VIDEO ANALYTICS COMES OF AGE

Enhanced surveillance: video analytics has often failed to live up to performance expectations in the field

©Getty Images

As a technology, video analytics has a slightly underwhelming history. The prevalence of CCTV cameras and the limited attention span of humans – typically reported at around 40 minutes – has driven the demand for intelligent video analysis (IVA) systems. Solutions that can process large volumes of video, in order to generate automated alerts and useful information, hold obvious benefits for operational monitoring, business intelligence and the protection of vulnerable or remote sites.

Yet organisations often face the aforementioned underwhelming choice when it comes to selecting a video analysis system to enhance their security, whether because of variable detection rates, the persistent false alarms of simple systems or the complex setup and cost of more advanced solutions. This has led to frustration for both installers and security operators in a market that already offers few definitive benchmarks across a confusing array of competing solutions.

It also hasn't been served well by the initial hype that surrounded analytics when the software debuted on the security market. Some vendors promised more than what the technology, at that stage, was capable of delivering. Buyers were invariably disappointed.

The introduction of video encoding, where a visual input is translated into a compressed format for transmission, was the catalyst for video analysis and video motion detection (VMD). Motion vectors are an outcome of the encoding process, providing information that can be readily used to detect motion across a scene. This avoids the need for further video analysis functions and processing overhead, which makes VMD easy to incorporate into devices – as well as low cost. In many cases, VMD is offered as a free utility.

VMD is incapable of distinguishing between a valid subject, such as a person or a vehicle, and other elements that introduce motion into a scene, however. This is most obvious when processing an external camera input, where the effects of weather and variable lighting render most VMD systems ineffective – as it becomes impractical for an operator to manually review numerous "nuisance alarms". Even systems offering more advanced video analysis can be prone to these issues.

The inherent limitations of VMD and less sophisticated video analysis systems has led to the development of more advanced IVA applications. These offer more analytical functionality, with software algorithms that are able to distinguish between subjects, such as people or vehicles, as well as types of behaviours, such as running or rolling. Most allow the definition of more sophisticated security scenarios, such as time-based loitering, and conditional rules involving movement between more than one zone.

Crucially, though, many IVA systems offer some form of nuisance alarm filtering to exclude the effects caused by weather, variable lighting and shadowing, car headlights (or other glare), and persistent movement in background vegetation or other scene elements. In some cases, this is based on background learning of a scene, where the software builds up an understanding of typical movement or activity in order to discern anomalies and therefore intrusion events. Processing such levels of scenario-based rules and nuisance alarm filtering, to the degree of accuracy required for a reliable operational tool, typically requires significant computational power.

So what does this mean in practice? Most large or highly-secure sites that have adopted server-based IVA across multiple cameras, running on dedicated IT infrastructure. In terms of installation, this type of analytics requires what, in essence, amounts to a significant IT project, with configuration of complex server and network infrastructure. As such, the upfront cost and ongoing maintenance costs involved in such a system are often substantial. This reliance on a costly and complex systems has limited the adoption of server-based IVA to larger or very secure sites, where higher performance levels are required to protect critical infrastructure. Smaller sites typically opt for the simpler VMD option, often embedded on a camera "at the edge", which is by its nature easier to deploy, particularly when installing this functionality with new IP cameras.

One of the other defining features of the majority of advanced IVA systems is the amount of effort involved in calibrating and configuring the system. Reliability depends on the correct siting of cameras and accurate modelling of the scene and the security scenario. Most applications provide manual tools and editors for these tasks, which often involve the installer having to add a number of details and measures to the scene – for every camera that is being configured. This can be a very effort-intensive and specialist task.

Such a combination of costly, specialist IT infrastructure and complex configuration has led many installers and integrators to question the value of server-based IVA. One noted that hardware issues, such as hard disk failures and the high cost of ensuring continuity of power, have led them to avoid server-based video analytics. Others cite complexity of setup (and ongoing fine tuning) as the main reason for server-based systems becoming increasingly unviable and uneconomic for installers and customers.

In fact, lack of affordability is a major factor in the growing frustration with server-based video analytics. The cost of a single software licence can run at upwards of £400 (€450) per channel, with an additional 10-20 per cent maintenance fee payable annually. Combine this with the cost of buying and maintaining server hardware, and the total bill for a server-based system can be significant – often exceeding the cost of purchasing the cameras themselves.

The trade-offs between performance, complexity and cost have defined the video analysis landscape. But there is an alternative approach – and it comes in the form of smart and simple VCA at the edge. We'll look at how this new generation of embedded applications stack up against their server-based counterparts shortly. But it's important to introduce another factor which has been a major factor in the uptake of analytics: credible performance benchmarking.

