White Paper

XENON High Frequency Trading Solutions

AI IN THE CLOUD FOR GOVERNMENT

Artificial intelligence (AI) has many and varied applications in the government sector, providing the potential to improve services for our citizens and reduce costs – sometimes to the extent that the previously unaffordable becomes a practical reality.

In particular, developments in cloud computing reduce the need for expensive hardware investments while still meeting security and sovereignty requirements.

AI AND DEEP LEARNING

The concept of artificial intelligence (AI) has been around for decades. It is the quest to have computers perform tasks that have traditionally required human judgment.

Attempts to explicitly program such systems met with limited success. The question has been whether computers could even be made to perform tasks like a human at all, let alone how well they could do it.

It has not been too difficult to write code to detect human faces within a picture – a task that even relatively low-end digital cameras can now perform. But being able to reliably recognise many different faces at once has been more challenging, especially if it needed to be done quickly and with many faces in the picture.

The development of machine learning has turbocharged the application of AI to these types of problems. Very loosely, the idea is that instead of codifying a task by writing a set of instructions specific to that task, you create a system that to some extent resembles the way a brain works, and then train it on a set of data set so it learns how to make perform the task.

These models fall into two broad classes. Supervised machine learning is used where the answers are already known. So you might feed thousands of photos into a model to train it to recognise pictures of cars. If you then present the trained model with another photo, it should be able to tell whether it is a car or not.

Unsupervised machine learning is useful when there is no explicit question. Instead, you are looking for commonalities in the data that may reveal something about the structure of the matter under consideration.

Deep learning goes a step further by significantly increasing the complexity of the neural networks used in these models – not just in terms of the number of neurons, but also the number of layers they are organised into – and also by massively increasing the amount of data used to train them.

In some areas, deep learning is already matching or exceeding the performance of human experts, for example in playing Go, recognising skin or colorectal cancers, and translating certain types of text from Chinese to English. Work in other areas - notably autonomous vehicles - has also advanced rapidly.

AI IN GOVERNMENT

AI has many current and potential applications in Government due to the breadth of its role in society and its large and varied data stores.

FACE RECOGNITION

Deep learning can greatly improve the accuracy and performance of face recognition systems.

While some of the most obvious applications are in the realm of security and public safety (many Australians are by now familiar with SmartGate self-service border control system), they can also be aimed at improving the customer experience. While SmartGate does that by significantly reducing the time spent in queues at border control, such systems can also be used to simplify other processes, as seen in the Sydney Airport/Qantas trial of face recognition to automate check-in, baggage drop, lounge access and boarding processes. The ability to quickly, consistently and accurately recognise individuals is key to the success of such systems.

VIDEO ANALYTICS

Humans soon become fatigued when monitoring surveillance video. Video analytics taking advantage of deep learning provides a way to bring out-of the-ordinary situations to the attention of human observers. Simple motion detection is unlikely to distinguish between a potential intruder walking along a fence line and a large plastic bag being blown along by the wind, but deep learning can, reducing the number of such false alarms.

Furthermore, given suitable video feeds, such systems can provide real-time insights into crowd behaviour, allowing police and protective services to respond more effectively to incidents at major events or in busy areas. The more efficient use of resources improves public safety in such situations, and increases the likelihood of a timely response to other unrelated events occurring at the same time.

NETWORK MONITORING

It is an archetypical 'needle in a haystack' situation. Its goal is to recognise patterns and identify anomalies among the millions of log entries and other records collected. Applying deep learning to this task can yield results more quickly than traditional analytics, potentially allowing the root cause to be identified and countermeasures applied before the significant damage is done.

One example is a system called PathScan that was developed at the Los Alamos National Laboratory in the US and is being commercialised by EY. In a nutshell, PathScan determines the normal flow of network traffic and then identifies exceptions that could indicate malicious activity.

The functions of government are now heavily dependent on IT, and the pool of data available to potential hackers has become an increasingly attractive target. Network security has become an increasingly important priority in maintaining public trust in these systems.

FRAUD DETECTION

It has much in common with network monitoring in that the objective is to spot a relatively small number of suspicious transactions within an ocean of legitimate activity. Deep learning can be usefully applied here thanks to its ability to identify patterns, whether that is to match new transactions with earlier ones that are now known to be fraudulent, or to uncover clusters of activity that warrant investigation.

Reducing fraud protects the public purse, helping citizens feel they are getting value for money from their government. Doing this effectively and efficiently with the aid of deep learning frees up resources that can be used to benefit law-abiding citizens.

CUSTOMER SERVICE

There are several ways to apply AI to improve customer service. For example, an AI-based system could dramatically reduce the duration of an interaction with Centrelink or the ATO – whether by collecting the required information while waiting for a customer service reporesentative to become available, or by being able to answer common questions and complete simple transactions.

