CALIPSA
EBOOK

A GUIDE TO ARTIFICIAL INTELLIGENCE IN VIDEO SURVEILLANCE

How the security industry can increase efficiency using AI technology

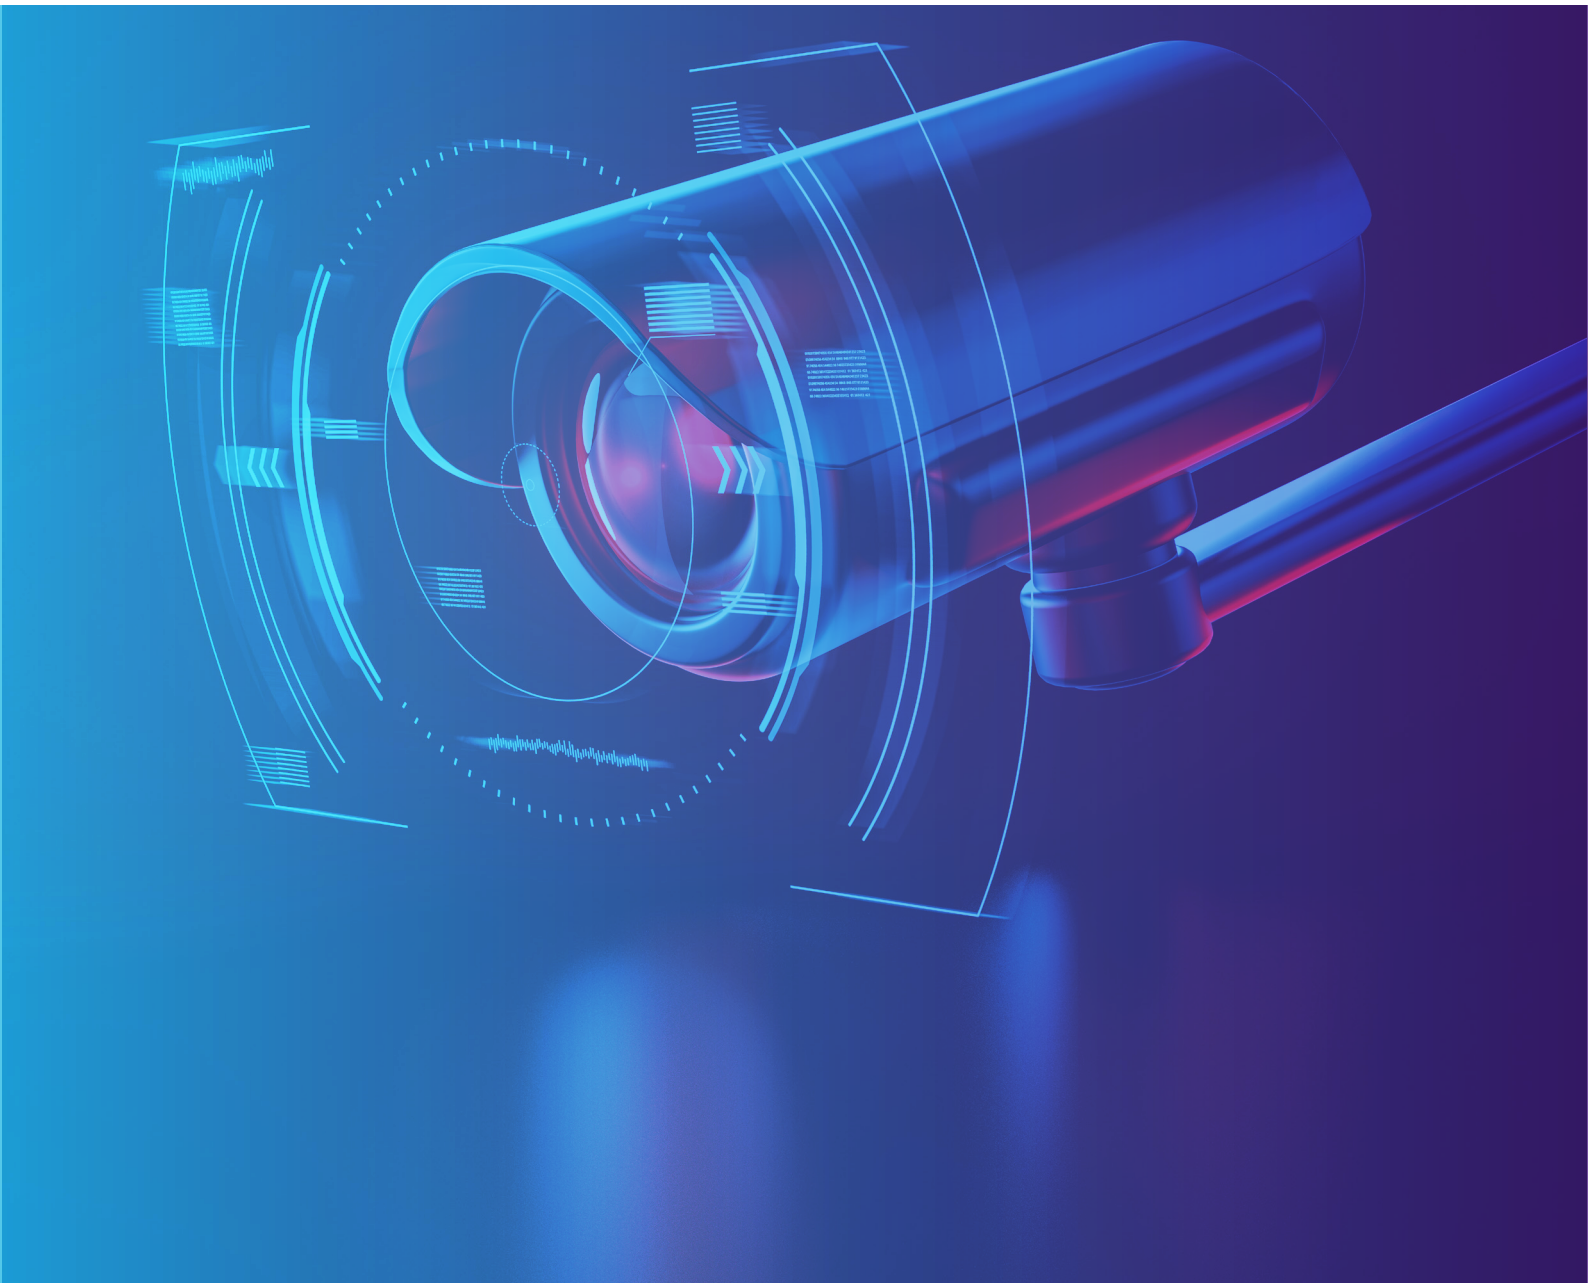

CONTENTS

INTRODUCTION 3

THE RISE OF AI IN SECURITY 4

TRUE OR FALSE? SOLVING THE CHALLENGE 5

AI 101: HOW DOES IT ACTUALLY WORK 6

RULE-BASED ALGORITHMS 7

THE NEURAL NETWORK APPROACH 8

THE BENEFITS OF AI FOR VIDEO SURVEILLANCE 9

