A Synergistic Exploitation Concept for Wide Area Search

Mark J. Carlotto[1]
General Dynamics, Advanced Information Systems

Abstract

Key to effective image exploitation is letting man and machine do what they each do best. Automated target recognition (ATR) systems rely on model-based approaches that operate in a highly structured predict-extract-match-search (PEMS) loop. Fundamental to wide area search (WAS) is the problem of detecting a large number of potential objects that are so diverse in nature that they cannot be enumerated let alone modeled. In addition there is the problem of unexpected objects (i.e., those that cannot be modeled a priori). A new approach to search based on a detect-organize-compile (DOC) paradigm is proposed and applied to wide area change detection (WACD). It combines an automated image screening algorithm for detecting manmade changes with an interactive web-based tool for ranking, reviewing, and compiling changes of interest. Performance of this approach against exhaustive manual search of the image shows about a 10X increase in throughput (changes detected per unit time) at a fixed level of performance.

1. Introduction

The performance of automated object and change detection techniques depend fundamentally on the statistical separation between significant change and false change in the background. Under typical operating conditions, achieving a high probability of detection typically results in a large number of false alarms. Requiring an image analyst to spatial scan detected changes in order to find those that are significant reduces the benefit of automation. For example, Fig. 1 shows the results of applying an automated change detection algorithm to relatively large Ikonos[2] pan image – about 10,000 pixels on a side. Approximately 50 significant changes are lost in over 2600 detections.

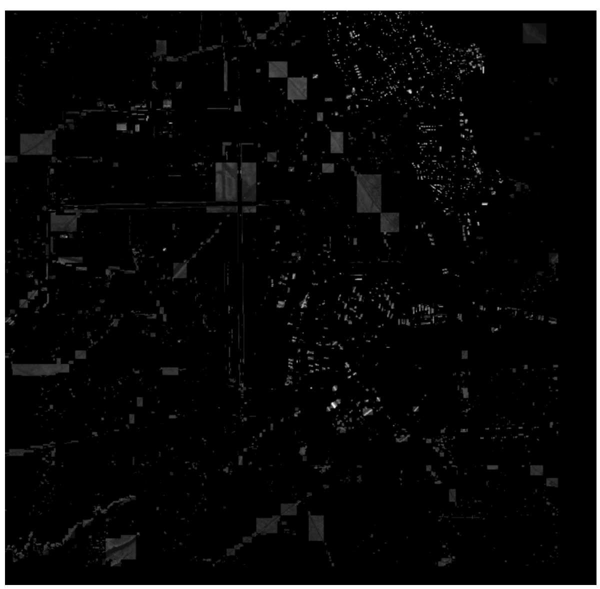

Derived from Ikonos © Space Imaging LLC

Fig. 1 Sample output from wide area search change detector (WACD)

---

[1] Mailing address: 8 Milne Way, Gloucester MA 01930. Email: mark.carlotto@gd-ais.com.
[2] All imagery courtesy Space Imaging (http://spaceimaging.com).