Perhaps the most important of the performance

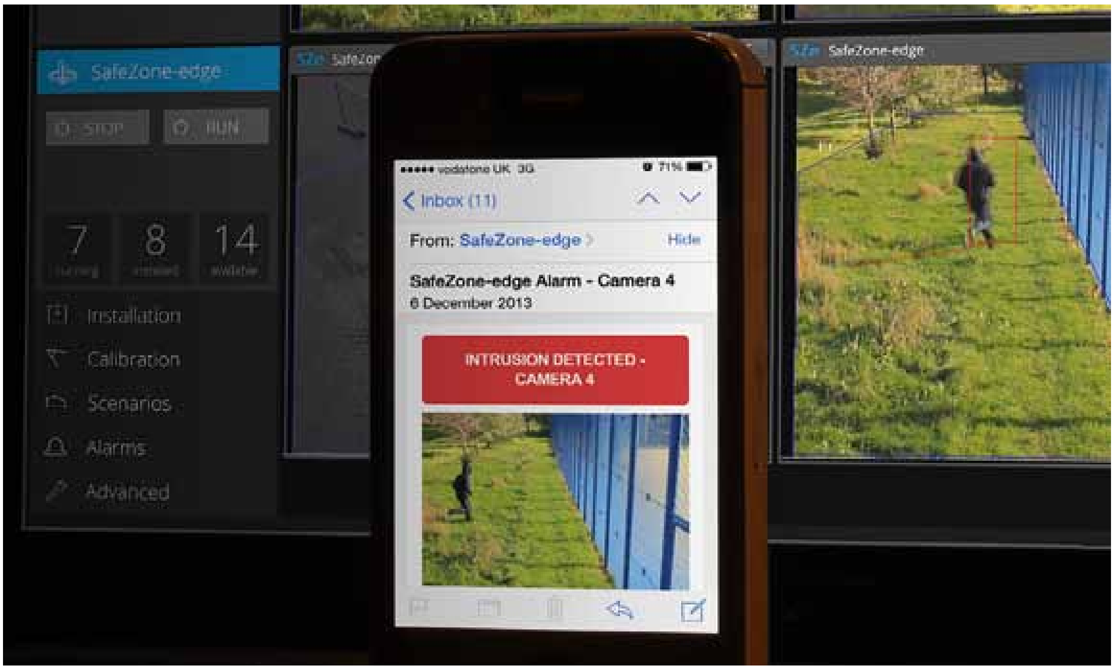

From: SafeZone-edge

SafeZone-edge Alarm - Camera 4
6 December 2013

INTRUSION DETECTED - CAMERA 4

©Digital Barriers

benchmarks that has driven the adoption of sophisticated IVA and server-based analytics is the Imagery Library for Intelligent Detection Systems (i-LIDS) scheme. i-LIDS is a UK Government initiative to facilitate development and selection of video analytics to meet Government requirements for security operations. It has been developed by the Centre for Applied Science and Technology (CAST) in partnership with the Centre for the Protection of National Infrastructure (CPNI).

i-LIDS offers a video library to benchmark video analysis systems against a number of security scenarios, based on accuracy. For example, sterile zone monitoring applications meeting i-LIDS performance criteria may be certified as a primary (sole) detection system or secondary (support) measure. Other scenarios include parked vehicle, abandoned baggage, doorway surveillance and new technologies.

The datasets for the event detection scenarios each contain approximately 24 hours of footage. Each of these datasets is filmed to represent all weather conditions, time of day and scene densities expected within the scenario, and consist of two or three camera views – referred to as stages – which are further segmented into shorter video clips of 30 to 60 minutes. The footage accurately represents real operating conditions and potential threats. In the case of sterile zone monitoring, analytics systems must detect the presence of persons in a restricted area or "sterile zone". CAST conducts regular trials of video analytics systems on each i-LIDS scenario. Systems demonstrating a sufficient level of performance are listed in a catalogue of approved products distributed to critical national infrastructure security managers.

The i-LIDS certification scheme began in 2006. In the US, the TRECVID evaluation – an initiative of the National Institute of Standards and Technology – aims to provide similar benchmarking. These schemes will prove crucial in ensuring that analytics systems are proven to be reliable and effective, and security managers at CNI facilities, both in the UK and around the world, should always look for these approved products as the choice of analytics software expands.

Only recently has the introduction of more powerful on-board camera and encoder microprocessors offered the ability to "piggyback" more sophisticated edge-based IVA. This gives the potential to deliver i-LIDS primary certification on an embedded application.

This has been proven by a very small number of IVA systems that can now provide government-approved performance at the edge. It is important to note that the majority of this generation of more powerful edge-based video analysis systems are designed around the same complex set-up and configuration approach as their server-based predecessors, however. This means that camera calibration remains a largely manual and time-consuming task, with configuration of security scenarios requiring a high level of installer expertise.

In addition, the cost of purchase can be comparable to that of a server-based licence – with some advanced edge-based (or distributed server and edge) applications priced per security scenario. So while edge-based IVA presents an opportunity to move away from the complex world of server-based analytics, the choice for installers and end customers is still less than straightforward.

©Digital Barriers

Credible performance benchmarking, such as the UK's i-LIDS scheme, inform users of the software's real-world capabilities

In recent times we've seen a different approach to automated intrusion detection emerging, with systems being developed to overcome the limitations of both server-based and conventional edge-based systems.

The best of these IVA applications offer performance and exceptional ease of use, and are priced to be affordable to smaller as well as larger sites. They focus on simplicity, in terms of setup, operation, purchasing and support. This means that for the first time we're seeing edge-embedded analytics which are suitable for deployment as a primary intrusion detection solution for high-security and CNI facilities.

Of course, while high-security sites and larger enterprises may find it easier to afford (and justify) the sophisticated IVA systems, the necessity for reliable intrusion detection, free from nuisance alarms, is equally acute for smaller sites and remote alarm receiving centres (ARCs). While a larger facility may have dedicated full-time security, remote monitoring services incur call-out costs in responding to alarm events. Thus the need for reliable and affordable automated intrusion detection, without the need for a dedicated IT back-end, is a universal requirement across industries.

In future, we are likely to see even more edge processing power, which will mean further features that had previously only been considered suitable for server-based applications are embedded on cameras – with reduced costs and complexity. These will provide even more reliable automated intrusion detection functions, all at the edge.

And as technology changes and improves, benchmarking schemes such as i-LIDS will prove even more important, providing an independent means of sorting those applications with performance that meets the requirements for CNI facilities from those that promise a huge range of features but come up some way short when it comes to robustness and reliability.

Dave Oliver is Director UK Enterprise Account Sales at Digital Barriers, which recently launched SafeZone-edge, a truly edge-based, truly intelligent video analysis system, which uniquely combines reliability, simplicity and affordability.