Critical Asset Protection, Perimeter Monitoring, and Threat Detection Using Automated Video Surveillance

Dr Alan J. Lipton, Craig H. Heartwell, Dr Niels Haering, and Donald Madden
ObjectVideo
11600 Sunrise Valley Dr, Suite 290
Reston, VA 20191
Email: {ajl | nhaering | don.madden}@objectvideo.com

Abstract

The latest generation of computer vision technology is revolutionizing concepts, applications, and products in video surveillance and CCTV. This is of prime relevance to security for large outdoor facilities such as commercial airfields, refineries, power plants, and office/industrial campuses. Most airfields, for example, have open (unfenced) perimeters, high volume heterogeneous traffic, are easily accessed on foot or by water, and exist in areas where regulations providing a safety buffer are difficult to legislate or enforce. And all airfields require 24x7 outdoor monitoring – snow, fog, rain or shine. Likewise, most high-value facilities appealing to criminals and terrorists are in close proximity to public areas (roads, residences, city, etc.).

The appeal of automated real-time surveillance is obvious – maximizing efficiency and effectiveness of security personnel and resources while increasing the probability of preventing a serious security breach. ***Computer vision*** based solutions have the potential for very discriminating detection and very low false alarms. The bottom line is that applied computer vision has the potential for the greatest return on investment (ROI), both short-term and long-term.

1. Introduction – why video?

Recent world events have prompted government and industry alike to rethink their approach to physical security. The threats we face are no longer large scale military attacks from known adversaries outside our borders. Our fears today derive from the possibility of a small group of individuals, perhaps already within our borders, having the ability to cause a large amount of damage. Such attacks could carry an extremely high cost in terms of economic and environmental damage, reduced national morale, and loss of human life. Not only has the nature of the threat changed, but recent events have redefined the nature of the targets. No longer are the prime targets military in nature – now public infrastructure and innocent civilians are facing attack. Organizations that control critical infrastructure and national assets such as airports, power production facilities, water supplies, and public transportation routes are feeling the pressure to increase their ability to detect "asymmetric threats" and respond to them in a timely manner.

These changes have forced a higher sense of vigilance upon many organizations previously unconcerned with major attack. Accordingly, we see an increase in the awareness of physical security issues and technologies along with increases in physical security budgets. The US Government has earmarked $37B for homeland security and created a new office to administer homeland security programs. A very large piece of the physical security pie is being devoted to video surveillance infrastructure and research.

Why video? People like video. It's one of the most ubiquitous sensing modalities available. It is real-time and highly intuitive (it's easy to understand what is happening in a video stream). Yet, curiously, video surveillance is not used primarily for real-time interdiction. It is used in two basic modes: as a deterrent and as a forensic tool. People are less likely to commit criminal activities if they believe they will be caught on camera; and if something does occur video is frequently used forensically to figure out what happened. Hence there is an apparent paradox: video is a ubiquitous, real-time, intuitive sensor that is *not* being used to provide real-time actionable intelligence.

Not using video surveillance to its full potential as a real-time threat detection system is unfortunate because video is an excellent tool in the fight to protect critical infrastructure. Most threatening activities begin with a prelude of hostile intelligence gathering – adversaries will often "case the joint" for a period of weeks or months before an attack. Appropriate video-based counter-measures can be used to detect these hostile patterns of activity. Furthermore, most hostile attacks begin with a perimeter breach, providing early opportunities for detection and interdiction. Again, video surveillance is an excellent tool to detect (in real-time) the nature and composition of a threat, its pattern of attack, whether it is a main force or merely a

diversion, and monitor the progress of an attack and the effect of counter-measures.

2. What's wrong with video today?

Figure 1 - "State of the Art" video surveillance system

Figure 1 shows a "state of the art" video surveillance system. Organizations often spend millions of dollars on video surveillance infrastructure consisting of hundreds or thousands of cameras. These camera feeds are usually backhauled to a central monitoring location where some of them are recorded for a period of time on local video storage media, and some of them are displayed in real-time to one or more security personnel on a bank of video monitors.

No matter how highly trained or how dedicated a human observer, it is impossible to provide full attention to more than one or two things at a time; and even then, only for a few minutes at a time. A vast majority of surveillance video is permanently lost without any useful intelligence being gained from it. The situation is analogous to an animal with hundreds of eyes, but no brain to process the information.

