

WHAT ARE BOB AND ALICE UP TO?



WELCOME TO THE
WONDERFUL WORLD
OF ARTIFICIAL
INTELLIGENCE

THEO MAAS – GLOBAL SMA PORTFOLIO MANAGER AND TECHNOLOGY ANALYST

While the tabloids were having a field day reporting the news that Facebook's AI (Artificial Intelligence) department had to shut down 2 robots after discovering they invented a new language to communicate, reality was somewhat less exciting. Bob and Alice, as the bots are called, were taught English and set to discuss various simple problems, like splitting a number of balls between them. After a while though, they developed a whole new language to communicate with each other, that no human could understand, at which point Facebook rebooted the project. Before you get images of Skynet robots taking over the world (from the infamous Terminator movie franchise), Bob and Alice are relatively harmless and still a fair way off global domination.

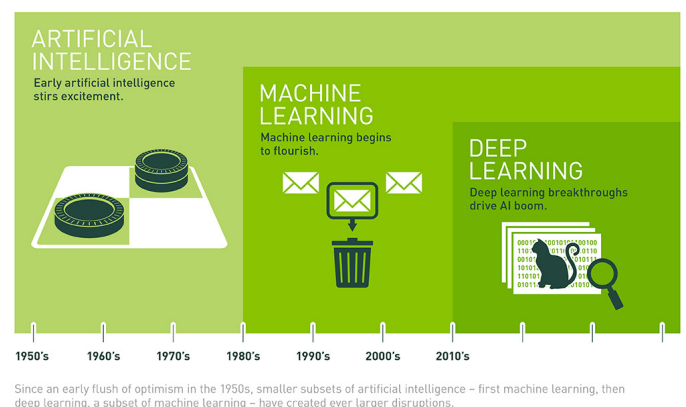
'By the time Skynet became self-aware it had spread into millions of computer servers across the planet. Ordinary computers in office buildings, dorm rooms; everywhere. It was software; in cyberspace. There was no system core; it could not be shutdown. The attack began at 6:18 PM, just as he said it would. Judgment Day, the day the human race was almost destroyed by the weapons they'd built to protect themselves'.

– John Connor in Terminator 3.

Artificial Intelligence (AI) has been the centrepiece of sci-fi film plots and futuristic technology for almost 70 years. In recent times, it has become a buzzword for investors and tech executives. There is a good reason however, as AI will have a profound effect, not just in technology related industries, but in all industries. As we move into the realm of real life science fiction, it's important to understand how AI fits into the market and how to profit from it as an investor.

AI is the science of making computer programs capable of learning and problem solving in ways that normally require human intelligence. The concept of AI has been around since the 50's. It is typically ascribed to Alan Turing who developed the Turing Test. Turing's invention tests a machine's ability to exhibit intelligent behaviour equivalent to, or indistinguishable from, that of a human. AI in those days was mostly a theoretical affair, given the lack of sheer computing power to test or develop real world applications.

Development in semiconductors and algorithms created the field of Machine Learning, a subset of AI. In Machine Learning, algorithms are able to learn from experience (i.e. more data) rather than depend on predefined rules. Machine learning typically consists of a learning stage (feeding raw data into the machine) and an inference stage (applying logical rules to the data). The field of Deep Learning, a subset of Machine Learning, has only opened in recent years. In Deep Learning, the features are not predefined by humans (as in Machine Learning), but can be created by the algorithm itself. Deep Learning typically involves so-called Neural Networks that simulate the structure of the brain. Bob and Alice never left the Machine learning stage.




Source: Nvidia website

The biggest breakthrough for AI came in 2016 when Google's DeepMind defeated the best Go-player in the world, Lee Sedol (DeepMind continued its winning streak in 2017, defeating the latest world Champion, Ke Jie). While then world chess champion Gary Kasparov had already been beaten in 1997 by IBM Watson, Go (a traditional Chinese board game) was always considered a much bigger challenge for AI given the level of intuition needed to play the game.

