

Ship Detection in Satellite Imagery Using Rank-Order Grayscale Hit-or-Miss Transforms

Neal R. Harvey*, Reid Porter, James Theiler
Space & Remote Sensing Sciences Group,
Los Alamos National Laboratory,
Los Alamos, NM 87545, USA

ABSTRACT

Ship detection from satellite imagery is something that has great utility in various communities. Knowing where ships are and their types provides useful intelligence information. However, detecting and recognizing ships is a difficult problem. Existing techniques suffer from too many false-alarms. We describe approaches we have taken in trying to build ship detection algorithms that have reduced false alarms. Our approach uses a version of the grayscale morphological Hit-or-Miss transform. While this is well known and used in its standard form, we use a version in which we use a rank-order selection for the dilation and erosion parts of the transform, instead of the standard maximum and minimum operators. This provides some slack in the fitting that the algorithm employs and provides a method for tuning the algorithm's performance for particular detection problems. We describe our algorithms, show the effect of the rank-order parameter on the algorithm's performance and illustrate the use of this approach for real ship detection problems with panchromatic satellite imagery.

Keywords: ship detection, panchromatic image, mathematical morphology, hit-or-miss transform, grayscale, rank-order

1. INTRODUCTION

Ship detection from remotely-sensed imagery has important applications that include maritime security and monitoring of fishing activities. However, detecting and recognizing ships in satellite imagery is a difficult problem. The most commonly-used ship detection techniques are based on the fact that ships on the ocean are generally brighter than their immediate surroundings, and therefore use local contrast in their detectors. However, these methods can produce unsatisfactory false alarm rates.

Template matching techniques are methods commonly found in target-detection applications. These techniques allow users to customize their detection algorithms to specific targets of interest. The basis of template matching is the comparison of a representative pattern (template) of the target of interest with an image, with the purpose of detecting occurrences of the reference pattern in the image. This ability to tailor a detector to specific targets holds the promise of being able to reduce false alarms in ship detection, when compared to the standard local contrast-based approaches.

Mathematical morphology is a branch of nonlinear signal and image processing that provides powerful tools for various image and signal processing applications. The Hit-or-Miss Transform (HMT) is a template matching method that exists within the mathematical morphology framework. The HMT has several benefits compared to other template-matching techniques, including the fact that it is based on rank-order operations which are generally faster than linear correlation techniques and that it is insensitive to DC variations of the template and the image.

In this paper we set out to investigate the utility of the Hit-or-Miss Transform to the ship detection problem. In section 2 we provide an overview of the Hit-or-Miss transform, its basis in the binary image processing domain and its extension to the grayscale domain. We focus on a particular grayscale extension and modifications to this, on which we based the work described in this paper. In section 3 we describe the data we use and the specific set of experiments we undertook. In section 4 we show and describe the results of these experiments. In section 5 we describe our conclusions from these experiments and further work that could improve upon the work described here.

Further author information: Neal R. Harvey: E-mail: harve@lanl.gov, Telephone: +1 505 667 9077

2. HIT-OR-MISS TRANSFORM

The Hit-or-Miss transform was first introduced by Serra¹ and was originally defined for binary images and signals.

The basic idea behind the hit-or-miss transform is that of extracting all pixels within an image that are matched by a given neighborhood configuration, consisting of some arrangement of foreground pixels and background pixels. The neighborhood configuration is therefore defined by a pair of disjoint sets: one for the foreground pixels and the other for the background pixels. If we represent our foreground and background sets with a pair of structuring elements (A, B) , (the ‘‘Hit’’ and ‘‘Miss’’ templates, respectively), then if we apply the HMT to an image X using (A, B) , we are looking for all positions where A can be fitted within the foreground X and simultaneously where B can be fitted within the background X^c (where X^c denotes the complement of X). This can be represented mathematically as :

$$X \odot (A, B) = \{x \mid A_x \subseteq X, B_x \subseteq X^c\} \quad (1)$$

Note that A_x (resp. B_x) denotes the translation of set A by x (resp. B by x). We can also describe the HMT in terms of an intersection between two erosions (where \ominus denotes the binary erosion operator):