3. Automated video surveillance

The solution to this problem is automated video surveillance (AVS)[3,4,5,6,7,8]. That is, computer software that watches video streams to determine activities, events or behaviors that might be considered suspicious and provide an appropriate response when such actions occur. The key technology is called ***Computer Vision***. This is a somewhat obscure branch of mainstream artificial intelligence research involving teaching machines to understand what they "see" through a camera. Traditionally, computer vision has had limited success in real-world commercial applications, but recent advances in technology and computational power along with a move of key talent from academia into industry have allowed computer vision to come out of the lab and into commercial video surveillance products.

Several years prior to the successes of applied computer vision, manufacturers of Digital Video Recorders (DVRs) and other video processing solutions began a dubious flirtation with a technology called Video Motion Detection (VMD)[2,9]. This technology analyses video imagery and determines where there is motion in the scene. The theory being that anything that moves is something interesting to users. What quickly emerged was the result that there is a lot of motion in the world that is highly annoying and something more was needed. Figure 2 illustrates the issue. In a situation such as detection of watercraft, the entire image is moving and VMD gets very confused. A more intelligent AVS system can accurately extract the relevant information from generic motion.

One of the main drivers behind the recent success of computer vision in surveillance applications has been the large amount of funding invested in the technology by organizations such as DARPA. For the last 25 years, DARPA has funded an Image Understanding program culminating in a 3-year project called Video Surveillance and Monitoring (VSAM) which concluded in 2000[1]. The goal of this program was to develop state-of-the-art algorithms for automated video surveillance – and in the last several years, this research has been bearing commercial fruit.

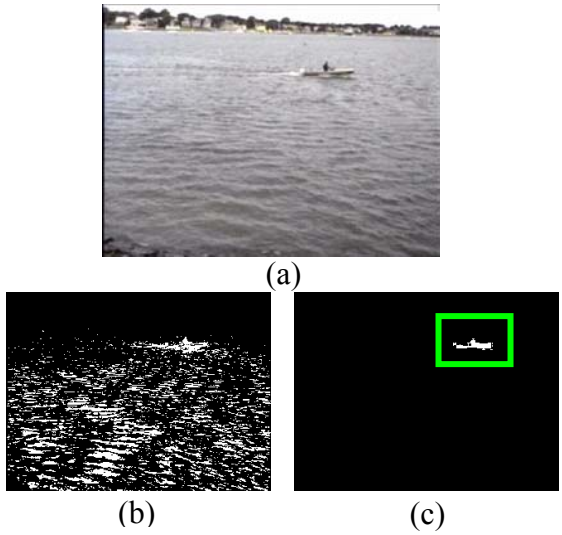

Figure 2 - VMD vs. AVS. (a) The source image - everything is moving. (b) VMD has trouble detecting the true object. (c) AVS accurately detects the object

ObjectVideo is one company that has successfully expanded and commercialized some of the

technologies developed by the VSAM program. We have used this technology as the basis of an AVS product that monitors video streams in real-time and detects activities that have been prescribed as interesting or suspicious.

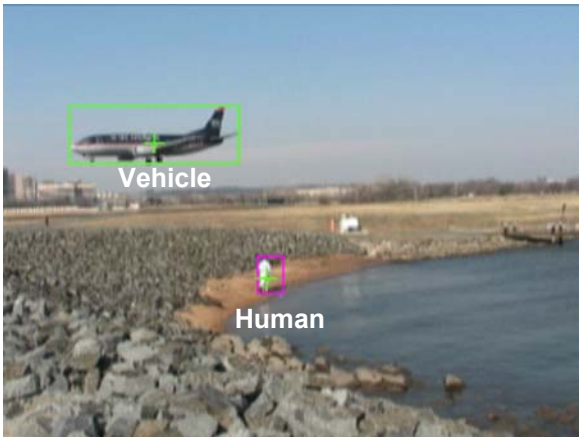

Figure 3 - Objects detected on the perimeter of Reagan National Airport

ObjectVideo's solution watches video streams and extracts descriptions of all relevant objects. It employs sophisticated algorithms for detection[3] and tracking of all relevant objects in the camera's view. It also contains algorithms for classification[10] of objects into specific types[1]. Figure 3 illustrates an image from ObjectVideo's application. Here an unauthorized human has been detected on the perimeter of an airport.

