

A New Design Tool for Feature Extraction in Noisy Images based on grayscale Hit-or-Miss Transforms

P. Murray* and S. Marshall

Abstract— The Hit-or-Miss transform (HMT) is a well known morphological transform capable of identifying features in digital images. When image features contain noise, texture or some other distortion, the HMT may fail. Various researchers have extended the HMT in different ways to make it more robust to noise. The most successful, and most recent extensions of the HMT for noise robustness, use rank order operators in place of standard morphological erosions and dilations. A major issue with the proposed methods is that no technique is provided for calculating the parameters that are introduced to generalize the HMT, and, in most cases, these parameters are determined empirically. We present here, a new conceptual interpretation of the HMT which uses a percentage occupancy (PO) function to implement the erosion and dilation operators in a single pass of the image. Further, we present a novel design tool, derived from this PO function that can be used to determine the only parameter for our routine and for other generalizations of the HMT proposed in the literature. We demonstrate the power of our technique using a set of very noisy images and draw a comparison between our method and the most recent extensions of the HMT.

Index Terms— Machine vision, Morphological operations, Object recognition, Segmentation

I. INTRODUCTION

Mathematical Morphology, first introduced by Matheron [1] and Serra [2] and later extended by Heijmans [3], provides an extremely powerful set of tools for image processing. Among these is the HMT [2] and [3], which is capable of identifying groups of connected pixels that comply with certain geometric properties. If there is noise in a given image, or if image features are extremely textured, the standard HMT will fail to detect objects which are of interest.

For the processing of binary images, the HMT is well

defined, [2]–[7], and involves searching an image for locations where predefined templates simultaneously fit the image. The templates, known as structuring elements (SEs) in morphology, are designed to match the geometry of objects of interest in the foreground and background of the image. If SEs are designed to closely match the geometry of the image features that are of interest, just one noisy pixel in the foreground or background can cause this transform to fail since the SEs will no longer fit as a result of the noise. To overcome these issues, various authors have proposed techniques that involve some pre-processing of the image or some modification of the SEs or of the HMT itself.

In [8], Zhao and Daut present a technique for the detection of imperfect shapes where the SEs are designed by first smoothing the original image using a morphological opening, before using the boundary of these smoothed shapes as SEs. In [9], the same authors present a technique which uses the skeletons of both the object to be recognized and its complement as SEs.

In [10], Bloomberg and Maragos introduce a Rank Hit-Miss Transform using rank order filters in place of erosions to improve the performance of the HMT and its robustness to noise. The same authors present a Blur Hit-Miss transform in [11] which uses “blur SEs” to dilate the foreground and background of the image prior to applying the respective erosions of the HMT. This helps remove noise and makes it easier for the SEs to match patterns by slightly modifying the geometry of features in a given image.

Various researchers have defined methods to extend the HMT for the processing of grayscale images and recently a unified theory for calculating a grayscale HMT has been presented in [12]. We review in this paper, the most prevalent extensions of the HMT for grayscale images and show the equivalences between these extensions. Despite the extensions of the HMT for processing grayscale images, the issues that cause the HMT to fail in the presence of noise in binary images, have the same effect in the grayscale case. Various techniques attempt to generalize the HMT for feature recognition in noisy images. In [13], Khosravi and Schafer present a formal definition of the grayscale HMT and analyze its performance in the presence of Gaussian and salt and pepper noise. To improve the performance of the HMT in noise, the authors generalize the HMT using rank order operators and also subsample the SEs used for

Manuscript received April 16, 2010. Paul Murray is with the Department of Electronic and Electrical Engineering, University of Strathclyde, Glasgow, G1 1XW, UK (phone: +44 141 548 2205; fax: +44 141 552 5781; e-mail: paul.murray@strath.ac.uk). Stephen Marshall, is with the Department of Electronic and Electrical Engineering University of Strathclyde, Glasgow, G1 1XW, UK (e-mail: [s.marshall@eee.strath.ac.uk]).
EDICS TEC BIP

template matching.

More recently, in [14], Perret et al present a Fuzzy Hit-or-Miss Transform which they use to detect features in very noisy astronomical images. Their technique uses rank order operations and a large set of SEs that are generated using a mathematical model. The authors also highlight some techniques that could be used to improve the robustness of other grayscale HMTs when operating in noise. We too discuss these extensions and show how their parameters may be calculated using our design tool.

Each of the aforementioned methods and extensions aim to improve the performance of the HMT such that it is more robust for feature detection in noisy images. Rank order operators feature heavily in this work, however, there appears to be a conceptual gap in that the authors fail to provide a method by which it is possible to select the appropriate rank or threshold for applying rank order filters in this way. This paper presents a Percentage Occupancy Hit-or-Miss transform (POHMT) which allows partial fitting of SEs in a similar fashion to the partial fitting allowed by rank order filters. The difference here however, is that we also provide a technique to accurately determine the rank of the filter that must be used for the detection of objects of interest. Furthermore, as a direct result of the plots that we use as a design tool, we show in this paper, how we can make the POHMT operate as a discriminatory filter which allows objects to be selectively marked or discarded by the transform.

II. MATHEMATICAL MORPHOLOGY

We first recall here the definition of the binary HMT before reviewing the work of various researchers who have extended the HMT such that it can be applied to grayscale images. We then present a novel, conceptual definition of the HMT in terms of SE occupancy and use this to explain the inability of the HMT to function where images are distorted by noise or when features exhibit internal texture.

A. The Hit-or-Miss-Transform

The HMT of a binary image X is the intersection of an erosion of X and erosion of the complement of X by a complementary pair of SEs B_{FG} and B_{BG} respectively where X , B_{FG} and B_{BG} are sets in 2D space, $E = \mathbb{Z}^2$. B_{FG} and B_{BG} are defined relative to a common origin in E where the composite SE $B = B_{FG} \cup B_{BG}$ and $B_{FG} \cap B_{BG} = \emptyset$. That is,

$$HMT_B(X) = \{ x \in E \mid (B_{FG})_x \subseteq X, (B_{BG})_x \subseteq X^c \} \quad (1)$$

where $(B)_x = \{b + x \mid b \in B\}$ and a feature is detected by the HMT if there is at least one point $x \in E$ such that the foreground SE $(B_{FG})_x$ is included in X whilst the background SE $(B_{BG})_x$ is simultaneously included in its complement, $X^c = E \setminus X$, see [2]-[7]. The HMT returns a “marker” consisting of single pixels or groups of pixels indicating the presence and locations of the objects that have the features specified by B . To recover the complete object after detection, an opening by reconstruction [7] may be applied.

Unlike other morphological transforms, extending the HMT for grayscale images is not a trivial task since the HMT is not an increasing transform [7], [12], [14]. Various researchers have proposed extensions of the binary HMT such that it can be applied to grayscale images. A thorough review of these techniques, as well as a unified theory for calculating the grayscale HMT is given in [12]. We present here a brief summary of the main extensions of the HMT, and further details can be found in [12] and [14].