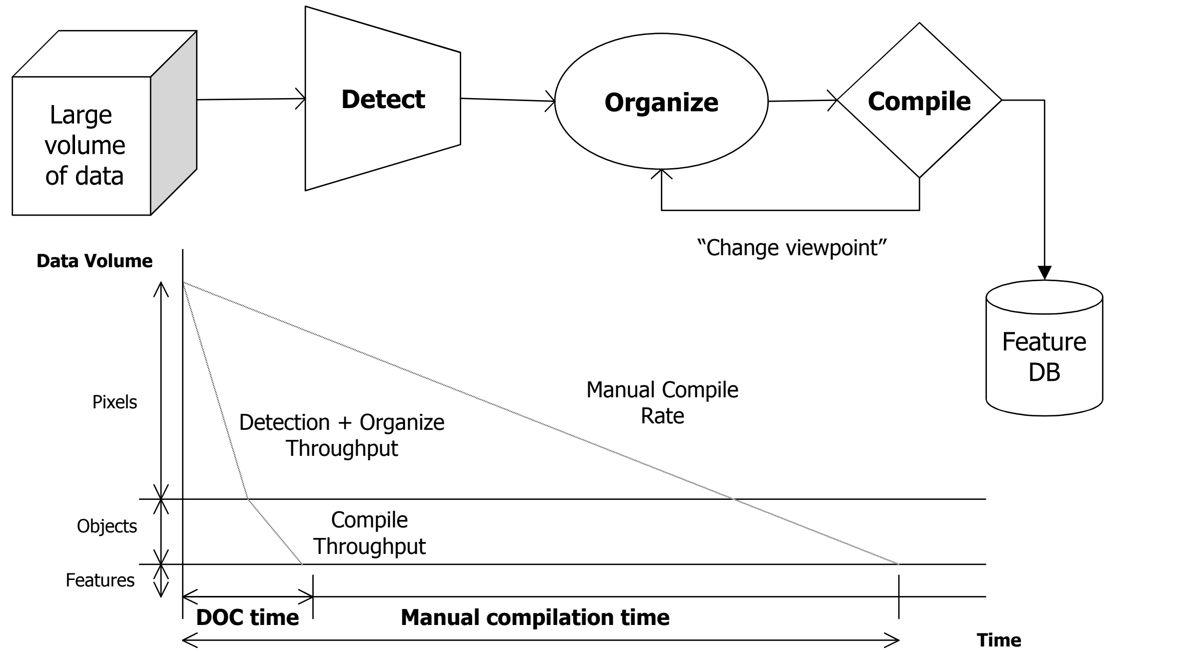

Fig. 2 Detect organize compile (DOC) paradigm

Model-based automatic target recognition (ATR) techniques are designed to find specific kinds of objects, and so are not directly applicable to general search problems. Fig. 2 proposes an alternative paradigm, known as detect-organize-compile (DOC), that addresses the problem of how to efficiently and accurately search a large volume of data for a relatively small number of unknown features (objects/changes) of interest. First, automated techniques detect candidate features to reduce the amount of data that needs to be searched by the analyst. Since the specific features are not known in advance, in order not to miss potential features of interest (i.e., to achieve a high Pd), a large number of FAs are generated. Second, to organize detected features for human interpretation, hypertext/hypermedia representations are created. These representations provide a means of presenting the analyst with alternative "views" of the data so that features of interest (the "needles in a haystack") can be quickly found. These changes are compiled and stored in a feature database. The organize/compile environment is designed to allow the analyst to find and compile one kind of feature (e.g., based on size and/or shape), change the "viewpoint", find and compile other (different) kinds features, etc.

Section 2 describes a wide area change detection algorithm (WACD) for search. It uses a cluster-based approach (Carlotto 2005) to detect manmade changes of any size or shape. Changes having a specific size and shape are found by filtering at the object-level. This occurs within an HTML/XML-based organize/compile environment (Section 3). Examples in Section 4 show how this environment can be used to quickly locate significant change, and how different types of change can be isolated and compiled. Comparison with the state-of-the-art and future trends are discussed in Section 5.

2. Wide Area Change Detection

Object-level change detection (OLCD) techniques have been used to detect compact objects in real time, wide area search systems like the Semi-Automated Image Processor (SAIP)[3]. Historically, the use of image level change detection (ILCD) techniques for WAS has been limited by their higher computational complexity, more stringent registration requirements, and background false alarm rates. WACD overcomes these limitations by: 1) dividing large images up into tiles for processing, 2) using a cluster-based anomaly detection (CBCD) approach for

[3] http://www.fas.org/irp/program/process/saip.htm

differentiating between manmade and natural changes, and 3) combining local registration before change detection with cross-correlation measurements after change detection to reduce false alarms caused by misregistration.