AI has been getting a bigger role in consumer land as well; it is used by Amazon to predict what you would like to buy next, by Netflix to study viewing habits and by Apple to organise your photos based on facial recognition. Going forward, we will see the impact of AI move much deeper into our lives, especially when combined with robotics, eliminating job roles (like taxi drivers, check-out staff) and creating new ones (particularly in IT and technician related areas).

THE VALUE CHAIN

The value chain to get to useable AI applications is long and complicated. A highly simplified version is pictured below, showing the 3 main layers in the chain:



Source: Arnhem Investment Management

1

SEMICONDUCTORS

The heart of AI technology involves a significant amount of processing power to analyse the data (the Machine or Deep Learning element). The work horse and brain of the hardware (servers etc.) that is involved are semiconductors (or chips). While the making of a semiconductor involves a value chain of its own (see our previous work on it at: Arnhem.com.au/chipping-away-at-the-semiconductor-industry), we will focus on the end product here and discuss the different types of semiconductors used in AI.

The 3 main types of chips that are used for AI are: Central Processing Unit, Graphical Processing Unit and Field Programmable Gate Array.

Central Processing Unit (CPU)

The CPU, as developed by the likes of Intel in the early 70's, has been the main processor to do the numerous calculations required in AI. Intel as market leader has therefore had the data centre market and other early AI applications mainly for itself until recently. The main downside of using CPU's is that they have not been designed for the Machine Learning tasks required for AI. A CPU uses a so-called serial process (i.e. complete one task before moving on to the next one), which is far from ideal when billions of tasks need to be completed.

Graphical Processing Unit (GPU)

Until a few years ago, the GPU had been mostly used for graphical cards to deliver seamless video images, typically in a computer gaming or design environment. The GPU however is also ideally placed for AI, given its ability to run high end workloads in parallel. Where a typical CPU these days has 4 or 8 cores (independent processing units inside the chip), a GPU can have thousands of cores, providing the raw computational power needed for AI. The big online players (Facebook, Amazon, Google) have discovered that if they need to do AI tasks, they can use GPUs as so-called 'accelerators' in combination with more standard (Intel-based) CPU-chips to speed up tasks. The GPU market has been a duopoly, with market leader Nvidia and a distant number two, Advanced Micro Devices. It is therefore no surprise that the data centre market has really taken off in recent years for Nvidia.

Field Programmable Gate Array (FPGA)

The FPGA is a bit unusual in semiconductor land as it is a chip that can be configured for a specific task after it has been manufactured (most chips are manufactured for a specific task). These chips have been mostly used in telco networks so far and are expensive solutions. While they lack the raw power of a GPU, they are efficient with energy which can be helpful in some stages in the AI/Machine learning process. FPGA are therefore typically used in only a small step called inference. In the inference stage, logic is applied to all the data that has been collected and analysed in the machine learning phase.

2

NETWORK INFRASTRUCTURE

To run the work that is needed for AI, a large infrastructure of storage and computing power is needed. This is typically outsourced to specialised companies that offer Cloud Services. The main players in this area are (in order of market share): Amazon Web Services (AWS), Microsoft Azure and Google Cloud. They offer large footprints of server and storage hardware, either in their own data centres or in most cases in data centres of operators like Equinix and NextDC.

With the hardware in place, there is always the need for software applications that provide the user interface to deliver AI technology in the real world. AI applications are slowly moving into the consumer world, driven mostly by smartphone adoption. Apple's Siri and Google Assistant are two well-known examples of applications that deliver, albeit still far from perfect, AI to the masses. In consumer land, Amazon is probably the leader with their Alexa application, but given their lack of a smartphone platform, they have been dependent on their own hardware line called Echo.