CASE STUDY: EYEQ MONITORING 10

CONCLUSION 11

INTRODUCTION

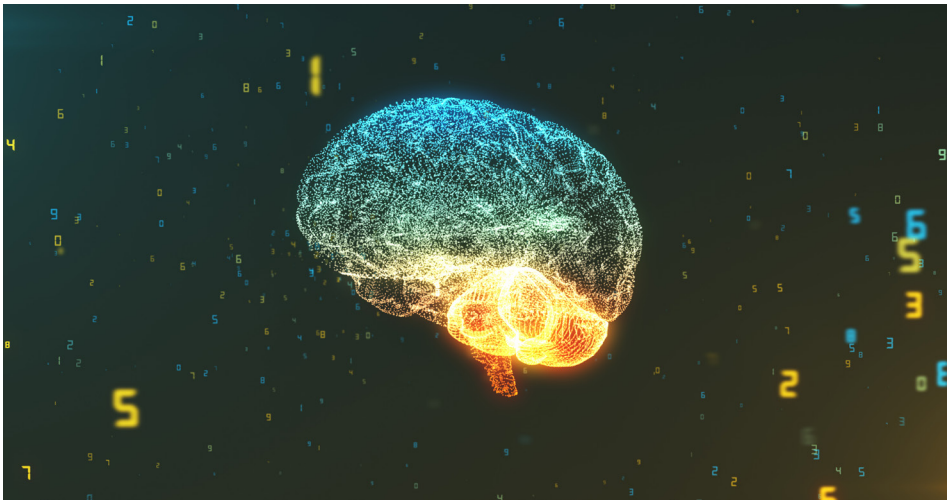

Machine learning and artificial intelligence. For the last couple of years they have been used everywhere. Some of the biggest companies are using them, sometimes without us even knowing. Everytime you search on Google, shop recommended products on Amazon or Facebook tags a friend for you; machine learning is being used.