These basic functions are nothing by themselves. They need to be part of a truly effective physical security system. To achieve this, the system needs to be

- **Scalable.** To effectively monitor a large enterprise, a system has to scale to hundreds, perhaps thousands, of sensors. It needs to be able to integrate and fuse information over both space and time
- **Integrated.** ObjectVideo realizes that video surveillance is a key part of a total solution, but not a solution unto itself. Any AVS system must be able to integrate seamlessly into other physical security systems. It must be able to physically connect to access control systems, data bases, and other sensor systems. It also has to be able to integrate with security personnel to provide actionable real alerts that increase their effectiveness.
- **Operationally effective.** The problem with VMD systems is that they are so annoying that end-users turn them off. Users have to be comfortable that an AVS system will provide a high level of relevant event detection with a low occurrence of nuisance false alarms.

There is a very broad spectrum of products that claim to be useful for AVS. Unfortunately very few of them are up to the challenges of a real-world 24x7 outdoor environment or the demands of operating environments where flexibility, accuracy, and usability are mandatory. Understanding the underlying science and technology behind computer vision and how to critically characterize target environments and operational requirements is the key to selecting the right solution(s) for specific applications.

4. Transitioning to the Real World

Computer Vision, although not a new science per se, is still very new as a commercially viable technology. Moreover, because applied computer vision products vary greatly and will evolve at the pace of *Moore's Law*[2], it is mandatory that adopters of computer vision technology have in place a solid technology insertion strategy as part of their overall security policy and strategic security planning. The following key strategy points are recommended to aid evaluators and adopters of computer vision:

1. Use in-house IT (information technology) talent – they are used to dealing with rapid technology changes.
2. Become educated about the science of computer vision – use consultants, vendors, and analysts to keep abreast of the market and use current literature to watch for trends indicating future directions and advances.
3. Look hard at new and emerging technology and products – these are going to be the best

[1] The basic product classifies objects into "human" and "vehicle" classes, however the product also provides a simple mechanism for custom classification algorithms to be developed for customers with more specific needs such as to distinguish humans from animals, or trucks from aircraft.

[2] Moore's Law and similar "rules of thumb" state that the "power" of computer-related technologies and applications double every 18 months. One can expect to see correspondingly dramatic changes and advances in applied computer vision products for the foreseeable future.

indicators of future capability and emerging thought-leaders and de facto standards.

4. **Try Before You Buy** – the wide range of computer vision applications will result in a large number of products claiming to be solutions to real-world security problems. The truth is that most computer vision products will have very limited application, and in tough (24/7/365/outdoor) environments, most will not be operationally effective.

5. **Vision Alignment** - Find experts (consultants, analysts, vendors) who share your vision, and use them; but own your vision, don't delegate it to outsiders. Also, take advantage of industry inertia; share your knowledge, experience, and strategies with others in your industry, and use the collective power of your industry to make computer vision technology market-driven to your benefit.

5. Case Studies

ObjectVideo promotes the strategy of *try before you buy*, and encourages adopters of video surveillance to evaluate candidate AVS solutions in the real environment in which it will be deployed. "Kicking the tires" is the only guaranteed way to be sure that a particular solution will be ***operationally effective*** – maximum detections with minimum false/annoyance alarms (few enough to not cause users to ignore or disable the system) and the right level of flexibility to be useful as part of an evolving security deployment or strategy.

ObjectVideo has conducted several pilot projects to demonstrate and evaluate the application of computer vision to protecting personnel, critical infrastructure and high value assets. ObjectVideo's powerful analysis tools enable us to rapidly qualify an environment, identify the cause of false alarms, and determine the effectiveness of different event and alert rules[3] and rule combinations. Our experience in field trials has convinced us that this capability is essential for effective and efficient deployment of computer vision solutions in complex environments.

Here we provide an overview of two case studies – an oil refinery and an airport. These case studies highlight the complexity of real-world environments but demonstrate the power of AVS and the recommended approaches for applying AVS and maximizing video surveillance technology for real-time critical asset protection, perimeter monitoring and threat detection.

5.1 Case Study: Oil Refinery

Our first case study is a pilot conducted at a typical oil refinery in which the host already had in place a particular camera and a particular view they were interested in evaluating. The host's goals were to gain an understanding of what can be achieved with AVS to automatically monitor their fenced perimeter, including understanding options for cameras, camera placement, illumination, and prescriptive threat detection rules.

