

A Colour Hit-or-Miss Transform Based on a Rank Ordered Distance Measure

Fraser Macfarlane*, Paul Murray*, Stephen Marshall*, Benjamin Perret†, Adrian Evans‡, Henry White§

* Department of Electronic & Electrical Engineering, University of Strathclyde, Glasgow, Scotland

† Université Paris-Est, Laboratoire d'Informatique Gaspard-Monge UMR 8049, UPEMLV, ESIEE Paris, ENPC, CNRS, F-93162 Noisy-le-Grand France.

‡ Department of Electronic & Electrical Engineering, University of Bath, Bath, England

§ BAE Systems, Air Sector, Filton, Bristol, UK

{fraser.macfarlane, paul.murray, stephen.marshall}@strath.ac.uk, benjamin.perret@esiee.fr, A.N.Evans@bath.ac.uk

Abstract—The Hit-or-Miss Transform (HMT) is a powerful morphological operation that can be utilised in many digital image analysis problems. Its original binary definition and its extension to grey-level images have seen it applied to various template matching and object detection tasks. However, further extending the transform to incorporate colour or multivariate images is problematic since there is no general or intuitive way of ordering data which allows the formal definition of morphological operations in the traditional manner. In this paper, instead of following the usual strategy for Mathematical Morphology, based on the definition of a total order in the colour space, we propose a transform that relies on a colour or multivariate distance measure. As with the traditional HMT operator, our proposed transform uses two structuring elements (SE) – one for the foreground and one for the background – and retains the idea that a good fitting is obtained when the foreground SE is a close match to the image and the background SE matches the image complement. This allows for both flat and non-flat structuring elements to be used in object detection. Furthermore, the use of ranking operations on the computed distances allows the operator to be robust to noise and partial occlusion of objects.

Index Terms—Image processing; Mathematical morphology; Hit-or-Miss Transform; Template matching; Object detection

I. INTRODUCTION

Mathematical Morphology (MM) is a fundamental set of techniques in the field of image processing, having first been introduced in the late 1960s and formalised by Matheron [1] and Serra [2] and later extended by Heijmans [3] in 1994. It applies the mathematical concepts of set theory, specifically lattice theory, to various image analysis tasks such as object or edge detection, segmentation and image de-noising [4], [5]. Initially, its use was limited to binary images however there has since been much effort to extend MM techniques for application to grey scale images and beyond [6]–[8].

The HMT [2] is just one of the many morphological operators that have been developed in the past 40 years. It uses a composite SE comprised of two sets, one to match the geometric properties of objects in the foreground and a second to match those of the background. It is commonly used for template matching in object detection applications. Its application to both binary [2] and greyscale images [4], like the rest of MM, are now well defined, however its extension to colour and multivariate images is not so straightforward. This is due to MM in general requiring a complete lattice,

or ordered set to allow the definition of its operators. While trivial for scalar, singular channel images, ordering a set of vectorial or multivariate data becomes challenging [6].

There have been multiple recent advances in the creation of a colour morphology framework using various techniques, each trying to overcome the multivariate data ordering challenge. For example, in [6] the authors use a distance measure to define a total order and define multiple MM techniques, although not the HMT. Another approach [8] defines a number of MM techniques, including the HMT using a similarity measure. Work has also been carried out to extend Barat's greyscale HMT (reviewed in [9]) into the colour domain, specifically in the CIELAB colour space [10].

Our proposed algorithm extends the HMT to colour images using distance measures to determine the fit of an SE to a query image. Our proposal differs from the work discussed previously in that the colour image data can be totally ordered using a reduced ordering approach [11] and thus MM operators, including the HMT, can be defined. Robustness to noise is included by utilising and extending the notion of a Percentage Occupancy HMT (POHMT) [12] to the colour domain.

II. BACKGROUND

A. Fundamental Morphological Operators

The fundamental operators of MM are erosion and dilation. The erosion by an SE, S , (ε_S) is the locus of all points in a binary image, $x \in X$, where the SE translated to point x , S_x , is a subset of X :

$$\varepsilon_S(X) = \{x | S_x \subseteq X\} \quad (1)$$

The dilation (δ_S), known as the dual of the erosion, is the locus of all points where the intersection of S and X is non-empty,

$$\delta_S(X) = \{x | S_x \cap X \neq \emptyset\} \quad (2)$$

B. The Hit-or-Miss Transform

The HMT differs from both erosion and dilation in the fact that it requires two non-intersecting SEs. The foreground SE (S_{FG}) is designed to fit the object under study in an image, X , and the background SE (S_{BG}) is designed to fit the




