BAE SYSTEMS

Executive Summary

Visual Profiler

Delivering actionable insights from Artificial Intelligence (AI)

GXP™
Geospatial eXploitation Products™

VIDEO inform
FOR A BRIGHT VISION

Table of Contents

Visual Profiler for GXP™

The currency and coverage of satellite and airborne imagery available today from a variety of sensors has become increasingly useful for disaster recovery, facility security, border protection, and intelligence workflows. However, when a quick turnaround is required, large volumes of imagery often overwhelm the human ability to interpret content. To address this challenge, BAE Systems Geospatial eXploitation Products™ (GXP) offers Visual Profiler for the automatic machine detection of objects of interest in imagery.

Integrated with GXP software, including GXP Xplorer® and SOCET GXP®, Visual Profiler seamlessly combines machine detections with human visualization and analysis. By providing examples of imaged objects of interest, such as vessels, vehicles, and infrastructure, a user can essentially train the Visual Profiler tool to detect what they are looking for. Visual Profiler then generates specific algorithms, Profiles, for these objects and autonomously scans images to detect and provide locations and counts of the objects.

As a geospatial data catalog and management system, GXP Xplorer automatically runs Visual Profiler as new imagery become available, and sends out cues when data is ready. Imagery can be further exploited by SOCET GXP, an advanced geospatial intelligence desktop application from GXP. Utilizing photogrammetry, imagery, and other geospatial data for users to visualize, identify, analyze, and extract ground features, SOCET GXP enables rapid situational awareness and critical decision-making.

Collectively leveraging the capabilities of Visual Profiler, GXP Xplorer, and SOCET GXP, human-machine automated workflows can deliver the rapid turnaround times needed for effective defense and intelligence, as well as disaster response.

> There is no need for any intervention by a programmer or computer vision expert in the Visual Profiler workflow

Visual Profiler benefits

- High accuracy performance within different object categories (both detection and identification)
- Requires only a small number of examples - tens to hundreds - to create an object profile with high detection performance
- Wide profiling flexibility empowering the analyst to define unlimited target objects
- Adjustable to almost any aerial/satellite imagery data from diverse sensors (including low-resolution imagery)
- Robustness for different image settings such as lighting and weather conditions
- Detects target objects in vertical imagery - from nadir up to 30-40° of nadir
- Rapid adjustment capabilities for new operational missions
- Investigative tools to explore ranked results
- Easily added to an existing GIS imagery management system
- Overall productivity improvements

Artificial Profiler technology

Visual Profiler is powered by Video Inform's proprietary AI-based technology called Artificial Profiler, which is designed to detect objects of interest in still visual data. The Artificial Profiler is composed of three main components:

- Profile Creator – Enables the end user to create a Profile by giving labeled examples and by teaching the algorithm to determine what to look for
- Profiler – Uses the Profile that was created by the Profile Creator to scan for objects it has been taught to detect
- Analyzer – Analyzes Profile performance by running the created Profile on a test area

For example, users can create a jet fighter Profile by training the system on aerial visual data, using the Profile Creator. Then, the system analyzes its performance on a new area, using the Analyzer. When the Profile is ready, the user can create a job that automatically searches new aerial imagery, seeking the appearance of jet fighters.

