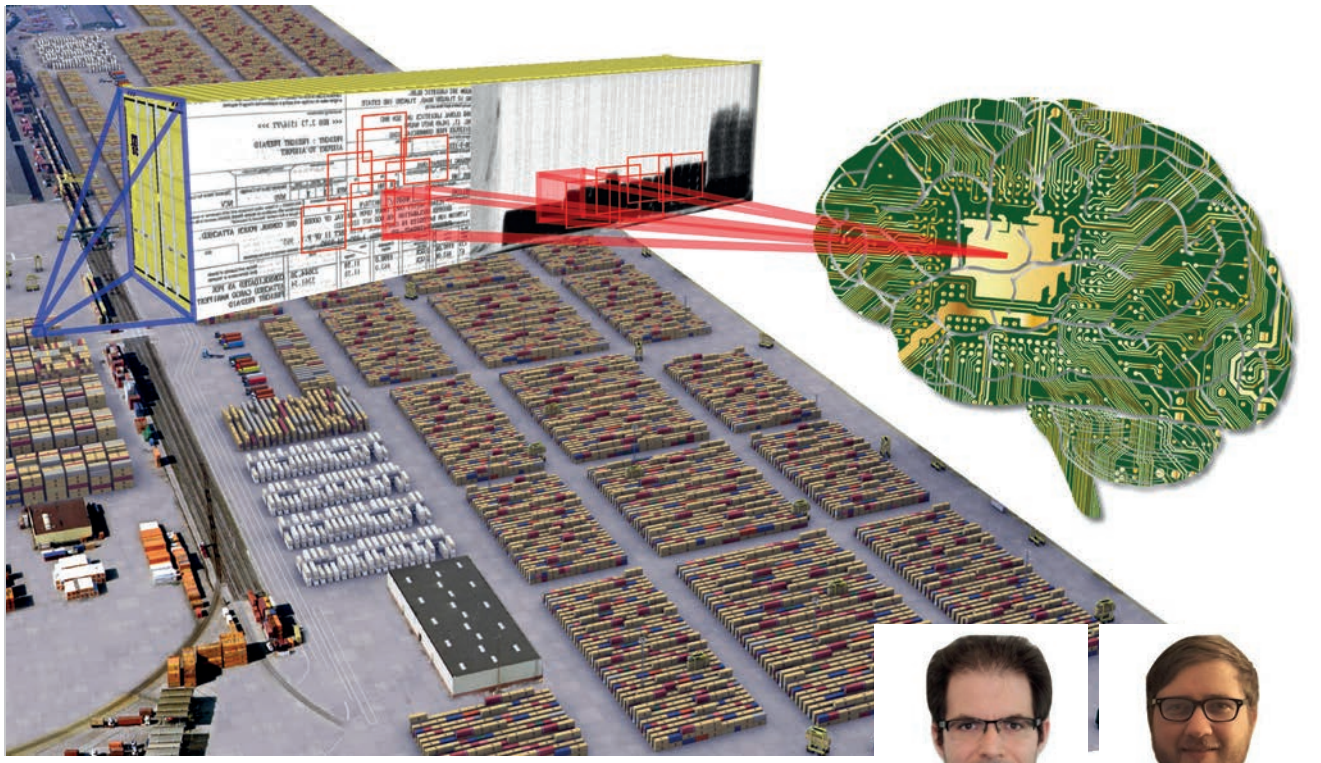

AUTOMATED INSPECTION

BY ARTIFICIAL INTELLIGENCE

Dr Nicolas Jaccard, and Thomas Rogers,
Visulytix, London, UK

The search for contraband, threats, and fraud in cargo containers is akin to searching for needles in an ever-growing field of many haystacks. The field represents the vast global container fleet that constitutes the global supply chain. The haystacks represent the diverse range of confusing legitimate items that fill a 40-foot shipping container. These are ever growing because the number of container transactions continues to grow each year.

To further complicate things, one doesn't know what the needles are until they find them. Is it narcotics? A weapon? Radiological material? Something that's never been seen before, or even conceived of?

CONTAINER SCREENING

Current screening protocols aim to address this issue by splitting the process into three parts:

1. Container selection based on a risk analysis, specific intelligence, or at random
2. Non-intrusive inspection of selected containers to form a radiographic image
3. Physical inspection.

There are a few problems with this approach. Firstly, only an estimated 4-5% of containers are selected for inspection. The unchecked 95% can, and sometimes do, contain threats or contraband. Secondly, false positive selections mean that human resources are wasted on benign containers. Thirdly, human errors in image inspection leads to unnecessary physical inspection: a time-consuming process in which the entire contents of a cargo container have to be carefully handled and documented.

HARDWARE INNOVATIONS

There has been a noticeable and recent innovation in the container screening market: X-ray scanners are now available that can automatically capture images of rail-hauled containers travelling at up to 60km/h. Such systems can scan a 40-foot container in less than a second. As such systems are further developed to accommodate all transport modalities, including trucks, it would, in principle, be possible to scan every container.