Vitals

Host	Oil Refinery
Application	24x7 monitoring of a 5-mile perimeter with an already in-place fence
Threat Scenarios	• Vehicles or people "watching" the facility from outside (warning) • Vehicles or people loitering outside fence (severe) • Vehicles or people approaching fence from roads running parallel to fence (severe) • Vehicles or people going over, under, or through fence (critical) • Vehicles or people near inside of fence (critical) • Abandoned objects near fence (inside or out),e.g., duffle bag of explosives (severe) • Camera disabling (moved, masked, lost power)

Physical Environment Issues

- Weather
- Insects on camera housing

Operating Environment Issues

- Low-level street lamp illumination and vehicle headlights at night
- Night time lighting phenomenon (refinery flare and flashing amber beacon on regular patrols by refinery security vehicles)

[3] Event rules specify the events to be detected. Alert rules specify the actions to be taken when an event has been detected.

- "Film" and dust accumulation in front of camera lens, rain, snow, wind (swaying grasses and trees)
- Manual PTZ use by security personnel

Set-up

The already in-place camera provided by the host simplified set-up. The only remaining set-up required was to split the camera video feed and run the signal into the AVS sensor.

Camera Type	Vicon Integrated Day Night (IDN); PTZ mount; all-weather housing
Camera Placement	Fixed placement; 60-foot industrial mast
Camera FOV	22 degree FOV

Figure 4 shows an aerial view and mark-up of the observed area.

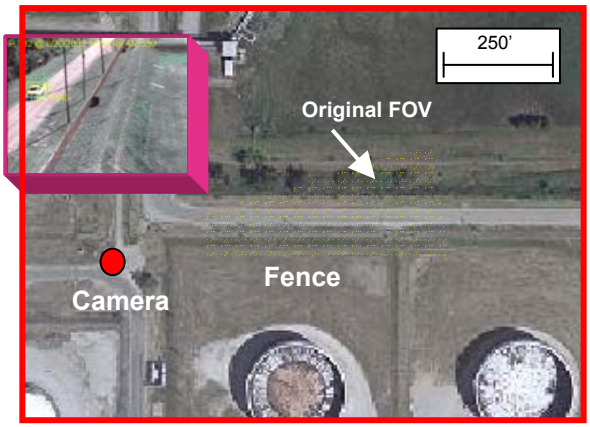

Figure 4 – Surveillance test area

Execution

Using the live camera input to the AVS sensor, 24x7 surveillance was conducted and recorded for a period of 10 days. During that time, ample null-hypothesis observation was conducted to analyse the normal activity and night time phenomena, some of which was initially responsible for a large number of false alarms. In addition, threat scenarios were staged in all categories at both day and night and several rule sets were used to evaluate the individual and combined effectiveness of different rule types and rule combinations.

Results

The day time performance of the out-of-the-box system was nearly ideal. Non-optimal camera placement limited the range of the camera somewhat, but within the predicted effective range, we achieved a better than 99% detection rate by simply using the right combination of rules to provide more than one detection opportunity per incident, see Figure 5, where False Alarm Rate is abbreviated FAR and Probability of detection is abbreviated Pr(d).

Night time performance was further limited by lack of adequate illumination in some areas of the camera's view, but there were no other fundamental problems caused by night time other than reducing the detection range. The night time issues that did cause problems were due to lighting phenomena, all of which were representative of the real world of the refinery. Lighting changes due to refinery flares caused no problems, but headlights and flashing lights did.

Scenario	Rule		Pr(d)	Pr(d) (Day)	Pr(d) (Night)	FAR /hour	FAR (Day)	FAR (Night)
Camera Disable	Masking		1.0	-	-	N/D	-	-
Camera Motion	Camera Motion		1.0	-	-	N/D	-	-
Hostile Surveillance	Abandoned Object		1.0	-	-	N/D	-	-
Perimeter Incursion	Combination of Directional Tripwires	Tripwire 1	.90	.94	.87	0.8	0.2	1.6
		Tripwire 2	.73	.70	.75	1.0	0.0	2.0
		Tripwire 3						
		Optimal combination ‡	.97	.98	.97	1.8	0.2	3.6

Figure 5 – Performance of Out-of-the-Box @ 450 feet

The unique visual phenomena created by the flashing amber lights atop the operations vehicles at the refinery were unexpected, and initially caused an unacceptable number of false alarms – as high as 2 per hour per camera would result in over 100 false alarms over a night shift. Although the "flashies" resulted in clustered false alarms (many immediately back-to-back), the total annoyance level of that many false alarms would be operationally unacceptable. In addition, there were a few false alarms caused by headlights from cars coming around an inside corner of the surrounding road. Using our analysis tools, we were able to analyze a week's worth of captured video and apply appropriate filters in a matter of hours, resulting in complete elimination of all observed false alarms. While we would never expect perfect operation, this dramatic level of improvement in a reasonable sample promises operational acceptance.