Fig. 1: **a)** Binary test image made up of multiple sizes of shapes. **b)** Binary composite SE. **c)** Output of the HMT of (a) using (b).

background of the object. In the binary case, S_{BG} will match the complement (X^c) of image X . The HMT of X using a composite SE, S is defined as:

$$\text{HMT}_S(X) = \{x | (S_{FG})_x \subseteq X, (S_{BG})_x \subseteq X^c\} \quad (3)$$

The notation in (3) can be thought of as the intersection of two separate erosions. Using (1), an alternative representation of the binary HMT can be defined:

$$\text{HMT}_S(X) = \varepsilon_{S_{FG}}(X) \cap \varepsilon_{S_{BG}}(X^c) \quad (4)$$

An example of the HMT on a synthetic binary image is shown in Fig. 1. A white square is used as the foreground SE and a hollow square as the background SE. This configuration is designed to detect 3×3 and 4×4 white squares on a black background.

C. Morphological Greyscale Extensions

The methods defined in traditional, binary, MM can be applied to greyscale images. The erosion of a greyscale image, g , by an SE, S , at point x in a given image is the minimum value of the image pixels coincident with S as defined in (5):

$$[\varepsilon_S(g)](x) = \min_{s \in S} g(x + s) \quad (5)$$

Conversely, the dilation can be thought of as the maximum value within the region defined by S :

$$[\delta_S(g)](x) = \max_{s \in S} g(x + s) \quad (6)$$

The HMT is not an increasing transform meaning its extension to greyscale is not a straightforward endeavour [4], [5]. There are various methods described in the literature which aim to address this, most notably the greyscale HMTs developed by Soille, Ronse, Barrat and Ducottet and Khosravi and Schafer, each of which is reviewed in [9]. An example of a greyscale HMT on a synthetic greyscale image can be seen in Fig. 2. Similarly to the binary image test shown in Fig. 1, the SE has been designed to detect 3×3 and 4×4 white squares on a non-white background.

D. The Percentage Occupancy HMT (POHMT)

While it performs well in many applications, the HMT is very susceptible to error given a noisy image. This is due to the underlying assumption that both SEs will fit the object under inspection perfectly, meaning just one noisy pixel in the foreground or background can thwart this operation.

There are multiple ways of relaxing the hard constraints of the HMT to detect these imperfect objects including the notion




Fig. 2: **a)** Greyscale test image made up of shapes with varying intensity and size. **b)** Greyscale composite SE. **c)** HMT of (a) using (b).

of a POHMT [12] which allows for a percentage error in the matching of an SE to an image. These Percentage Occupancy (PO) measures can be implemented by replacing the hard limits of erosions (minimum filters) and dilations (maximum filters) with rank k filters [12]. The output of a rank order filter with rank k within the region of an SE, S , centred at point x in an image I is calculated using the following:

$$[\zeta_{S,k}(I)](x) = k\text{th order statistic}_{s \in S} \{I(x + s)\} \quad (7)$$

A rank order filter with rank k will return the k^{th} smallest element, or k^{th} order statistic, of an ordered set of data. When $k = 1$, an erosion takes place and when $k = n$ a dilation takes place, where n is the sample size of the input set. The desired rank given some PO value, P , is denoted as k_P and allows for the unbiased specification of rank regardless of the size and shape of an SE, S , and the region it describes within an image [12]:

$$k_p = \text{floor} \left(\frac{P}{100} \times \text{card}(S) \right) \quad (8)$$

The POHMT can then be defined in terms of the values of two rank order filters and the desired value of k given some value for P , k_p . The first of these operations is equivalent to a relaxed erosion of the foreground, $[\zeta_{S_{FG},k_{100-p}}(I)](x)$, and the second is the relaxed background dilation, $[\zeta_{S_{BG},k_p}(I)](x)$, such that:

$$[\text{POHMT}_S(g)](x) = \begin{cases} 1 & \text{if } [\zeta_{S_{FG},k_{100-p}}(I)](x) > [\zeta_{S_{BG},k_p}(I)](x) \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

In (9) the output is the set of points in the image where the intensity at rank k_{100-p} of S_{FG} , when coincident with some image pixel x , is greater than the intensity at rank k_p of S_{BG} at that same point, x , in I . An example of the POHMT is shown in Fig. 3 where the images from Fig. 1 and 2 were input into a POHMT with $P = 90\%$ and the result is a more accurate object detection in the presence of noisy pixels. This can be seen in the results shown in Fig. 3 as the squares with non-white centre pixels are detected despite the non-ideal shape or intensity.

