditional providers of system design tools, software, IP and services for applications such as vision, deep learning, audio, communication and connectivity Focus is on offering DSPs for neural network acceleration 	7	
Start-ups	<ul style="list-style-type: none"> Typically, small companies with focus on a single deep learning chipset Often supported by VCs or funding provided by IC vendors or tech giants 	47	

Source: TMT Analytics, Medium

There are only 14 Deep Learning chipset manufacturers, including Intel, NVIDIA, IBM and Xilinx, focused on both data center and edge applications, whereas the majority of the manufacturers are focused on edge applications only. We believe this further corroborates the growth anticipated in the edge computing market. Furthermore, Intel has the most diversified portfolio of deep learning chipsets, which include CPUs, FPGAs, ASICs and neuromorphic chipsets.

When it comes to market share, NVIDIA is the clear leader in the GPU market segment. It has close to 70% share in global GPU sales. NVIDIA is also the market leader in the Deep Learning training market. However, in the Deep Learning inference market, Intel claims to have more than 80% global market share.

While the competition in CPUs, FPGAs and ASICs is quite strong and intensifying, we believe BRN, with its IP for Neuromorphic Computing chipsets, has very few competitors in the

neuromorphic computing market (Figure 21). Also, the majority of neuromorphic systems available today are at the early prototype stage.

However, the companies that are active in this space, such as Intel and IBM, are all established and very large.

FIGURE 21: KEY COMPETITORS IN THE NEUROMORPHIC COMPUTING MARKET


Company	Headquarters	Product Portfolio	Applications
IBM	US	IBM TrueNorth Neuromorphic Chips	<ul style="list-style-type: none"> • Embedded deep learning applications on smartphones and robots • Defect detection in 3D Printing • Deep learning inferencing in HPC data centers
Hewlett Packard	US	Dot Product Engine	<ul style="list-style-type: none"> • Suitable for HPC applications
Intel	US	Loihi Chips	<ul style="list-style-type: none"> • Suitable for devices that need to learn in real-time such as drones and autonomous cars • Embedded vision for cameras
General Vision	US	CMIK	<ul style="list-style-type: none"> • Suitable for solving pattern recognition problems, text and data analytics, and vision applications

Source: TMT Analytics

Rising consolidation in AI and deep learning markets

Overall, the AI market is witnessing increasing consolidation. Big corporations are trying to integrate machine learning and deep learning into their products. However, lack of AI talent has led these companies to acquire AI start-ups to build their capabilities. According to CB Insights, 115 AI start-up acquisitions were reported in 2017, an increase of 44% compared with 80 acquisitions in 2016 (Figure 22).

FIGURE 22: NUMBER OF AI START-UP ACQUISITIONS



Source: CB Insights

Google is the leading acquirer of AI start-ups with 14 acquisitions since 2013. Intel has also expanded its capabilities in the deep learning chipsets market by acquiring key players such as Nervana Systems, Movidius and Mobileye (see below). Other top AI start-up acquirers include Apple, Facebook, Microsoft and Amazon.

Given the company's attractive IP portfolio, with 10 patents granted or pending, and the fact that the commercialization of Akida has already commenced, we wouldn't rule out BRN becoming an acquisition target at some point.

M&A provides guidance to BRN's potential valuation

BRN currently generates modest revenues from product sales (BrainChip Studio), technology licenses and development work, i.e. US\$ 505k in 1HY18 (through June 2018). However, forecasting revenues for 2019 and beyond would be speculative at this stage, given that the timing of new client wins is impossible to predict, both for BrainChip Studio and for Akida, which is only in the early stages of its commercialization process.

And given that we believe Akida will be the key driver of BRN's long term value, we will need more commercial data points before we can model revenues from Akida properly, and hence model BRN's long term valuation on a reasonable basis.

Intel has paid ~US\$ 400M for comparable technologies on multiple occasions

However, we believe looking at the semiconductor industry's willingness to pay for unique IP around neuromorphic processing provides a real-world gauge of BRN's potential valuation.

In August 2016, Intel announced the acquisition of US-based Nervana Systems for US\$ 408M (source: Recode). Nervana is developing a hardware-based neuromorphic computing solution optimized for unsupervised learning. The company was established in 2014 and received US\$ 24.4M in Venture Capital (VC) funding prior to being acquired by Intel.