The tuned system was run for an additional 18 hours to ensure that there were no residual false alarms and that the filtering had not reduced the detection rate.

Scenario	Rule		Pr(d)	Pr(d) (Day)	Pr(d) (Night)	FAR /hour	FAR (Day)	FAR (Night)
Camera Disable	Masking		1.0	-	-	N/D	-	-
Camera Motion	Camera Motion		1.0	-	-	N/D	-	-
Hostile Surveillance	Abandoned Object		1.0	-	-	N/D	-	-
Perimeter Incursion	Combination of Directional Tripwires	Tripwire 1	.98	1.0	.95	0	0	0
		Tripwire 2						
		Tripwire 3	.90	.98	.81	0	0	0
		Optimal combination	.99	1.0	.99	~0	~0	~0

Figure 6 – Performance of Tuned System @ 450 feet

Recommended Surveillance Configuration(s)

Despite the excellent results achieved with the original configuration, the placement and FOV of the camera limited the detection range somewhat. An optimal placement and configuration were provided to the host for consideration in future deployment of new cameras (see Figure 7). This recommended configuration can be repeated across an array of cameras to provide 100% coverage of the entire facility perimeter.

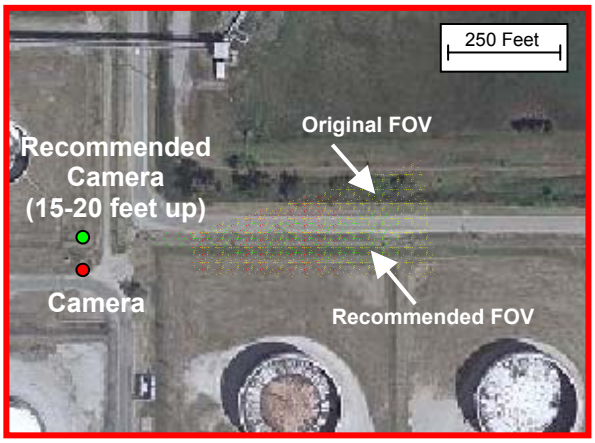

Figure 7 – Recommended placement and configuration for maximum effective monitoring

Specific recommendations in the following areas were provided in a final report to the host:

- Cameras
  - Static cameras or pre-sets
  - High-quality IDN cameras w/ illumination
- Camera View and Placement
  - Lower placement (15-20 ft)
  - Narrower Field of View
  - Overlapping Fields of View
- Further Testing
  - Usability
  - Seasons/weather
  - As part of a phased deployment approach

Expected Operational Effectiveness

Application	Effective Detection Rate	Observed False Alarms
Camera Disabling	100%	0
Camera Motion	100%	0
Counter-Surveillance	100%	0
Perimeter Threats	99%	< 0.01 / hour

Table 1 – Cumulative performance of tuned system

With the recommended adjustment and deployment configuration, we would expect to achieve nearly ideal performance around the entire perimeter of this facility (and, indeed, others like it). During the period of this pilot we did not experience a full range of weather conditions, however, based on our experience adjusting such systems up to ideal performance, we would expect to achieve and maintain ideal performance on a year-round basis with only occasional adjustments that should be covered by any standard system maintenance contract.

5.2 Case Study: Boston's Logan Airport

Our second case study is a pilot conducted at Boston's Logan airport. Logan's goals were to gain an understanding of what can be achieved with AVS to automatically monitor the airfield's waterfront perimeter and a newly established safety zone on the water (0-500') in which public boat traffic is highly restricted. The specific results and recommendations will be used to help design and specify a perimeter surveillance system.

Vitals

Host	Boston's Logan Airport
Applications	• 24x7 monitoring of an open 8-mile waterfront perimeter • 24x7 monitoring of water safety zones (0-250' no-boat zone and 250-500' no-stop zone)
Threat Scenarios	• Water craft passing through 250-500' safety zone (interest) • Water craft stopping/ loitering in 250-500' safety zone (warning) • Water craft stopping/ loitering at ILS piers (severe) • Water craft passing through 0-250' no-boat zone (severe) • Water craft stopping/ loitering in 0-250' no-boat zone (severe) • Water craft approaching shoreline in 0-250' no-boat zone (critical) • People or vehicles on mud flats or marsh land within 0-500' safety zone (warning) • Vehicles or persons on land approaching airfield from shore (critical)

Physical Environment Issues

- Water (constantly moving)
- Tall dense grass to support endangered marsh wildlife (constantly waving in the wind)
- Tides (edge and height of water constantly changing in view of camera)
- Weather (rain, snow, fog, rapid shifts from sunny to cloudy)
- High winds
- Heavy insect population (some attracted to infrared illumination)
- Shore birds (primarily sea gulls)
- Nocturnal fauna (skunks, opossums, etc.)