To remain consistent with the literature, we define the notation used throughout this paper. Let E represent a two dimensional digital space ($E = \mathbb{Z}^2$) and T^E be the set of all graylevel functions from a subspace of E to T where $T = \mathbb{R} \cup \{+\infty, -\infty\}$ or $T = \mathbb{Z} \cup \{+\infty, -\infty\}$ such that T is a complete lattice with respect to the order “ \leq ”. Let $I \in T^E$, denote a grayscale image, and $B \in T^E$ denote a grayscale SE. We can then define the grayscale erosion and dilation of image I by the SE, B , where the erosion of image I is denoted, $(I \ominus B)$, and the dilation of I is denoted, $(I \oplus B)$: for all $x \in E$,

$$(I \ominus B_{FG})(x) = \min_{b_{FG} \in \text{supp}(B_{FG})} (I(x + b_{FG}) - B_{FG}(b_{FG})), \quad (2)$$

$$(I \oplus B_{FG})(x) = \max_{b_{FG} \in \text{supp}(B_{FG})} (I(x - b_{FG}) + B_{FG}(b_{FG})). \quad (3)$$

The grayscale HMT uses a pair of foreground and background SEs, $B_{FG}, B_{BG} \in T^E$, where the grayscale HMT defined by Ronse in [15], (denoted RHMT in [14]) using our notation becomes,

$$RHMT_{[B_{FG}, B_{BG}]}(I)(x) = \begin{cases} (I \ominus B_{FG})(x) & \text{if } (I \ominus B_{FG})(x) \geq (I \oplus B_{BG}^*)(x) \neq +\infty \\ -\infty & \text{otherwise} \end{cases} \quad (4)$$

where $B_{BG}^* : E \mapsto T$ and $B_{BG}^*(b) \rightarrow -B_{BG}(-b)$, i.e. the dual of B_{BG} .

In [7] and [16], Soille defines an unconstrained HMT (UHMT) using flat SEs, which $\forall x \in E$, returns the number of cross sections of a grayscale image, I , where $(B_{FG})_x$ fits the cross section, $CS_t(I)$, and $(B_{BG})_x$ simultaneously fits the complement of this cross section, $\complement CS_t(I)$,

$$UHMT_{[B_{FG}, B_{BG}]}(I)(x) = \text{card} \{ t \mid (B_{FG})_x \subseteq CS_t(I), (B_{BG})_x \subseteq \complement CS_t(I) \}. \quad (5)$$

In [12], the authors extend the UHMT as written in (5), to allow grayscale SEs, such that,

$$UHMT_{[B_{FG}, B_{BG}]}(I)(x) = \max \{ (I \ominus B_{FG})(x) - (I \oplus B_{BG}^*)(x), 0 \}. \quad (6)$$

By comparing (4) and (6), the similarity between the RHMT and the UHMT is clear. Both of these extensions mark features in a grayscale image when the result of the foreground erosion is greater than or equal to that of the background erosion. The output of the RHMT is an image

containing the result of foreground erosions when this condition is satisfied. Soille's UHMT produces a graylevel image where the intensity of each pixel indicates the number of cross sections where both SEs fit the image i.e. the difference between the foreground and background erosions.

In [14], the similarity between Soille's and Barrat's grayscale HMT defined in [17] is shown, where Barrat et al.'s HMT (denoted $BHMT$ in [14]) is written,

$$BHMT_{[B_{FG}, B_{BG}]}(I)(x) = (I \oplus B_{BG}^*)(x) - (I \ominus B_{FG})(x). \quad (7)$$

Clearly, this grayscale HMT closely resembles that of Soille's given in [12]. The two differ however in that Soille's HMT returns the difference between the *foreground erosion* and the *background erosion* where Barrat et al. return the difference between the *background erosion* and the *foreground erosion*. In Soille's HMT, the higher the output value, the better the fit of the SEs. The opposite is true for the BHMT, and the equivalence of these two operators is shown in [14] as,

$$UHMT_{[B_{FG}, B_{BG}]}(I)(x) = -\min\{BHMT(I)(x), 0\} \quad (8)$$

Khosravi and Schafer, in [10], present their grayscale HMT which requires only one SE, B_{FG} , as,

$$KHMT_{[B_{FG}]}(I)(x) = (I \ominus B_{FG})(x) - (-I \ominus (-B_{FG}))(x). \quad (9)$$

This definition is discussed in [14], where it is shown that the KHMT is in fact equivalent to the BHMT as shown in (10) and (11),

$$KHMT_{[B_{FG}]}(I)(x) = (I \ominus B_{FG})(x) - (I \oplus (B_{FG}^*))(x), \quad (10)$$

$$= -BHMT_{[B_{FG}, B_{FG}]}(I)(x). \quad (11)$$

Each of the proposed methods can be used to extract the features of a grayscale scale image that match the geometry of both B_{FG} and B_{BG} , however, all of these techniques fail in the presence of noise unless further modifications are made. All of these techniques are discussed further in [14] and a thorough review of the grayscale HMTs proposed by Ronse and Soille is given in [12]. In this paper we extend the definition of the grayscale HMT given by Soille in [7] since it is consistent with our conceptual description of the operation of the HMT.

B. A New Conceptual View of the Hit-or-Miss-Transform

Interpreting a grayscale image as a topographic surface allows the HMT to be considered as a translation of the two SEs in this 3D image space searching for places where they simultaneously fit the image to detect objects. In the various definitions of the grayscale HMTs given in Section II.A, the SEs are translated in the 2D space, and standard morphological operations are used to probe the image where these operators interact with the graylevel at each

pixel as shown in (2) and (3). In the 3D space, we must still translate the SEs to all points $x \in E$ in the 2D space, however, we may also consider the translation of SEs in the vertical direction. From an implementation perspective we may first translate the SE by a vector x such that the origin of the SE is coincident with an image pixel x , $\forall x \in E$. Then, at each point $x \in E$, the concept of a vertical translation of the SE may be implemented by interrogating in some way, (dependent on the operation) the image pixels that are coincident with the elements of the SE, $\forall t \in T$.

The grayscale erosion of the image foreground can be described conceptually as a process of translating B_{FG} to a point $x \in E$ and raising the SE to the highest level t for which it is entirely beneath or fully occupied by the signal. By this interpretation, the erosion of an image at any point $x \in E$ is equivalent to calculating the maximum level, t , for which the foreground SE, B_{FG} , is fully occupied in the image. For the foreground erosion, the SE is fully occupied if the intensity of all image pixels that are coincident with the elements of B_{FG} is greater than or equal to t . We denote by $l_{FG}(x)$, the maximum level, t , for which the SE is fully occupied when its origin is at any point $x \in E$,

$$l_{FG}(x) = \max_{b_{FG} \in B_{FG}} \{t \mid \forall b_{FG} \in B, I(x + b_{FG}) \geq t\}. \quad (12)$$

By definition, the HMT uses a foreground erosion to match patterns from below the topographic surface and a background erosion to match patterns from above. The background erosion can be described by a similar process of translating B_{BG} to a point $x \in E$ and lowering the SE to the lowest level t for which it is entirely above or fully occupied in the image. We therefore define the background erosion at a point $x \in E$ as the minimum level, t , for which the background SE, B_{BG} , is fully occupied in the image. For the background erosion, the SE is fully occupied if the intensity of all the image pixels that are coincident with the elements of B_{BG} is less than t . We denote by $l_{BG}(x)$, the minimum level, t , for which the SE is fully occupied when its origin is centered at a point $x \in E$.

$$l_{BG}(x) = \min_{b_{BG} \in B_{BG}} \{t \mid \forall b_{BG} \in B_{BG}, I(x + b_{BG}) < t\} \quad (13)$$