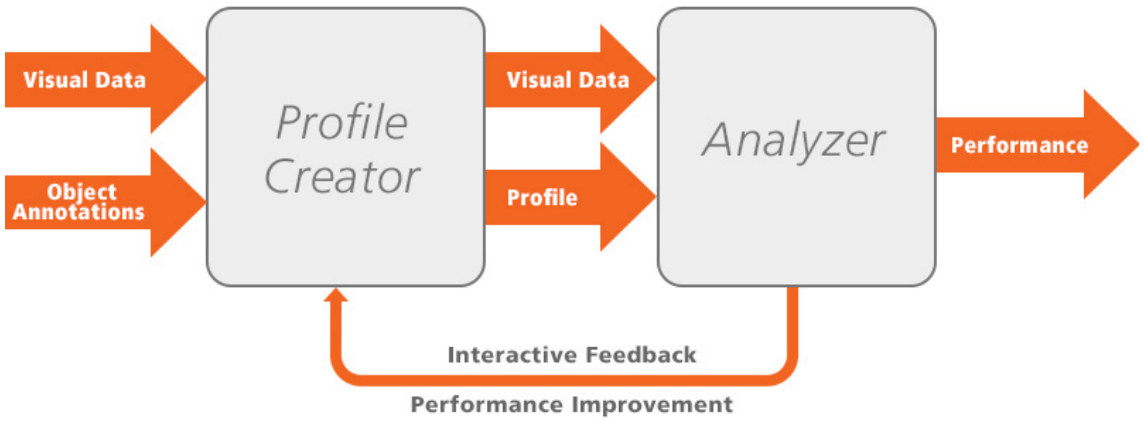

Figure 1

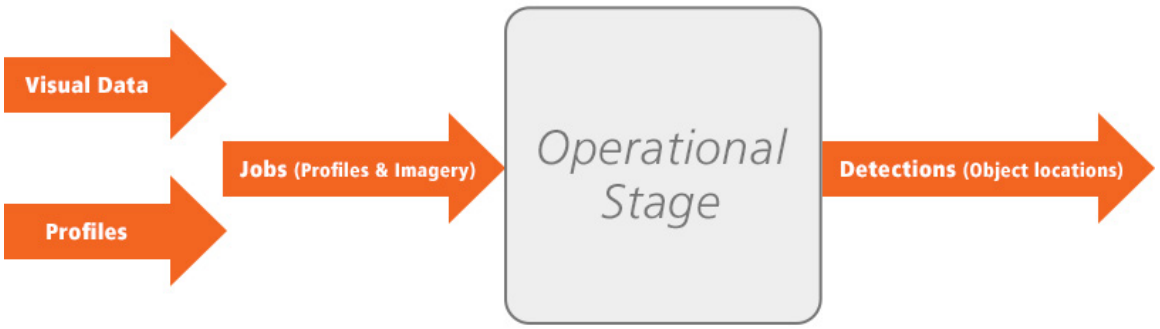

Figure 2

The dynamic architecture of the Artificial Profiler is researched and enhanced constantly by Video Inform

One of the unique characteristics of the system is the iterative nature of the training process. The solution continuously gathers user-supervised examples and learns from them; the more examples supplied, the better the performance.

Viewing and exploring results and detections can be done in two modes: as a standard geographic layer on the imagery in SOCET GXP, or by using the Visual Profiler album of detections, which gives direct access from the detections in the album to the imagery.

The Artificial Profiler algorithm is based on the architecture of an ensemble of state-of-the-art deep learning algorithms with unique machine learning approaches. These ensemble methods create multiple models and combine them to produce improved results and more accurate solutions than a single model on its own. The Profiler algorithm is implemented in a highly efficient manner to achieve rapid search times in scanning new data.

The Profile Creator engine is designed to achieve a high detection rate with a small number of training examples - from tens to several hundred - depending on the relative difficulty (i.e. how the objects differ from their surroundings, the size of the objects, etc.)

In the most recent Enhanced Visual Profiler AI core, additional quantum leaps in detection accuracy and performance have been achieved, and proprietary ResNet-based Convolutional Neural Network (CNN) and context VInets have been introduced. This customized architecture uses balanced computations through the Neural Network (NN) for fast inference while maintaining high accuracy, and is also adaptable to different input image sizes.

Visual Profiler: Empowering the analyst

- Any analyst can build a new Profile quickly and easily for any target object
- Few examples are required to create a Profile with high detection performance
- Analysts can provide feedback to the system, enabling it to continue to train
- Simple and intuitive user interface
- Results can be filtered according to rank

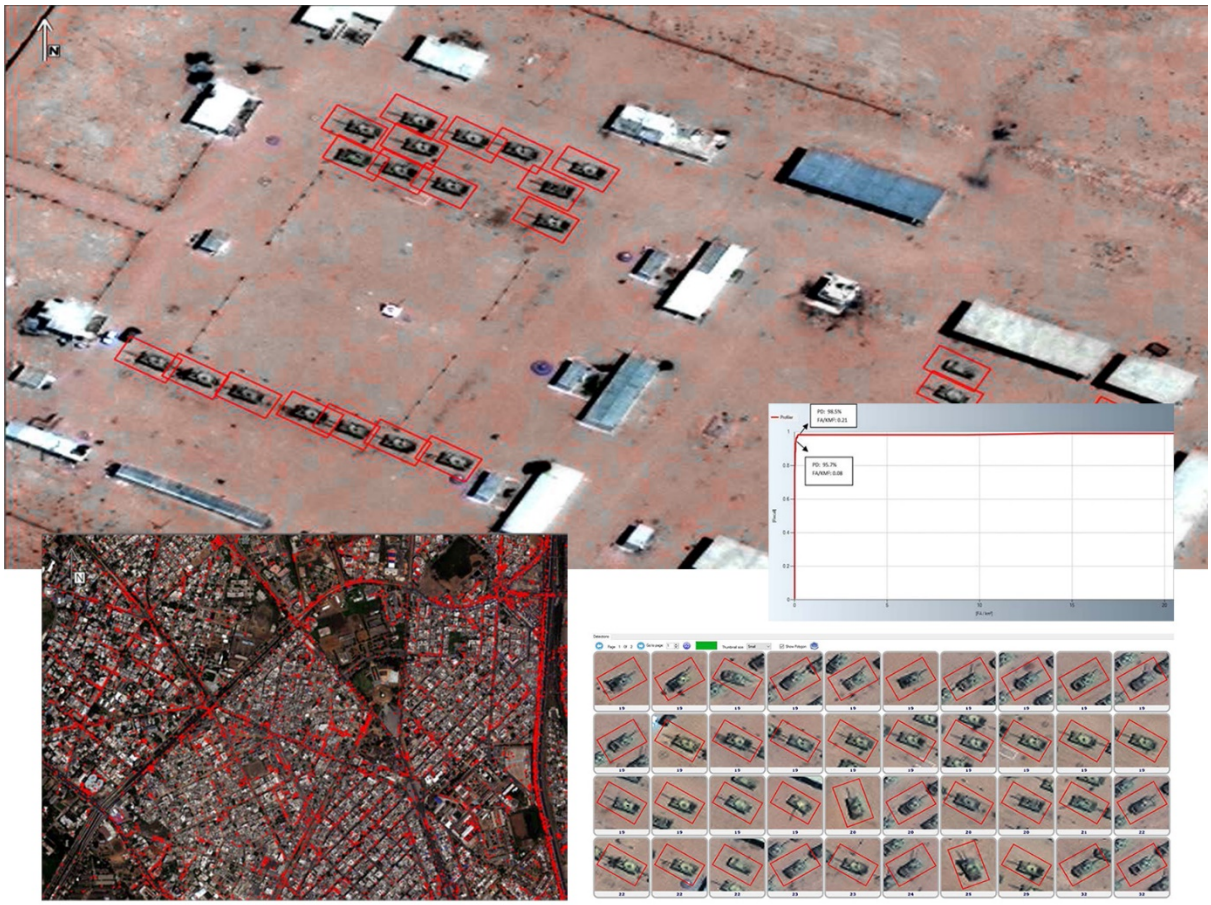

The analyst can create their own Profiles, maintaining confidentiality about what they are looking for