$$X \odot (A, B) = (X \ominus A) \cap (X^c \ominus B) \quad (2)$$

2.1 Hit-or-Miss Transform: Extension to Grayscale

While the binary HMT has been widely used, its extension to gray-level images is not straightforward, due to the fact that the HMT is not an increasing transform (a transform is increasing when it preserves inclusion - i.e. containment relationships are maintained. For an increasing transform Ψ , $X \subset Y \Rightarrow \Psi(X) \subset \Psi(Y)$). However, there have been several definitions of the grey-level hit-or-miss transform (GHMT).²⁻⁵ Naegel et al. brought several of the main definitions of GHMT together, reviewing them and providing a unified theory.⁶ Since then, other authors have extended the definition of the GHMT further to include a *fuzzy* GHMT.^{7,8}

The grayscale hit-or-miss transform that we use in this work is based on that described by Barat et al.⁴ Barat et al. described an operator that they called Single Object Matching using Probing (SOMP), and an extension to this, which they called Multiple Objects Matching using Probing (MOMP).

2.1.1 The SOMP Transform

The fundamental SOMP operator is based on the principle of mechanical profiling in which surface height variations of an object are measured using a probe, as it is passed across the surface of the object. This probing technique can be transposed to image analysis using the elementary morphological operators of dilation and erosion to probe the topographical surface of the grayscale image being analyzed. The key idea is to detect a target of interest by analyzing the relevant position of two probes brought into contact with the image surface, above and beneath simultaneously. Similar to the Khosravi and Schafer’s GHMT,³ the SOMP uses identical structuring elements for the ‘‘Hit’’ and ‘‘Miss’’ templates. The SOMP approach corresponds to a topological approach to pattern matching and has the properties of a metric. As such, it is a more efficient alternative to correlation algorithms (requiring no multiplications, only comparisons). In addition, compared to the Hausdorff metric, SOMP has the advantage that it is invariant to changes of intensity of the image being searched and automatically compensates for intensity offsets between the target object and the image being searched.

Let us represent our grayscale image to be probed as f , our target template as t and t_v as a vertical translation of t by an amount v . Probing the image f from above and below involves translating t upwards by values $v1$ and $v2$ respectively, so that the translated versions of t touch the surface of f from above and below. $v1$ corresponds to the result of a grayscale morphological erosion of f by t and $v2$ is equal to a grayscale morphological dilation of f by $-t$ (sometimes called an anti-dilation). The SOMP transform is equal to the difference between the vertical positions of the inferior and superior probes, $(v2 - v1)$. If we denote the grayscale erosion of f by t as $\varepsilon_t(f)$ and the grayscale (anti-) dilation of f by $-t$ as $\delta_{-t}(f)$, then we can define SOMP as follows:

$$SOMP_t(f) = \delta_{-t}(f) - \varepsilon_t(f) \quad (3)$$

2.1.2 The MOMP Transform

While the SOMP algorithm is designed to search for a single instance of a target of interest, some applications require the ability to detect objects with varying shapes in a image. The variations in these objects could be due to different types of targets or a specific target type with varying size or noisy distortion. One approach to this problem would be to apply the SOMP algorithm iteratively, using all possible variations of the target template. However, this approach would be far too time-consuming. In order to better address this problem, Barat et al. proposed an extension to the SOMP algorithm - the MOMP algorithm, that allows the search for variations of the target of interest, using a unique pair of adapted probes. These probes are distinct from one another (compared to the SOMP algorithm which uses identical probes) and define a template which contains all the variations in the targets to be detected.