In September 2016 Intel also acquired the Irish company Movidius, which develops ultra-low power vision processors for computational imaging and vision processing for use in industrial and consumer electronics, including Virtual Reality (VR) devices and drones. Movidius was established in 2008 and prior to being acquired, the company received a total of US\$ 86.5M in VC funding. According to CrunchBase, Intel paid US\$ 400M for Movidius.

While no revenue data was disclosed for either Nervana or Movidius, we believe that revenues for either company at the time of acquisition would have been very small, if in fact they generated revenues at all.

Nervana was only established in 2014 and was still in early-stage development mode. While Movidius had been in operation for eight years up to the moment of acquisition, it seems their deals with customers/partners were mostly early-stage pilot projects, which typically do not generate substantial revenues.

In other words, given both companies' development profiles for hardware-based unsupervised learning and our assumption of low levels of revenues for both companies, we believe the take-out prices for Nervana and Movidius two years ago provide a reasonable proxy for BRN's potential value.

Applying the approximate ~US\$ 400M (~A\$ 550M) average equity valuation for Nervana and Movidius to BRN would yield a value of A\$ 0.52 per share (A\$ 0.43 on a fully diluted basis). While we are not suggesting BRN will be acquired anytime soon, we do believe that the company's Neuromorphic Computing IP could potentially be valued at these levels.

AudioPixels is valued at ~A\$ 560M, implying a value of A\$ 0.44 per BRN share

ASX-listed Audio Pixels Holdings Limited (ASX:AKP) is developing new techniques to generate sound waves directly from digital audio streams using micro electromechanical structures (MEMS) rather than conventional loudspeaker elements. In other words, AKP aims to generate sound by using tiny devices that are manufactured using standard semiconductor

manufacturing processes. We believe AKP is some time away from generating revenues, whereas BRN is already generating revenues, albeit modest revenues at this time.

Similar to BRN, AKP is currently in the process of transferring its technology to silicon. The company is currently valued at approximately A\$ 560M, which implies a potential value for BRN of A\$ 0.53 per share (A\$ 0.44 per share fully diluted).

Standalone commercialization would create highest value for shareholders

Even though several comparable Neuromorphic Computing companies have been acquired by large industry players in recent years, we believe it is BRN's intention to commercialize the Akida IP rather than pursue a near to medium term exit through M&A. In that standalone scenario, in which BRN should be able to attract multiple licensees for the Akida IP in various industry verticals, we can see substantial revenue potential.

To illustrate this point, assume an individual licensee of Akida IP pays 2% royalties per chip sold at an Average Selling Price of US\$ 25. Based on this customer's unit sales in Figure 23, BRN would generate ~US\$ 37.5M (A\$ 51M) in royalty income during the first five years of the customer's production of this particular chip using Akida technology.

FIGURE 23: REVENUE EXAMPLE FOR INDIVIDUAL CUSTOMER (US\$M)

Royalty %	2%	BRN
ASP (US\$)	25	revenues
Unit sales Y1 (M)	5	2.5
Unit sales Y2 (M)	10	5.0
Unit sales Y3 (M)	15	7.5
Unit sales Y4 (M)	20	10.0
Unit sales Y5 (M)	25	12.5
Revenues to BRN over 5 year-period (US\$ M)		37.5

Source: TMT Analytics

Please note we haven't included income from NRE and license fees. And while all the variables in this example are arbitrary, this exercise does illustrate BRN's scalability with royalty income for a marginal customer falling straight to the bottom line.

Successful commercialization of Akida in the next five years could potentially see BRN win dozens of customers in a range of different verticals. Hence, we believe that the valuations seen in recent M&A transactions are an intermediate step towards higher valuations longer term, which we expect to be driven by BRN successfully executing on its commercialization strategy.


BRN is undervalued at current share price levels

At BRN's current share price, the company's valuation of approximately A\$ 50M is substantially lower than valuations seen in M&A transactions as well as the valuation of its ASX-listed peer AKP.

Outstanding options could bring in A\$ 40M in fresh capital

BRN currently has 190.55M options outstanding at various exercise prices (Figure 24). These options could potentially generate A\$ 40.2M in fresh capital if and when exercised.

FIGURE 24: EXERCISE PRICES (Y AXIS) FOR BRN'S 190.5M OUTSTANDING OPTIONS (X AXIS)



Source: ASX, TMT Analytics

Conclusion

In November 2018, BRN indicated that it was involved in 21 active or committed pilot projects for Akida and BCS combined, in addition to 17 design wins, 55 qualified sales opportunities and more than 500 leads.