This description of the HMT resembles the one given by Soille in [7] where he states that the HMT is equivalent to the number of intersections of the intervals $[0, t_{FG}]$ and $[t_{BG}, t_{max}]$. Soille defines that t_{FG} is the highest level where B_{FG} fits the foreground, t_{BG} is the lowest level at which B_{BG} fits the background, and t_{max} is the highest intensity in the image as determined by the bitdepth. In our definition of the HMT, l_{FG} and l_{BG} are equivalent to t_{FG} and t_{BG} , however, our HMT differs significantly in that it may be calculated in a single pass of the image with erosion and dilation implemented simultaneously using a composite SE. In (6), the definition of UHMT states that for a point to be marked by this grayscale HMT, the result of I eroded by B_{FG} must

be greater than or equal to that of I dilated by B_{BG}^* . By analogy, we define that our HMT will mark an object as a “hit”, iff $I_{FG}(x) \geq I_{BG}(x)$. We express our grayscale HMT in terms of $I_{FG}(x)$ and $I_{BG}(x)$, $\forall x \in E$,

$$[HMT_B(I)](x) = \begin{cases} 2^n - 1 & \text{if } I_{FG}(x) \geq I_{BG}(x) \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

The concept of this HMT is consistent with the grayscale HMTs defined by Ronse and Soille, however, in our definition the result is a binary marker containing the location of any objects that have been detected. This is illustrated in Fig.1.




Fig.1 Grayscale HMT operating on a topographic surface. (a) A synthetic grayscale image. (b) Topographic representation of (a) with the HMT detecting the middle feature when $I_{FG}(x) > I_{BG}(x)$ (ticks represent objects that are detected and crosses denote points that have not been marked).

By considering the HMT in terms of SE occupancy, the traditional implementation of the HMT (which requires two erosions) may be simplified by combining B_{FG} and B_{BG} into a unified, composite, SE (B , shown in Fig.2a). The composite SE, B , is translated to each point x in the image, $(B)_x = \{b + x | b \in B\}$. A point $x \in E$ is marked in the result if there exists a level, $t \in T$, which for all of the elements, $b_{FG} \in (B_{FG})_x$, $t \leq I(b_{FG})$, while, simultaneously, for all of the elements, $b_{BG} \in (B_{BG})_x$, $t > I(b_{BG})$, i.e.

$$[HMT_B(I)](x) = \begin{cases} 2^n - 1 & \text{if } \exists t \in T, [\forall b_{FG} \in B_{FG} | t \leq I(x + b_{FG}) \text{ and } \forall b_{BG} \in B_{BG} | t > I(x + b_{BG})] \\ 0 & \text{otherwise} \end{cases} \quad (15)$$

This technique is illustrated in Fig.2 where a combined, composite SE is shown in Fig.2(a) and an example of this SE discriminating between similar objects is shown in Fig.2(b).




Fig.2 The HMT implemented using a composite SE. (a) The composite SE where elements of B_{FG} are shown in dark gray and elements of B_{BG} shown in light gray. (b) The SE searching for places where it is fully occupied in the image.

Our definition of the grayscale HMT allows this operator to be calculated in one pass of the image instead of the common two pass method followed by an intersection, summation or comparison of the two resultant images. As a

result, the transform is faster and simpler than the standard method.

C. The Hit-or-Miss-Transform in noise

Fundamentally, a “hit” i.e. an object which is detected and marked by the HMT, is one which satisfies the conditions stated in Section II.B. This strict definition of the HMT requires that the composite SE must be fully occupied in both the foreground and background of the image for successful detection of an image feature. Often, when features are distorted by noise or if image features contain a large amount of internal texture, it is not possible for B to be fully occupied in the image even if its geometry matches that of the feature. This causes the HMT to miss objects that should be detected as illustrated in Fig.3.




Fig. 3 Operation of the HMT in noise (a) Composite SE that can be used to detect a circle (b) Fully occupied, composite SE detecting the object of interest (c) Composite SE cannot be fully occupied due to noise

Fig. 3(a) shows a composite SE which can be used to detect circular objects using the standard HMT. B_{FG} is a solid disk and B_{BG} is a solid ring. The black line in Fig.3(b) represents a noise free shape which is to be detected using the SE shown in Fig.3(a). In this case, the elements of B corresponding to B_{FG} are fully occupied by the shape and the elements corresponding to B_{BG} are simultaneously fully occupied by its background at all levels t , until t is greater than the intensity of the shape. This feature and any feature that has not been corrupted by noise and whose dimensions are greater than B_{FG} and less than that of B_{BG} will be detected by the HMT when using this SE.

In the case that an object of interest, its edges, or both are corrupted by noise, the elements of B corresponding to B_{FG} and B_{BG} may never be simultaneously fully occupied by the object. This is illustrated in Fig.3(c) where both the foreground and background regions of the object shown in Fig.3(b) have been perturbed by noise. Since some of the foreground pixels within the object are at a level t that is lower than the level of its noisy background, there is no level t at which B_{FG} and B_{BG} can be simultaneously fully occupied. As a result, the HMT will fail to recognize this feature as an object of interest. For the same reasons stated here, objects which have internal texture, such as biological cells, may fail to be detected by the standard HMT.

III. A PERCENTAGE OCCUPANCY HIT-OR-MISS-TRANSFORM

In the standard HMT the foreground structuring element, B_{FG} , must fit entirely within the foreground of the object and the background SE, B_{BG} , must fit entirely within the background surrounding the object. In other words, they must be fully occupied by the foreground and background respectively. Any noise, even just one pixel, in either the foreground or background of the object can prevent an

otherwise legitimate hit occurring.

The idea behind the POHMT is to make the detection process less sensitive to moderate amounts of noise (or texture) in the image. We propose relaxing the constraint that the SEs must be 100% occupied and allow them to be only partially occupied and to still record a ‘hit’. Attempts at relaxing these strict constraints have been proposed in [7], [14], [19] and [20], however, in this paper we introduce a new design tool in order to set the appropriate level of partial (or percentage) occupancy. This tool may also be used in order to set the parameters for equivalent methods in an objective rather than an empirical way.

This paper builds upon a Percentage Occupancy Hit or Miss Transform (POHMT) introduced in [18], which allows a percentage of the SE to be “punctured” by noise or texture in a signal and still detect a “hit”. This section first defines a method that can be used to calculate the extent to which B_{FG} and B_{BG} are occupied by a signal for all levels $t \in T$ when their origin is coincident with any point $x \in E$. The POHMT is then defined using this approach, to allow objects to be detected in places where the SE is only partially occupied by the signal.

A. Calculating the occupancy of Structuring Elements

We have shown that a grayscale HMT, using flat SEs, can be implemented with a single, composite SE, which searches the image to identify places where its foreground and background elements are simultaneously 100% occupied. By designing the SE to match the geometry of an object in both the foreground and background and measuring the extent to which the object occupies the SE when coincident with an image feature, we can estimate how far we must relax the 100% occupancy requirement in order to detect this object in the presence of noise.

To facilitate this explanation and its comprehension, we first consider B_{FG} and B_{BG} separately before showing how the two can be combined (as in Section II.B) into a single operator, capable of processing an image in one pass. Separating the two SEs and measuring the extent to which features occupy B_{FG} and B_{BG} allows us to plot this data to determine a minimum occupancy requirement so that the SEs can detect objects in the presence of noise or texture.