Defining the MOMP Probes (“Hit” and “Miss” Templates): Let N be the number of examples of the targets of interest and $t_{i,i \in [1,N]}$ by their representative functions (grayscale images). Contrary to the identical nature of the two SOMP probes, the two MOMP probes are distinct. The definition of the probes is based on that described by Doh et al. for their extension of the HMT to the problem of detecting different binary shapes,⁹ but extended to the grayscale case. The mathematical definition of the probes is established using set theory. The inferior (resp. superior) probe, denoted by t_{inf} (resp. t_{sup}), corresponds to the function obtained by taking the minimum (resp. maximum) of all functions t_i on the union of their spatial supports (let us define T_i as the spatial support of t_i). Thus, they are defined as follows, $\forall y \in \bigcup T_i$:

$$t_{inf}(y) = \min(t_i(y)) \quad (4)$$

$$t_{sup}(y) = \max(t_i(y)) \quad (5)$$

Defining the MOMP Transform: The MOMP transform is expressed in the same way as the SOMP transform, only using distinct “hit” and “miss” templates (structuring elements t_{sup} and t_{inf}):

$$MOMP_{(t_{inf}, t_{sup})}(f) = \delta_{-t_{sup}}(f) - \varepsilon_{t_{inf}}(f) \quad (6)$$

Figure 1 provides an illustration of the principle behind the MOMP transform. For simplicity it is shown for a 1-D signal. An example pair of Hit and Miss) probes is illustrated, together with how they are brought into contact with the signal being probed and the resulting distance metric produced.




Figure 1. An illustration of the principle behind the MOMP transform.

2.1.3 Extensions with Rank-Order Filters

The commonly used definitions of the GHMT use the grayscale morphological operations of dilation and erosion. These operators are based, respectively, on maximum and minimum operators. Using these operators can lead to the GHMT to be sensitive to impulsive noise.³ In order to reduce such sensitivity, one can consider using other order statistics in place of the maximum and minimum. Therefore, we can replace the grayscale dilation and erosion operators with the corresponding grayscale rank-order equivalents in the SOMP and MOMP transform. By doing so, we not only reduce the transforms' sensitivity to noise, but also provide some additional "slack" in the fitting of the hit and miss templates (probes) to the image being searched. While this additional slack may allow fitting to targets that are slightly outside the variability encompassed within the hit and miss templates and improve the transform's detection ability, it may also allow a fitting to non-target regions of the image and therefore has the possibility of increasing false alarms.

2.1.4 Our Implementation of the GHMT

As we have previously stated, we have based our implementation of the GHMT on the MOMP described by Barat et al. We modified the implementation to allow rank-order operators in the definition of the transform, instead of the standard dilation and erosion operators. Note that, by specifying a rank-order parameter of 1 the rank-order operators become the maximum and minimum operators, which are the dilation and erosion operators and thus is equivalent to the original definition of the MOMP transform. The output of the MOMP transform, at each pixel in the image being searched, is a distance measure of how well the neighborhood of the pixel fits with the hit and miss templates. Having a distance measure output for each pixel position, produces a multi-valued output image. The implementation of the MOMP described by Barat et al. used a threshold value defined as the difference between grayscale values of the hit and miss templates at their center point to define their detection criterion. Our approach is that, with a multi-valued output image, we can threshold the output image over a range of values and, for each threshold value, calculate a detection rate and false alarm rate. In this way, by varying the threshold, we can trace a Receiver-Operator Characteristic (ROC) curve, to show the performance of the GHMT on the target detection problem at hand.

Our specific software implementation of the GHMT was based on the functionality provided by Kitware's Open Source Insight Segmentation and Registration Toolkit (ITK)¹⁰ software.

3. EXPERIMENTS AND DATA SETS

3.1 Experiments

Template matching methods in general have been proven to work extremely well for problems where the target appearance is well characterized and there is little variability or noise likely to be present. For the GHMT it has been shown in the literature that the transform can perform well for some problems, but often these problems are rather simplistic or synthetic. There are relatively few examples that show the performance of the GHMT to real world detection problems, with the inherent problems this entails, such as large target variability and noise.

We wanted to test the GHMT on a difficult real-world problem. To this end, we chose ship detection in panchromatic imagery. Panchromatic imagery only has one (grayscale) channel and there is therefore no spectral information to help in the detection problem. Ships, even for one specific type, can exhibit an enormous variability in their appearance. There can be differences in the cargo and other equipment placed on deck. There can be differences in the shade of the vessel, depending on how it has been painted. Because we are using satellite imagery, there can be great differences in the illumination of the vessels, due to position of the satellite or the sun or orientation of the vessel.