Automatic Object Detection with GXP

The GXP solution for Automatic Object Detection provides for:

a) Initial training of the machine to detect objects of interest
b) Steady state autonomous detection of objects on new images
c) Human machine teaming to use object detections to perform analyses that support courses of action

Automatic object detection is useful for a wide variety of tasks. Two example use cases for illustration purposes include detecting illegal fishing vessels and monitoring activity levels in an airfield of interest. With the GXP solution for object detection, users first train the machine learning application to detect relevant objects; in these examples, vessels and aircraft. Users then import or create geospatial boundaries of regions, for these cases, fishing zones and airfields. GXP Xplorer receives and catalogs imagery that covers the regions and automatically triggers the machine learning application to process incoming images and locate objects within the regions.

Upon completion of processing, the object locations are stored and a completion notification is generated. Once users are alerted to the notification, they can retrieve the object detections to confirm, analyze, and take action. In the illegal fishing scenario, when a suspicious vessel is detected, the action may be for maritime law enforcement to deploy and investigate. For the airfield monitoring scenario, aircraft detections may be aggregated and analyzed to identify patterns such as busiest times of the month.

Initial training of machine

Visual Profiler Creator and SOCET GXP are used to create and train an object profile for the machine to accurately detect objects within imagery. Machine training may only need to be done once if detection success rates are proven sufficient, however, sample objects taken from images in the steady state can be used to re-train and further improve accuracy if required. The machine training flow is depicted as follows (see Figure 3):

1. Images are opened in SOCET GXP (with the Visual Profiler add–on)
2. User creates a new Object Profile
3. User identifies imaged objects to use as positive samples for training the Profile
4. User or software identifies negative samples, which are things in the image that are not the objects of interest
5. Visual Profiler Creator trains the machine and generates a candidate Object Profile for testing
6. Visual Profiler Creator runs the Object Profile for testing
7. User reviews test results for mistakes, which may be false positives (wrong object), and false negatives (missed objects)
8. If too many mistakes are discovered, user identifies additional samples and repeats the training and testing
9. Iterating stops when desired success rate is achieved
10. Object Profile is stored in profile database, ready to apply on new images

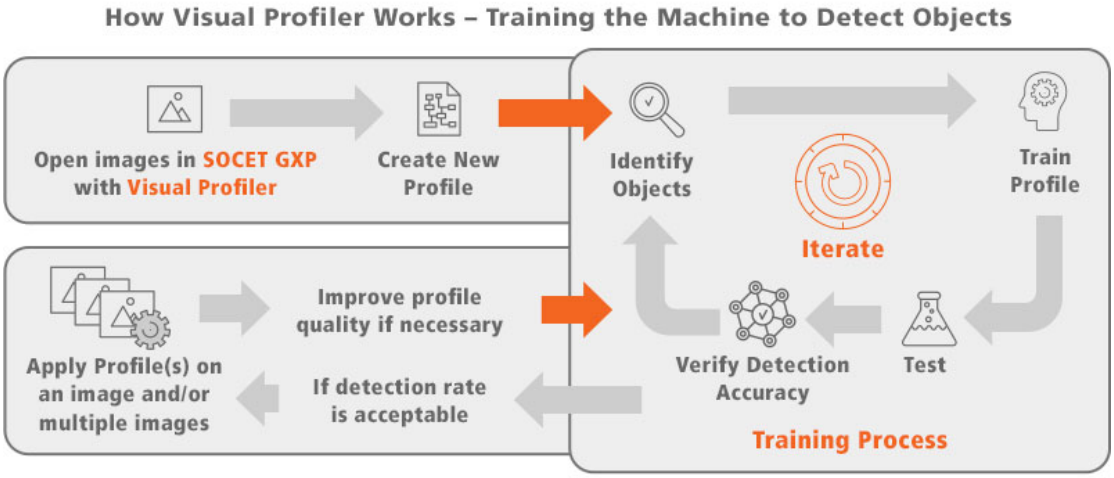

Figure 3

Steady state – Automatically detect objects on new images

With a trained object profile, the steps to autonomously detect objects on new images in a steady state are as follows (see Figure 4):

1. User configures workflow rules in GXP Xplorer that include:
   a) Trigger condition to run the machine object detection process automatically, such as whenever a new image is added to the GXP Xplorer catalog
   b) The types of objects to try to find, specified by selecting one or more object profiles that are trained for specific object types
   c) The geospatial regions in which to look for objects, which may be the entire geographic region covered by a full image, or within specific pre-defined boundaries, such as airfields or border crossing regions
2. When the trigger condition set in the workflow rules is met, GXP Xplorer causes Visual Profiler Server to run the object detection process using criteria configured in the workflow
3. Visual Profiler stores detections and attributes such as latitude, longitude, heading, length, object type, image source name, detection confidence, and image chip in GXP Xplorer
4. GXP Xplorer sends an alert to a user
5. User views the detected objects in the GXP Xplorer map or SOCET GXP

