Stereo Image Matching

Breakthroughs in Automated Change Detection

*By Bingcai Zhang and Mike Walter, BAE Systems (*www.baesystems.com/gxp*).*

In a good stereo pair, humans readily fuse the two images and perceive a 3-D scene. The relief may be exaggerated, but our brains are comfortable with the presentation. Similarly, stereo correlation algorithms used for automatic terrain extraction operate nicely on good stereo pairs. But localized differences between stereo image pairs can cause headaches for humans and correlation software alike.

For human observers, these differences confound our natural stereoscopic vision. It is immediately apparent that something isn't right. The differences also confuse image-matching terrain extraction software. The results include spikes, wells and other elevation anomalies that previously required manual editing to correct.

BAE Systems recently developed an algorithm to automatically detect and remove false matches caused by moving vehicles as part of Next-Generation Automatic Terrain Extraction (NGATE) enhancements for the company's SOCET SET and SOCET GXP photogrammetry software. The algorithm also provides an automated way to detect change between image pairs, including change caused by moving vehicles. Importantly, the method works well with either panchromatic or multispectral images. The two accompanying case studies illustrate the use of the NGATE stereo image matcher in moving vehicle change detection, and removal of digital terrain model (DTM) defects caused by change or motion.

Continued on page 18.

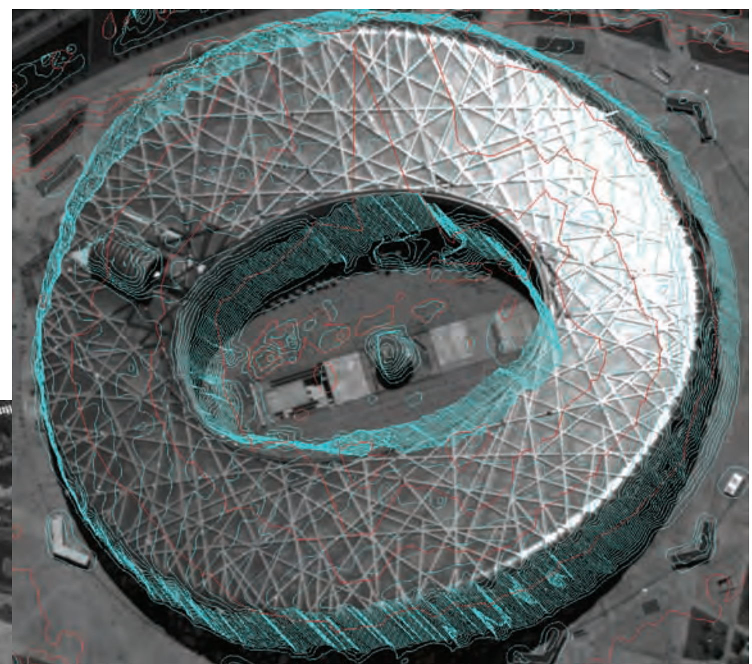

BAE Systems

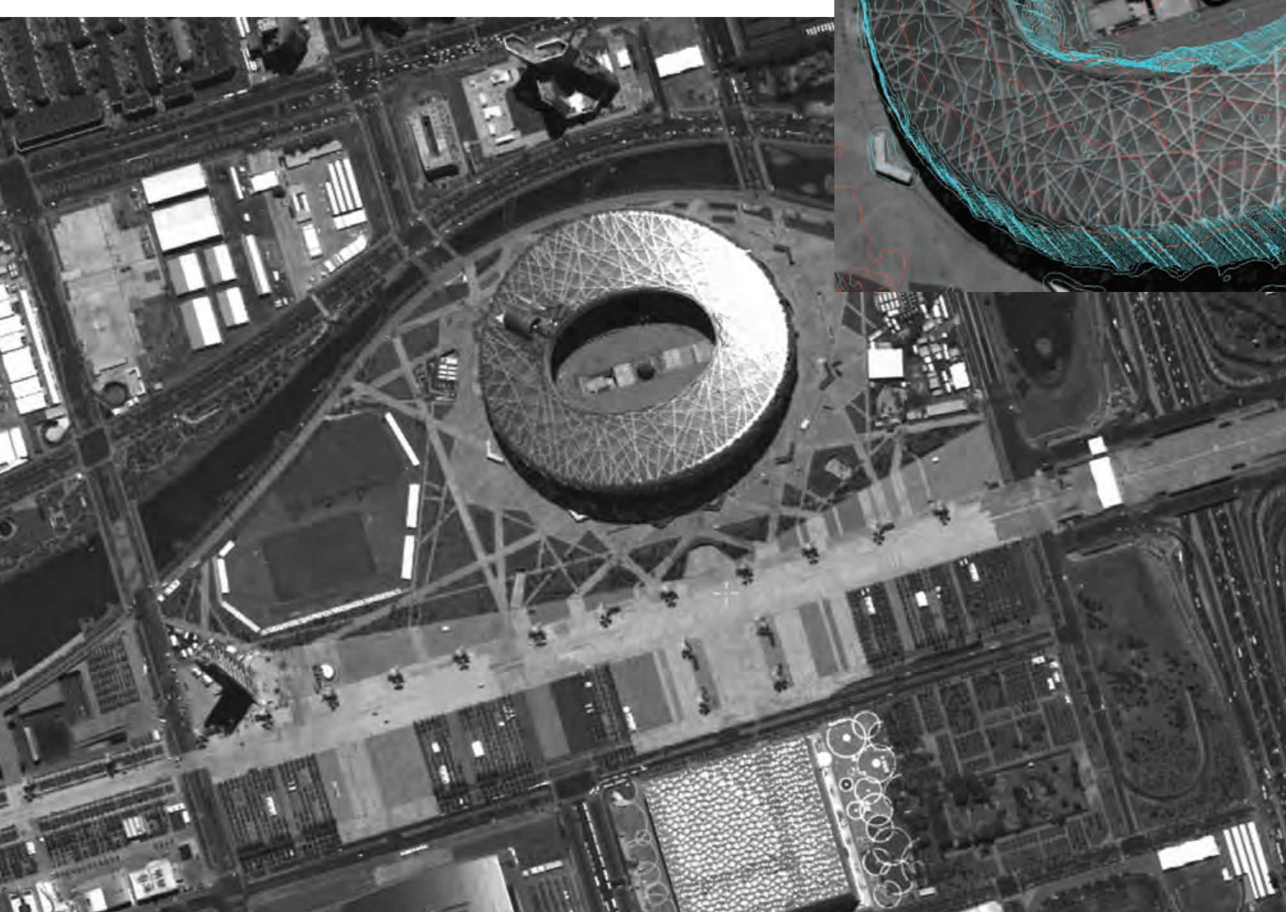

DigitalGlobe