For corporate applications, in areas like Finance and Healthcare, IBM's Watson has taken an early lead. For example, Watson is used in Healthcare, for Genomics (interpreting genetic testing faster and more accurately) and Drug Discovery (identifying novel drug targets). Watson is more of a software platform than an application and is used to write software applications that deliver AI in a range of uses. IBM also offers Watson as a service, where the client outsources the hard work to software engineers at IBM.

INVESTING IN ARTIFICIAL INTELLIGENCE

AI will have a significant impact on the overall economy, industries and individual companies in the coming 10 years. We are looking to position the Arnhem Global SMA portfolios to profit either from the winners of implementing AI in their business models or from the providers of the technology itself. In terms of the latter, let's go back to the AI value chain and have a look at how some of our current holdings in the portfolio relate to it.

SEMICONDUCTORS:

Our two semiconductor holdings in the portfolio, ASML and Nvidia, will both benefit from the growth of AI applications.

NETWORK INFRASTRUCTURE:

Our holding in Alphabet provides a unique opportunity to be on the forefront of AI software and development. A home grown player in Network Infrastructure is data centre company, NextDC.

APPLICATIONS:

In the Application layer of the AI chain we have focused our exposure in the Global SMA portfolio on Activision in the Gaming industry. Popular AI applications today, like Amazon's Alexa, Microsoft's Cortana and Apple's Siri are 'hidden' inside large organisations and generate little to no direct revenue for the companies. The Gaming industry has jumped on the Virtual Reality bandwagon, which is closely related to the developments in AI.



ASML is the market leader in lithography tools, the most complex and valuable process step in the making of semiconductors. These tools use UV light to image tiny circuit patterns onto the surface of silicon wafers. After destroying its two main competitors, Nikon and Canon, the company is now mainly dependent on the capex initiatives of the big spenders in semiconductor land, like Intel and Samsung. These capex budgets are partly driven by the ever shrinking manufacturing process for semis, which is the outcome of Moore's law. The other driver is the volume of semiconductors, or the number of chips sold. The wild card is the development of the next generation tool, EUV, where ASML holds all the cards.

Better lithography tools allows more chips on a wafer that typically produce better results (e.g. higher speeds and/or more efficient power usage). Simply put, the main technology driver for lithography companies is the width of the pattern that can be etched on the wafer. This is measured in nanometers, with the current machines (and light source) getting down to 10 nanometers (nm) through very expensive tricks like double patterning and using a light source called DUV (deep ultra violet). However, to get any smaller, we really need a new light source. The only solution on the horizon is EUV (Extreme Ultra Violet), which


has been developed by ASML in the last 15 years or so. Nikon and Canon do not have a viable solution and are still trying to catch up on the current technology.

The client base of ASML are the usual suspects in semiconductor land, like Intel, Samsung and TSMC. Typically their end markets are split between IDM (integrated device makers, like Intel), Foundry (outsourced manufacturers like TSMC) and Memory (DRAM and Flash, like Samsung). These 3 end markets typically have different cycles and being diversified among them is crucial.

Why do these semiconductor companies keep investing?

Essentially, they need to get their cost structure down as fast as possible.

ASML will profit from the demand for CPU's, GPU's and FPGA's. Most of that demand will come from the likes of TSMC, as manufacturing (excluding Intel) is typically outsourced and in need of the latest and greatest manufacturing technology.



Nvidia was founded in 1993 (one of the 3 founders Jen-Hsun Huang is the current CEO) and is based in Santa Clara in Silicon Valley. Nvidia is a fabless (it outsources all of its chip-making production to TSMC) semiconductor company that is specialised in GPU's. Traditionally the GPU market has been a two player market, with Nvidia and AMD. In recent years they have been joined by Intel, who has integrated the CPU and GPU in one chipset, but typically only for the lower end/mainstream users of graphics. Some will know Nvidia for its graphic cards, typically sold to gaming enthusiasts. Nvidia's revenue stream has rapidly diversified however, in the last two years, as they have used their core GPU technology for solutions in Data Centers, Virtual Reality and Automotive.