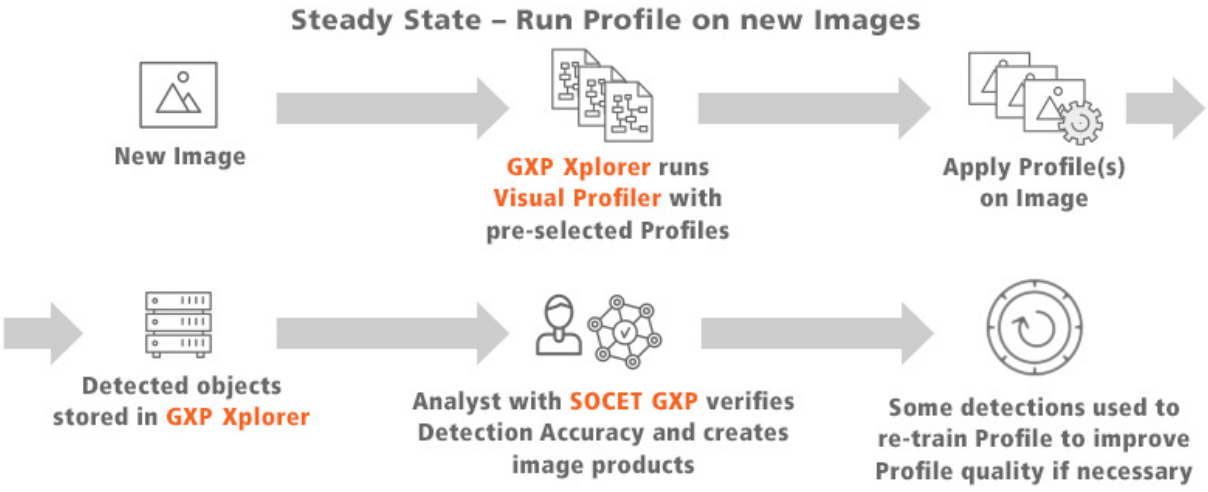

Figure 4

Human machine teaming in the analyst's production workflow

Machine detections are used in streamlined GXP workflows to save analysts the time of looking for objects manually without adding undue work burden to validate the machine results. Machine detections that are auto-generated after imagery ingest are available immediately for import by analysts. The detections are imported into day-to-day analyst production flows so that validation occurs as part of product generation tasks rather than in a separate validation task. Validated detections are then fed back to improve the machine profile success rates for future images. The workflow is shown with the following steps (see Figure 5):

1. GXP Xplorer triggers a Visual Profiler task to run on a new image using pre-trained profiles to detect objects of interest
2. When Visual Profiler completes this task, a work package is added to a queue and analysts are alerted
3. Analyst accepts the task from the queue, loads the image from the task work package into the SOCET GXP Multiport™, and imports and displays the detected object locations as candidate observations using the Observation toolbar in the SOCET GXP Ribbon
4. Analyst validates and fixes any mistakes in the candidate observations, compares to previous work products to analyze changes, and creates a set of validated observations in GXP Xplorer and a new work product
5. SOCET GXP software collects information about the mistakes corrected by the analyst as Candidate Labeled Data and stores this data where it can be used by the human-machine training team to re-train profiles and enhance their success rate for future processing of imagery

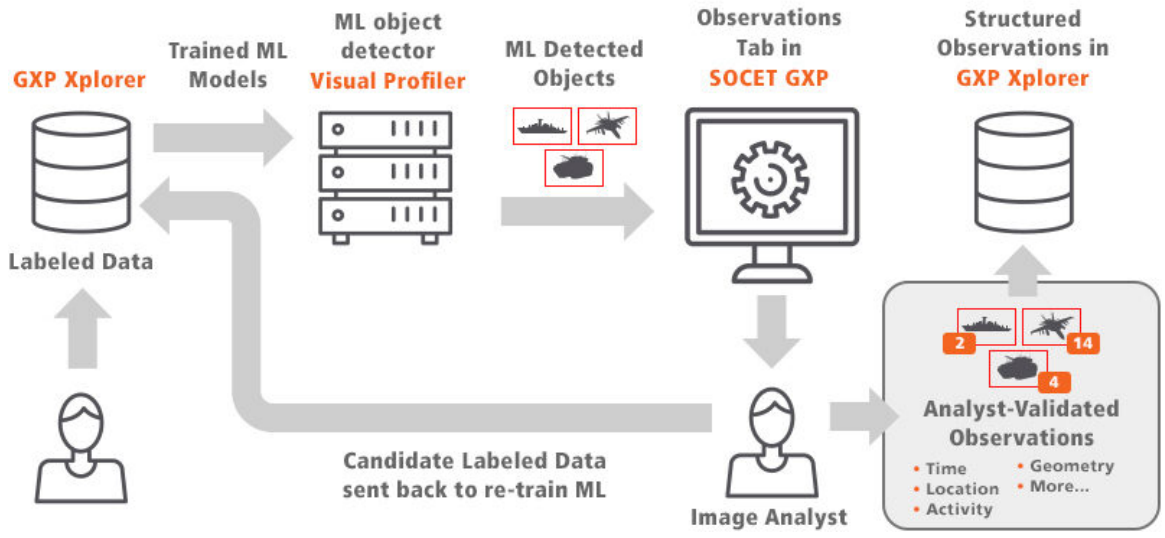

Figure 5

Latest enhancements

Better Detection Rate

Major detection accuracy improvements

Up to 15% better Detection/False Alarm Rate

Faster Profiling

Accelerated profiling performance

Up to 2.5x faster

Faster Training Time

Up to 35% faster

Enterprise Capabilities

Multi-user Annotations

Multi-user Review of Detected Objects

Use cases and processing throughput rates

Visual Profiler has been utilized for numerous use cases over the last five years to detect many different types of objects on a large variety of image resolutions and modalities. The table below summarizes the processing speeds for several of these use cases (additional details are provided in the following pages).