One of the difficulties of applying AI and deep learning to government services is that critics will almost inevitably play the 'Big Brother' card. The technology that improves the identification of cancers and other anomalies in

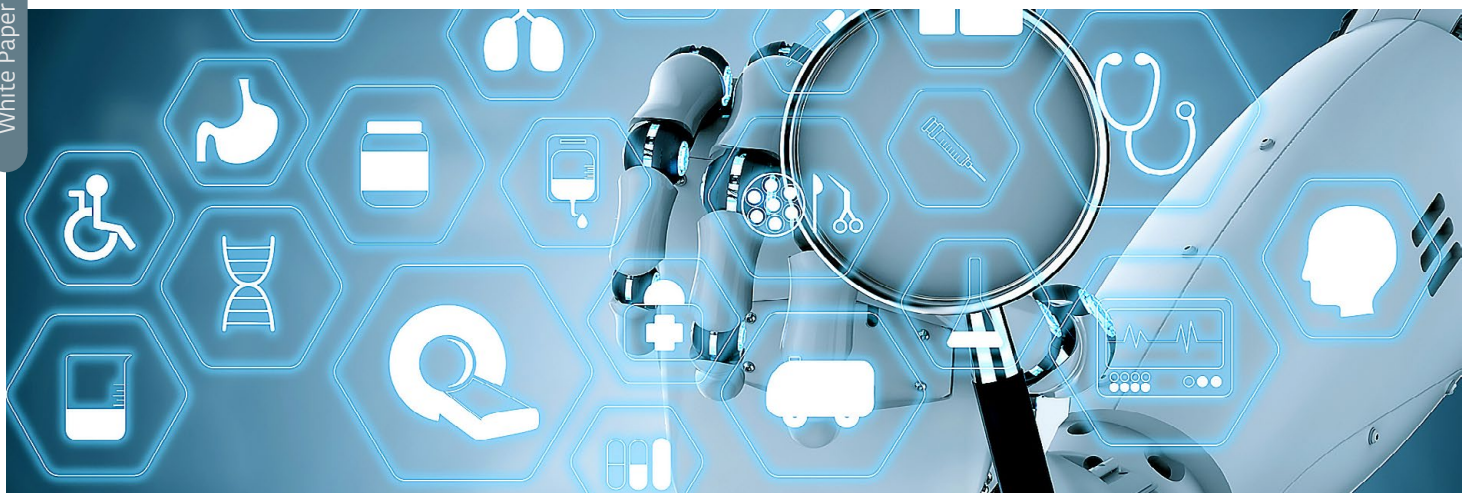

medical scans also helps recognise the faces of people captured in grainy security videos as they walk down streets at night, whether they are going about their lawful business or up to no good.

So it is important for government services to show that the adoption of these technologies can improve citizens' lives.

As noted above, one of the characteristics of deep learning is that it may help identify patterns that aren't immediately obvious. There may be a temptation to simply use the technology for fraud detection when it comes to welfare benefits and taxation, but it also could also identify citizens who may be short-changing themselves: "Many people in your situation are entitled to claim..." or "You have reported an abnormally low amount at label X given your type of business and its revenue - please check your figures, but remember you must be able to substantiate all amounts shown in your return."

OVERSEAS EXAMPLES

Deep learning is being used to give everyone from the military to financial analysts the performance to convert large amounts of data into actionable information – all while reducing costs and the power necessary to complete the mission. Today, aerospace, defence, and intelligence industries are taking advantage of deep learning for better, more data-informed decision making.

For example, the US Army has been looking at possible applications of machine learning in a wide variety of areas, including predictive maintenance, medicine, armour design, cybersecurity, robotics, and propaganda detection and targeting.

CLOUD INFRASTRUCTURE FOR AI AND DEEP LEARNING

Cloud Infrastructure as a Service (IaaS) has several advantages for AI/deep learning projects, but they all revolve around its inherent scalability.

AI/deep learning makes substantial demands on computing hardware, but the exact requirements can fluctuate widely both during development and ongoing operation. A conventional approach to infrastructure would involve sizing the system to meet the expected peak load, and either letting it idle during quiet periods or trying to find another application that can take up the slack. Infrastructure as a Service (IaaS) allows resources to be brought into play as required, and the client only pays for what they use - it is the provider's responsibility to reallocate any spare resources.

Current approaches to deep learning revolve around GPU (graphics processing unit) computing. Both training and using deep learning systems are computationally intensive tasks, and therefore unsuited to general-purpose CPUs. Fortunately, the way the algorithms work means many of these calculations can be performed in parallel, so a way around this is to throw large numbers of processor cores at the problem - this can most affordably be done by taking advantage of GPUs that each contain hundreds of cores.

Much of the recent interest in deep learning has been the result of the way GPUs have dramatically improved the price/performance ratio of the systems that can run deep learning models sufficiently quickly for them to be useful rather than merely being of academic interest.

Not only does Vault, a secure cloud provider, make GPUs a part of its IaaS offering, it has also partnered with XENON to take advantage of that company's extensive experience and expertise in GPU computing.

XENON introduced GPU computing to Australia in 2008, and through its relationship with GPU vendor NVIDIA it became the first to offer turnkey AI/deep learning appliances in the region.

XENON is a supplier to many leading universities and research institutes – including CSIRO, Data61, the Australian National University, the University of Melbourne, Monash University and the University of Sydney – and to the Department of Defence.