E. Multivariate Morphology and Ordering

While the HMT has been extended to greyscale and made robust to noise, one of the main obstacles in extending morphology to incorporate multivariate data is the necessity to define a vectorial ordering. As discussed in [11], there




Fig. 3: **a)** Binary test image from Fig. 1a **b)** POHMT of (a) using Fig 1b and $P = 90\%$. **c)** Greyscale test image from Fig. 1a **d)** POHMT of (c) using Fig 2b and $P = 90\%$.

exist multiple ways of ordering multivariate data, including: marginal (m-ordering), reduced (r-ordering), partial (p-ordering) and conditional ordering (c-ordering). The r-ordering approach, which our proposed algorithm is based on, seeks to reduce the dimensionality of multivariate data to a scalar value by measuring its similarity to some other observation. R-ordering is commonly carried out using various distance measures or projections into reduced spaces. The combination of an r-ordering scheme and the robustness to noise provided by a POHMT gives rise to our proposed new algorithm - the Multivariate Distance Percentage Occupancy Hit-or-Miss Transform (MDPOHMT).

III. MULTIVARIATE DISTANCE POHMT (MDPOHMT)

The MDPOHMT uses r-ordering based on point-wise distances between two multivariate SEs (S_{FG}, S_{BG}) and an image or, more generally, any multivariate set of data.

We can think of each colour or multivariate observation as a single point in multivariate space. By designing SEs of the same dimensionality and translating them to some point x in a given image, b , we can compare each element in the SE with its corresponding image pixel at each $x \in b$ to compute a measure of distance d_{FG} for foreground and d_{BG} for background.

The simplest similarity metric and the one used here is the Euclidean Distance in the RGB colour space. The definition for an n -dimensional Euclidean Distance, d , between two points in space a and b is defined by:

$$d(a, b) = \sqrt{(b_1 - a_1)^2 + (b_2 - a_2)^2 + \dots + (b_n - a_n)^2} \quad (10)$$

Although the Euclidean Distance is used for representation purposes, our algorithm can be easily generalised to incorporate alternative distance metrics and other colour and multivariate representations.

The notion that the SEs match, or fit, the image when the foreground matches the image and the background SE matches its complement is retained from both the standard HMT and POHMT. However, as these measures are now defined in terms of distances, the polarity of standard definitions for erosions and dilations are reversed. With the erosion pertaining to a maximal distance and the dilation as a minimal distance as opposed to minimal and maximal intensity, respectively, for the traditional definitions. Each of the foreground and background point-wise distances d_{FG} and d_{BG} can be collated into a single *distance image*, D , which can be based on an erosion, dilation or rank k filter giving the following definitions, respectively:

$$[\varepsilon_S(D)](x) = \min_{s \in S} d(I(x+s), S(s)) \quad (11)$$

$$[\delta_S(D)](x) = \max_{s \in S} d(I(x+s), S(s)) \quad (12)$$

$$[\xi_{S,k}(D)](x) = k\text{th order statistic}\{d(I(x+s), S(s))\}_{s \in S} \quad (13)$$

By combining the distance based rank order filter from (13) and the POHMT from (9), a definition for the proposed MDPOHMT for a multivariate image, I , composite multivariate SE centred at x , S , and desired PO, p , can be formalised as follows:

$$\text{MDPOHMT}_S(I)](x) = \begin{cases} 1 & \text{if } [\xi_{S_{FG}, k_{100-p}}(D)](x) < [\xi_{S_{BG}, k_p}(D)](x) \\ 0, & \text{otherwise} \end{cases} \quad (14)$$

The output of (14) will be the points where the distance between the foreground SE and the image is less than that between the background SE and the image. This represents a good “fit” and is in keeping with the fundamental concept of the HMT. A demonstration of applying our proposed MDPOHMT to a synthetic colour image using three different colour composite SEs is shown in Fig. 4 using both full rank and rank k operations.

IV. RESULTS

The proposed algorithm introduces a new colour HMT which also allows the transform to compensate for severe noise in an image in order to accurately detect objects of interest. Thus far we have used only synthetic, ideal test images for demonstration of the proposed techniques and the key concepts. In this section we apply our proposed algorithm to




Fig. 4: **a)** Colour test image made up of Red, Green and Blue shapes on an orange background. **b)** Red composite SE. **c)** Green composite SE. **d)** Blue composite SE. **e)** Output of the HMT of (a) using (b). **f)** Output of the HMT of (a) using (c). **g)** Output of the HMT of (a) using (d). **h)** Output of the MDPOHMT of (a) using (b) with $P = 90\%$. **i)** Output of the MDPOHMT of (a) using (c) with $P = 90\%$. **j)** Output of the MDPOHMT of (a) using (d) with $P = 90\%$.