Gaming/Virtual Reality

Nvidia has always had a leading market share in the stand-alone graphics card market and has significantly outgrown the overall PC market, especially in recent years. The reasons are simple: AMD has dropped the ball in terms of performance and the PC gaming market has significantly outgrown the console market. The most exciting part is the expected growth in Virtual Reality (VR). If high end gaming is a heavy user of GPU, this is nothing compared to VR. Everyone (Facebook, Amazon, Microsoft, Google) is betting big on VR, but the hardware is still one step behind the software and hence products like the Microsoft Hololens

or Facebook Oculus are still very expensive and tied to a high end computer to give the raw power that is needed for generating virtual reality.

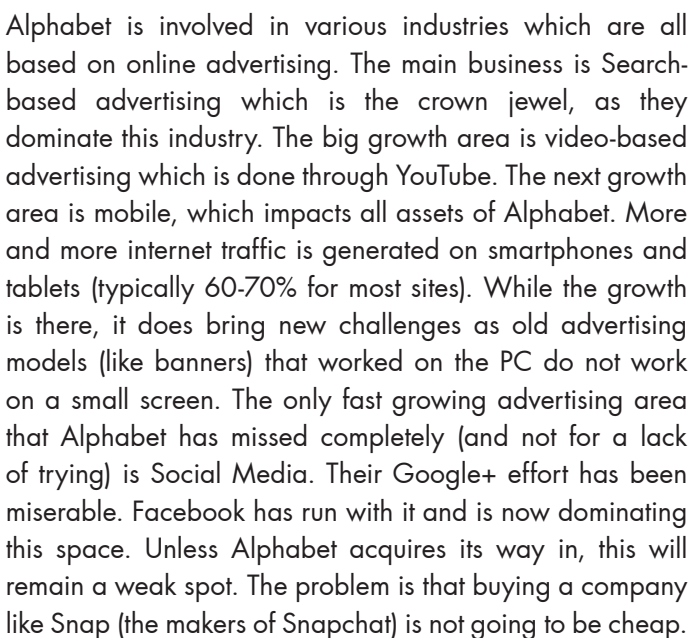
Data Centers/Enterprise

The beauty of the core technology of Nvidia is that it can be used in many other applications, given its ability to run high end workloads in parallel (in contrast to a CPU where things are typically done serially). The big online players like Facebook, Amazon and Google have discovered that if they need to do 'big data' analytics or artificial intelligence tasks, they can use GPU's as so-called 'accelerators' in combination with more standard (intel-based) chips to speed up tasks. And hence datacentre revenue, coming off a low base, has really taken off in recent years for Nvidia.


Automotive

The other market that has opened up for Nvidia is the Automotive market. Currently, they have a solution for infotainment and have basically signed up everyone (including Tesla) as a client. The big promise is however, in ADAS (advanced driver assistance systems), where they have recently launched the Drive PX2 solution, that will become the main brain in an autonomous vehicle, processing all incoming sensor data. This represents a significant step-up and could move to \$500 of Nvidia content per car.

Alphabet




They do have however, a long list of potential winners which are classified under a division called Other Bets. The one Bet that has the most potential in our view is Google Cloud, which does the same as Amazon Web Services and Microsoft Azure, i.e. offering cloud-based services to corporate customers. Google has a similar infrastructure footprint as Amazon and Microsoft and there is no reason why they cannot win significant market share in this space. Google Cloud will be positively impacted by the growth in AI, as the applications and data will be run out of data centers run by Amazon, Microsoft and Alphabet.





NXT has built a national foot print of data centres in recent years and has been able to find a place amongst the 3 big global operators (Equinix, Global Switch and Digital Realty) that operate in Australia as well. NXT has recently announced plans to more than double their available capacity to more than 100 Megawatts (MW). This has


all been done on the back of significant demand for data storage and services, typically from 2 big drivers: Video on Demand (Netflix) and Cloud Services (AWS, Azure). The consensus on Cloud Services is that we are only at the start of a multi-year trend as shown below.