An NGATE analysis of a WorldView-1 satellite image of Beijing's Bird's Nest Stadium shows a DSM as cyan contours and a DEM as red contours (inset). Both are two-meter intervals. In the middle of the stadium, the DSM and DEM have blunders due to moving objects, which makes it difficult for a human operator to place an extraction cursor on the ground. In the upper left corner, the water body also causes DSM and DEM blunders. BAE Systems' forthcoming automatic water body extraction will address this problem.

Change Detection for Moving Vehicles

The first case study is in an urban area. Image resolution or ground sample distance is 0.144 meter. Image pixel depth is 16 bits. Forward overlap is 90 percent. Images have one panchromatic band and dimensions of 9,420 by 14,432 pixels. These are high-quality, high-resolution and accurately registered stereo images. In the test area the stereo image matching algorithm detected 20 vehicles out of 22 visible in the scene. In one case, a similar vehicle moved to the same location previously occupied by another vehicle, and NGATE processed this as the same object. In the second case, a vehicle is somewhat indistinct from its surroundings. The two failed cases might be solved if special logic, such as velocity gating, was added to NGATE in the future.

Figure 1 shows the effect of moving vehicles on terrain extraction algorithms. The uppermost image pair shows the displacement of vehicles in the interval between the capture of two images. The middle image pair shows raw elevation contours with blunders. The lower image pair shows the removal of elevation blunders by improved NGATE processing.

Figure 2 is the same imagery, but with posts instead of contours. Figure 2 shows the bad posts detected and removed by the improved NGATE processing. In the upper image pair, moving cars have caused three clusters of posts where false matches have occurred. The three false matches turn into elevation blunders, which were shown as contours in Figure 1. Figures 1 and 2 demonstrate that the algorithm can detect and remove elevation blunders caused by moving objects.