Even if it seems like quite a new concept, the premise of this technique was invented as early as the 1950s by a man named Alan Turing. Already famous for having created the enigma machine which was used to decode the Nazi encrypted communication during the Second World War, he was one of the first to realise that one day a machine could 'think'. He came up with the famous 'Turing test' which is intended to understand whether an algorithm is truly intelligent.

For a very long time, this research stayed in academia and those techniques were only theoretical. It was only more recently that artificial intelligence and machine learning began to be used in the real world. The massive increase of computational power made it possible to train machine learning systems at scale for real-world application.

The security industry is one area that has already started to benefit from the increased capabilities of artificial intelligence in recent years. But how does it work and what are the benefits for the industry? In this guide, we explore the role of artificial intelligence in video surveillance.

THE RISE OF AI IN SECURITY

From home surveillance to perimeter protection, and access control to site security; artificial intelligence has a big role to play. And the past few years have seen the exponential rise of security technology, including AI-powered products.

While the likes of Google and Amazon have been using deep learning technology since the early 2010s, paving the way for new ideas and enhancements, the security industry has more recently started to capitalise on the advancements.

For an industry traditionally heavily dependent on manpower and human resources, AI opens new doors. For example, as skilled labour becomes harder to come by, and more expensive to employ, companies can now look towards technology to fill the gaps. AI has been introduced to take on repetitive tasks, allowing man hours to be focused on more complex tasks and adding value to a business or customer.

The rise of artificial intelligence in the security industry, video surveillance included, is evident. And it's only just beginning.

> **In 2019, research valued the 'artificial Intelligence in security' market at $5.08 billion. This is expected to grow even further, reaching an estimated $14.18 billion by 2025, representing a CAGR of 18.64%**
> *https://www.mordorintelligence.com/industry-reports/artificial-intelligence-in-security-market*

Security companies have started to realise the value technology can offer their business - and their people. They need to ensure they are investing in the most accurate, reliable and intelligent technology to grow their business.

TRUE OR FALSE? SOLVING THE CHALLENGE

For security technology to be truly effective, the product must be designed to tackle, answer and solve a clear use-case or objective. When it comes to technology, one size doesn't fit all, so finding the right technology to solve the challenge your business is facing is key. In CCTV, AI is particularly powerful in reducing false alarms.

False alarms are a serious burden for the security industry. It's currently estimated that more than 95 percent of all alarms triggered by video surveillance cameras are false. This can be caused by anything from a change in lighting, to a spider web across the lense.

As a result, tackling the false alarm challenge has become one of the most common use cases of AI in the security industry. Instead of relying on human operators to review every single false alarm, AI is able to identify and filter alarms that aren't considered genuine activations. In order to teach AI algorithms to perform this task, we need to distinguish between two categories.

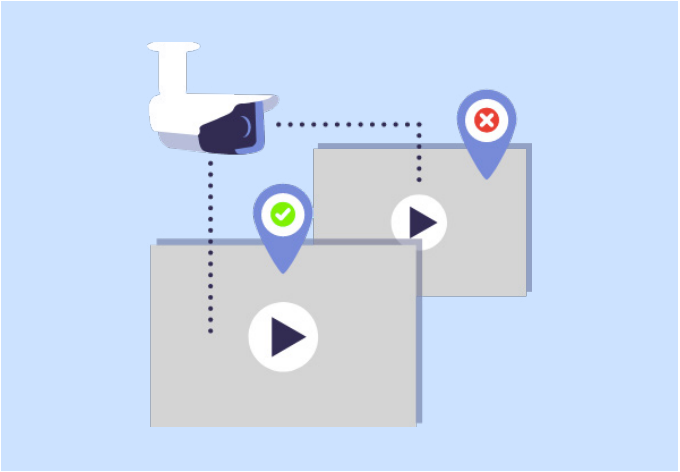

A true alarm

is an event that contains human activity in a scene where no one is supposed to be there.