The number of elements of B_{FG} that are occupied by a signal in the foreground can be calculated by translating B_{FG} to a point x in the image, and $\forall t \in T$, calculating the cardinality ($Card$) of the set, $(B_{FG})_{x,t} = \{b_{FG} \in B_{FG} \mid I(x + b_{FG}) \geq t\}$, of image pixels which are coincident with B_{FG} and have intensity greater than or equal to t . For all $t \in T$, we calculate the foreground occupancy, $O_{FG,t}$,

$$O_{FG,t} = Card\{(B_{FG})_{x,t}\}. \quad (16)$$

By an equivalent technique, it is possible to measure the extent to which a feature occupies the background SE, B_{BG} ,

$$O_{BG,t} = Card\{(B_{BG})_{x,t}\}. \quad (17)$$

In this case, B_{BG} is translated to a point $x \in E$ and $\forall t \in T$, we calculate the background occupancy, $O_{BG,t}$, i.e. the cardinality of the set, $(B_{BG})_{x,t} = \{b_{BG} \in B_{BG} \mid I(x + b_{BG}) < t\}$, of image pixels, coincident with B_{BG} , that have intensity less than t .

Using (16) and (17), we obtain two, one dimensional arrays, $O_{FG,t}$ and $O_{BG,t}$, of length 2^n which contain the number of elements that are occupied by the signal in B_{FG} and B_{BG} respectively at each level t . The elements of both arrays can be converted to percentages such that we obtain the percentage occupancy of B_{FG} and B_{BG} for all $t \in T$ when their origin is coincident with a point $x \in E$. We denote the percentage occupancy of the foreground and background SEs, PO_{FG} and PO_{BG} respectively where, $\forall t \in T$,

$$PO_{FG,t} = \frac{O_{FG,t}}{Card(B_{FG})} \times 100 \quad (18)$$

$$PO_{BG,t} = \frac{O_{BG,t}}{Card(B_{BG})} \times 100. \quad (19)$$

Since the cardinality of B_{FG} and B_{BG} is generally known, PO_{FG} and PO_{BG} may be calculated directly, $\forall t \in T$, using,

$$PO_{FG,t} = \left[\frac{Card\{b_{FG} \in B_{FG} \mid I(x + b_{FG}) \geq t\}}{Card(B_{FG})} \right] \times 100 \quad (20)$$

$$PO_{BG,t} = \left[\frac{Card\{b_{BG} \in B_{BG} \mid I(x + b_{BG}) < t\}}{Card(B_{BG})} \right] \times 100 \quad (21)$$

The advantage of working in a relative measure such as percentages is that when calculating the POHMT, being able to specify a minimum percentage of the SE that must be occupied for successful detection rather than the actual number of SE elements makes the transform more general. Calculating PO_{FG} and PO_{BG} also allows these quantities to be plotted against each other in the form of a percentage occupancy (PO) plot such as those shown in Fig.4.

Fig.4(a) shows a noise free, synthetic, 8 bit grayscale image, containing a homogeneous circle on a uniform, dark background. By designing B_{FG} such that it can be contained entirely within this circle and the elements of B_{BG} to form a ring to encompass the disk, PO_{FG} and PO_{BG} may be calculated using (20) and (21). We have plotted PO_{FG} and PO_{BG} against intensity for the noise free feature in Fig.4(b) to illustrate how these quantities vary with t . By observation of Fig.4(b), it is clear that B_{FG} is 100% occupied until $t = 150$ i.e. until B_{FG} is above the signal and B_{FG} is 0% occupied for $t > 150$. We can also see that B_{BG} is 0% occupied until $t = 50$ i.e. it is completely below the background and is 100% occupied for $t > 50$ when the SE is completely above the background level of the image. We show in fig4.(c) the PO plot that can be generated by plotting the points of PO_{FG} and PO_{BG} against each other for

each level t . By interpolating these discrete points we obtain a profile, which, in this case, takes the form of a right angle.

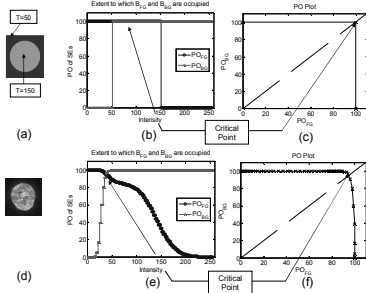


Fig. 4 Images and their PO plot (a) Synthetic image (b) Plot of PO_{FG} and PO_{BG} against intensity (t) for (a). (c) PO plot indicating that the standard HMT will detect the noise free object. (d) Noisy cell image (e) Plot of PO_{FG} and PO_{BG} against intensity (t) for (d). (f) PO plot indicating the HMT will not detect the cell. N.B. If the HMT will not be affected by noise ((a), (b) and (c)), the critical point may be a set of points, the cardinality of which gives the number of times that the SEs fit the feature.

The image shown in Fig.4(a) is not perturbed by noise and the feature of interest does not exhibit internal texture, hence the standard HMT, using the SEs described could be used to detect this object. This is reflected in the PO plot since it shows that there is at least one level, t , such that when B_{FG} and B_{BG} are centered at $x \in E$, B_{FG} and B_{BG} are simultaneously, 100% occupied. This is indicated in Fig 4(c) by the line forming a right angle which intersects the point on the 45° diagonal where, $PO_{FG} = PO_{BG} = 100$. However, if the image is corrupted by noise, the PO plot will not form the ideal right angle but will instead tend towards a curve. This is demonstrated using the image of a noisy cell shown in Fig.4(d). Again, we have plotted PO_{FG} and PO_{BG} against intensity in Fig.4(e), and the corresponding PO plot, generated by plotting PO_{FG} vs PO_{BG} and interpolating, is shown in Fig.4(f). In this case, by examining Fig.4(e), it is clear that there is no level t for which $PO_{FG} = PO_{BG} = 100\%$ and hence the HMT will fail to detect this feature using the SEs described. This is reflected in the PO plot shown in Fig.4(f) since instead of forming the right angled profile shown for the noise free shape, the PO plot in the case of noise, tends towards a curve which crosses the 45° line where $(PO_{FG} = PO_{BG}) < 100$. It should be noted that the critical point on the PO profile (Fig.4(c) and Fig.4(f)) is the point where the curve crosses the 45° line. This point is equivalent to the point at which PO_{FG} and PO_{BG} intersect in Fig.4(b) and Fig.4(e). The critical points and their equivalences in the plots are highlighted in Fig.4.

As the noise and texture is increased, the distance between the curve and this ideal right angle increases. This effect is demonstrated in Fig.5 where zero mean, Additive White Gaussian Noise (AWGN) of increasing power has been added to a synthetic image similar to the one shown in Fig. 4(a).

The PO plot shown in Fig.5 can be used to set a minimum percentage occupancy requirement for B such that the circle may be detected, using this SE, even in very noisy conditions. Since there is only one object (the grey circle) in each image on the left of Fig.5, the minimum occupancy

requirement can be set for the image with the highest noise power.




Fig.5 The effect of noise on the PO plot. (left) Images corrupted by AWGN of zero mean and increasing power. (right) Corresponding PO plots for the object in each image with increasing noise power.