Object type	Average GSD (m)	Image sensor	Processing rates for object detection		
Use Case			**Square km per hour**	**Gigapixels per hour**	**Images per hour**
Helicopters	0.4	Satellite WV-2,WV-3,GeoEye-1	680	4.3	3.5
Oil tanks	0.4	Satellite WV-2,WV-3,GeoEye-1	3,900	24.4	19.8
Tennis courts	0.4	Satellite QB02,WV-2,WV-3,GeoEye-1	365	2.3	1.9
Large vessels	0.4	Satellite WV-2,WV-3,GeoEye-1	1,200	7.5	6.1
Cement mixer trucks	0.09	UAV VisionMap	500	61.7	50.2
Commercial airplanes	0.4	Satellite QB02,WV-2, WV-3, GeoEye-1	1,120	7.0	5.7

Performance is measured on 64 CPU Cores with Visual Profiler version 6.2.28 from July 2020. Throughput is approximately linear with CPU Cores so rates for 128 Cores will be about 2x the rate for 64 Cores.

Regarding GPU use, the Visual Profiler object training processing currently uses CPU+GPU. The object detection processing shown in the table currently uses CPU. Use of GPU for object detection processing is planned for a future Visual Profiler release.

Images per hour rate assumes 1.23 Gigapixel/image, based on the WorldView Pan 35Kx35K pixel "scene" size. WorldView images are also delivered in other sizes.

Machine learning can be tuned to either maximize detection success or run faster. Maximizing detections can find all objects but could produce false alarms. If some mistakes and missing detections are acceptable, say for 80% success, speeds up to 4x faster can be achieved.

Use case 1: Helicopters

- Maxar Imagery – WorldView-2 (WV-2), WorldView-3 (WV-3), GeoEye1
- 0.3-0.5 cm Ground Sample Distance (GSD)
- 1,403 training examples
- Tested area - 62 KM^2
- For a 95.7% detection rate (135/141), we had 0.016 false alarms per KM^2
- For a 97.2% detection rate (137/141), we had 0.096 false alarms per KM^2
- Running time - 680 square kilometers per hour (64 cores CPU cluster)

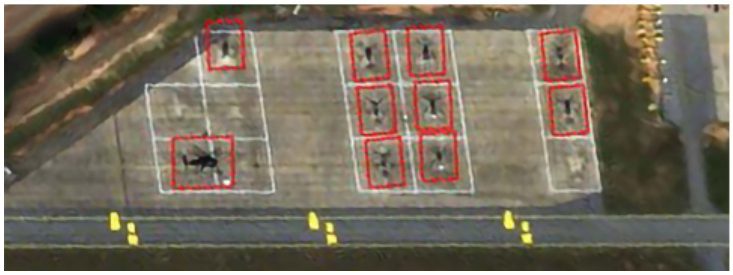

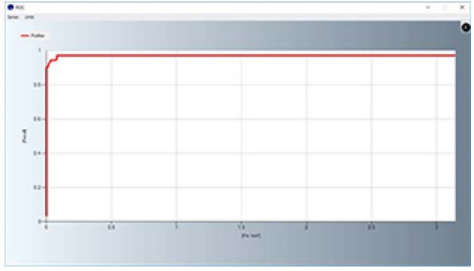

Use case 2: Oil tanks

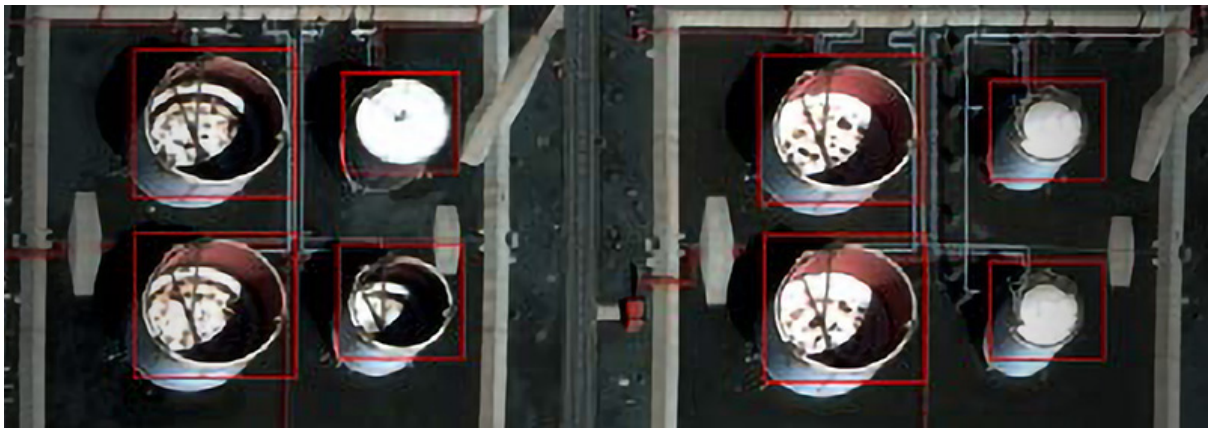

- Maxar Imagery – WV-2, WV-3, GeoEye1
- 0.25-0.6 cm GSD
- Greater than 20m long
- 4,400 training examples
- Tested area - 40 KM^2
- For a 98.1% detection rate (214/218), we had 0 false alarms per KM^2
- Running time - 3,900 square kilometers per hour (64 cores CPU cluster)