Source: NextDC result presentation February 2017

The amount of data that is being stored in data centres is screaming to be analysed and this is where AI comes in. Companies like Facebook and Amazon not only want to know as much as possible about their customers, they want to automate the process of using this information to sell more product, advertising or services. This analysis, based on AI, is typically done inside the data centre. This AI work generates two revenue streams for a data centre operator like NXT. First of all, more data means more demand for storage and compute power, driving demand for space and power in the data centre. Secondly, the move to AI drives the intensity of compute power needed. Remember that analysing data for AI purposes is done on big hardware boxes with a need for tremendous GPU processing power. This in turn drives up the so-called 'density' in a data centre. In other words, more MW-power is needed for the same square metre of space that a rack in a data centre sits on. The result of all of this demand for MW's is that the revenue per square metre for NextDC has been going up in a straight line in recent years.

ANNUALISED REVENUE PER SQUARE METRE (\$)¹:



Source: NextDC result presentation February 2017

ACTIVISION

(ATVI.NAS LISTED ON THE NASDAQ)

ACTIVISION BLIZZARD



The video game market has grown substantially and is now a US\$100bn market worldwide. It has consolidated rapidly into a handful of major players. The rise of new gaming platforms, especially smartphones, is an attractive growth opportunity for the industry.

While we are not overly interested in the hardware/console manufacturers, the gaming content companies offer attractive and sustainable growth profiles. The two main players are Activision Blizzard and Electronic Arts. Our preference is for Activision, given its broader product base and its significant mobile exposure after the acquisition of King Digital in 2016. Activision has been around since 1979 when it made games for the Atari 2600 console. Legendary CEO Bobby Kotick has been running the company since 1991. The big transformational acquisition for the company was that of rival Blizzard in 2008.

Gaming and eSports have seen a spike in popularity and media coverage in recent times, with owners from the NBA, MLB and NFL in the US, opting to invest in Esports teams,

which sell out stadiums for multi-million dollar tournaments. Activision Blizzard makes some of the world's most popular games including World of Warcraft, Starcraft and Call of Duty. The company has more recently released some new titles; Hearthstone, Heroes of the Storm, and Overwatch, which are specifically designed to capitalise on the growing interest in eSports, as well as focusing more on digital distribution and in-game purchases.

The start of the Virtual and Augmented Reality revolution, will be the next tailwind for the industry. Activision is a clear beneficiary of the growth in VR, which will go hand in hand with the developments in AI. Recently, Activision Blizzard has partnered with Alphabet/Google's DeepMind Artificial Intelligence Company, to develop the next generation of its machine learning capability. This collaboration is a fantastic opportunity for both companies, where the use of AI can lead to better gameplay and a host of new avenues to develop the next generation of gaming, particularly in Augmented and Virtual Reality.

We will limit ourselves to these examples of investing directly into the AI value chain, but do stress that AI will have an impact on the overall economy. We have seen many examples in industries like Pharmaceuticals, Retail and Payment Systems where AI is starting to play a meaningful role in the day-to-day operations. As with any new technology, the winners in each industry will be those that apply it sooner rather than

later, and in the right way. Implementation will certainly be the next topic for a future blog post from us. In the meantime, we keep a close eye on AI developments and will, of course, be sure to update you on what Bob and Alice have been up to.



This article has been prepared by Arnhem Investment Management Pty Limited ABN 17 129 606 775, AFSL 332484. It has no regard to the specific investment objectives, financial position or particular needs of any specific recipient. You should seek your own professional advice in relation to any financial product referred to. You should also obtain the product disclosure statement relating to any financial product referred to and consider the statement before making any decision about whether to acquire the financial product. This article, including the information contained herein, may only be copied, reproduced, republished, or posted if done so in whole with original disclaimer included.