Continued advances in computer hardware developed for the gaming industry[4] should enable real-time implementations of the more compute/storage-intensive ILCD techniques in the near future. To facilitate parallel processing, and to improve registration locally, input images are divided up into tiles. Input images must be ortho-rectified and geo-registered in order to assign geographic coordinates to detections.

The WACD process flow is depicted in Fig. 3. The first step is to align the tile data. Since registration differences are largely translational, a Fast Fourier transform (FFT) approach is used to find the optimal shift between the image pair. Next, CBCD finds possible manmade changes. This involves dividing the reference image into a set of clusters. Automatic clustering techniques are used to determine the best number of clusters to use tile-to-tile. Each cluster represents a homogenous population of pixels in the reference image. Over the set of pixel locations in a reference image cluster, a different set of pixel values are observed in the test image. If there are no manmade changes affecting this cluster, the pixel values in the test image will be related to those in the first; i.e., if there are correlated changes in the background over the cluster, the test image mean and covariance may be different from the reference image, but the pixels in the cluster will remain clustered around the mean. But if the cluster is affected by manmade change, new (uncorrelated) values will be introduced tending to produce clusters with a mixed population of pixels, having a wider spread of values. These manmade change pixels are detected by their higher Mahalanobis distance relative to the background.

Like the RX algorithm (Reed and Yu 1990) the CBCD test statistic (the Mahalanobis distance) has the property of having a constant false alarm rate (CFAR) independent of background type. CFAR thresholding is thus automatic with the threshold value determined as a function of the level of significance of the test (typically 0.99). Detection thresholding is performed in the forward and backward directions. New objects (appearances) are detected as forward changes, while disappearances are detected as backward changes.

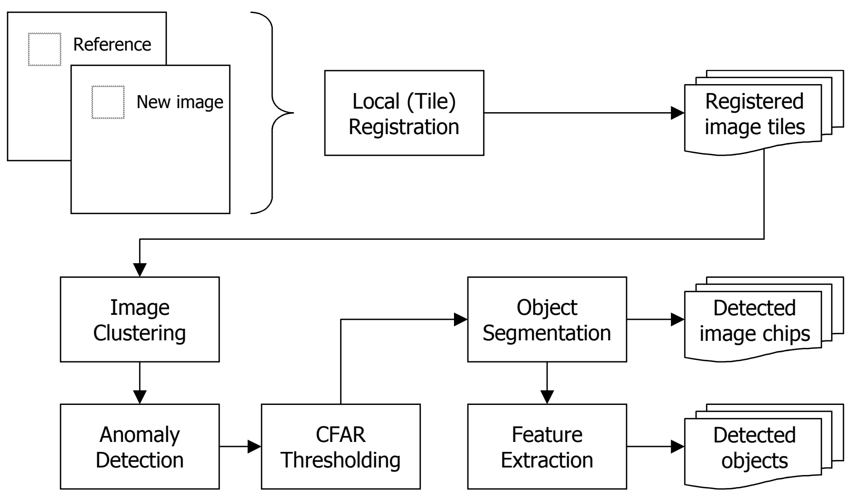

Fig. 3 Wide area change detector (WACD)

Object segmentation labels connected change regions above threshold and converts the thresholded image into an object-level representation. Regions outside a given size range are eliminated and a set of features extracted (geometrical, textural, spectral, and user-defined). The regions are then scored and rank-ordered. Areas having significant terrain relief and vertical objects like buildings pose a problem for ILCD techniques. To mitigate the FAs caused by local misregistration effects, a correlation test is performed over each region. It involves computing the maximum value of the normalized correlation coefficient between the test and reference images within the area of the detection over a specified search area. Changes caused by misregistration have a higher correlation coefficient than object appearances/disappearances and are assigned lower scores.