Moving vehicles may also cause no matches as shown in Figures 3, 4 and 5. In Figure 3 there are 11 areas on a highway that NGATE can't match because of vehicle movement. Red polygons with blue dots inside indicate the 11 areas. Each blue dot represents a group of no-match pixels. All of the no-match areas are due to moving vehicles, and NGATE successfully detected all 11.

Continued on page 18.

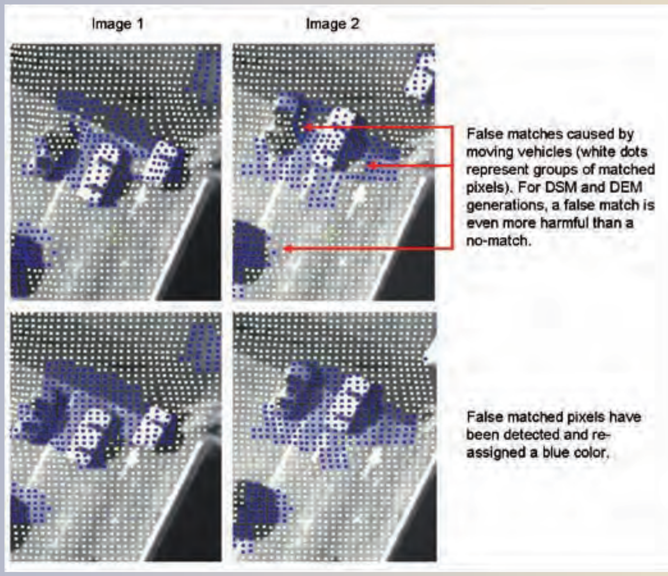

Figure 2. Three groups of false matched pixels have been detected. Moving vehicles cause either false match or no match as indicated by blue dots. In most cases, elevations in no-match areas can be interpolated by using adjacent matched pixels as shown in Figure 1.

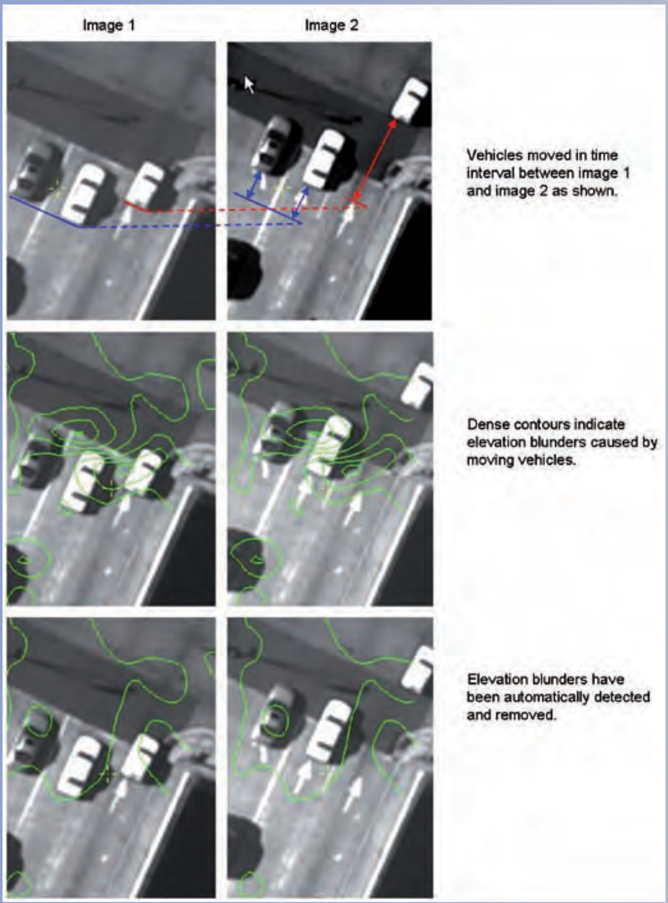

Figure 1. Moving vehicles cause three false matches on the middle images. The three false matches result in elevation blunders. Note how in the lower images, the algorithm has detected and removed the blunders.

