Surveying Coastal Ship Traffic With LANDSAT

Ron Abileah
jOmegak
San Carlos, CA 94070 USA

***Abstract*- A semi-automated algorithm was developed to detect ships in LANDSAT 7 images. The algorithm combines multispectral and pattern recognition methods to discriminate ships from ocean clutter. Automated processing enables us to process a large number of images and gather a statistical picture of ship traffic patterns. As a test case we applied the algorithm on 54 LANDSAT images in the area of Jacksonville, FL, from the period 1999-2003. The area and time period are the same as an earlier ship traffic study by Ward-Geiger et al. using ship reports in the Mandatory Ship Reporting System (MSRS). The similarities between the two studies suggest that LANDSAT is a good alternative for surveying nearshore ship traffic.**

I. Introduction

The applications for remote sensing of ships include maritime safety and security, marine pollution, protection of marine mammals, and regulating fishing in EEZs. Existing and proven methods include coastal radars, Synthetic Aperture Radars on satellites, underwater acoustics, and others. We were interested in testing the utility of LANDSAT images as another alternative. This paper reports our initial findings and experiences with using LANDSAT images for this purpose.

The following section is a brief discussion of existing systems and techniques and the motivation for looking into LANDSAT as an alternative. Section III discusses the computer algorithms developed for processing LANDSAT imagery into ship detections. Section IV presents a comparison between LANDSAT detections and an independent analysis using ship position reports collected by the Mandatory Ship Reporting Systems (MSRS), as reported by Ward-Geiger [1], henceforth referred to as WG.

II. Motivation for using LANDSAT IMAGERY

Ship detection has been demonstrated with a variety of remote sensing technologies. They fall into two broad categories: systems that can in theory monitor ship traffic continuously and systems that provide snapshots of ship traffic at discrete times. In the continuous monitoring category are the following

1) Underwater acoustics. The US Navy's SOund SUrveillance System (SOSUS) was originally intended for deep ocean monitoring of submarines and military ships. The author demonstrated that the system could detect driftnet fishing in the North Pacific [2]. Detection of illegal driftnet fishing became one of the first non-military missions for SOSUS. The application of SOSUS to near-shore ship traffic is less well established.

2) The over-the-horizon HF radars. This is another cold-war era technology originally intended for long range aircraft detection. A team at SRI International proved that ships could also be detected with a different radar waveform and algorithms [3].

3) Coastal HF Surface radars. A large number of these systems are now deployed world-wide for continuous mapping of near-shore currents. The US network of nearly 100 radars is providing real time coverage of the entire West and Northeast coasts, with coverage in the Gulf and Southeast coasts growing. With appropriate reprocessing of the raw radar data such radars can also detect ships, sometime even beyond the radar horizon. Detections up to ~100 km range have been demonstrated by Vesecky and others [4,5].

In the second category are radar and optical imaging systems on low-earth orbit satellites which by their very nature are limited to snapshot looks. The revisit intervals range from days to months. In this category are the following already proven methods

1) *Synthetic Aperture Radar satellites.* The SAR satellites detect both the ship hulls and wakes. It is a day/night, and almost all-weather sensor. The Canadian RADARSAT [6] is the best known SAR satellite in use for ship detections. Several companies offer SAR ship detections as a commercial service [7,8].
2) *Ship clouds.* Ship stack emissions produce long cloud-like streaks that can be detected with the MODIS and AVHRR satellites [9,10].

LANDSAT falls into the this category. But with the SAR and ship clouds methods already proven and SAR now available as a commercial service, what is the motivation for using LANDSAT? There are two advantages to LANDSAT imagery that may be of interest in certain applications.

First, the LANDSAT image archive is longest and most complete time history of any of the sensors or systems mentioned above. The LANDSAT series of satellites began over 30 years ago. The most recent, LANDSAT 7, began operations in 1999 and is still acquiring image data. There are plans to continue the series with LANDSAT 8 to be launched December 2012. The *entire* surface of the Earth is revisited every 16 days. (Although as explained later the actual effective revisit rate is two months.)

The second advantage is cost. The initial $4500 cost of one image scene (185 km x 185 km) was eventually reduced to about $300, and starting in late 2008 the entire LANDSAT image archive became available free of charge with orthorectified Level 1G processing. The images can be obtained by downloads from the LANDSAT portal at http://edcsns17.cr.usgs.gov/EarthExplorer.