Thus, for the experiments described here, we chose panchromatic satellite images from a region that has a great deal of ship traffic and chose to attempt to develop a detection algorithm, using a specific GHMT (rank-order MOMP), for a specific type of ship. We wanted to investigate how well the rank-order MOMP transform could perform in detecting this specific ship type, and not other types of ship. We also wanted to investigate the effect of varying the transform's parameters on its performance.

3.2 Data

The data we used for the experiments described here was Quickbird¹¹ panchromatic imagery. The particular set of Quickbird data used was segments of two larger images taken over Hong Kong in January 2004 and January 2007. We had a total of 9 image segments of approximately 1000 x 1000 pixels. Figure 2 (a) shows an example of one of the image segments used in the training set. Note that the contrast has been stretched to show the details in the image more clearly.




Figure 2. Examples of Original Panchromatic Image Data and Target Masks for “Training” Data

3.2.1 Targets of Interest

As our target of interest we chose a kind of cargo vessel with a distinctive “round” bow. For each of the 9 image segments we marked up the targets of interest that were present in the images. A total of 40 targets of interest were marked in these images. Figure 2 (b) shows an example of a mask indicating where these targets of interest are in one the image segments. The images were then divided into two sets, with each set containing 20 targets of interest. One set was used to create the hit and miss templates, i.e. it was the “training” set and the other set was used as the testing or validation set.

Figure 3 shows the 20 examples of the target of interest that were extracted from the training set data and used to create the hit and miss templates. It can be seen from Figure 3 that there is considerable variation in the examples of the target of interest. What is particularly obvious is the variation in lighting in the examples, as well as the variation in the cargo content in the vessels.

Figure 4 shows the hit and miss templates created from the training set, using the maximum and minimum operators, as described previously.

3.2.2 Applying the GHMT to the data

Once the hit and miss templates were created from the analyst mark-up, we then applied our GHMT implementation to the training and testing images. Because we wanted to be able to detect the targets of interest at any rotation in the image, we created versions of the hit and miss templates at 64 different rotations. We then applied the GHMT algorithm using all 64 rotated hit and miss combinations separately. Each application of the GHMT algorithm with one of the rotated hit and miss combinations produced an output image with a distance measure at each pixel. To combine the output images from all the rotations and create the final output image we simply found the minimum value at each pixel over all the rotations. Thus we were finding the best match to the hit and miss templates over all the rotations.





Figure 3. Examples of target of interest from the training set, which were used to create the hit and miss templates.



(a) Minimum (Hit) Template (b) Maximum (Miss) Template

Figure 4. Examples of (a) Hit and (b) Miss templates. Note that the contrast has been stretched for the Hit template in order to make it easier to see.

3.2.3 Varying the rank-order parameter

In order to determine what effect varying the rank order parameter had on the performance of the GHMT, we applied the GHMT as described previously, but multiple times, using a different value for the rank-order parameter each time. We used rank-order parameters ranging from 1 to 96 in increments of 5.

3.2.4 Calculating ROC Curves for each rank-order parameter

For each output from each GHMT application we created a ROC curve by varying the threshold on the output distance measure and calculating a detection rate and false alarm rate for each threshold level. For each target of interest in the image (that had been marked up by the analyst), we determined a true detection to have occurred if, after thresholding, there were any pixels that intersected with the target mask. Thus our detection rate was calculated at an object level.

In addition to the mark-up for the targets of interest, the analyst had also provided mark-up for regions that he was certain did not contain any targets of interest. These regions represented our background masks. We calculated our false alarm rate as the number of pixels, after thresholding, that intersected with the background mask, divided by the total number of pixels in the background mask. Thus, our false alarm rate was calculated at a pixel level.