Fig. 5: **a)** Original colour image of a traffic light. **b)** Circular non-flat SE used to find the red light. **c)** Circular non-flat SE used to find the amber light. **d)** Results of taking the MDPOHMT of (a) using (b) with $P = 100\%$. **e)** Results of taking the MDPOHMT of (a) using (c) with $P = 100\%$.

real images first using various forms of SE and then applying multiple noise models to an image at different density.

Fig. 5 shows the first experiment undertaken, the aim of which was to detect the red and amber lights present in the image of a traffic light shown in Fig. 5a. The red and amber non-flat SEs used are shown in Fig. 5b and 5c respectively. They are comprised of circular regions cropped directly from the image which form the foreground component of each SE. The background SEs are both hollow squares with colours which are the average of those contained in each of the foreground SEs. Taking the MDPOHMT of the image using each of these SEs in turn yields the images shown in Fig. 5d and 5e where there are single points at which the distance between each SE and the image is exactly zero - denoted as a blue crosses superimposed onto the original image. This is to be expected as, in this case, the SEs are regions taken directly from the image under test.

The previous experiment was repeated with two flat (single-coloured), circular SEs rather than using the cropped regions from the image itself. The colour of the red SE was chosen as the average of the red non-flat SE used previously and likewise for the amber SE. The flat red and amber SEs are shown in Fig. 6b and Fig. 6c respectively, with the results of the MDPOHMT shown in Fig. 6d and 6e for each SE. The results show a small set of points where each SE matched the image well, again detecting the presence of round red and amber objects. In each case, the PO measure was decreased to allow for the non-exact match between the image and the SE, caused by averaging.

In order to assess the improvement of the MDPOHMT by incorporating colour information the previous experiment was repeated after transforming the colour image into a greyscale image and using the unconstrained HMT (UHMT) [5] as a comparison. The greyscale conversion is achieved by discarding the hue and saturation information present in the image while keeping only the luminance values, resulting in a single channel greyscale image. The UHMT fails to distinguish between the two traffic lights due to the similarity in greyscale intensity between the two as seen in Figure 7. The greyscale image of the traffic lights is shown in Figure 7a with the red




Fig. 6: **a)** Original colour image of a traffic light. **b)** Circular SE used to find the red light. **c)** Circular SE used to find the amber light. **d)** Results of taking the MDPOHMT of (a) using (b) with $P = 90\%$. **e)** Results of taking the MDPOHMT of (a) using (c) with $P = 90\%$.

and amber greyscale equivalent SEs shown in Figures 7b and 7c respectively. The results of the UHMT using each SE are shown in Figures 7d and 7e. The results display a small set of points where each SE matched both illuminated traffic lights in each test. As mentioned previously, this is due to the small difference in intensity compared with a comparatively large difference in colour and distance.

In order to assess the algorithm's robustness to noise, multiple correlated noise models [13] were applied to the image from the previous tests. Adding varying strengths of correlated impulsive, impulsive bipolar and Gaussian noise to the image allowed for the validation of the PO component of the algorithm. Specifying a probability gives the images shown in Fig. 8, using the cropped red and amber SEs, along with varying ranks, yields the same results even in the presence of large amounts of noise. By varying the PO tolerance, P , in our algorithm the transform identified our region of interest regardless of the type or intensity of noise in each image.

Further experiments were undertaken with the aim of using our algorithm to detect particular objects from satellite imagery. One of these such experiments is shown in Fig. 9 where swimming pools in the Beverly Hills area of Los Angeles were

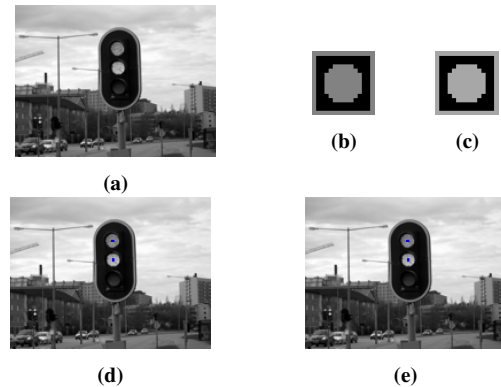


Fig. 7: **a)** Original image of a traffic light converted to greyscale. **b)** Circular SE used to find the red light. **c)** Circular SE used to find the amber light. **d)** Results of taking the greyscale UHMT of (a) using (b). **e)** Results of taking the greyscale UHMT of (a) using (c).