The portal has a user friendly interface where the user enters an area of interest with latitude and longitude coordinates, a place name, or using an interactive world map. Dates and other selection criteria may also be entered. The user is then presented with a list of archived images and image thumbnails. The thumbnail images are very low resolution but sufficient to judge if the image quality is suitable for ship surveying. For example one can easily established with thumbnail images whether a particular image has too much cloud cover to bother with. The user places the selected images in a shopping cart and then, like shopping at any Internet store, submits the order. The portal promises that orders will be filled in three days, but lately the turnaround has been same day. When notified by email the user returns to the LANDSAT portal to download the desired images. Each image is a bundle of 9 image bands packed in a zip file of about 250 MB (which expands to 600 MB unzipped); the download time is ~3 minutes. A typical ship traffic study such as the one we will describe in Section IV can be accomplished in one or two days.

III. ALGORITHM FOR SEMI-AUTOMATED PROCESSING

Fig. 1 and 2 are typical LANDSAT images. The two figures are of the same 100 km stretch of the Georgia-Florida coast, from Brunswick, GA in the North to Jacksonville, FL in the South, but taken on different dates. Land and coastline are on the left side. Ships are easy to detect when one knows exactly where to look. A few vessels can usually be found by zooming in on the inlets into the three ports. However a manual search of the entire image to find all ships is a very tedious task. Automated computer processing is essential to searching the entire area. In this section we describe the algorithm developed for automated detection.

It is first useful to categorize images in three level of quality from cloud-free and excellent visibility to heavy overcast and poor visibility of the ocean surface. Fig. 1 is and example of a Category 1 image. It is cloud free and the bottom reflection for shallow depth areas just off the coast indicates excellent visibility conditions. The right panel of Fig. 1 is a zoom on one of the ships found with automated processing. In this case, as in all Category 1 images, ships stand out as very bright objects against a black background . Large ships are resolved into multiple image pixels that outline a distinct and recognizable ship hull. The algorithm to automate this detection is fairly simple.

Fig. 2 is a Category 2 image where there is partial cloud cover and brighter ocean surface due to sunglint from wind roughened surface. Long ocean wave crests are apparent (right panel of Fig. 2). The wave crests are about 100 m apart, which is typical of ocean waves. Ships are not as distinct from the background and more sophisticated processing is required to automate detections.

A category 3 image (not illustrated) has greater cloud overcast and possibly also cirrus clouds and haze. Images can be processed into ship detections with still greater algorithm sophistication, and then require a manual quality check.

Algorithm for clutter free cases. First we describe the algorithm for a clutter free case (such as Fig. 1). This is also the preliminary step for the full algorithm that deals with clutter images. The intensity values in the original LANDSAT image are integer counts in the range of 0 to 255. First, the ocean background radiance in each band is estimated with the median in a sliding ½ km x ½ km area and subtracted pixel by pixel. Henceforth radiance refers to the residual after this subtraction. The background now has an average value of zero and a specially uniform noise of ~1 count, which is probably the sensor noise (i.e., not in the real scene). A typical ship is 20-30 counts in *each* pixel in *each* of the three LANDSAT visible bands, the blue (B), green (G), and red (R), bands also known in LANDSAT parlance as the B10, B20, and B30 bands, respectively. Thus a one pixel ship is a "20 sigma" events, even before we integrate the radiance over the three visible bands. Large ships fill several pixels. It is easy to see why detection is so easy. The probability of a string of pixels, each >20 sigma relative to background noise, being an accidental occurrence in noise is negligible.

At this point the algorithm requires only the visible bands. There is more information in the IR, but there is no justification for using IR when ships are already detected with very high probability. The ship wakes are additional confirmation of a ship, but also not needed.