That is, by reference of the PO plot in Fig.5, setting the minimum occupancy requirement of B to 75% guarantees that the gray circle will be detected by the POHMT in all five images. This is clear from the PO plot which indicates that 75% is the lowest occupancy of the SE in all of the images. If there were other objects in the image, setting the minimum PO requirement so low, may invoke erroneous hits in the images that are distorted by noise of lower power. Increasing the minimum PO appropriately, using the PO plot as noise power decreases, will reduce the likelihood of erroneous detection.

It is possible for B_{FG} and B_{BG} to be combined to form a composite SE, B , as in Section II.B. This allows the extent to which the elements of B corresponding to the foreground, $b_{FG} \in B$, and background, $b_{BG} \in B$, of an image to be calculated simultaneously for all $t \in T$, in a single pass. The highest percentage of B , denoted PO_B that is occupied by a signal for all $t \in T$, when its origin is coincident with a point $x \in E$ may be calculated by finding the critical point that was introduced in Fig.4. In the PO plot, the critical point is the one that crosses the 45° line which occurs when $PO_{FG} = PO_{BG}$. However, the values on the curve are discrete points at finite integer values of t and hence it is unlikely that a data point belonging to this profile would lie exactly on the 45° diagonal. We must therefore find the point in the PO plot which is closest to the one that intersects the line to determine the value of this critical point. Given that the critical points shown in Fig.4(b) and Fig.4(e) are equivalent to those shown in Fig.4(c) and Fig.4(f) respectively, it can be seen that the critical point may be most conveniently computed using,

$$PO_B(x) = \max_{t \in T} \left[\min \{ PO_{FG,t}, PO_{BG,t} \} \right]. \quad (22)$$

In (22), the quantities PO_{FG} and PO_{BG} may be calculated $\forall x \in E$ using (20) and (21). The POHMT, which is introduced in the next section, uses PO_B to determine whether or not a pixel at position $x \in E$ should be marked, as a hit, in the output of this transform, based on the extent to which PO_B is occupied at each position in the image.

B. The POHMT

The POHMT uses a composite SE to detect features in

noisy images by allowing objects that occupy only a percentage of the SE to be marked in the output of this transform. The POHMT can be calculated $\forall x \in E$,

$$POHMT_{[B \in B_{FG} \cup B_{BG}]}(x) = \begin{cases} 2^n - 1 & \text{if } PO_B(x) \geq P \\ 0 & \text{otherwise} \end{cases}, \quad (23)$$

where PO_B is calculated using (22) and P is the minimum percentage of the SE that must be occupied for successful detection of an image feature. The value of P can be set by trial and error, or, if the power and distribution of noise that corrupts a signal is known, then an accurate value for P may be calculated using noise models. A third technique, which we propose in this paper, is to measure an appropriate value for P using the PO plot and a set of test images that is representative of the real data. By designing SEs to best match the geometry of image features that are to be detected, the method described in the previous section can be used to generate a PO plot for each object of interest in the set. The PO plot for each object can then be used to find a minimum value for P such that all features of interest will be detected in all images using the POHMT. This can be done by observation of the PO plot, or, alternatively, P can be determined automatically by calculating the critical point using (22). An example of generating a PO plot, using it to set P and the result of applying the POHMT to a noisy image containing cancer cells is shown in Fig.6, where the images are inverted for convenience when printing.




Fig.6 Example of POHMT using a PO plot to set P . (a) Original, noisy image. (b) PO plot showing the PO profile for each cell using the same composite SE. (c) Binary marker produced by the POHMT (d) Result of applying the POHMT and performing an opening by reconstruction.

By observation of Fig.6(a), it is clear that there are four cells in the image. Each cell has different characteristics in terms of shape, intensity and noise. To detect these cells using the POHMT, the geometry of B was set using a priori knowledge of the shapes and sizes of the cells. The elements of B used to probe the foreground of the image were designed as a flat disk measuring 90 pixels in diameter such that it could fit inside that smallest cell in the image (bottom left cell in Fig.6(a)). The elements of B used to probe the background of the image formed a ring with an inner diameter of 110 pixels which was designed to encompass the largest cell (bottom right of Fig.6(a)). The extent to which the cells occupied B was measured by centering B on each cell and in turn using (20) and (21) to calculate PO_{FG} and PO_{BG} and generate the PO plot for each cell as shown in Fig.6(b). By interpretation of the plot, it is clear that setting P to any value less than or equal to 70% is sufficient to ensure that all four cells will be detected by the

POHMT using B .

The POHMT was calculated $\forall x \in E$ with P set to 70%. The POHMT produced a binary marker, as shown in Fig.6(c) which contains four groups of marker pixels in the same locations as each of the four cells in the Fig.6(a). Performing an opening by reconstruction using the original image as the “mask” and the result of the POHMT as the “marker” produced the image shown in Fig.6(d). It should be noted that the standard HMT did not detect any of the cells as predicted by the PO plot. We also point out here, that the POHMT is an extension of the HMT and hence the standard HMT can be implemented as a special case of the POHMT. Setting P to 100% in (23) and calculating the POHMT of an image will give the same result as any of the grayscale HMTs discussed in Section II of this paper.

IV. EXTENSIONS OF THE HMT AND THE PO PLOT

We have demonstrated that the PO plot can be used as a design tool to set the only parameter P for the POHMT. In addition to this, we show here that the PO plot may be used as a tool by other researchers to set corresponding parameters for their own routines. Further, we show that by exploiting the information contained within the PO plot, it is possible to discriminate between image features, having similar properties, using a single composite SE.

A. A design tool for existing grayscale HMTs in noise

In addition to the grayscale extensions of the HMT that have been presented by various researchers (as discussed in Section II of this paper), numerous methods have been proposed which aim to generalize the HMT to make it more robust to noise. We review here a few of these techniques and provide a method which exploits the properties of the PO plot in order to set parameters for these methods.

In [14], the authors present a “Generic solution to improve noise robustness” where they indicate that the grayscale HMTs proposed by Ronse and Soille can be made more robust to noise if the distance between the two SEs, B_{FG} and B_{BG} , is increased. In [14], an example of how to modify this distance is given as, $B'_{FG} = B_{FG} - l$ and $B'_{BG} = B_{BG}$, or, $B'_{FG} = B_{FG}$ and $B'_{BG} = B_{BG} + l$. However, no formula or method is provided that can be used to calculate an appropriate value for l .

It is possible to use the PO plot to determine an appropriate value for this parameter, l , by forcing the PO plot to form the ideal right angle. By calculating a distance, d , from the PO arrays, and shifting the elements of either PO_{FG} or PO_{BG} , by this distance, we can force the PO plot to form a right angle despite any noise or texture in the image. We calculate d , as the difference between the highest level, t , for which B_{FG} is 100% occupied and the lowest level, t , for which B_{BG} is 100% occupied. More formally,

$$d = \max_{t \in I} (t | PO_{FG,t} = 100) - \min_{t \in I} (t | PO_{BG,t} = 100). \quad (24)$$

We can then force the plot to form a right angle by shifting the elements of either PO_{FG} or PO_{BG} by this distance d , to

the right or left respectively, to obtain either PO'_{FG} or PO'_{BG} . By plotting PO'_{FG} vs. PO_{BG} or PO_{FG} vs. PO'_{BG} , we obtain the right angled plot which implies that by setting $l = d$, we can accurately set the distance between SEs using the method described in [14] to improve the noise robustness of the RHMT or the UHMT. That is, the level d that is calculated in order to force the plot to form the right angle, is equivalent to the minimum distance that must be allowed between B_{FG} and B_{BG} such that the feature of interest may be detected by either of these HMTs in the presence of noise. To demonstrate this technique we use the synthetic image shown in Fig.7(a) which has been corrupted by zero mean, AWGN, of variance 9.22.