Upon successful completion of transitioning the Akida IP from lab to fab, expected in the second half of 2019, we believe BRN should be in a position to secure a variety of different customer types for Akida, including cell phone manufacturers, semiconductor foundries, Automotive OEMs for ADAS and Autonomous Vehicles, third-party semiconductor IP providers, IDMs (Integrated Device Manufacturers) and companies in the Imaging space.

Given that BRN will be commercializing Akida through a manufacturing partnership with Socionext and an asset-light licensing model with recurring royalty revenues, we anticipate high gross margins once Akida sales ramp up. Additionally, in the near to medium term we expect BrainChip Studio to increase its traction through direct sales and through BRN's channel partnerships, such as with GPI.

Valuation range of A\$ 0.40-0.45 if Akida NSoC shows commercial viability

Looking at where Akida currently is in its development/commercialization process, we are of the opinion that BRN's share price does not accurately reflect the company's commercial potential described above. In our view, the company's valuation could move towards levels seen for comparable companies, such as Nervana and Movidius that were acquired by Intel Corp (NASDAQ:INTC), i.e. equivalent to A\$ 0.43 per fully diluted BRN share. Similarly, we believe AudioPixels' (ASX:AKP) A\$ 560M valuation, equivalent to A\$ 0.44 per fully diluted BRN share, also provides a gauge to BRN's potential valuation, given where AKP is positioned in its development process.

In the near to medium term, though, we believe BRN will first need to demonstrate commercial viability of the Akida NSoC by signing one or more commercial agreements that involve integration of Akida and/or components of the Akida IP stack into customers' designs.

Given the strong upside we can potentially see for the BRN's share price if Akida is commercially validated, we start our coverage with a Speculative Buy recommendation.

Near-term share price catalysts

- Conversion of current discussion partners into paying customers and development partners, e.g. to develop and deploy Akida for specific applications such as for cell phones and ADAS.
- Updates on the Akida commercialization roadmap, specifically around first silicon test results.
- Additional customers for BrainChip Studio, especially in the Gaming and Law Enforcement sectors.
- One of the Top-20 shareholders, Metals X Limited, seems to have large sold its position on market recently. The selling pressure from the sale of these 11.9M shares may have exacerbated the downward pressure on BRN's shares. With selling pressure from Metals X now gone, we may see a bounce in BRN in the short term.

SWOT Analysis

Strengths

- The Akida IP is unique in that it combines fast and low power Neuromorphic Processing with a targeted market price point that is expected to be well below similar technologies.
- The asset-light IP licensing model should make for very high gross margins once sales of Akida IP ramp up.
- BRN is already working with a large number of prospective Akida customers on the testing and integration of Akida NSoCs..

Weaknesses

- The Akida IP has yet to be transferred into an actual chip. This transfer process may encounter setbacks, which would push out the timeframe for delivery of first sample chips, currently planned for the second half of 2019.
- BRN will be competing with larger industry players that may be able to more easily fund IP development and run pilot projects with customers, which may inhibit BRN's Akida commercialization process.
- Given the average quarterly cash burn of US\$ 1.7M in FY18, with US\$ 10M in cash on the balance sheet per the end of the September quarter, BRN may need to raise additional capital to fund the company until it reaches cash flow break even, diluting current shareholders.
- Capital restrictions may limit BRN in developing IP, in addition to what the company is already working on, potentially inhibiting future growth from new products.

Opportunities

- The market for Deep Learning chipsets is potentially very large with Tractica estimating a market size of US\$ 66BN by 2025, compared to US\$ 5BN today. This implies a CAGR of nearly 45%.
- Additionally, most of the Deep Learning chips sold today are not dedicated Neuromorphic chips like Akida but general-purpose chips, such as GPU's, tailored to perform certain tasks such as image recognition. Hence, we see a very large displacement potential in the Deep Learning chipset market for Akida.

Threats

- Global tensions around ownership and theft of semiconductor IP between the United States and China may result in Western semiconductor companies, including BRN, being restricted in the IP they are allowed to sell to Chinese companies. This would potentially restrict BRN's growth in one of the world's largest markets for semiconductors.
- In generating revenues from its Neuromorphic System-on-a-Chip, BRN will be competing with some very large players, such as Intel that has acquired a number of companies similar to BRN in the last few years.