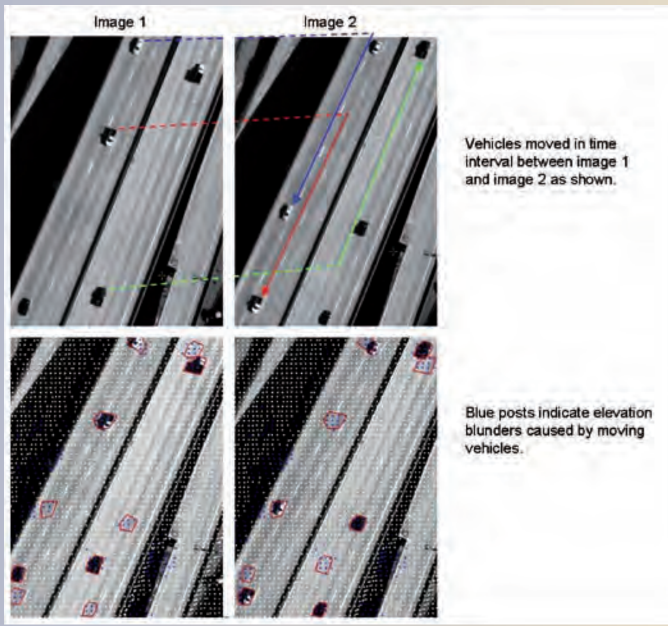

Figure 3. Moving vehicles result in no-matches. Each blue dot represents a group of no-match pixels. Stereo image matching is based on pixel similarities. Each moving vehicle results in two no-match areas (red polygons with blue dots inside).

Continued from page 17.

In Figure 4, there are another 10 areas of no matches on the highway, indicating the presence of vehicles. NGATE detected 10 of 11 vehicles present in the two frames. The detected vehicles in the scene exhibit some contrasting details, including sun glint, contrasting windows and other facets in shadow. NGATE may ignore objects containing few pixels, or with limited contrast. Note that light poles and low-contrast grassy areas also exhibit poor correlation results.

In Figure 5, NGATE detects 11 areas of no matches caused by moving vehicles. In one case a vehicle moved to the same location occupied previously by a similar vehicle resulting in good correlation but for the wrong reason.

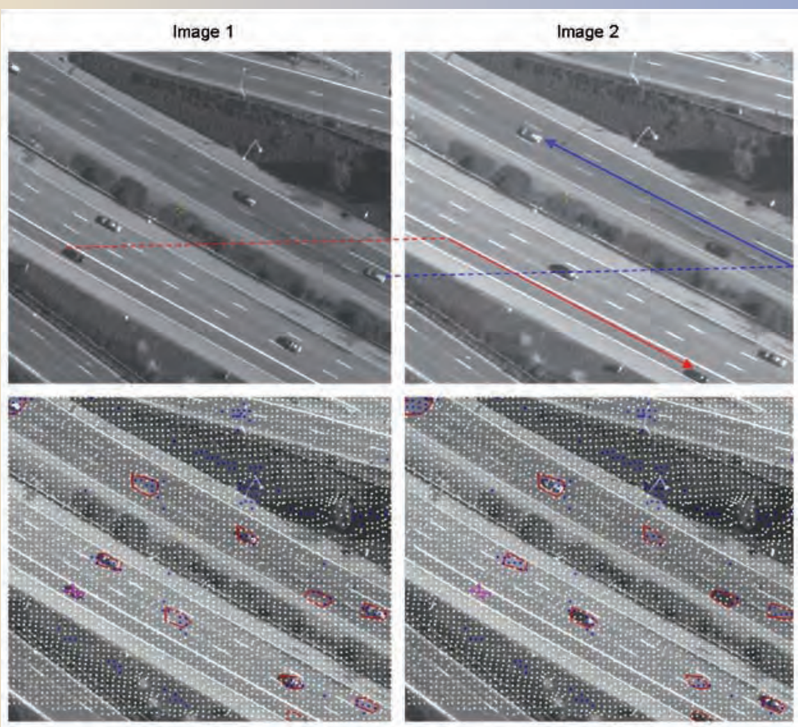

Figure 4. In this stereo pair, vehicles have moved too much for NGATE to correlate, so it reports posts with poor FOMs in the unmatched regions. Ten of these regions are shown as red polygons with blue dots inside. The blue dots are groups of pixels that NGATE can't match.