Operating Environment Issues

- Non-optimal camera placement
- No artificial illumination on perimeter
- Plane landing and take-off lights at night
- Headlights and flashing lights on operations vehicles (infrequent patrols at night)
- Water reflections of lights from planes, vehicles on opposite shore, water craft, city skyline, and residential areas
- Clam diggers who will be allowed access to adjacent mudflats within 0-500' safety zone
- Camera height and camera placement restrictions and around airfield

Set-up

Because Logan's airfield currently has very few in-place surveillance cameras, and power is intermittently accessible, ObjectVideo used a portable, self-contained rig to allow sampling at any location and in any configuration. We also provided infrared illumination for night time testing.

Camera Type	Sanyo Digital IDN; Cosmicar/Pentax zoom lens; Pelco outdoor housing; fixed mount and PTZ mount
Camera Placements	Variable placements: 8 foot height at edge of airfield (water incursion) 10 foot height slightly off edge of airfield (perimeter incursion) 12-14 foot height well off airfield (marsh and mudflats)
Camera FOV	Testing was done at 3 settings: • 10-14 degree FOV • 20-25 degree FOV • 45 degree FOV

Figure 8 – Locations selected for perimeter incursion staging

Figure 11 – Camera view of IR-illuminated perimeter incursion

Figure 9 – Locations selected for water incursions

FOV: 21°
Expected Range: 300'

Figure 12 – Camera coverage and view of marsh area

FOV: 10-14°
Expected Range: ~450'

Day view (with tripwires)

Night view (with tripwires)
(Illuminated)

Figure 10 – Camera coverage and view of perimeter

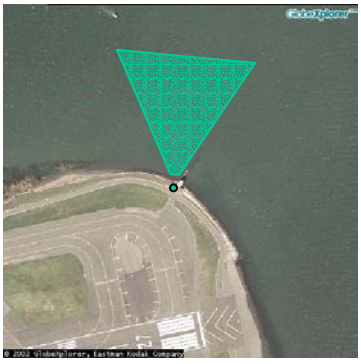

FOV: ~45° Expected Range: ~500'

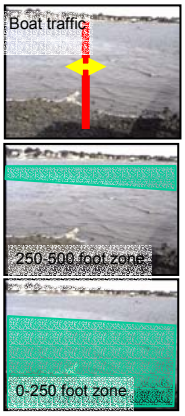

Figure 13 – Camera coverage and view for day time water incursions

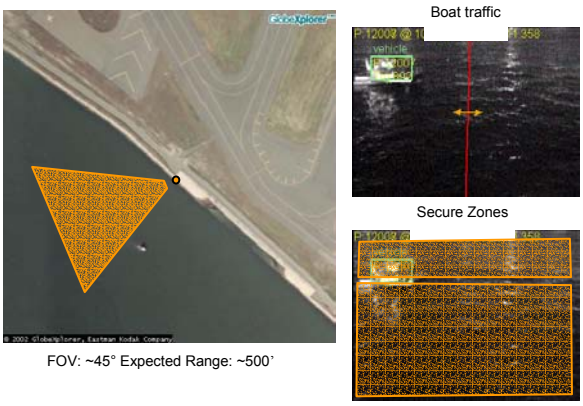

Figure 14 – Camera coverage and view of night time water incursions

Execution

Because there were no installed surveillance cameras to support testing of perimeter incursions or water front incursions, field trials were conducted in stages. Each trial consisted of design, set-up, filming, processing and analysis. The mobile system was used for filming – both null hypothesis and staged scenarios. Film was played back into the AVS sensor (housed at the MA State Police garage at the airport), simulating real-time video input, resulting in equivalent data collection and video capture. Analysis was performed as usual on the collected data, and rule sets were tested on the "live" playback.