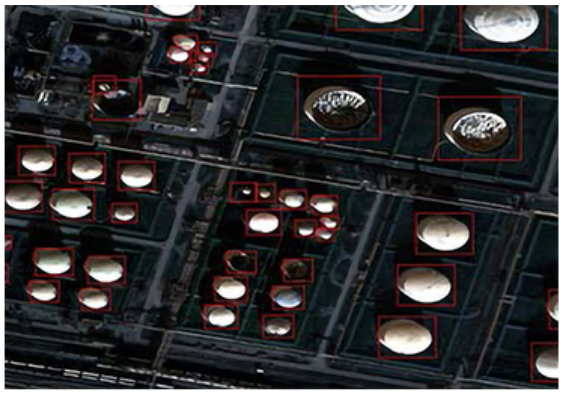

Use case 3: Tennis courts

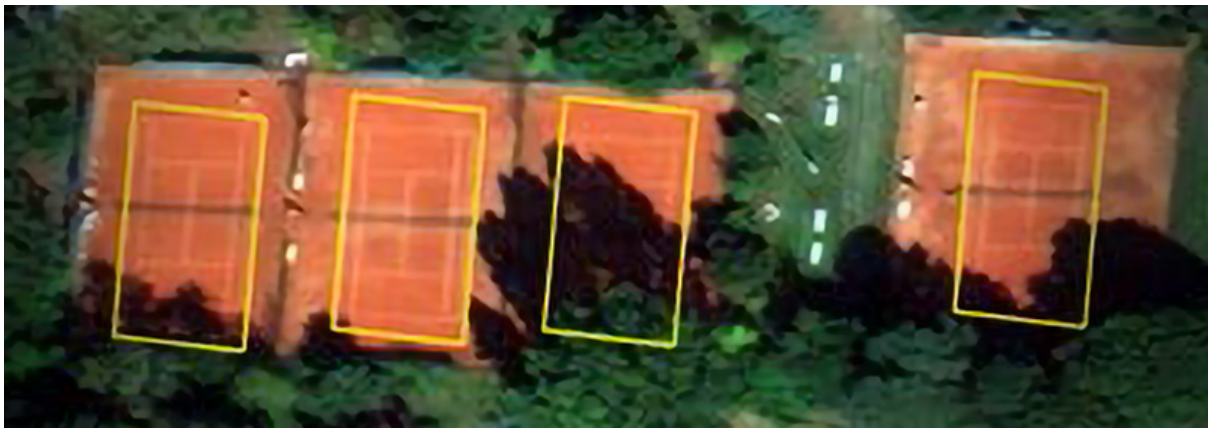

- Maxar Imagery – QuickBird-2 (QB02), WV-2, WV-3, GeoEye1
- 0.3-0.6 cm GSD
- 1,014 training examples
- Tested area - 105 KM^2
- For a 95.8% detection rate (70/73), we had 0.01 false alarms per KM^2
- For a 97.2% detection rate (71/73), we had 0.037 false alarms per KM^2
- Running time - 365 square kilometers per hour (64 cores CPU cluster)