A false alarm,

is an alarm that does not have this human activity and is considered as 'noise' for the security company who is monitoring the camera.

Once a clear objective has been outlined, you can begin to train and use artificial intelligence to achieve the desired outcome.

Example of a true alarm

Example of a false alarm

AI 101: HOW DOES IT ACTUALLY WORK?

When you hear the words artificial intelligence, they will often be accompanied by a mention of terms such as machine learning and deep learning. The first step to understanding how AI works is to distinguish between the different approaches being used by companies to power their software. As AI technology continues to become more prevalent in the security industry, the more you will be exposed to the different techniques.

The definition below outlines the difference between the terminology being used in the market:

Artificial intelligence
Any technique that enables computers to mimic human intelligence using logic, if-then rules, decision trees, and machine learning (including deep learning).

Machine learning
A subset of AI that includes abstruse statistical techniques that enable machines to improve at tasks with experience. The category includes deep learning.

Deep learning
The subset of machine learning composed of algorithms that permit software to train itself to perform tasks, like speech and image recognition, by exposing multilayered neural networks to vast amounts of data.

Within these different subsets of the technology, there are two main ways of working: rule-based algorithms and the neural network.

Rule-based algorithms

Most video analytics software that has been developed to solve the issue of false alarms, such as motion detection, line triggering etc. use what is known as a rule-based algorithm. This means they rely on a decision tree to make their final decision, where the algorithm will test multiple hypotheses until making a prediction.

When using a rule-based algorithm to analyse an alarm, the algorithm will ask multiple questions of the image, testing for an outcome before it is able to determine whether an alarm was indeed true or false. In this particular case, the algorithm will ask multiple questions such as whether there was movement and then how big that movement is etc. At each step, the algorithm will go down this decision tree in order to make a prediction.

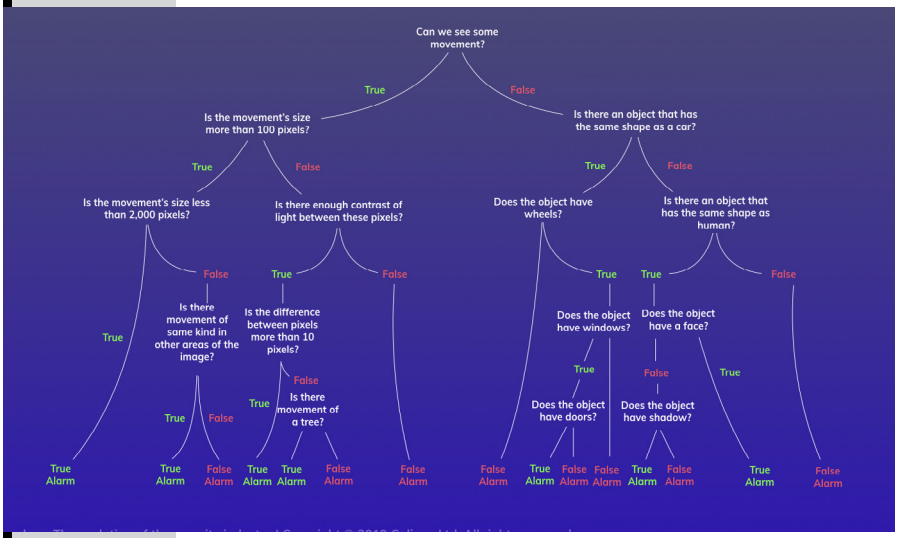

This technique has serious limitations:

1. It doesn't matter how smart the engineers behind those algorithms are, there will always be an edge case where the techniques will not work.
2. To make sure that nothing is being missed, the parameters are often tweaked at the expense of the overall accuracy.
3. Since most alarms are very complex, these more simple algorithms often fail at detecting them.

Calipsa does not use this technique but rather uses neural networks to solve the problem of false alarms.

The neural network approach

Even though it is scientifically inaccurate to state that a neural network works like the brain, their invention was largely inspired by it. Instead of designing complex algorithms like the decision tree described above to understand whether a particular image contains a specific element, we can build a neural network where layers of neurons will be stacked on top of each other between the input image and the decision.