Perimeter incursion scenarios were staged by the ObjectVideo crew; the Massachusetts State Police Marine Unit staged water incursions with police boats (1 small and 1 large). Several local boats were also included serendipitously. Logan operations personnel supported filming by identifying danger zones around runways, ensuring that height restrictions were adhered to, and transporting equipment and crew safely onto the airfield. All scenarios were filmed both day and night, with additional filming in some areas to test alternative camera configurations and illumination.

Results

Similar to the refinery pilot, performance of the out-of-the-box system was nearly ideal. Since 500 feet was the recurring monitoring distance, all cumulative detection rates were calculated to 500 feet. At 500 feet, we achieved a 85% detection rate on perimeter incursions with a single "tripwire" detection rule; at ranges less than 500 feet, detection rates reached greater than 97%. By combining rules, detection rates greater than 90% were achieved at 500 feet.

High winds (a residual effect of a hurricane) were initially a factor in lost accuracy. This was overcome by using a sturdier mounting mechanism for the camera housing instead of the lightweight PTZ head and mobile tripod originally deployed. This information was useful in the overall effort to design and specify a final system from scratch including all necessary hardware. Where possible, it is important to use static rigid mounting and wind-tolerant housings (e.g., domes).

Day time perimeter performance alone achieved 99% detection rate with no observed false alarms. Night time performance was limited by less than optimal illumination, but still incurred no false alarms. The effective range of the 500W IR illuminators was 300 feet. For complete night time coverage, illuminators placed every 300 feet are necessary.

Day time boat incursion detection at 500 feet was 97% with an observed false alarm rate of 5 per hour due entirely to the incoming tide's moving edge. Detection of stopping boats was 75% due to inadequate mooring simulation; discounting the bad simulations yields a potential detection rate of 100%.

Night time performance was limited by less than optimal illumination, with ~4 false alarms per hour observed due to reflections from plane lights. Detection of moving boats was 82% due to over-filtering; detection of stopped boats was 100%.

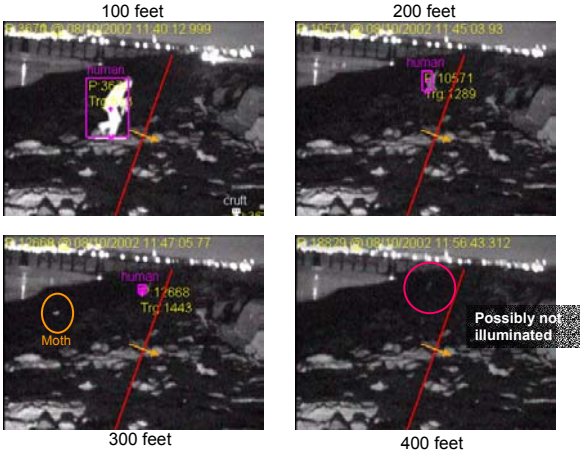

Figure 15 – IR Illuminator Performance

Recommended Surveillance Configuration(s)

The optimal camera array should contain a mix of IDN cameras, IR illuminators, and thermal cameras. If the

cost of thermal cameras is too high, a good solution would consist of only IDN cameras and illuminators; such a system would be susceptible to some bad weather such as heavy snow and fog through which the IR illumination would not adequately penetrate.

- **Option 1**: IDN sensors with appropriate IR illumination
  - Pros: moderately priced; effective 24/7 in conditions from good weather to light snow and rain
  - Cons: not effective in heavy snow, rain, or fog
- Option 2: Thermal sensors
  - Pros: effective 24/7 in all weather
  - Cons: expensive; does not give "intuitive" imagery; higher maintenance cost
- **Option 3**: Combination of IDN and Thermal sensors
  - Pros: 24/7 all weather coverage in sensitive areas; 24/7 "most weather" coverage in other areas; readily optimized for effectiveness vs. cost
  - Cons: additional expense of thermal sensors

Figure 16 – Camera recommendations

Cameras every 500 feet with partially overlapping fields of view
Sensor FOV: ~10°-15°

Figure 17 – Recommended Perimeter Configuration

Dual cameras every 500 feet with partially overlapping fields of view
Sensor FOV: ~45°-60°

Figure 18 – Recommended Water Approach Configuration

Cameras every 500 feet with partially overlapping fields of view
Sensor FOV: ~10°-15°

+

Dual cameras every 500 feet with partially overlapping fields of view
Sensor FOV: ~45°-60°

Figure 19 – Combined Configuration for 100% Coverage

Expected Operational Effectiveness

No significant adjustments were made to the baseline AVS sensor to achieve good performance in this environment. The false alarms and missed detections that occurred can be easily accommodated. We expect to be able to achieve ideal performance with minor adjustments for tidal motion and transient water reflections. Within 500 feet, all detections should achieve greater than 97% with no significant false alarms.