APPENDICES

Board members

Stephe Wilks (Non-Executive Chairman): Mr. Wilks joined the board in February of 2019 and currently serves as Non-Executive Director of BluGlass Limited (ASX:BLG), Non-Executive Director of DataDot Technology Limited (ASX:DDT), Non-Executive Director and Chairman of Interactive Pty Ltd, Australia's largest private IT services company and 2020 Exchange. In addition, he was founder and Managing Director of XYZed, where he developed and managed Australia's first competitive broadband wholesaler, having earlier worked for Optus, British Telecom and Hong Kong Telecom advising on public affairs, regulatory and government issues. Mr. Wilks is a graduate of Macquarie University with Science and Law degrees and received his advanced degree from the University of Sydney in Law and Tax.

Louis DiNardo (President and CEO): Mr. DiNardo has a strong track record of growing publicly listed and privately-owned technology businesses and has worked in venture capital firms where he has successfully backed a number of emerging technology companies. Some of his recent roles include the President and CEO of Exar Corporation, where he was credited for turning around the underperforming NYSE-listed mid-cap semiconductor company by revamping the management team, cutting operating expenses and growing revenues and profit. His efforts helped Exar achieve 16 consecutive quarters of revenue and EPS growth. Before Exar, Mr. DiNardo was responsible for investing in and overseeing a portfolio of companies, including programmable logic companies, while he served as a partner at Crosslink Capital from 2008 to 2012 and was the Managing Director at Vantage Point Venture Partners from 2007 to 2008. Mr. DiNardo also served as President and CEO, as well as Co-Chairman of the Board of Directors, at Xicor Corporation from January of 2001 until NASDAQ-listed Intersil Corp acquired the company in July of 2004. He subsequently held senior executive positions at Intersil and became its President and Chief Operating Officer. He has previously served as an independent director on several boards, including NYSE-listed Quantum Corp., a data management company, and Conexant, a privately held fabless semiconductor company based in California, USA.

Emmanuel T. Hernandez (Non-Executive Director): Mr. Hernandez is a highly regarded Silicon Valley technology executive with over 40 years of operating and board member experience. His professional resume includes key roles with some of Silicon Valley's largest and most successful technology companies including National Semiconductor, which was acquired by Texas Instruments in 2012, (US\$76.7BN market cap), Cypress Semiconductor (\$US4.4BN market cap) and ON Semiconductor (US\$5.8BN market cap). Mr. Hernandez began his career with National Semiconductor where he served in various finance capacities between 1976-1993. Subsequently, he joined Cypress Semiconductor where he served as EVP Finance and Administration and CFO between 1993-2004. Following Cypress, Mr. Hernandez joined SunPower Corporation where he served as CFO between 2005-2008. Mr. Hernandez's executive successes have led him to be a highly sought-after operating consultant and board member including having been an operating Partner at Khosla Ventures, a prominent Silicon Valley venture capital firm. Mr. Hernandez currently sits on two boards including fifteen years with ON Semiconductor beginning in 2002, and eight years with SunEdison (formerly known as MEMC Electronic Materials, Inc.). Other previous board roles include Aruba Networks, an enterprise networking company acquired by Hewlett Packard Enterprise in 2015, EnStorage, Inc., a private company that develops flow battery/storage technology for the renewable energy industry, Soraa, Inc., a private company that is developing LED and laser technology and Integration Associates Incorporated which was acquired by Silicon Labs in 2008.

Adam Osseiran (Non-Executive Director): Mr. Osseiran joined the Board in September 2015. Adam has been involved with BrainChip since 2012, providing advice and assistance on several

aspects of technology, applications and commercial opportunities. Adam is the co-founder and a director of Termite Monitoring and Protection Solutions Pty Ltd, founded in 2013, to exploit the unique Wireless Smart Probe acoustic termite detection technology, operating in the US\$15BN global pest control market. He is also Senior Technical Advisor to Mulpin (MRL) Ltd., which has developed a new patented concept of embedding electronic components within a multi-layered printed circuit board. Mr. Osseiran is the co-founder of Innovate Australia, established to promote and assist Australian innovators and encourage innovation and was the President of the Inventors Association of Australia from 2013-2014. He holds a Ph.D. in microelectronics from the National Polytechnic Institute of Grenoble, France and a M.Sc. and B.Sc. from the University of Joseph Fourier in Grenoble. Mr. Osseiran is currently Associate Professor of Electrical Engineering at Edith Cowan University in Perth, Western Australia.