In this total-scanning paradigm, and if Artificial Intelligence (AI) is indeed good enough, it would make sense for an AI system to analyse each image, and in combination with risk indicators, determine which images should be sent to humans for inspection. This would significantly reduce the number of false alarm images that human resources are wasted on, in addition to increasing the number of accurate detections over the entire container population. This way, 100% container screening could be realised; a long-standing, but never met, goal of the Container Security Initiative introduced after 9/11.

Figure 1: Automated container mover

Figure 2: Physical inspection

Figure 3: Cargo train

THE AI REVOLUTION

We are undergoing a revolution in AI. It is not difficult for one to imagine a future, though perhaps distant, in which supply chains are fully automated. Signs of it are already emerging, as companies such as Amazon and Tesla are investing in methods to automate logistics through drone delivery and driverless trucks. Indeed, in recent times we have even seen the opening of AI-controlled seaports, such as London Gateway, where containers can be picked and stacked with minimal human intervention.

As automation becomes more widely adopted, the volume of trade is likely to catapult to levels far surpassing those observed today. As such, it will inevitably add further weight onto the shoulders of security officers as they wade through this now much larger field of many haystacks. Thus, it is also inevitable that help will be sought from AI for screening too.

BENEFITS AND CHALLENGES OF AI

There are many benefits to AI inspection. AI systems cannot be bribed. They do not get tired or lose concentration. When they break, they are simple to replace by installing new hardware. They are easy to scale; only extra hardware is required and there is no need for expensive hiring and training of staff. They are also more consistent; a given size of cargo container takes the same time for AI to inspect regardless of its complexity, making resource requirements predictable and allocation straightforward. Finally, in an operational context, the performance of an AI system is easier to quantify, and the levels of detection and false positives are easily tuned to the requirements of a given port.

The main challenge for AI is achieving high-accuracy with limited training data. Contraband data is rare, and modern AI systems, such as Deep Learning, need to be trained on vast amounts of both benign and contraband data. Already, in the academic literature, several approaches have been proposed to address this problem. This has resulted in systems that seemingly match or exceed human accuracy and speed. This super-human performance is definitely evident in mainstream applications of AI, where Deep Learning has been able to beat the world's best GO player, and outperform humans at a range of tasks including voice and image recognition.

INSPECTION ALGORITHMS

The current market for container inspection algorithms can be split into three applications:

1. Object detection - is there a known threat or contraband item present?
2. Anomaly Detection - is there anything unusual in the image?
3. Manifest Verification - does the image content perfectly match shipping manifest?

An AI system should perform each of these at least to human levels of accuracy in order to fully replicate manual operation.

Object Detection and Manifest Verification should be achievable with some domain-specific tweaks to current state-of-the-art Deep Learning methods. Anomaly Detection is more challenging, both practically and intellectually. For example, what is an anomaly? How do we evaluate performance? It will be some time before an AI system can detect anomalies as well as humans.

But Anomaly Detection is essential in cases where contraband is so well concealed that it cannot be detected by a specific detector. In such scenarios, the concealment often leads to an anomalous change in texture or density. Anecdotally, operators detecting these rare and unusual events are often what leads to discovery of contraband.

OBJECT DETECTION

There are a number of algorithms on the market for Object Detection that are integrated into proprietary scanning software. These assist operators by highlighting detected objects in the image. Current capabilities include cigarette and high-Z detection. These applications often rely on well-established but deprecated machine vision approaches; it is likely that the adoption of Deep Learning, or similar recent AI advancements, would lead to significant improvements and pave the way for image-based automated container selection.

In academia, big strides have been made toward modernising Object Detection algorithms. Milestones include the development of methods for the detection of concealed cars (e.g. cars involved in trafficking), and so-called small metallic threats (SMTs) that may be used by terrorists and organised crime networks. These methods use bespoke-trained Convolutional Neural networks (CNNs) to achieve very high detection performance with low false alarm rates. Jaccard et al. [10] show that cars concealed by large amounts of dense and confusing clutter can be detected, even with limited training data.

In addition, Rogers et al. [11] and Jaccard et al. [12], show that by using Threat Image Projection (TIP), which is used in aviation security to train and test human operators, a large contraband dataset can be generated, making it possible to train highly accurate CNN models for contraband detection. While these approaches yield great promise, they have yet to be commercialised and made available to the market.

MANIFEST VERIFICATION

Fully-fledged Manifest Verification tools are not yet commercially available, and attempts described in the academic literature are underwhelming. However, tools for Empty Container Verification (ECV), a sub-problem of Manifest Verification, are available. ECV is important since an estimated 20% of the container fleet is declared-as-empty. Detection of cargo inside such containers can be a sign of duty-evasion, fraud or smuggling. Commercial ECV implementations can spot the most obvious cases but might miss smuggling attempts where only small loads are present (e.g. small quantities of drugs). Recent research suggest that

machine learning helps in those cases, hinting that modern approaches such as Deep Learning could potentially achieve close-to-perfect ECV performance. The ability of achieving fully-fledged Manifest Verification is contingent on the availability of suitably large and diverse image datasets complete with manifest information.