Algorithm for detecting in clutter. The visible bands are not sufficient for automated detection in clutter cases, such as Fig. 2. In addition to the three visible bands we incorporate data on the near-IR (N), short-wave IR (S), and mid-wave IR (M)), also in 30 m resolution; the two thermal IR bands with 60 m resolution, which for the purpose of this paper are summed into one thermal (T) image; and a panchromatic (P) image with 15 m resolution

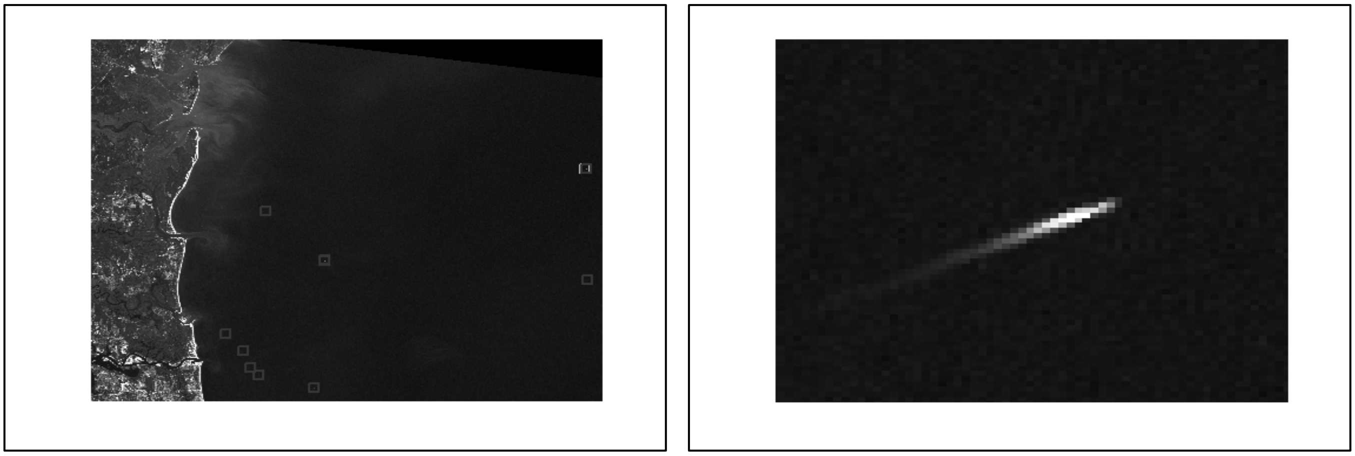

Figure 1 A clear day image. Left: 100 km x 100 km area with squares marking candidate ships detected by automated processing; Right: zoom-in on one of the candidates showing a bright ship hull and trailing wake.

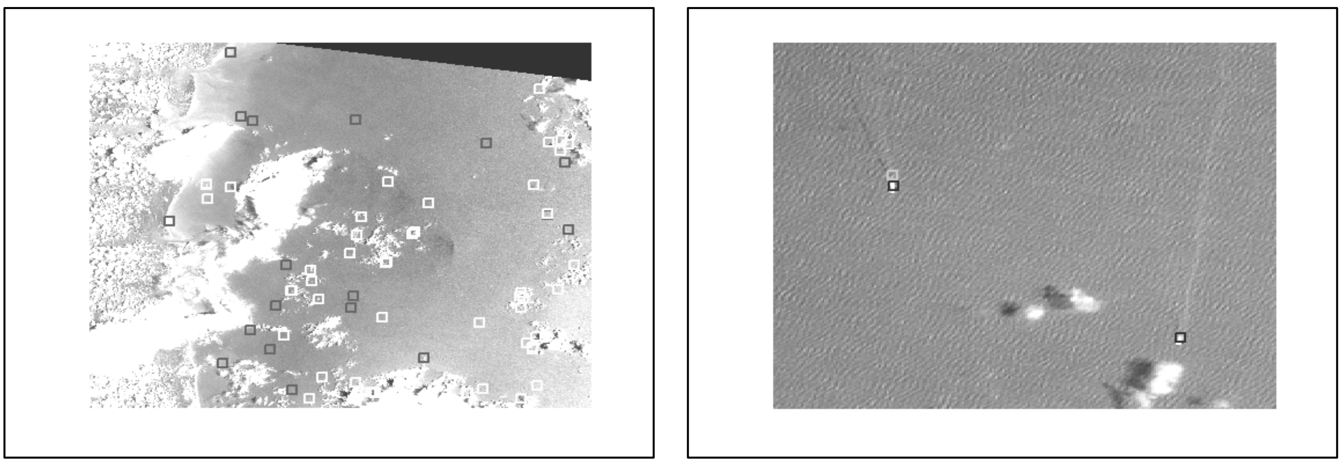

Figure 2 The same scene as in Fig. 1 on another day with clouds and ocean wave. Right panel zooms in on an area with small clouds, waves, and two obvious ships. Ships are detected by various criteria discussed in the text. Ship wakes are a further confirmation that the detections are actual ships.