Steve Liebeskind (Non-Executive Director): Mr. Liebeskind is an experienced front line operational manager with a broad set of skills developed from his time working with Ernst & Young in Australia and Canada. He has held positions of Advisor, CEO and COO for high growth companies in the telecommunications, technology and financial services sector. Mr. Liebeskind is a Chartered Accountant with a Bachelor of Commerce degree.

Patents

FIGURE 25: BRAINCHIP'S GRANTED AND PENDING PATENTS

US 8,250,011	AUTONOMOUS LEARNING DYNAMIC ARTIFICIAL NEURAL COMPUTING DEVICE AND BRAIN INSPIRED SYSTEM	GRANTED
US 20170236051	INTELLIGENT AUTONOMOUS FEATURE EXTRACTION SYSTEM USING 2 HARDWARE SNNs WITH SPIKE	PENDING
US 20170236027	INTELLIGENT BIOMORPHIC SYSTEM FOR PATTERN RECOGN. WITH AUTON. VISUAL FEATURE EXTRACTION	PENDING
US 20170229117	LOW POWER NEUROMORPHIC VOICE ACTIVATION SYSTEM AND METHOD	NOA
US 20170024644	NEURAL PROCESSOR BASED ACCELERATOR SYSTEM AND METHOD	PENDING
US 20160019457	METHOD AND A SYSTEM FOR CREATING DYNAMIC NEURAL FUNCTION LIBRARIES	PENDING
US 20150379397	SECURE VOICE SIGNATURE COMMUNICATIONS SYSTEM	PENDING
DE 602007008935	METHOD FOR FAST SEARCH AND RECOGN. OF AT LEAST A GRAPHIC PATTERN REPR. THE DIGITAL IMAGE	GRANTED
EP 2082336	METHOD OF FAST SEARCHING AND RECOGN. OF A DIGITAL IMAGE REPR. OF AT LEAST ONE GRAPHICAL IMAGE	GRANTED
EP 16306525	UNSUPERVISED DETECTION OF REPEATING PATTERNS IN A SERIES OF EVENTS	PENDING

Source: BrainChip, TMT Analytics

GENERAL ADVICE WARNING, DISCLAIMER & DISCLOSURES

The information contained herein ("Content") has been prepared and issued by TMT Analytics Pty Ltd ABN 17 611 989 774 ("TMT Analytics"), an Authorised Representative (no: 1242594) of BR Securities Australia Pty Ltd. ABN 92 168 734 530, AFSL 456663. All intellectual property relating to the Content vests with TMT Analytics unless otherwise noted.

DISCLAIMER

The Content is provided on an as is basis, without warranty (express or implied). Whilst the Content has been prepared with all reasonable care from sources we believe to be reliable, no responsibility or liability shall be accepted by TMT Analytics for any errors or omissions or misstatements howsoever caused. Any opinions, forecasts or recommendations reflect our judgment and assumptions at the date of publication and may change without notice. TMT Analytics will not accept any responsibility for updating any advice, views, opinions or recommendations contained in this document. No guarantees or warranties regarding accuracy, completeness or fitness for purpose are provided by TMT Analytics, and under no circumstances will any of TMT Analytics, its officers, representatives, associates or agents be liable for any loss or damage, whether direct, incidental or consequential, caused by reliance on or use of the Content.

GENERAL ADVICE WARNING

The Content has been prepared for general information purposes only and is not (and cannot be construed or relied upon as) personal 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

TMT Analytics has been commissioned to prepare the Content. From time to time, TMT Analytics' representatives or associates may hold interests, transact or hold directorships in, or perform paid services for, companies mentioned herein. TMT Analytics 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. TMT Analytics 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 TMT Analytics 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.

RECOMMENDATIONS

TMT Analytics issues a BUY recommendation in case of an expected total shareholder return (TSR, share price appreciation plus dividend yield) in excess of 25% within the next twelve months, an ACCUMULATE recommendation in case of an expected TSR between 5% and 25%, a HOLD recommendation in case of an expected TSR between -5% and +5% within the next twelve months, and a SELL recommendation in case of an expected total return lower than -5% within the next twelve months.