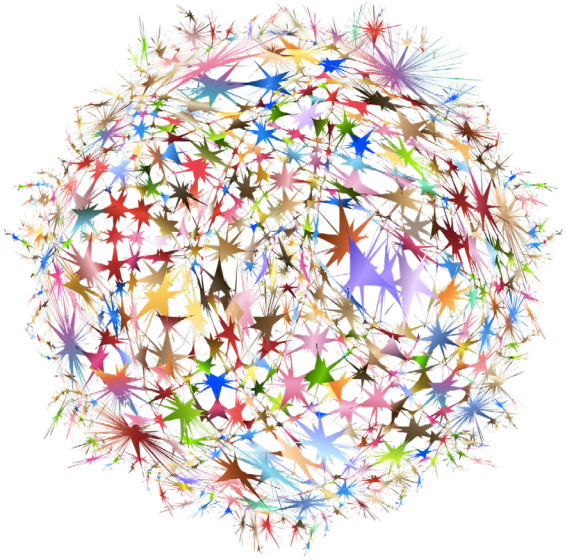

> **Neural network**
> Artificial neural networks are computing systems vaguely inspired by the biological neural networks that form the human brain. These systems 'learn' to perform tasks by considering examples, generally without being programmed with task-specific rules.

Let's dive a little deeper and try to understand how everything works by doing a 'brain scan'. The figure below represents the output of the activations at each layer for a specific example. In this case, the the objective is to know whether or not the image contains a cat. This process and the explanation is transferable - in Calipsa's case we use it to identify human or vehicle motion in video surveillance footage.

In the first layers of the neural network, the outputs look quite similar to the original image. However the deeper you go in the neural network, the more abstract the representation becomes. At each layer, the data is transformed little by little in order to have a deeper representation of what is going on. In the first layers, the neural network will typically recognise some basic features such as edges and colors. Then in the later layers some more advanced features will be recognised, such as faces and legs. This knowledge is built up through the neural network in order to make a final prediction - in this case, is this image a cat or not?