Fig.7 Setting the SE separation for Ronse's and Soille's HMTs. (a) Noisy synthetic image. (b) Histogram equalization of (a) for visualization of noise. (c) Intensity profile of image center row, no noise (blue), noise corrupted signal (green). (d) PO plot obtained before (green) and after (blue) shifting the elements of PO_{FG} by d as well as intermediate plots for increasing d . (e) PO Plots obtained before (green) and after (blue) setting the SE separation to d and calculating PO_{FG} and PO_{BG} .

For the purpose of illustrating the effect that the noise has on this signal, we have shown, in Fig.7(b), the image after histogram equalization, and in Fig.7(c), we show a 1D intensity profile taken from the center row of the image before and after the noise has been added. By generating the PO plot for the case where the distance between the SEs is initially zero, we can use (24) to calculate the distance d that should be set between the SEs to allow this feature to be detected. By shifting the elements of PO_{FG} to the right by d and plotting PO'_{FG} vs PO_{BG} , we obtain the right angle as shown in Fig.7(d). To demonstrate the way in which the plot is forced to form the right angle, we have shown three additional curves in the PO plot in Fig.7(d). These curves have been generated purely for example by setting d to values that lie between zero and the critical distance of 20 graylevels that was calculated using (24).

By reference of Fig.7(d), it is clear that as d increases, the plot gradually approaches the desired right angle before attaining this profile when $d = 20$. If the distance, d , between the SEs is sufficient to allow the HMTs defined by Ronse or Soille to detect the noisy feature, the PO plot that is generated after fixing this distance between the SEs and recalculating PO_{FG} and PO_{BG} , also forms the right angle. This case is shown in Fig.7(e) where the distance between the SEs was set to 20 graylevels before calculating PO_{FG} and PO_{BG} .

A drawback with the method in [14] of increasing the distance between the SEs is that it is prone to erroneous detections since any group of pixels lying between the SEs will be marked as a "hit". When applying the RHMT or the UHMT and setting the distance between the SEs to be d as described, the disk in the noisy image is successfully detected. There are however, as expected, many erroneous "hits" in the resultant image. The POHMT on the other hand marks only the feature of interest.

The authors [14] also discuss the performance of the HMTs proposed by Barrat et al. and Khosravi and Schaefer when images are corrupted by noise. They conclude that since these HMTs already evaluate a distance between the SEs, the problem of finding a suitable distance is transformed into a problem of thresholding the result of their HMTs. It is possible to use the PO plot in the same fashion as before to determine this threshold, where the threshold is equivalent and can therefore be equated to the distance d that was previously calculated using (24).

By performing the grayscale HMTs proposed by Barrat and Khosravi and thresholding the results at $d=20$, we successfully detect the disk in the center of the image. However, as is the case with the previous example, a large number of erroneous detections appear in the result when using this technique. Thresholding the result of these HMTs at a level less than d does not allow successful detection of the circle in noise, however, the result still contains a high number of false positives. The same is true for the RHMT and the UHMT, where setting the distance between the SEs to be less than 20 graylevels results in erroneous "hits" while the feature of interest is not detected.

In [14], the authors state that if an image is corrupted by AWGN, and the variance of this noise is known, then the distance between B_{FG} and B_{BG} can be set to equal twice that of the standard deviation of the noise. This theory can be verified by setting this distance between the SEs and generating a PO plot to form the ideal right angle as has been done in previous examples. This method is reliable if the power and distribution of the noise is known. Usually, this is not the case and hence the PO plot could be used in such situations, to calculate this parameter.

We have shown here, that the problem of defining a suitable distance between the SEs and finding a suitable threshold to apply to the result of the HMTs are equivalent as stated in [14]. We also show that these parameters can be estimated using the PO plot by finding the minimum distance, d , which forces the result to form the right angle. However, an issue with relaxing the conditions of the HMT using these techniques is that the transform becomes more susceptible to producing erroneous hits in the output image.

Perret et al present their solution to overcome the difficulties faced by the HMT in the presence of noise in [14]. The authors introduce a Fuzzy Hit-or-Miss Transform (FHMT) which they use to detect features in very noisy astronomical images and provide an impressive set of results. Their technique involves generating a large set of SEs using a mathematical model that incorporates the characteristics of the features they aim to detect for various

scale lengths, orientations and elongations. A measure of fitness is obtained for all patterns in the set of SEs at each point in the image, and a record of the best fitting SE at each pixel is stored as well as a measure of how well this SE fits the image. This data is used to form a so called "Scoremap" which is thresholded at a particular level to produce a binary marker image from which detected features can be reconstructed.

For the FHMT, the PO plot could be used to set the ideal distance between the SEs, or to provide an indication of a suitable threshold that can be used on the output of this transform. Additionally, the PO plot and a suitable set of training data could be used to set a minimum occupancy requirement for a single SE, or at least a small subset taken from the large set of SEs that are used currently. This would allow the algorithm to execute, in the same way as the POHMT, in a fraction of the time taken by the current routine (2 minutes per image) described in [14].

B. A discriminatory filter

Often, features that are to be detected in an image are not geometrically identical. If, therefore, we wish to extract from an image, a number of features, which differ from each other in terms of shape and size, we can design a number of composite SEs i.e. one to match the geometry of each object that we seek in the image. We can then perform a grayscale HMT using each of the composite SEs in turn before calculating the union of all the resulting binary images to obtain a single image that contains markers for each image feature that has been detected. That is of course assuming that the HMT will not fail to detect these features due to noise or texture in the image.

As was illustrated in Fig.6, the POHMT allows multiple objects which are geometrically very different to be detected using just one composite SE in a single pass of the image. This can be achieved by exploiting the information contained within the PO plot in order to determine an appropriate level for P such that we can guarantee to detect all the features in this image. The PO plot does however provide a further advantage in that we may set P in such a way that we can discriminate between image features using just one composite SE. The simplest case of discriminating between features using the POHMT is to set the value of P high enough to eliminate objects which simultaneously occupy a maximum percentage of B that is always less than P . An example of selectively detecting cells in the image by varying P using the information contained in the PO plot is shown in Fig.8.

Clearly, by reference of Fig.8(a), setting the level of P to any value that lies between the curve representing the cell in the bottom right and left of the image, we can eliminate the cell in the bottom left while successfully detecting the other three cells. In Fig.8(c) we have extracted only three of the four cells shown in Fig.8(b) by setting $P = 75\%$ in order to eliminate the cell in the bottom left of the image for which the maximum, simultaneous occupancy of the SE when coincident with this cell is 70%. By raising the level of P to 90% and then 96% in accordance with the PO plot shown in

Fig.8(a), we can extract respectively two cells (Fig.8(d)) then one cell (Fig.8(e)). Evidently, the PO plot is an extremely powerful design tool, as it provides information that allows objects to be detected selectively using one composite SE in a single pass of the image.