ANOMALY DETECTION

To date, little work has been published on Anomaly Detection. The ACXIS project aims to provide functions to assist operators with anomaly detection. The project proposes to search a database for similar images (e.g. of the same truck) and allow operators to compare a fresh image with a database image to help spot any irregularities.

In academia, Andrews et al. [18] have sought to develop Deep Learning based methods for anomaly detection in an unsupervised fashion. In this, the algorithm is only trained on a large dataset of normal data that contains no contraband, threats, or anomalies. The system learns a model of the distribution of normal images, from which outliers (anomalies) can be detected. The method offers promise, however, is not currently at levels required for commercial use. This is expected to change in the future, as more advances are made in mainstream Deep Learning for unsupervised learning of feature representations.

NEW ATTACK VECTORS?

Whilst AI can reduce the burden on human operators, could it also introduce new attack vectors for organised crime and terrorist networks? Potentially the number of cyber attacks on ports and security systems could increase. Attacks could be aimed at disrupting security by denial of service, or at tricking or overriding AI systems to make incorrect decisions. For example, an attack could increase the number of false alarms, leading to increased distrust of AI systems, perhaps so much that they are switched off altogether.

Cyber attacks could also be aimed at theft of AI models. With model access, criminals could use AI to generate adversarial examples capable of evading detection. In turn this could be used to develop new methods of concealment. In any case, careful thought should be applied when developing AI inspection systems, to account for potential novel attacks.

CONCLUSION

It is not difficult for one to imagine a future where the global supply chain is entirely automated. The thirst for automation is driven by the need to cope with the increasing throughput demands. On the hardware side, commercial scanners capable of imaging containers in less than a second have been developed.

The main limitation, in terms of both security and throughput, is the use of human operators who are overburdened by the ever-growing flow of trade. This burden will only worsen as automation in logistics increases.

It is inevitable that AI will be given a deeper and more important role in inspection to cope with these demands. However, the current market lags far behind what is possible with recent innovations in AI such as Deep Learning, and efforts to provide automated inspection need to catch up with the efforts to automate trade.

With the help of AI, the operator's search for needles in the ever-growing field of many haystacks can be reduced to a much smaller, more tractable problem. And the extra human resources that would have been required can be freed to tackle tasks that AI is not suitable for, such as intelligence, policing operations, and physical inspection.

REFERENCES

[1] http://data.worldbank.org/indicator/IS.SHP.GOOD.TU/countries
[2] https://www.gpo.gov/fdsys/pkg/CHRG-112hhrg76511/html/CHRG-112hhrg76511.htm
[3] https://www.rapiscansystems.com/en/products/cvi/rapiscan_eagle_r60
[4] http://www.sds.l-3com.com/cargoscreen/CX-Rail.htm
[5] https://arstechnica.co.uk/information-technology/2017/05/deepmind-alphago-go-ke-jie-china/
[6] https://arxiv.org/abs/1610.05256
[7] https://arxiv.org/abs/1501.02876
[8] https://www.rapiscansystems.com/en/products/radiation_detection/rapiscan_auto_z_detection_software
[9] https://www.smithsdetection.com/index.php?option=com_k2&view=item&id=546&Itemid=101&lang=en
[10] https://doi.org/10.3233/XST-16199
[11] http://dx.doi.org/10.1117/12.2262662
[12] https://arxiv.org/abs/1609.02805
[13] http://dx.doi.org/10.1007/s12198-013-0121-3
[14] http://dx.doi.org/10.1117/12.706272
[15] https://doi.org/10.1016/j.apradiso.2005.05.033
[16] https://doi.org/10.1049/cp.2015.1762
[17] http://www.casra.ch/uploads/tx_tvpublications/Paper_ICCST2016_20160730_01_WiVi.pdf
[18] http://dx.doi.org/10.1117/12.2261101

ABOUT THE AUTHORS

Dr Nicolas Jaccard obtained his PhD from University College London (UCL) in 2014, where he developed novel approaches for the automated analysis of microscopy images. Before joining Visulytix as CTO, Nicolas worked as a post-doctoral research associate and helped establish the Computational Security Science group at UCL. There, he pioneered the use of Deep Learning for the detection of security threats in X-ray security imagery.

Thomas is a data scientist at Visulytix specialising in Security and Defence. He holds an MSci Theoretical Physics (Imperial College), an MRes Security Science (UCL), and is nearing completion of his PhD (UCL). He previously held research posts at DSTL and the LHCb project. An amalgamation of these experiences has lead to Thomas' interest in developing real-world Deep Learning solutions for the Security and Defence sectors.

ABOUT THE ORGANISATION

Visulytix is a London-based artificial intelligence company that provides decision support systems across a diverse range of sectors. Our goal is to develop deployable, cutting edge machine learning techniques to improve outcomes while reducing costs and time.

ENQUIRIES

media@visulytix.com
http://www.visulytix.com/