[4] http://www.itnews.com.au/newsstory.aspx?CIaNID=19234

3. Detection Compilation Environment

WACD output files include registered tile images, detected image chips, and detected region objects. A virtual detection compilation environment (DCE) is created by organizing these files into a set of web pages (Fig. 4). Two kinds of web pages are created: image maps, which highlight the detections in the image, and tables, which list detections in ranked order. Detections can be accessed spatially by clicking on the detection in the image map. Tables allow the user to quickly scan the detections in order of their significance (rank). Image maps and tables are defined at the scene- and tile-level and are cross-referenced. During the compilation process, certain types of changes may become important. At that point the detections can be scored against a more specific set of criteria, re-ranked, and a new set of web pages generated.

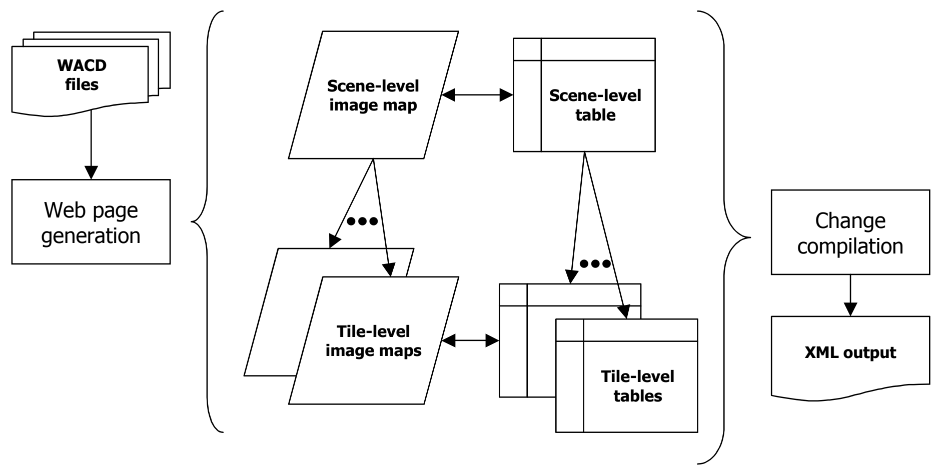

Fig. 4 Web-page generation and compilation

A user reviews and selects objects of interest in the DCE by specifying their class (e.g., construction, demolition, vehicles, aircraft, etc.) in an HTML list structure. An optional text field is also provided to describe the object. The information is submitted via an HTML form (client side) and processed on the server side using a CGI script. Class, description, feature information (including location), links to image chips, and other data are encoded as XML and stored in a file (Fig. 5). Storing results in XML facilitates their use in other applications.


```
- <Changes Creator="WACD" Dtg="June03_Feb04">
- <Detection>
  <Class>aob</Class>
  <Description />
  <GIFnew>14_17n.gif</GIFnew>
  <GIFref>14_17r.gif</GIFref>
  <Latitude>63.993217</Latitude>
  <Length>70.000000</Length>
  <Longitude>-22.624449</Longitude>
  <PixelXbar>219.435760</PixelXbar>
  <PixelXmax>252.000000</PixelXmax>
  <PixelXmin>182.000000</PixelXmin>
  <PixelYbar>989.756958</PixelYbar>
  <PixelYmax>996.000000</PixelYmax>
  <PixelYmin>986.000000</PixelYmin>
  <Pose>88.360367</Pose>
  <Rank>1.000000</Rank>
  <Score>873.258423</Score>
  <Size>358.000000</Size>
  <Width>10.000000</Width>
  </Detection>
  </Changes>
```


Fig. 5 XML representation of a WACD detection

4. Ikonos Pan Change Detection Example

To illustrate our concept a change detection experiment was performed on a registered pair of Ikonos[5] Pan images over Iceland. The images were divided into 1024x1024 pixel tiles for processing. Changes less than 100 sq. meters were eliminated. Fig. 6 shows scene-level image maps for the new image (with and without no-change areas masked out). Image maps for the reference image are accessed by clicking on the links at the top of the page.