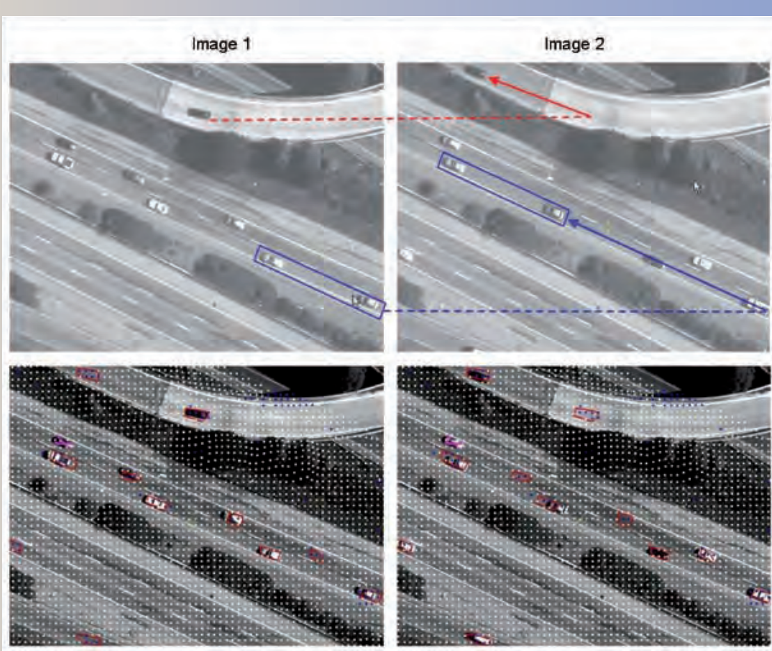

Figure 5. Moving vehicles prevented NGATE from correlating pixels in 11 locations temporarily occupied by moving cars in the two frames. There was one false match in the upper left, marked with a red X. The reason for the false match is that a similar vehicle moved to the same location. To NGATE, the vehicle didn't move. NGATE detected 11 of 12 moving vehicles, eliminated blunders and produced an accurate DEM.

Continued from page 16.

Change Detection Considerations

Automatic change detection and moving vehicle detection are longstanding priorities in remote sensing. Various algorithms use image radiometric properties to detect changes between two or more images. The more successful algorithms tend to operate on multispectral imagery, because the signatures are more distinct. These algorithms usually operate on one image at a time, producing an intermediate product like a terrain categorization (TERCAT) file. Two TERCAT files then are compared, often by subtracting pixel values in corresponding positions of the images.

The factors that make stereo change detection effective can have deleterious effects on automatic terrain extraction algorithms. This fact can be exploited for change detection.

To ensure reasonable alignment of the compared images, the TERCAT files should be orthorectified before comparison. Thus, the entire process requires several steps, including triangulation, orthorectification, TERCAT and finally comparison. The method is only robust for multispectral imagery, which means that a wealth of high-resolution panchromatic imagery can't be treated this way.

Typically, analysts use a side-by-side method to compare images, but many analysts prefer stereo viewing. In stereo, change can be glaringly evident. As anyone knows who has looked at ad hoc stereo with significant time differences between images, change areas don't fuse stereoscopically. A vehicle parking lot, for example, filled in one image and empty in the other, literally pokes you in the eye, even when viewed at reduced resolution. Construction (or destruction) is another good example, but stereo viewing also detects more subtle changes. A single car out of place from one frame to the next will immediately attract an observer's attention. You can search two images at the same time, and, because change is apparent even at reduced resolution, you can gener-

Removing Digital Terrain Model Defects

The accompanying commercial satellite images illustrate some example uses of NGATE to identify changes at military installations in Iraq during the early phase of the Second Gulf War. In Figure 1, image 1 is from Jan. 22, 2002, and image 2 is from Feb. 5, 2003.

The red dots in Figure 1 indicate posts with poor FOM values, while the green dots are good FOM values. The image GSD is about 1 meter. In addition to the obvious absence of vehicles in image 2, a new dirt road, an earthen wall, some new revetments and other soil disturbances are seen.

Close inspection of the red dot areas in Figures 2 and 3 show several changes. The most interesting are the addition and removal of revetments used to protect military equipment. There are also changes in the soil caused by vehicle traffic, and there are some new earthen walls in the second image.