Fig.8 Example of POHMT operating as a discriminatory filter. (a) PO plot for the four cells in (b). (b) Noisy image containing four cells. (c) Three of the four cells detected by setting $P = 75\%$. (d) Two of the cells detected by setting $P = 90\%$. (e) One of the cells detected by setting $P = 96\%$.

The case demonstrated here is a powerful yet trivial one since it is obvious that increasing the level P or in other words increasing the strictness of the transform results in objects being discarded in the detection process.

What is more interesting, is that by a similar technique to the one described above, it is possible to isolate any of the four cells in the image shown in Fig.8(b) and hence we can segment any combination of the image features using just one composite SE. In this case, the PO plot can be used as a shape descriptor which allows us to use one composite SE to discriminate between objects of interest in an image and objects which may have very similar geometrical properties in the spatial domain, for example, the two cells at the top of the image. Fig.9 shows each of the cells being extracted on their own using the same composite SE and the POHMT.




Fig.9 POHMT operating as a discriminatory filter (a) Image containing four cells of different shape and size. (b) Bottom left cell isolated. (c) Top right cell isolated. (d) Bottom right cell isolated. (e) Top left cell isolated.

The results shown in Fig.9 can be easily achieved by firstly detecting and reconstructing all of the four cells in the image. Then, by detecting the objects that are not desired and reconstructing this image, the difference image can be calculated such that only the features specified by a user etc. are picked out by the POHMT.

V. EXPERIMENTAL RESULTS

In this section, we test the performance of the POHMT by implementing each step described in Section III to detect a biological cell in a series of very noisy images. We also test our method on the same images as Perret et al. in [14] to determine if our routine is capable of detecting the LSB galaxies in less time. Our technique is simpler firstly as it processes the image in a single pass. Secondly, a fast POHMT has been implemented using techniques similar to

those used to optimize median filtering [21] and morphological operators [22]. However, instead of searching for the min, median or max value in the window, the POHMT searches for the rank specified by P .

An example of the noisy biological images (of an immune system cell) is shown in Fig. 10 where we have chosen three of the images (Fig.10(a)) to be used as a training set in order that P can be determined and used to detect the features of interest in our test set (Fig.10(b)). We can see a small group of pixels in each image in Fig.10 that represents the feature of interest, while the rest of the image contains noise and other features that are not of interest. The images are extremely noisy, and, by observation of the data, it is evident that the shape and orientation of the cell changes between the images.




Fig. 10 The image set containing the cell of interest and some other features where the entire image is submersed in noise. (a) Training set to determine an appropriate value for P . (b) The set of test images in which we seek the feature after P has been fixed using the PO plot for the test set. (c) PO plot obtained for the training set shown in (a).

The first stage in the process is to generate a PO plot for each feature of interest in the test set in order to determine an appropriate level for P . Although the cell is not a constant shape and size in all images, we can design B such that its elements corresponding to B_{FG} will fit inside it in each image. Similarly, B_{BG} was designed to encompass all of the features of interest in each image to guarantee that we can detect the cell in all possible orientations and variations of shape and size. By increasing the spatial distance between the SEs, as we are here, it can be argued that the transform may produce erroneous hits. If a problem occurs, this issue can easily be overcome by exploiting the discriminatory property of the POHMT shown in Section IV B. We also note that although automatic techniques are available for SE design, we have used a manual method here to compare our method with the one presented in [14]. We have used square SEs for processing simplicity, however, this may be readily extended to arbitrarily shaped SEs using the method described in [22]. B was used to generate a PO plot for each image in the training set in order to obtain a suitable level for P , such that the feature could be detected in the test set, without picking up erroneous hits. The PO plot, generated for the training set, is shown in Fig. 10(c). Clearly, by reference of the PO plot, setting $P=81\%$ is sufficient to ensure that this feature may be detected using one composite SE for the entire test set. The POHMT was calculated for each image in the test set where the results of applying this transform and reconstructing the features of interest that have been marked

are shown in Fig.11.

We also calculated the average processing time using the described SEs when analyzing this image on a PC with a Pentium IV processor. The image is 512 x 512 pixels in size and the average time taken to process an image was measured to be 0.87s. This does not include the opening by reconstruction which is performed largely for illustration and is not normally required in a practical detection or feature recognition system. To improve visibility in Fig.11, we have dilated the each image resulting from the opening by reconstruction.




Fig.11 Result of applying opening by reconstruction to the result of the POHMT for each image in Fig.10(b).

We have also tested our routine on the images used in [14] to compare the performance and efficiency with the one described by Perret et al. We show an example of the POHMT detecting an LSB galaxy in Fig.12 where the contrast of the image shown in Fig.12(a) has been enhanced to make the LSB clearly visible. The average processing time of our routine was 4 seconds per 512 x 512 image which is a substantial improvement compared to the 2 minute execution time of the optimized routine presented in [14].

For the example shown, our method is faster and simpler to implement than the method proposed by Perret et al [14] since it only requires one composite SE. Further, we point out here that we do not optimize our routine by subsampling the image, the SE, or by using any potentially lossy, heuristic techniques.




Fig.12 The POHMT detecting a LSB galaxy. (a) Original noisy image containing a LSB galaxy in the lower right quadrant of the image. (b) The output of the POHMT when processing the image in (a).

Clearly, our proposed routine provides a significant improvement to the optimized method presented in [14] while achieving accurate results. It should be noted however, that the method proposed in [14], is extremely well suited to the authors' application since the visual appearance of LSB galaxies varies significantly between images. For this reason, the authors in [14] use an accurate mathematical model to generate a large set of SEs for

processing each image, however, when the sought patterns do not vary substantially between images our technique is optimal and requires that only a few frames (as demonstrated) be used to train the routine and set the only parameter, P . In addition to the simplicity of our method, the benefits of the PO plot in terms of discriminating between image features and allowing other researchers to set parameters for their routines, provides an element of novelty that is not yet found in similar techniques.

VI. CONCLUSIONS

In this paper we have presented an extension of the grayscale HMT following the definition given by Soille, although various definitions exist in the literature. We have shown the equivalences between these extensions and we have highlighted the reasons that noise in images may cause these transforms to fail. We have presented a solution that offers improved robustness to noise, in the form of a Percentage Occupancy Hit-or-miss Transform, which relaxes the fitting criteria of the SEs making them more flexible such that they can successfully detect image features even in very noisy conditions.

In addition to the POHMT, we have shown in this paper, a novel design tool in the form of a PO plot. The plot can be used to set the only parameter required by the POHMT and can be used by other researchers to set parameters for their own routines. We have given various examples of how this can be achieved and we have used the PO plot to incorporate some suggested modifications by other researchers to make the grayscale HMT perform better in noise. When performing the HMT and using the suggested modifications, we have shown that image features can be detected in noise, but unlike the POHMT, there are also a large number of false positives in the result.

Further to the PO plot being used to set parameters for grayscale HMTs in noise, we have shown that this tool provides some additional benefits. The discriminatory filter aspect of the POHMT which is a direct result of analyzing the PO plot, allows us to differentiate between objects in the image that we wish to detect and others which may appear visually similar in the spatial domain but that are not of interest. We have also implemented a fast algorithm that can be used to calculate the POHMT. On the set of images containing the noisy biological cell, the fast POHMT executed in less than one second while detecting the image features of interest. We have shown that our method outperforms all of the grayscale HMTs that have been discussed in this paper when images are noisy and that even using the suggested techniques for improved robustness to noise, we still achieve better results. We have also shown that our method achieves better results than the most recent extension of the HMT presented by Perret et al. and we have verified this using their images. We have also shown that by exploiting the properties of the PO plot, we obtain our results in a fraction of the time, using just one composite SE, in a single pass of the image.