4. RESULTS

Figure 5 shows how the ROC curves vary with different rank-order parameters. Figure 5 (a) shows this for the training data and Figure 5 (b) for the testing data. It can be seen that as you increase the rank-order parameter there is an increase in performance (i.e. the ROC curves get pushed toward the top-left corner), up to a certain limit, at which point the performance begins to decrease. We see very much the same behavior for both the training and the testing data. What is very obvious is the vast difference in performance between the application of “standard” MOMP definition, for which the rank-order parameter is 1 and the MOMP application with the best-performing rank-order parameter. This is especially evident for the training data (Figure 5 (a)). Here the ROC curve for a rank-order parameter of 1 is pushed way down to the bottom right - exactly the opposite of what we want to see, whereas for the best-performing rank-order parameter, the ROC curve is pushed way up to the top left - exactly what we do want to see. This behavior is still clearly evident for the testing data (Figure 5 (b)), but it is not quite as marked as for the training data.

As we are ultimately interested in obtaining low false alarm rates and high detection rates, we are interested in the parts of the ROC curves that are towards the left-hand side. In order to compare ROC curves for various rank-order parameters objectively, we can pick a low false-alarm rate (e.g. 0.025) and see what detection rate we can get with that false-alarm rate for a particular ROC curve and compare these values for the different rank-order parameters.

Figure 6 shows how the probability of detection attainable for a probability of false alarm of 0.025 varies with rank-order parameter, for both training and testing data. As we saw with the ROC curves themselves, we can see that there is a rapid improvement in performance as one increase the rank-order parameter, but that, at some point (in this case at a rank-order of about 80) we reach a maximum performance and any further increase in the rank-order parameter causes a drop in performance. We can also see that the behavior is the same for both the training and testing data, with the optimal rank-order parameter for both training and testing being around 80. It should be noted that the size of the support of the Hit template was 146 pixels and the size of the support of the Miss template was 293 pixels. Therefore, given that the same rank-order parameter is used for the erosion with the hit template and the (anti-) dilation with the miss template, the maximum rank-order parameter possible would be 146 for this example.

Figure 7 shows the result of applying the GHMT transform to one of the testing images. Here we took the distance measure output of the image and picked a threshold that produced an optimal output in terms of a combination of minimizing false alarms and maximizing true detections. The resulting detection “blobs” are overlaid onto the original image. In addition, due to the fact that some of the detection blobs are quite small, to further highlight the detection locations, we have also overlaid 100x100 boxes, centered on the centroids of the detection blobs. Within this image there was one target of interest that the analyst was certain was of the target type. This ship is the one that has been highlighted towards the lower-right part of the image. The algorithm has also detected another ship in the image - the one highlighted towards the upper-right part of the image. This is not of the same type as the target of interest and is therefore a false alarm. There are several other false-alarms highlighted within the image, but these are all on the land, in the upper-left part of the image. Looking at the areas on the land that are highlighted, they are predominantly buildings that have some similarity

in size, shape and texture to the targets of interest and it is therefore not so surprising that they have been detected. For a real-world ship detection application, it would be a simple procedure to mask out all land regions and these land-based false-alarms could be easily discounted. There exist several shoreline databases that could be used for masking out the land in the image, including NOAA's Shoreline Database and NGA's Prototype Global Shoreline Data, both of which can be accessed via links available at NOAA's National Geophysical Data Center (NGDC) website.¹² What is quite interesting in this result is that the false-alarm vessel highlighted by the GHMT detection algorithm seems to be less similar in appearance to the correctly-identified target in the image than other vessels in the image - for example the vessel immediately to the right of the false-alarm. The appearance of the false-alarm vessel must lie within the range of variation allowed by the hit and miss templates together with the slack provided by the rank-order parameter value. However, it still remains surprising that the vessel to the right of the false-alarm didn't have a close match. This highlights the fact that the question of how much variability this method allows in the target type is still an area that needs to be addressed.