The most prevalent sources of clutter, in order of difficulty, are clouds, ocean waves, large breaking waves, and surface slicks. (Ocean bottom and water turbidity can also be a source of clutter, but only in shallow water, which we do not address in this paper.) When applied to a clutter image such as Fig. 2 the simple algorithm detects many more false targets than real ships. To appreciate the magnitude of the problem consider the example in Fig. 3 which shows a close-up on an image with clouds, ocean waves, and surface streaks. The brain can process this image and quickly determine that there is one ship in the lower left quadrant. A trailing wake indicates that the ship is heading left. But automating the brain process is challenging. The implementation must be very accurate. There is little room for error in deciding whether an ocean feature is clutter or ship. To illustrate this point consider the following hypothetical ship detection problem. The study area is 100 km x 100 km area . The scene is predominated by wave crests at 100 m spacing. The detector is a sliding 300 m x50 m box (a template for a typical large ship) integrator, with the length of the box aligned with the wave crests (the most stressing case). There are ~10^6 independent, non-overlapping, detector outputs. Half will be on bright wave crests and produce false positives.

Automation of detections in this case requires a technique that will differentiate between clutter and ship with 99.9999% accuracy. Anything less precise will produce many more false positives than real ships and thus produce meaningless statistics. Note that shape and size would not discriminate because most wave crests have similar dimensions to the ship template.

Automatic rejection of clouds is just as difficult as waves. In Fig. 3 note there are several small, elongated cloud streaks that are individually indistinguishable in size, shape, and intensity from ships.

There is no silver bullet to the clutter problem. No one technique is sufficient to separate ships from clutter. Several techniques must be applied on the data in parallel. Table I is a summary of the techniques we applied. The fully automated techniques are implemented algorithms. The manual techniques are more difficult to implement in a computer program. They are presently applied in a visual quality check on the output of automated processing. A brief explanation of the methods follows.

Spectral discrimination is based on the idea that most ships are brightest in B-G-R bands, and much less bright in S-M band. In other words, ships reflect visible light more than IR (as compared to land where the opposite is true). By comparison, clouds, waves, and breaking waves have a “white” (flat) spectrum from blue to midwave IR (recalling that radiance is now defined as excess over an area median value). So one step in cleaning up the clutter in Fig. 3 is to mask all features where the visible and IR radiances are about equal. This removes virtually all the clutter, but it is only a “99%” solution, which is not good enough and needs to be followed with other discriminators.


Figure 3 Closeup on ship with ocean waves, streaks, and clouds

The thermal band is an additional discriminator for clouds. Clouds are generally cooler (by 5-10 counts in the T channel) than the background. So a threshold on the pixel counts in T is used as another flag for cloud pixels. It would have also been helpful if the thermal band could be used to detect “hot” ships. Unfortunately the thermal band resolution , 60 m, tends to dilute the a ship’s thermal signal with the surrounding water.

Pattern recognition methods have also been suggested for ship detection [11]. In our experience the spectral discrimination methods just described rejects more clutter. However, pattern recognition is useful as an additional step to clean clutter features left after spectral discrimination.

Pattern recognition is useful only if the objects of interest are resolved into multiple pixels. As a general rule of thumb pattern recognition algorithms require six or more pixels on the target, although fewer pixel may be sufficient for very simple pattern recognition problems (for example: a linear hull vs jagged cloud). This is where the LANDSAT panchromatic (P) 15 m resolution, which is twice the resolution of the multispectral bands, becomes useful.

We implement the “slender shape test.” This test looks at the P radiance of pixels in the port and starboard sides of the hull. In a case of a slender ship we expect the neighboring cells to be the same as the general water background. If higher than nominal ocean background the detection is suspect. The test also requires that the hull length is at least three times the width. A typical ship length to beam ratio is ~5, so this test is lenient. But it does an excellent job of rejecting breaking waves since they are generally less than one pixel in the P image. This test also eliminates clouds and waves that somehow made it through the spectra and thermal IR tests.