Although the applications of this method have been demonstrated for images of a biological and astronomical

nature, our method may be applied to any feature recognition problem. The only requirement of this routine is that we must know the spatial characteristics of the patterns that we seek, however, this requirement is consistent with most morphological operations.

VII. ACKNOWLEDGMENTS

The authors would like to thank Eric Bullinger, now at the University of Liege, for his guidance at an early stage in this work. Thanks also to Monica Schliemann and Peter Scheurich at the Institute of Cell Biology and Immunology, University of Stuttgart for supplying the cancer cell images used in this paper. We would also like to thank Stefano Fumagalli, Patrick Ejlerskov and Jonathan A. Coles in the Centre for Biophotonics, University of Strathclyde and James Brewer at the University of Glasgow, for allowing us to use the images shown in Fig.10. Thanks must also go to Benjamin Perret, at the University of Strasbourg for providing us with the set of astronomical images and for granting us permission to use this data to produce Fig.12 in our manuscript. Finally, we would like to thank Des McGhee at the University of Strathclyde for his advice and guidance on various aspects of set theory and notation.

VIII. REFERENCES

- [1] G. Matheron, *Random Sets and Integral in Geometry*. New York: Wiley, 1975.
- [2] J.Serra, *Image Analysis and Mathematical Morphology*. London: Academic Press, 1982.
- [3] H. J. Heijmans, *Morphological Operators*. New York: Academic Press, 1994.
- [4] R.M. Haralick and L.G. Shapiro, *Computer and Robot Vision*, New York: Addison-Wesley, 1992.
- [5] R.C. Gonzalez and R.E. Woods, *Digital Image Processing*, Third Edition, New Jersey: Pearson Prentice Hall, 2008.
- [6] M. Sonka, V. Hlavac and R. Boyle, *Image Processing, Analysis and Machine Vision*, London: Chapman & Hall, 1993.
- [7] P. Soille, *Morphological Image Analysis: Principles and Applications*, New York: Springer, 2003.
- [8] D. Zhao, D. Daut, "Shape recognition using morphological transformations," *Proc. IEEE 16th International Conference on Acoustics, Speech, and Signal Processing (ICASSP'91)*, pp. 2565–2568, Apr. 1991, doi: 10.1109/ICASSP.1991.150925
- [9] D. Zhao and D. G. Daut, "Morphological Hit-or-Miss Transformation for shape recognition," *J. Visual Communication and Image Representation*, vol. 2, no. 3, pp. 230–243, Sept. 1991.
- [10] D.S. Bloomberg, P. Maragos, "Generalized hit-miss operations," *Proc. SPIE*, vol. 1350, *Conf. Image Algebra and Morphological Image Processing*, pp. 116–128, 1990.
- [11] D.S. Bloomberg, L.M. Vincent, "Pattern matching using the blur hit-miss transform," *J. Electronic Imaging*, vol. 9, no. 2, pp. 140–150, Feb. 2000.
- [12] B. Naegel, N. Passat, C. Ronse, "Grey-level hit-or-miss transforms-part I: unified theory," *J. Pattern Recognition*, vol. 40, no. 2, pp. 635–647, Apr. 2007.
- [13] M. Khosravi, R. Schafer, "Template matching based on a grayscale hit-or-miss transform," *IEEE Trans. Image Processing*, vol. 5, no. 6, pp. 1060–1066, June 1996, doi: 10.1109/83.503921.
- [14] B. Perret, S.Lefèvre, Ch.Collet, "A robust hit-or-miss transform for template matching applied to very noisy astronomical images," *J. Pattern Recognition*, Vol. 42, no. 11, pp. 2470–2480, Nov. 2009.
- [15] C. Ronse, "A Lattice-Theoretical Morphological View on Template Extraction in Images," *J. Visual Communication and Image Representation*, Vol.7, No.3, pp 273 – 295, Sept.1996.
- [16] P. Soille, "Advances in the analysis of topographic features on discrete images," *Proc. Lecture Notes in Computer Science*, 10th

International Conference on Discrete Geometry for Computer Imagery (DGCI'02), pp. 175–186, 2002.

- [17] C. Barat, C. Ducottet, M. Jourlin, "Pattern matching using morphological probing," *Proc. IEEE Conf. International Conference on Image Processing (ICIP'03)*, pp. 369–372, Sept. 2003, doi: 10.1109/ICIP.2003.1246975
- [18] P. Murray S. Marshall, E. Bullinger, "The Percentage Occupancy Hit-or-Miss Transform", *Proc. The 17th European Signal Processing Conference (EUSIPCO '09)*, pp.253-257, Aug.2009.
- [19] N. R. Harvey and S. Marshall, "Rank-order morphological filters: A new class of filters", *Proc. IEEE Workshop, Nonlinear Signal and Image Processing*, pp. 975-978, 1995.
- [20] P. Soille, "On morphological operators based on rank filters", *J. Pattern Recognition*, vol. 35 no. 2, pp. 527–535, Feb. 2002
- [21] T.S. Huang, G.J. Yang, and G.Y. Tang, "A fast two dimensional median filtering algorithm," *IEEE Trans. Acoustics, Speech and Signal Processing*, vol 27, no. 1, pp. 13-18, Feb 1979.
- [22] M. Van Droogenbroeck, H. Talbot, "Fast computation of morphological operations with arbitrary structuring elements," *J. Pattern Recognition Lett.*, vol. 17, no. 14, pp. 1451-1460, 1996.



Paul Murray received the degree of MEng (with Distinction) in Electronic and Electrical Engineering from the University of Strathclyde, Glasgow, Scotland in 2008. He is currently pursuing a PhD in Image Processing at the University of Strathclyde where he is in his third year of study. His

research interests include extending and developing new image processing algorithms in mathematical morphology, segmentation, tracking and machine learning for biological applications. His main focus has involved automating the analysis of large sets of biological images using new and existing methods and techniques.



Stephen Marshall received a first class honours degree in Electrical and Electronic Engineering from the University of Nottingham in 1979 and a PhD in Image Processing from the University of Strathclyde in 1989. In between he worked at Plessey Office Systems, Nottingham, University of

Paisley and the University of Rhode Island, USA.

In recent years, his research activities have been focussed in the area of Non Linear Image Processing. In this time, he has pioneered new design techniques for morphological filters based on a class of iterative search techniques known as genetic algorithms. The work is summarised in his book *Logic Based Image Processing* published by SPIE.

He has published over 150 conference and journal papers on these topics including IEE, IEEE, SPIE, SIAM, ICASSP, VIE and EUSIPCO. He has also been a reviewer for these and other journals and conferences. He has also been successful in obtaining research funding from National, International and Industrial sources. These sources include EPSRC, EU, Rolls Royce, BT, DERA, the BBC and Scottish Enterprise.

He is a Fellow of the IET and a founder of the IET Professional Network on Visual Information Engineering (VIE) and Chair of the Scottish Biosystems Modelling Network SBMN. He is a Professor in the Department of Electronic and Electrical Engineering and Director of the Centre for excellence in Signal and Image Processing.