Fig. 8: **a)** Image with Gaussian noise - $\sigma = 0.5, \rho = 10$. **b)** Detected red light with $P = 70\%$. **c)** Detected amber light $P = 70\%$. **d)** Image with 10% impulsive noise. **e)** Detected red light with $P = 90\%$. **f)** Detected amber light with $P = 90\%$. **g)** Image with 30% impulsive noise. **h)** Detected red light with $P = 70\%$. **i)** Detected amber light with $P = 70\%$. **j)** Image with 10% bipolar impulsive noise. **k)** Detected red light with $P = 90\%$. **l)** Detected amber light with $P = 90\%$. **m)** Image with 30% bipolar impulsive noise. **n)** Detected red light with $P = 70\%$. **o)** Detected amber light with $P = 70\%$.


detected using the MDPOHMT. Fig. 9a shows the cropped image with various objects of interest in the scene. Fig. 9b shows the flat blue SE used to produce the results shown in Fig. 9c where detected regions are highlighted.

V. CONCLUSIONS

We have presented a novel algorithm for extending the Hit-or-Miss transform to colour images whilst also incorporating techniques which offer robustness to noise. Our algorithm is an extension of the POHMT presented in [12] that overcomes the multivariate ordering problem by incorporating r-ordering techniques described in [11]. The comparison between the greyscale UHMT and our MDPOHMT has also been investigated and we offer improvements that working in colour spaces can bring. We have tested the robustness of our new algorithm to noise by using various forms of impulsive and Gaussian noise models with increasing intensity. In each case the algorithm performs well, including cases of high noise. While this algorithm is demonstrated using a Euclidean distance measure in the RGB colour space, this can be changed in order to support any distance metric or colour space.

ACKNOWLEDGEMENTS

The authors would like to thank the EPSRC as well as BAE Systems for the financial support for this work.



(a) Imagery ©2018 Google, Map data ©2018 Google



(c)

Fig. 9: **a)** Satellite imagery of the Beverly Hills Area (34.077566, -118.407321). **b)** Square SE used to detect swimming pools in the image. **c)** Detected swimming pools with $P = 70\%$ overlaid onto the original.

REFERENCES

- [1] G. Matheron, *Random sets and integral geometry*. New York, NY: Wiley, 1975.
- [2] J. Serra, *Image analysis and mathematical morphology*. New York, NY: Academic Press, 1982, vol. 1.
- [3] H. Heijmans, *Morphological Image Operators*. Academic Press, 1994.
- [4] B. Perret, S. Lefèvre, and C. Collet, "A robust hit-or-miss transform for template matching applied to very noisy astronomical images," *Pattern Recognition*, vol. 42, no. 11, pp. 2470–2480, 2009.
- [5] P. Soille, *Morphological image analysis: principles and applications*, 2nd ed. Springer Science & Business Media, 2013.
- [6] J. Angulo, "Morphological colour operators in totally ordered lattices based on distances: Application to image filtering, enhancement and analysis," *Computer Vision and Image Understanding*, vol. 107, no. 1, pp. 56–73, 2007.
- [7] S. Lefèvre, E. Aptoula, B. Perret, and J. Weber, "Morphological template matching in color images," in *Advances in Low-Level Color Image Processing*. Springer, 2014, pp. 241–277.
- [8] C. W. Yeh and D. Pycocock, "Similarity colour morphology," in *5th Computer Science and Electronic Engineering Conference*, 2013, pp. 71–76.
- [9] P. Murray and S. Marshall, "A review of recent advances in the hit-or-miss transform," in *Advances in Imaging and Electron Physics*, P. W. Hawkes, Ed. Elsevier, 2013, vol. 175, ch. 5, pp. 221–282.
- [10] A. Ledoux, N. Richard, and A.-S. Capelle-Laizé, "Color hit-or-miss transform (CMOMP)," in *Proceedings of the 20th European Signal Processing Conference*. IEEE, 2012, pp. 2248–2252.
- [11] V. Barnett, "The ordering of multivariate data," *Journal of the Royal Statistical Society*, vol. 139, no. 3, pp. 318–355, 1976.
- [12] P. Murray, S. Marshall, and E. Bullinger, "The percentage occupancy hit or miss transform," in *Proceedings of the 7th European Signal Processing Conference*. IEEE, 2009, pp. 253–257.
- [13] A. N. Evans and X. U. Liu, "A morphological gradient approach to color edge detection," *IEEE Transactions on Image Processing*, vol. 15, no. 6, pp. 1454–1463, 2006.