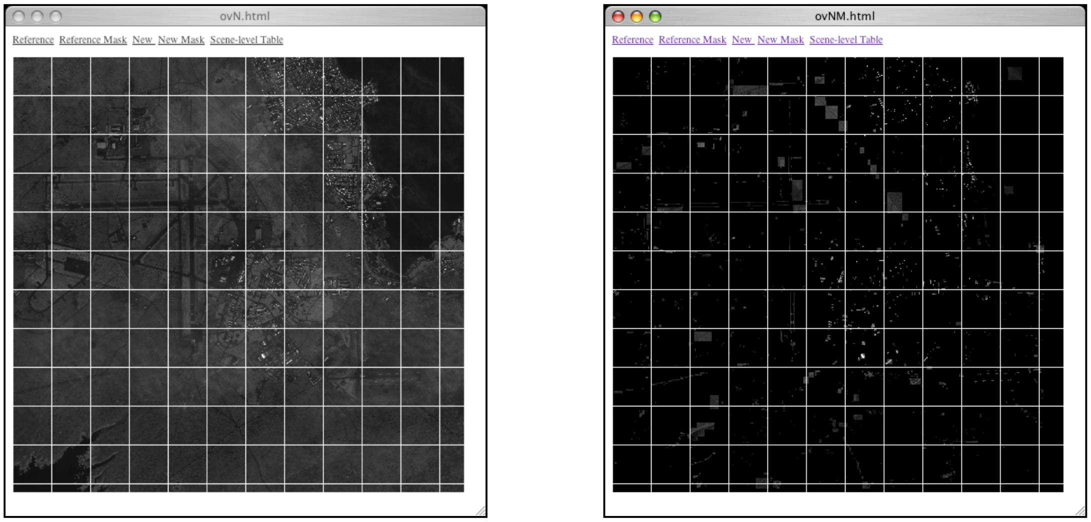

Derived from Ikonos © Space Imaging LLC   Derived from Ikonos © Space Imaging LLC

Fig. 6 Scene-level image maps for the new image

Clicking on a tile (delimited by the overlaid grid), changes the page to a tile-level image map (Fig. 7). Clicking on the toggle link goes back and forth between the new and reference image. One large change due to the appearance of a building is evident at the bottom. Clicking on that area in the image map takes us to a table of detections (Fig. 8a).

Derived from Ikonos © Space Imaging LLC   Derived from Ikonos © Space Imaging LLC

Fig. 7 Tile-level image maps for the reference (left) and new (right) images

---

[5] http://www.spaceimaging.com

In this table, all of the detections in the tile are rank-ordered. The selected detection is the 4-th out of 32 detections in that tile. The top ranked detections are other (smaller but brighter) structures that have appeared. As we move down the list, the significance of changes decreases, with those at the bottom of the list mostly false alarms caused by lighting differences and layover changes. Fig. 7b is the scene-level table of detections, which shows the top detection in each tile. To see all of the ranked detections in a tile, one clicks on the link which leads to the tile-level table. Clicking on the image chips in the tables flickers between the new and reference images. Key features are also listed in the tables with links back to the image maps.

Tables provide a fast way of accessing the top-ranked detections in the image without having to scan the entire image. To quantify the speed at which one can compile significant change different rows of tiles (assumed to have about the same number of changes on average) were processed three ways. First, significant changes were found by flickering the tile-level image maps (manual). Links at the top of the page allow a user to quickly navigate left/right (and top/bottom). Second, changes were found by flickering the tile-level image maps with no-change areas de-emphasized (e.g., Fig. 7). Third, changes were found by scrolling through the scene- and tile- level tables from top to bottom.