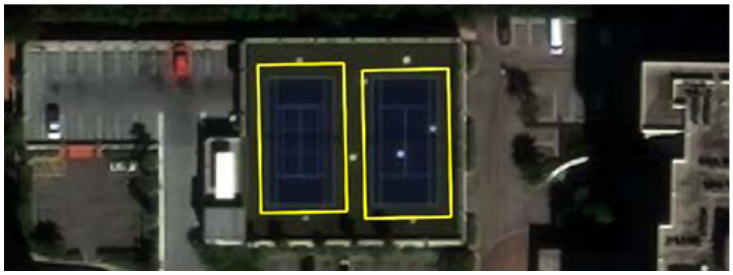

Use case 4: Large vessels

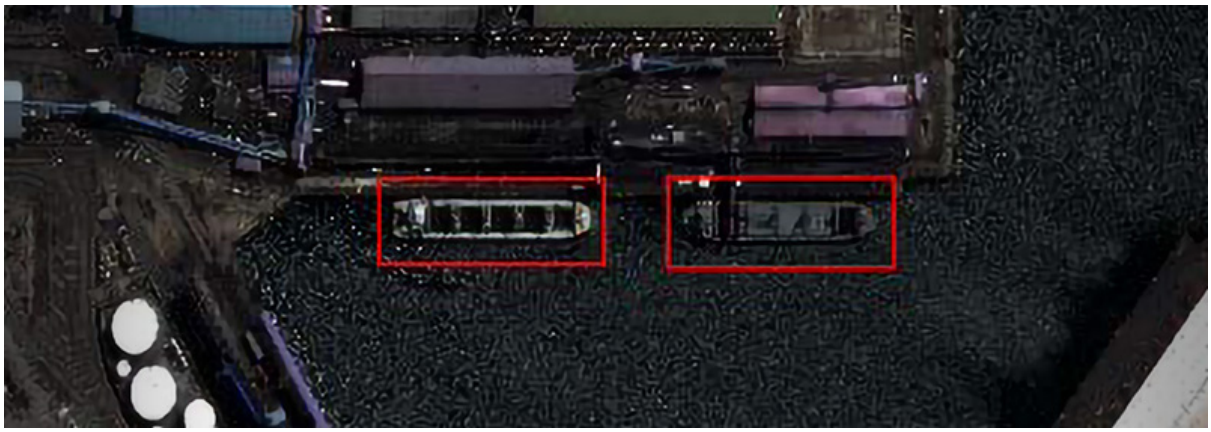

- Maxar Imagery – WV-2, WV-3, GeoEye1
- 0.3-0.6 cm GSD
- 1,014 training examples
- Tested area - 1,217 KM^2
- Open ocean - For a 100% detection rate (121/121), we had 0.004 false alarms per KM^2
- Another test area that includes Port areas - For a 96.6% detection rate (69/72), we had 0.3 false alarms per KM^2
- Running time – 1,200 square kilometers per hour (64 cores CPU cluster)

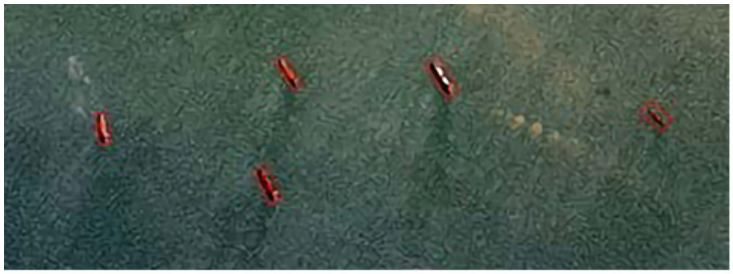

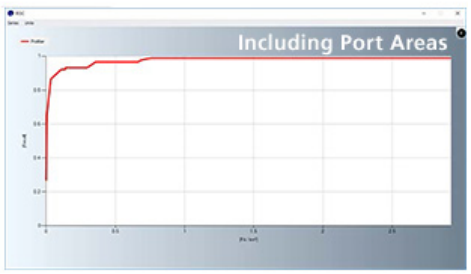

Use case 5: Cement mixer trucks

- VisionMap Unmanned Aerial Vehicle (UAV) Imagery
- Small number of examples for training
- ~9 cm GSD
- 130 training examples
- Urban area - For a 98.5% detection rate (65/66), we had 0.88 false alarms per KM^2
- Running time - 500 square kilometers per hour (64 cores CPU cluster)

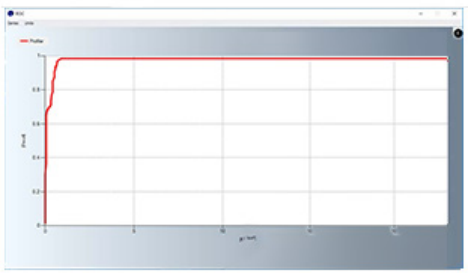

Use case 6: Commercial airplanes

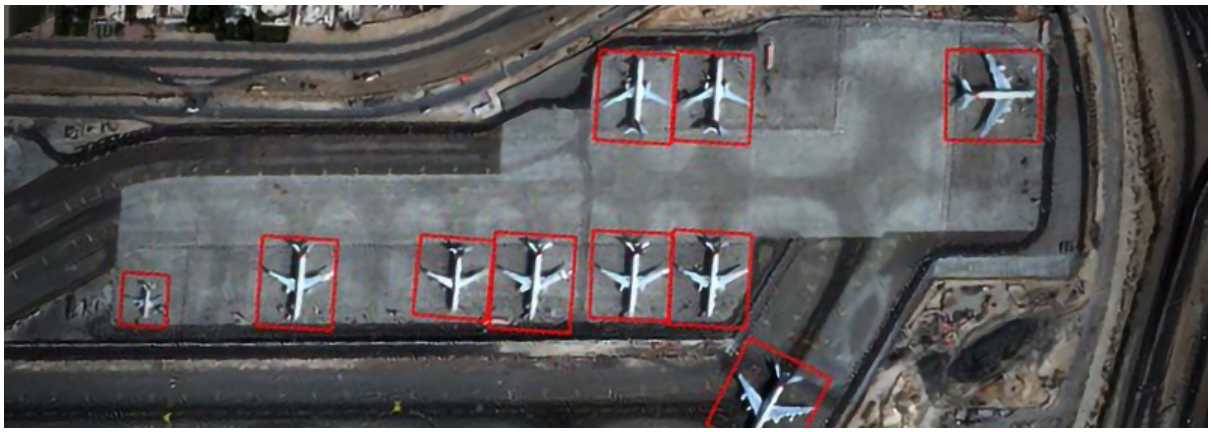

- Maxar Imagery - QB02, WV-2, WV-3, GeoEye1
- 0.3-0.65cm GSD
- 1,339 training examples
- Tested area - 300 KM^2
- For a 96.5% detection rate (340/352), we get 0.04 false alarms per KM^2
- For a 98.0% detection rate (345/352), we get 0.12 false alarms per KM^2
- Running time - 1,120 square kilometers per hour (64 cores CPU cluster)