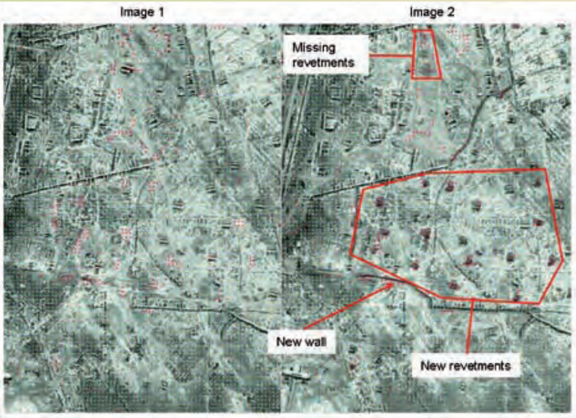

Figure 2. Changes in revetments and defensive walls are clearly visible.

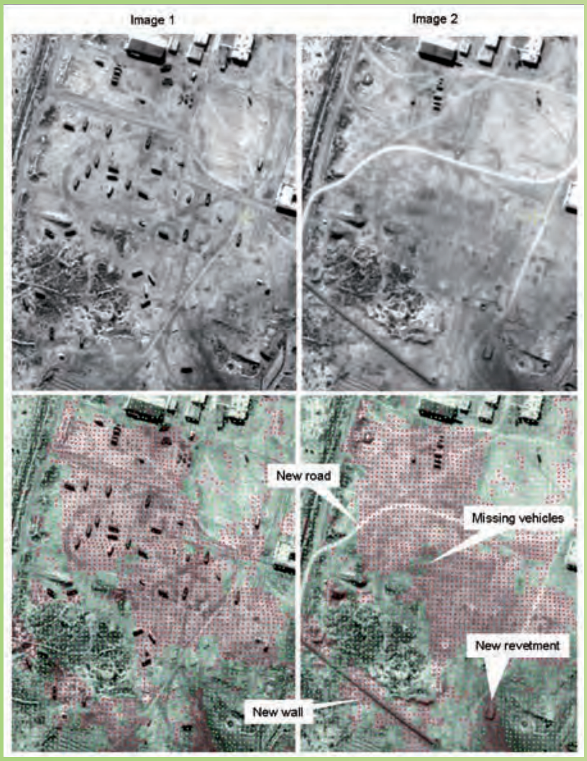

Figure 1. In image 1, enemy vehicles are in garrison. In figure 2 the vehicles have departed. NGATE detects the change as no matches. New roads, earth walls and revetments are also detected.

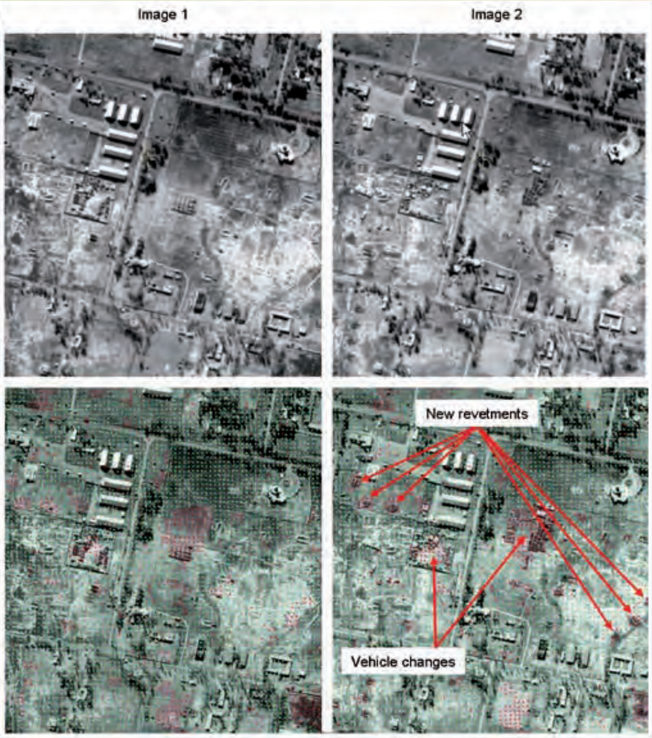

Figure 3. Movement of vehicles and construction of new revetments can be detected.