TABLE I
CLUTTER MITIGATION METHODS

	Phenomena			
	Clouds	Waves	Breaking waves	Streaks
Full automated methods	Spectral Thermal Pattern recognition Size	Spectral Radiance threshold	Spectral Size	Pattern recognition Size
Manual methods	Proximity to other clouds Ship wakes Cloud shadows	Orientation parallel to wave crests Ship wakes		

Once all the above have been applied to the data there may still be a few false targets. To eliminate them we now use a manual inspection of the output of automated detections. The human brain does the additional processing. For instance clouds are distinguished by the shadows cast on the ocean surface (see for example Fig. 2 and 3) and by their proximity to other clouds. Ship wakes are a further help in cases where there is remaining doubt about whether a detection is a real ship or not.

It is important to appreciate that there is a penalty in lost ships with each clutter mitigation technique. For instance, some ships (10% in our sample) have the same flat spectrum from blue to midwave IR that we use to identify clouds and wave crests. If the image is free of clouds and waves we could skip this discriminator and avoid losing ships. The length-to-width ratio test eliminates ships <50 m in length, which is unfortunate because LANDSAT can in principle detect ships smaller than a pixel. Clutter mitigation should be used discretionally only where it is needed.

IV. Test case: survey in the South East Right Whale Critical Habitat

The area shown in Fig. 1 and 2 is also known as the South East US (SEUS) critical habitat. This is Northern right whale's winter calving area. Unfortunately it is also an area of high ship traffic, both commercial and military, in and out of ports in Jacksonville, FL in the south end of the SEUS, Fenandina, FL in the middle, and Brunswick, GA in the north end. Vessel strikes have become a major concern in the survival of this endangered specie. NOAA is very active in research, advisories, and regulations aimed at reducing risks for the right whale. Vessels exceeding 300 gross tons have been required to report through the MSRS when entering the SEUS zone.

WG used the MSRS data base for 1999-2003 in a comprehensive analysis of the ship traffic in the SEUS. Upon first hearing of the WG study we immediately saw it as an opportunity to test the LANDSAT ship detection. We downloaded 54 LANDSAT images for the SEUS area, spanning 47 months in the same time period as the WG study.

There is about one image per month, but the image quality varies. Only 20 images are in Category 1. The rest are Category 2 (13) and Category 3 (21). Weighting each image by the fraction of water area that is visible and usable for ship detection leads to an effective 27 revisits, or one per two month.

To make the comparison of LANDSAT with WG most meaningful we attempt to include the same ship population. We used a 150 m threshold for ship hull length. The larger ships are more likely to be reporting to MSRS. The length threshold also gives greater assurance that all the detections are real ships, not false positives. There were 132 ship detections with this length criterion, which works out to 132/27 ~ 5 ships per image, where an image is weighted as previously described. The 132 detections are plotted in the left panel of Figure 4 as vectors. The vectors indicate the ship hull's orientation. The left panel of Fig. 4, taken directly from WG, is a color coded ship density, based on 1004 vessel reports.

There are striking similarities bewteen the LANDSAT and WG plots. Both show concentrations in the vicinity of the pilot buoys where ships meet the pilot boats. In both surveys ships are absent from the immediate offshore areas between the ports. This makes sense: the large ships are either going in and out of port or in transit. Transits would be well off shore in deep water. There is less agreement further out to sea where LANDSAT ships appear to be spatially homogenous and WG indicates concentration of ships in narrow corridors. There are several possible explanations for this difference. Some WG tracks are straight-line from the point where ships enters the SEUS to the buoy they indicate as their destination. The actual ship track may have diverted from a straight line In fact the hull directions (in Fig. 4) suggest that ships meander somewhat from straight-line tracks. We also see evidence of turns and jaggy tracks in ship wakes. Another difference is that WG includes only inbound traffic. LANDSAT ships are both in- and outbound. Outbound tracks would not aim at a pilot buoy and would thus be more dispersed. Yet another explanation is that LANDSAT detections include military vessels (several detections are actually multiple ships in military formation) which are not in the MSRS data base.

V. Summary & Conclusions

We successfully demonstrate that ship traffic patterns can be derived with LANDSAT images. In our test case LANDSAT was compared with an earlier study [1] which used MSRS ship reports. There is good agreement between the two approaches. What differences exist can be attributed to the different sampling and data reduction procedures used in each of the survey methods.