5. DISCUSSION AND CONCLUSIONS

We have demonstrated the application of a grayscale hit-or-miss transform to the problem of ship detection in panchromatic satellite imagery. We have also shown how, by modifying the transform to allow the use of rank-order operators instead of the standard dilation and erosion operators in the implementation of the transform, we can greatly improve performance on the ship detection problem. The choice of optimal rank-order parameter is clearly problem dependent.

The results of our experiments do seem to indicate that the GHMT has potential for addressing the ship-detection problem. While we have demonstrated its potential and shown how we can improve performance of the GHMT on ship detection problems through rank-order operators, the results are not perfect. Some false alarms and missed detections still occur and this motivates further work.

For real world applications it would be necessary to have a hit and miss template for each target type. However, how broad or narrow the definition of target type would need to be has yet to be investigated. How much variation in the target type can be tolerated?

In our experiments we demonstrated that choosing the right rank-order parameter can have a dramatic effect on the algorithm's performance. We performed an exhaustive search to find the optimal rank-order parameter. We did not investigate possible methods for efficient search for the optimal rank-order parameter.

While the morphological template matching techniques are known to be more efficient than the standard techniques, for real-world applications such as those described here, it is still necessary to perform moving window operations with large windows and this requires considerable computing effort. Investigations into efficient implementations of the GHMT for large hit and miss templates could prove beneficial in terms of making these techniques more useful in real-world scenarios.


ACKNOWLEDGMENTS

This work was funded by the US Department of Defense. We also gratefully acknowledge the assistance and advice provided by Luis Ibanez of Kitware Inc. and Richard Beare of the ITK users group that enabled us make the ITK modifications necessary to implement our GHMT software.


REFERENCES

- [1] Serra, J., [*Image Analysis and Mathematical Morphology*], Academic Press (1983).
- [2] Soille, P., [*Morphological Image Analysis: Principles and Applications*], Springer, 2 ed. (2003).
- [3] Khosravi, M. and Schafer, R., "Template matching based on a grayscale hit-or-miss transform," *IEEE Trans. on Image Processing* **5**, 1060–1066 (June 1996).
- [4] Barat, C., Ducottet, C., and Jourlin, M., "Pattern matching using morphological probing," in [*Proc. International Conference on Image Processing (ICIP)*], I: 369–372 (2003).
- [5] Weeks, A., Sartor, L., and Richie, S., "New grayscale hit-miss operator," *Journal of Electronic Imaging* **13**(1), 169–171 (2004).

- [6] Naegel, B., Passat, N., and Ronse, C., “Grey-level hit-or-miss transforms - part i: Unified theory,” *Pattern Recognition* **40**, 635–647 (2007).
- [7] Perret, B., Lefevre, S., Collet, C., and Vollmer, B., “Astronomical object detection with a robust hit-or-miss transform,” in [*Proc. European Signal Processing Conference (EUSIPCO)*], (2008).
- [8] Perret, B., Lefevre, S., and Collet, C., “A robust hit-or-miss transform for template matching applied to very noisy astronomical images,” *Pattern Recognition* **42**, 2470–2480 (2009).
- [9] Doh, Y., Kim, J., Kim, J., Kim, S., and Alam, M., “New morphological detection algorithm based on the hit-miss transform,” *Optical Engineering* **41**, 26–31 (January 2002).
- [10] “Insight segmentation and registration toolkit (itk).” <http://www.itk.org>.
- [11] “Digitalglobe — an imagery and information company.” <http://www.digitalglobe.com/>.
- [12] “NOAA’s National Geophysical Data Center shoreline/coastal data web page.” <http://www.ngdc.noaa.gov/mgg/shorelines/shorelines.html>.



(a) Comparison of ROC curves for various rank-order parameters for training data.



(b) Comparison of ROC curves for various rank-order parameters for testing data.

Figure 5. Comparison of ROC curves for different rank-order parameters used for the GHMT: (a) For training data; (b) For testing data.




Figure 6. Variation of maximum probability of detection attainable for a probability of false alarm of 0.025 with rank order parameter.



Figure 7. Result of applying the GHMT algorithm to one of the testing images, using the Hit and Miss templates created from the training images, as shown in Figure 4.