Minimizing Undesirable Effects

Changes and moving vehicles aren't the only cause of no matches or false matches in stereo image pairs. One must be careful not to classify all poor matches as vehicles. Near roads, streets and highways, there are three other common causes that may lead to no matches or false matches:

1. Vertical structures. Narrow vertical structures such as light poles, traffic signal lights and billboards can result in poor matches. Stereo image matching technology assumes that the underlying surface is continuous and reasonably smooth (in a mathematical sense). Narrow vertical structures violate this assumption and may lead to false matches or no matches. Although a bare earth algorithm can detect and remove these errors from a digital elevation model (DEM), it can't currently distinguish such objects from vehicles. This capability needs to be refined.
2. Trees under windy conditions. When wind blows, trees move and cause dissimilarity between the left image and the right image. Image intensity variations may allow users to differentiate them from moving vehicles.
3. Smooth, low-contrast surfaces. When image radiometric quality isn't high enough to differentiate pixels on a smooth surface, an image match may fail due to low signal power. Users may be able to apply a localized histogram to differentiate these correlation failures from moving vehicles.

To minimize these effects, stereo images should have four characteristics:

1. Smaller convergence angles. The smaller the convergence angle is, the more similar the stereo image pairs are. Therefore, more pixels can be matched as long as the underlying objects or surface don't change or move. Small convergence is good for change detection or vehicle detection, but it isn't good for digital surface model (DSM) generation because the accuracy of the DSM is inversely proportional to the convergence angle. A balance needs to be struck.
2. High resolution. Image matching uses an 11 x 11 window (121 pixels), and every pixel contributes to the matching result. Image resolution must be high enough that a vehicle covers at least the number of pixels of a matching window.
3. High radiometric quality. The higher the radiometric quality, the higher the success rate of image matching. Most modern digital sensors can acquire much higher quality images than old scanned films. Digital images shouldn't be compressed, as compressed images may lose valuable data for the matching process.
4. High registration accuracy. Stereo image pairs are epipolar-rectified so each row in the left image has a corresponding row in the right image. Epipolar rectification requires stereo images to be accurately registered to less than one pixel of Y parallax. The advantage is faster correlation with fewer errors.

ally work faster than with side-by-side comparison.

Certainly, there are acceptable limits for forming an ad hoc stereo pair. Overly wide convergence angles, prominently disparate shadows, poor radiometric matches, or gross seasonal differences (snow, foliage on or off) will defeat the method. One can say the same for any change detection approach.

The factors that make stereo change detection effective can have deleterious effects on automatic terrain extraction algorithms. This fact can be exploited for change detection. Areas of poor correlation (due to image differences) get a low Figure of Merit (FOM). The terrain posts within these areas can be color-coded to indicate areas where one suspects change has occurred. The same suggestion was pursued to develop NGATE.

The theory of stereo image matching is that conjugate pixels are from the same fixed ground object. NGATE performs stereo image matching on every pixel. With sufficiently high resolution and accurately registered stereo images, users can match most pixels even on comparatively patternless surfaces such as highways. To compute the elevation of a pixel, a user multiplies a constant by the amount of X parallax. The amount of X parallax is the shift of a conjugate pixel between stereo image pairs.

When a vehicle moves, the conjugate pixels are no longer from the same location. If the stereo image matcher is able to match pixels from moving vehicles in two scenes, the amount of X parallax is wrong and results in an elevation blunder. Depending on how much a vehicle moves between two images, the stereo image matcher will make either "no match" or a "false match" in the area near the vehicles. In the no-match case, the vehicle has moved too far for the stereo image matcher to find matching pixels in both frames. In the false-match case, the vehicle is correlated in both frames, but has exaggerated X parallax due to displacement. An elevation blunder results. Over fast-moving highways, moving vehicles generally result in no-matches. Over slow-moving traffic, false matches are common. There are also other phenomena that can produce false- and no-match results (see "Minimizing Undesirable Effects" at left).

As part of NGATE enhancements, BAE Systems developed an algorithm to automatically detect and remove elevation blunders caused by moving vehicles.

Combining algorithms that operate on radiometric properties and image geometric properties may lead to further success in change detection and vehicle detection.