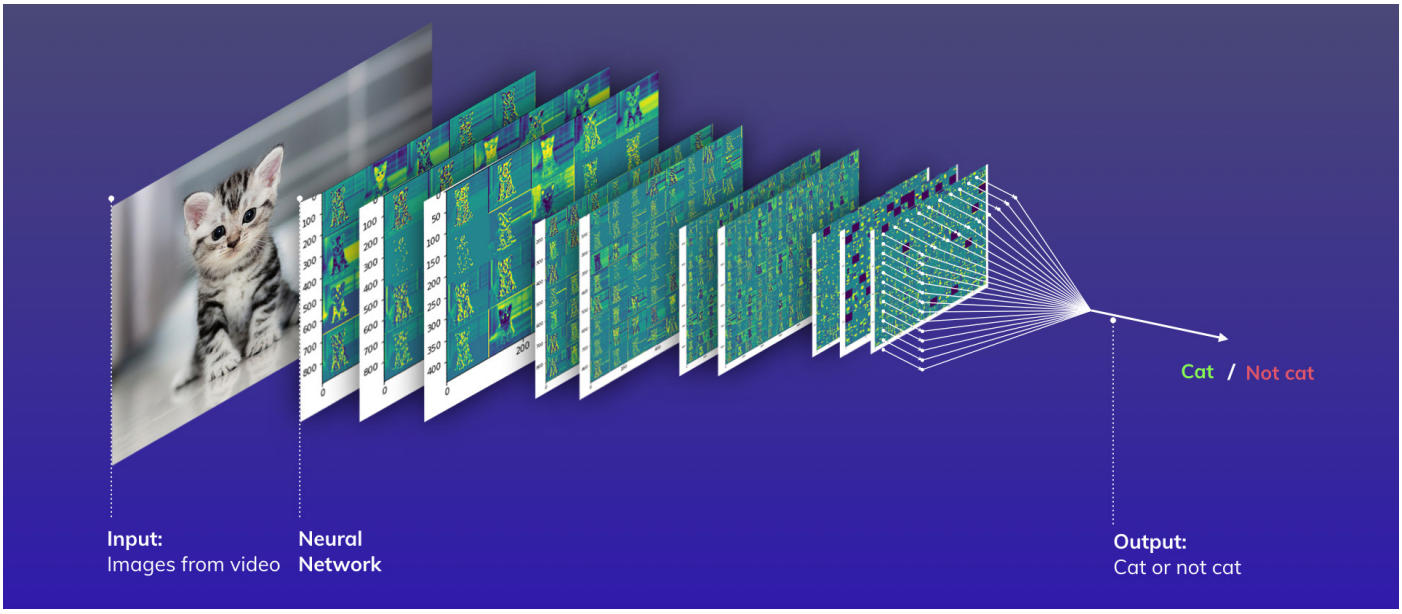

The process of getting the right representation for different images in order to make the correct prediction is done via a training process. Instead of changing each of the parameters of each neuron manually, we expose a neural network initialised with a random state to millions of images. At each step, we let the neural network make a prediction and update its parameters when it provides a wrong answer. This process is done via the backpropagation algorithm, and lets the neural network learn little by little from its mistakes until it outperforms all other methods.

Backpropagation
Short for "backward propagation of errors," it is an algorithm for the supervised learning of artificial neural networks using gradient descent. Given an artificial neural network and an error function, the method calculates the gradient of the error function with respect to the neural network's weights

Because we always start from a random state, the internal representation that is being built by training will be different every time we train a new neural network. It is the same with humans. Ask somebody to visualise a cat in their mind and the chances are their visualisation will be different from that of the person next to them.

If you pay enough attention to the above neural network, it is possible to notice that some neurons are completely black on the last layer. That simply means that some of the expected features were not present but others compensate to make the decision. For example, while you may think of a cat being a small animal with four legs and a tail, even if you don't see a tail in the image, you can still be pretty sure that it is a cat.
This example of a neural network for cats can be extended to pretty much anything and provides a lot higher accuracy than other methods, such as the rule-based approach.

THE BENEFITS OF AI FOR VIDEO SURVEILLANCE

Increased speed and efficiency

It is widely reported that control room operators will review a minimum of 3 alarms every single minute. With 95% of these estimated to be false, the result is a significant amount of time wasted on false activations and limited time available to action genuine events. Artificial intelligence allows operators more time to focus on the alarms that matter. By removing up to 90% of false alarms, the technology works hand in hand with operators to help them react quicker and more efficiently to security threats.

Highly accurate decisions

When developing deep learning algorithms for crime detection, there are two key metrics to consider: recall and reduction. Achieving both the right balance, and highest accuracy, in both of these areas ensures that the technology will perform at it's best - delivering on the objective to identify true and false alarms. Recall refers to how accurate the technology is at capturing true alarms. Reduction refers to the percentage of false alarms filtered out. At Calipsa, our recall rate is 98.7% while our reduction averages at 90% [Feb, 2020], meaning the technology is extremely accurate in identifying security threats and nuisance alarms.

Continuous learning and improvement

Artificial intelligence is constantly learning. As AI algorithms are exposed to more video footage, the better they become at identifying true and false alarms. This means that companies that invest in AI technology aren't just purchasing a static product, but a product that will continue to get better and better over time. For example, through continuous learning and iterations of Calipsa's algorithms, we have successfully transformed the reduction of our product from 30% accuracy to 90% in a matter of years. Those highly accurate algorithms are now used to detect false and true alarms and will continue to develop as more they receive more footage to learn from.

CASE STUDY:
USING VIDEO ANALYTICS TO IMPROVE OPERATIONAL MONITORING STRUCTURE

By implementing Calipsa's AI-powered False Alarm Filtering Platform, EyeQ Monitoring have been able to efficiently scale their monitoring operations