Derived from Ikonos © Space Imaging LLC   Derived from Ikonos © Space Imaging LLC

Fig. 8 Tile-level (left) and scene-level (right) tables of ranked detections

Results for this image show that, on average, compiling changes from masked-out registered tiles is about 3x faster than from just registered tiles (i.e., no visual cueing). An additional 3-4x increase in speed over visual cueing is achieved by using the tables. Thus, it is our claim that about an order of magnitude increase in compilation speed can be achieved using DCE over manually scanning for change in registered images. Fig. 9a shows all of the detections generated by WACD; compiled changes are plotted in Fig. 9b. WACD processed the full image (11789 x 11397 pixels) in about 3 hours and 25 minutes on a Macintosh G4. If tiles are processed in parallel using the same hardware technology, the run time would be about 1 minute and forty seconds. The organize and compilation time was about 17 minutes (without interruptions or coffee breaks).

5. Discussion

WACD detects potential changes of interest differing sizes and shapes at the pixel level so that they can be filtered to find specific features of interest at the object level. This is to be constrasted with the use of template matching, spatial anomaly detectors, and ATR techniques, which do not provide the required flexibility for search (i.e., looking for different kinds of objects without having to re-run up-front pixel-based anomaly detectors).

WACD detects possible objects of interest via change detection. Alternatively a single image (spectral) anomaly detector (e.g., based on CBAD) could be used instead to find possible objects of interest. The DCE is a prototype implementation of the organize and compile processes in DOC. It provides both HTML and XML representations, which makes it ideal for interfacing with other applications. The feedback loop back from compile in the DOC process flow represents the ability to dynamically re-organize the detections as one searches for different kinds of objects. This is a new feature that does not exist in current systems. Future work will explore ways of further automating the compilation process.

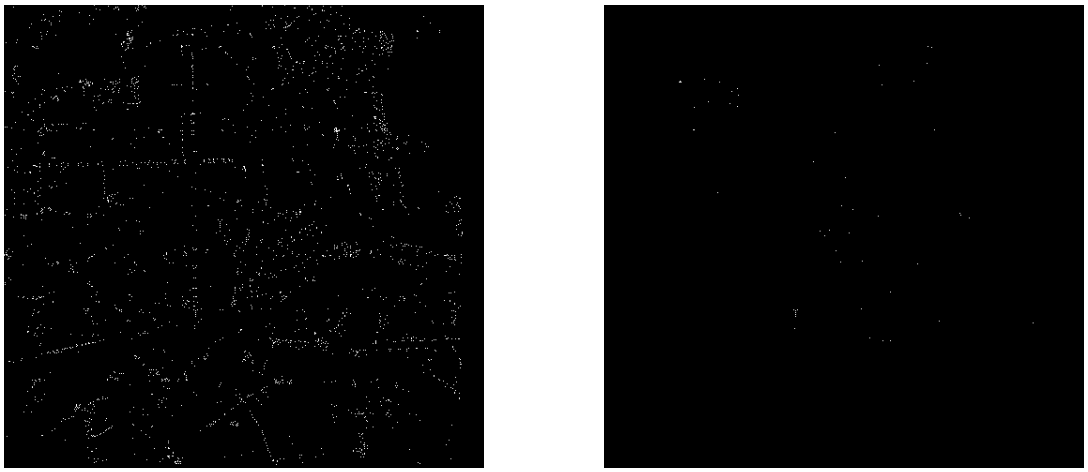

Fig. 9 Raw detections (left) and compiled changes (right).

Ackowledgement – Ikonos imagery provided by Space Imaging (http://spaceimaging.com)

References

M. J. Carlotto, "A Cluster-based Approach for Detecting Manmade Objects and Changes in Imagery," *IEEE Trans. Geoscience and Remote Sensing*, Vol. 43, No. 2, Feb. 2005.

Reed and X. Yu, "Adaptive multi-band CFAR detection of an optical pattern with unknown spectral distribution," *IEEE Trans. Acoustics, Speech, and Signal Processing*, Vol. 38, pp 293-305, March 1990.