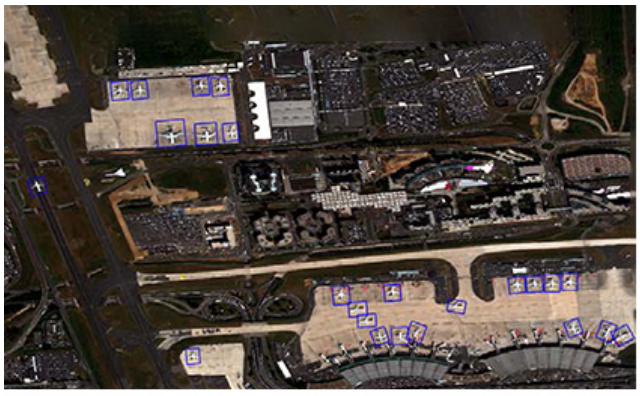

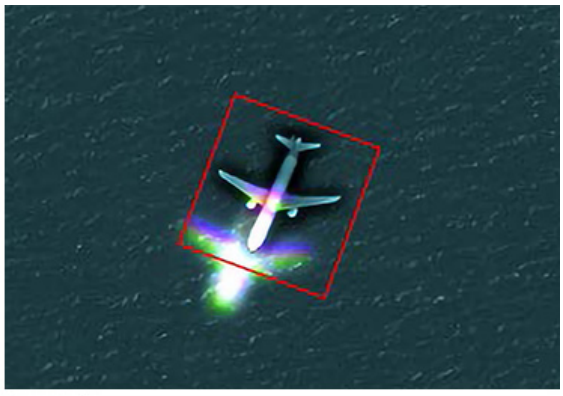

System overview

Visual Profiler requires the use of SOCET GXP software and its exploitation capabilities. A REST Application Programming Interface (API) is available in Visual Profiler for easy integration to any GIS management systems. Visual Profiler runs on Microsoft® Windows® operating systems. Alternatively, it can be supplied as a DOCKER container and run on a Linux cloud-based system.

The Profile Creator is processed on both CPU and GPU, while the Profiler runs on standard computer CPU cores, working in a distributed manner and utilizing every free CPU that it is defined to run on.

High level system diagram

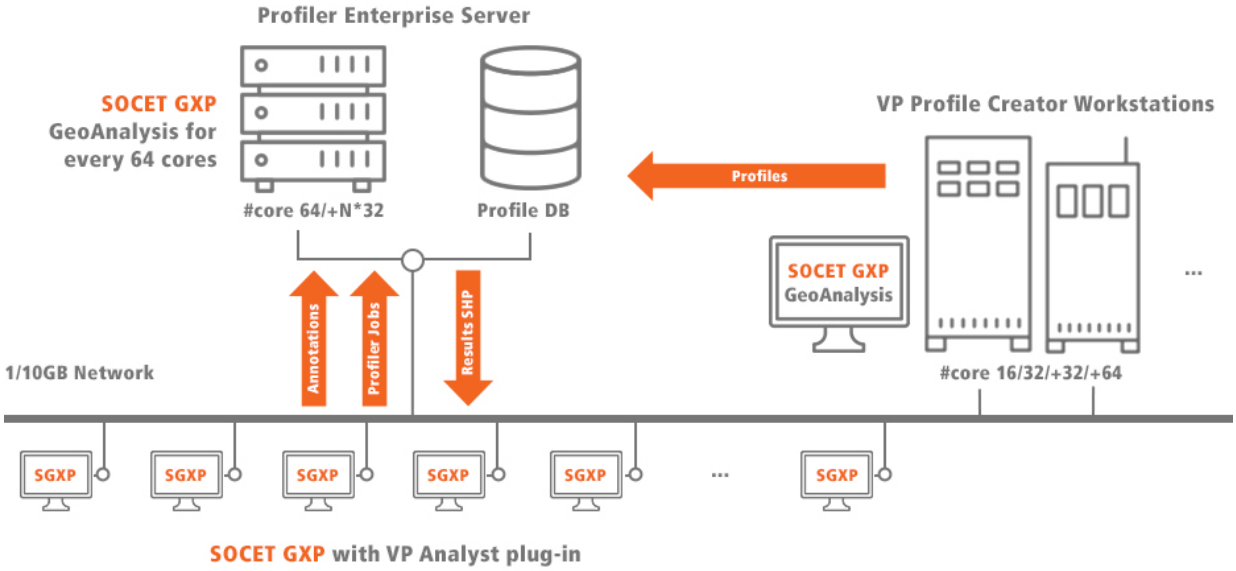

Figure 7

BAE Systems

BAE Systems, Inc. is the U.S. subsidiary of BAE Systems plc, a global defense, security and aerospace company which delivers a full range of products and services for air, land, and naval forces, as well as advanced electronics, security, information technology solutions, and customer support services.

Supporting development of the most advanced geospatial intelligence, BAE Systems Geospatial eXploitation Products (GXP) software enables rapid discovery, exploitation, and dissemination of mission-critical geospatial data. From key military, security, and incident response operations, to a variety of commercial development and research initiatives, GXP provides a comprehensive suite of solutions to inform effective decision-making and ensure a safer world.

Video Inform

Video Inform, a leading provider of large-scale Geospatial Analytics, introduces the new version of the Visual Profiler – a state of the art Deep Learning scalable solution, to enable detection, interpretation, and classification of objects of interest in Airborne and Satellite imagery. Our solutions deliver new levels of actionable intelligence insights.