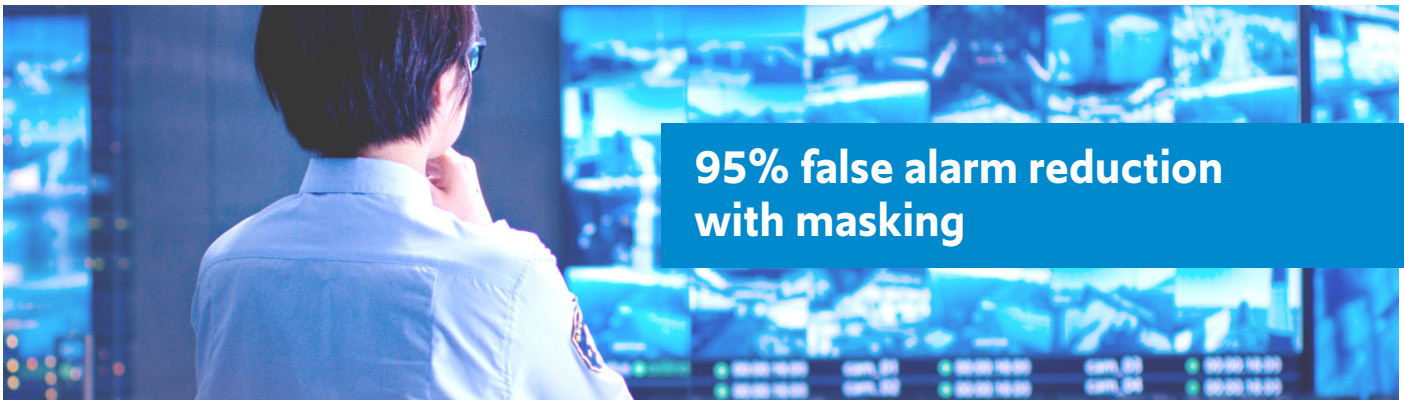

EyeQ Monitoring, formerly NVMC Solutions, is one of the largest focused-providers of live video monitoring in the United States. Established in 2007, EyeQ Monitoring installs and maintains the latest video surveillance technology, and also employs and trains agents based in the United States who then monitor those systems remotely after hours - responding to threat issues in real-time.

EyeQ Monitoring had been using analytics solutions for 2 years and decided to implement Calipsa's reliable true cloud solution to allow them to scale their video monitoring business effectively. With Calipsa, EyeQ Monitoring have been able to:

- ✓ Filter up to 95% of false alarms using deep learning algorithms
- ✓ Provide true alarm output to monitoring agents from live monitoring sites
- ✓ Reduce the number of agents required to monitor the same number of sites
- ✓ Improve the deployment time for integrating analytics

> *Calipsa does a great job of classifying false alarms. We compared competing vendors to get a valid comparison of performance and after properly setting masking, we saw the accuracy rate of Calipsa average 95%. My department evaluates many innovative leading edge products each year and I have to say that Calipsa was one of those products that delivered.*
>
> **WARREN NEUBURGER**
> Chief Technology Officer

CONCLUSION

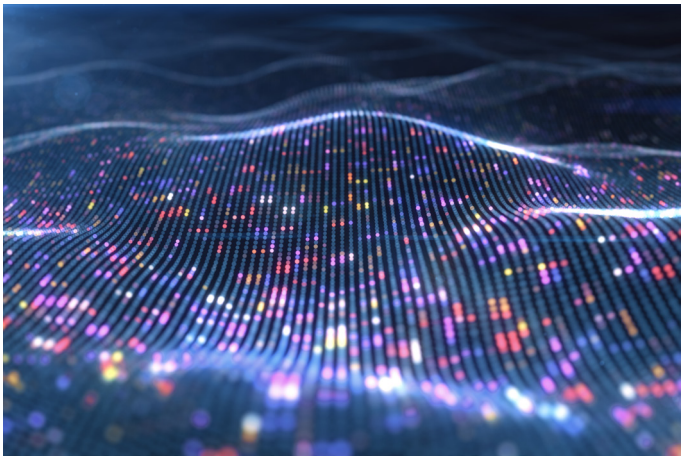

In the context of security, deep learning and artificial intelligence have huge potential, with the power to transform performance. It's hard to deny that when used correctly, the technology brings enormous value to video monitoring. However it's important to remember that AI should never be seen as a complete replacement for existing processes. Security professionals should not be looking to artificial intelligence to run their monitoring operations, rather it should be viewed as an extension and enhancement to current their existing operations. By building a collaboration between humans and technology, security companies will be able provide more value to customers than they could possibly do by using either in isolation.

By integrating technology such as Calipsa into monitoring operations as an extension to human verification, security businesses can have confidence that they are providing their customers with an accurate and constantly improving service.

ABOUT AUTHOR

Boris Ploix

CTO and Co-Founder - Calipsa

Boris co-founded Calipsa in 2016 with a mission to make the world a safer place. As Calipsa's CTO, drawing upon 10 years' machine learning and applied mathematics experience including Masters degrees from both Imperial College London and UCL, Boris is instrumental in driving Calipsa's machine learning technology and product development.