NGATE is able to judge the legitimacy of pixel matches, using rule-based logic. Suspect elevations get a low FOM. During terrain extraction, NGATE then will apply a special filter to eliminate implausible spikes. In addition to being an automatic terrain extraction algorithm, NGATE is, in effect, an automatic moving vehicle detector.

What's on the Drawing Board?

The idea of using stereo image matching technology to detect change, including moving vehicles, is worth further investigation. BAE Systems' research indicates this idea may lead to technology that can successfully detect moving vehicles in accurately registered stereo images of suitable resolution. Stationary vehicles can also be detected using a combination of bare earth algorithms and stereo image matching algorithms inside NGATE. More work must be done to reduce false alarms, because change and motion aren't the only reasons no-match or false-match results occur.

Combining algorithms that operate on radiometric properties and image geometric properties may lead to further success in change detection and vehicle detection. NGATE doesn't currently attempt to measure the velocity vector of moving vehicles, so this should be an area of future research. Nor does NGATE have the ability to definitively classify objects as vehicles. For now, it simply recognizes that a vehicle-sized object has caused correlation to fail and repairs the DEM in that vicinity.

As part of an alerting system, NGATE would send the polygonal boundaries of the presumed moving vehicles to analysts for assessment. In an orchestrated workflow, the analyst's image exploitation application would automatically load the imagery at each candidate target location. The analyst would make the final judgment as to whether the object is, in fact, a vehicle, and record the information accordingly. NGATE processes images quickly and could prove valuable as an automated change detection algorithm.

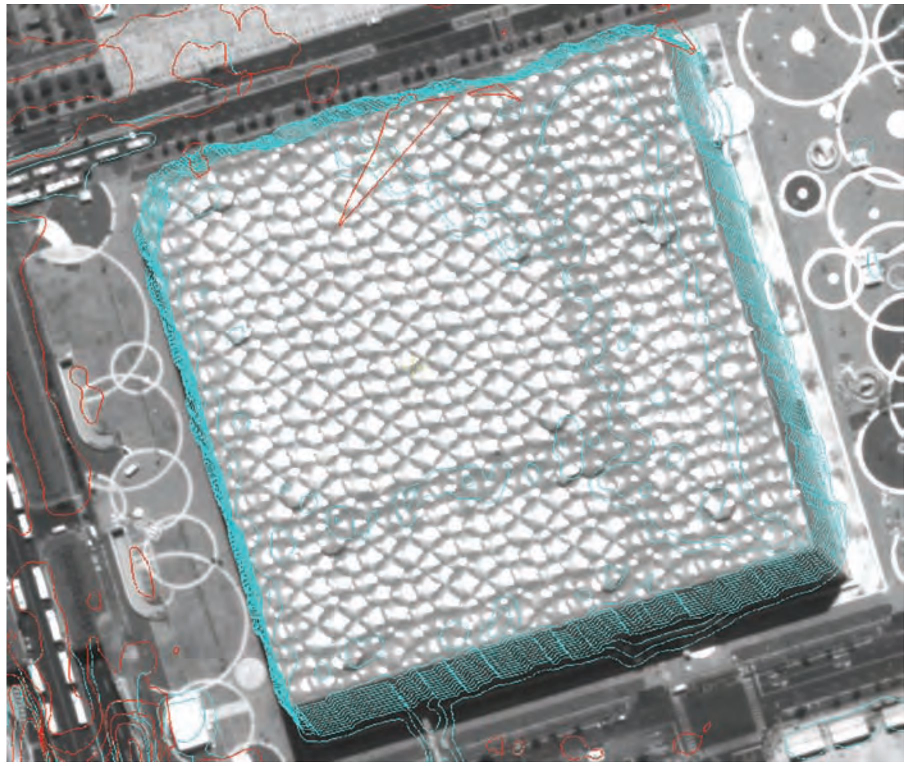

NGATE analysis of the Water Cube aquatics center at the site of the 2008 Beijing Olympic Games shows a DEM in red and a DSM in cyan. An underground tunnel on the lower-left corner alters the DEM contours. There's also a bridge on the lower middle edge. NGATE's simultaneous DSM and DEM generation successfully identified and removed the bridge from the DEM. Future NGATE enhancements for automatic building extraction include algorithms for creating right-angle building corners.