6. Where Do We Go From Here?

What ObjectVideo has achieved so far is only the beginning of what is possible for automated, integrated video surveillance solutions. New technologies will be available over the course of the next several years that will provide far more advanced performance – ObjectVideo will be at the forefront of that development.

The increase in computational power and the integration of IT infrastructure with physical security infrastructure indicates the emergence of several trends in AVS:

- **Enterprise-wide surveillance systems**. The surveillance systems of the future will truly span the entire enterprise and every element within.
- **Fully integrated.** It will integrate seamlessly with all physical security systems (access control, other sensors, and security personnel). It will also integrate with appropriate data bases (such as corporate records, INS, DMV, Law enforcement, etc)
- **Real-time interdiction.** ObjectVideo's system can already provide real-time alerting

mechanisms. Future systems will be able to perform real-time interdiction.

- **Intuitive user interface.** ObjectVideo has already made a start at building a simple prescriptive graphical language to describe threats. Future versions will move from a sensor-centric view of a site to a model-centric view. System tasking can be performed by relating actions to geographical locations rather than camera views. And situational awareness will be provided using a map-based or model-based paradigm.
- **Increased "intelligence".** The system will have increased computer vision functionality to perform automatic object recognition and human ID via video biometrics. In addition, future systems will be able to detect patterns of normal behavior and be proactive about alerting a user to any abnormal events.

7. Conclusions

In this paper, we have outlined two examples indicating that AVS based on computer vision technology is a useful piece of the solution for asset protection, perimeter monitoring and threat detection. The Logan airport example demonstrates that this technology is desirable over other technologies because it is passive, relatively inexpensive, operationally effective, and provides real-time, actionable intelligence.

This technology, however, comes with the caveat that the customer has to become educated about its underlying technology and its applicability. Many proponents of computer vision technology are advocating commercial systems that do not perform adequately in real-world environments – they are subject to poor detection rates and high false alarms rates in realistic, unstructured environments. At ObjectVideo, we strongly recommend that potential customers trial the technology in their own unique environments to determine the utility of this technology and its adaptability to environmental pressures. Our example shows that the ObjectVideo system is, in general, extremely effective as a turnkey system – and in cases with unique environmental phenomena, our system is rapidly adaptable to overcome operational concerns.

8. References

[1] www.cs.cmu.edu/~VSAM

[2] H.H. Nagel, "Formation of an Object Concept by Analysis of Systematic Time Variations in the Optically Perceptible Environment", Computer(14), No. 8, Aug. 1978, pp. 29-39

[3] A.J. Lipton, M. Allmen, N. Haering, W. Severson, T.M. Strat, "Video Segmentation Using Statistical Pixel Modeling", US Patent #20020159634 Pending

[4] I.Haritaoglu, D. Harwood, and L.S. Davis, "W4s: A Real-Time System for Detecting and Tracking People in 2 ½ D", in 5th European Conference on Computer Vision, 1998, Freiburg, Germany: Springer Verlag.

[5] C. Wren, A. Azarbayejani, T. Darrel, and A. Pentland, "Pfinder: Real-time tracking of the human body", in IEEE Transactions on Pattern Analysis and Machine Intelligence, 1997, 19(7): pp. 780-785.

[6] M.A. Isard and A. Blake, "ICondensation: Unifying low-level and high-level tracking in stochastic framework", in Proceedings of the 5th European Conference on Computer Vision, 1998, pp. 893-908.

[7] K. Toyama, J. Krumm, B. Brumitt, and B. Meyers, "Wallflower: Principles and Practice of Background Maintenance", in Proceedings of the International Conference on Computer Vision, 1999, pp. 255-261.

[8] A. Lipton, H. Fujiyoshi, and R. Patel, "Moving Target Detection and Classification from Real-Time Video", in Proceedings of the IEEE Workshop on Applications of Computer Vision, 1998.

[9] M. Onoe, N. Hamano, K. Ohba, "Computer Analysis of Traffic Flow Observed by Subtractive Television", in the Journal of Computer Graphics and Image Processing(2), 1973, pp. 377-392.

[10] R.O. Duda, P.E. Hart, D.G. Stork, "Pattern Classification", 2nd Edition, 2000, published by John Wiley and Sons.