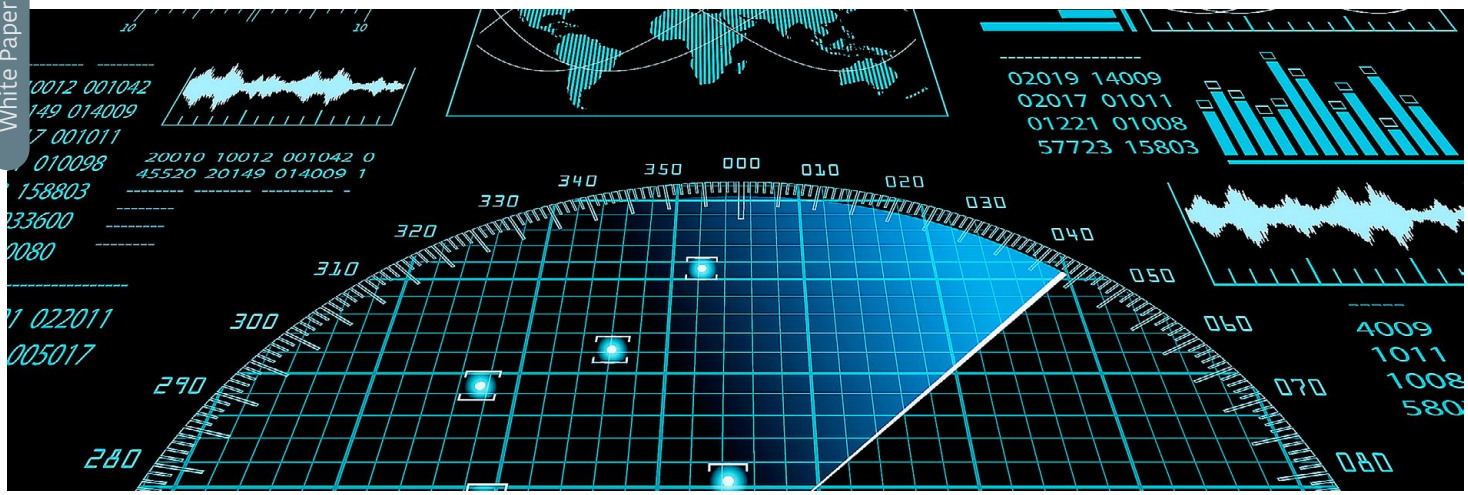

GPU-accelerated machines in the cloud will be the foundation of most AI-based applications. Such systems will be well-placed to take advantage of the growing number of datasets from the billions of sensors that will be coming online in the next decade.

"While some people still regard AI as hype, it is becoming clear to most observers that artificial intelligence utilising neural networks and deep learning will have a significant – and in some areas game-changing – effect on the way governments go about their business," said XENON Technology Group CEO and Founder Dragan Dimitrovici.

"But it is very unlikely that these benefits will be achieved without the use of GPUs, which are currently the most cost-effective way of performing the massive number of calculations required within a realistically short time," he explained.

XENON's consultancy services help clients accelerate their development of AI systems. One example from the resources sector is the application of machine vision in order to reduce exploration times while simultaneously improving the accuracy of predicted yields.

Storage is another issue for deep learning. Such projects often involve large and growing data sets, but the use of IaaS means that rather than struggling with capacity planning and complex upgrades, the client can simply dial up the amount of storage as the data set grows - or down again when a project ends or a temporary bulge in the data is aged out of the current set.

While storage capacity is rarely a problem in an IaaS environment, it is sometimes impractical to maintain a second copy of a large data set in that provider's cloud if the original is located elsewhere. That means the IaaS provider must be physically colocated with the primary data, or able to access it via fast and secure network links.

Vault's data centre in the ACT is on multiple networks, including ICON for the secure communication of information with a classification up to PROTECTED to other sites on the network. This provides convenient access to large data sets that are stored outside the Vault cloud.

Security is a consideration for all IaaS users, but government departments are subject to particular mandatory requirements, including onshore storage and processing.

Vault Cloud was designed from the outset to meet Australian Government security requirements. It is fully certified by ASD for use for OFFICIAL (UNCLASSIFIED DLM) and PROTECTED data, and was one of the first cloud providers to gain PROTECTED certification. It is already used by Australian Government departments and agencies including the Department of Human Services, the Department of Jobs and Small Business, and the Department of Defence.

Vault Systems is an Australian owned and operated business. All of its infrastructure is located inside Australia with multiple data centres in Canberra and Sydney to provide geographical redundancy. Its employees are all located in Australia and have security clearances appropriate to the roles they perform.

CONCLUSION

Deep learning is a powerful technology with many applications in the government sector.

Vault's association with XENON provides access to specialist knowledge in the areas of GPU computing and deep learning. This expertise is also available to Vault's clients.

Vault Cloud provides a technologically suitable, secure and certified home for agencies' deep learning projects, combining the cost-effectiveness of IaaS with the specialist hardware required for deep learning.