The task of finding ships in images is mostly automated. In clutter-free images (about 30% in our test case) ships are easily detected as objects of one or more pixels with

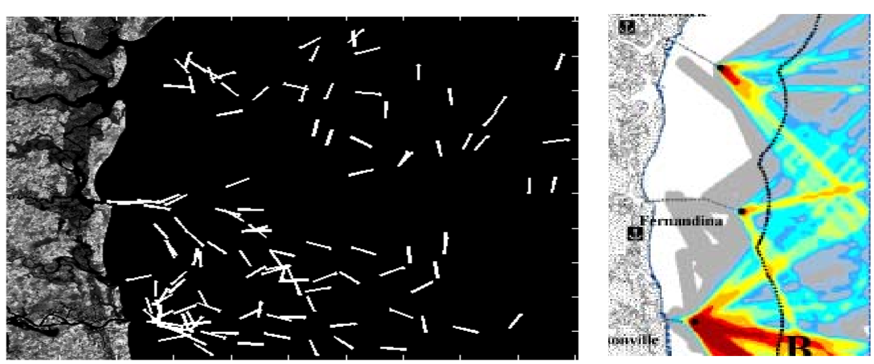

Figure 4 Left: LANDSAT ships. Right: ship traffic density extracted from a figure in Geiger-Ward et al. [1]

radiance levels that are significantly above the background. Only the visible bands (blue, green, red) are required to make a reliable detection. For more cluttered images (i.e., with clouds, ocean waves, etc.) we require all the LANDSAT visible and IR bands (blue to midwave IR, and thermal IR) to discriminate ships from clutter. Pattern recognition methods are applied on the panchromatic image for further discrimination. The algorithm is very effective for clutter mitigation but at the expense of sacrificing some actual ships. The clutter mitigation should therefore be used with discretion, applied only on images with clutter.

ACKNOWLEDGMENT

The author thanks Mark Mueller, Florida Fish & Wildlife Conservation Commission, for helpful discussions pointing to the Ward-Geiger et al. study. The LANDSAT data used in this study was downloaded from the USGS EarthExploerer portal.

REFERENCES

[1] L.I. Ward-Geiger, G.K. Silber, R.D. Baumstark and T.L. Pulfer, "Characterization of ship traffic in right whale critical habitat," *Coastal Management,* v. 33, p. 263-278, 2005

[2] R. Abileah and D. Lewis, "Monitoring high-seas fisheries with long-range passive acoustic sensors," *Oceans'96. MTS/IEEE Conference Proceedings*, Volume 1, 23-26 Sept. 1996 pp:378 - 382 vol.1

[3] J. Barnum*, "Ship detection with high-resolution HF skywave radar", IEEE Journal Oceanic Engineering*, vol. OE-11, no. 2, pp. 196-209, April 1986.

[4] J.F. Vesecky, K. Laws, and J.D. Paduan, "Monitoring of Coastal Vessels Using Surface Wave HF Radars: Multiple Frequency, Multiple Site and Multiple Antenna Considerations," *Geoscience and Remote Sensing Symposium (IGARSS) 2008,* Volume 1, 7-11 July 2008 pp: I-405 - I-408

[5] T. M. Blake, "Ship detection and tracking using high frequency surface wave radar," *Seventh International Conference on HF Radio Systems and Techniques*, (Conf. Publ. No. 441) Volume , Issue , 7-10 Jul 1997 pp:291 – 295

[6] P.W. Vachon , J.W.M. Campbell, C. Bjerkelund, F.W. Dobson, and M.T. Rey, "Ship detection by the RADARSAT SAR: Validation of detection model predictions", *Can. J. Remote Sensing,* Vol. 23, No. 1, 1997 pp. 48-59

[7] Kongsberg web site: http://www.ksat.no/Products/ShipDetection.htm

[8] H Greidanus, "Benchmarking operational SAR ship detection," *IGARSS 2004,* 20-24 Sept 2004, Anchorage, Alaska

[9] J.M. Weiss, R. Luo, R.M. Welch, , "Automatic detection of ship tracks in satellite imagery," *IGARSS '97,* 3-8 Aug. 1997, Vol 1, pp. 160 - 162

[10] N. Srinivasa Rao, M. M. Ali, M. V. Rao, and I. V. Ramana, Estimation of Ship Velocities From MODIS and OCM , *IEEE Geoscience and Remote Sensing Letters*, VOL. 2, NO. 4, October 2005, pp: 437-439

[11] D.W. Burgess, "Automatic ship detection in satellite multispectral imagery," *Photogrammetric engineering and remote sensing,* 1993, vol. 59, no2, pp. 229-237