HAL open science

Robust Object Detection in Colour Images Using a Multivariate Percentage Occupancy Hit-or-Miss Transform

Fraser Macfarlane, Paul Murray, Stephen Marshall, Benjamin Perret, Adrian Evans, Henry White

To cite this version:

Fraser Macfarlane, Paul Murray, Stephen Marshall, Benjamin Perret, Adrian Evans, et al.. Robust Object Detection in Colour Images Using a Multivariate Percentage Occupancy Hit-or-Miss Transform. Mathematical Morphology - Theory and Applications, 2021, 5 (1), pp.128-152. 10.1515/mathm-2020-0111 . hal-04340893

HAL Id: hal-04340893

https://hal.science/hal-04340893v1

Submitted on 13 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Math. Morphol. Theory Appl. 2021; 5: 128–152 DE GRUYTER

Research Article Open Access

Fraser Macfarlane*, Paul Murray, Stephen Marshall, Benjamin Perret, Adrian Evans, and Henry White

Robust Object Detection in Colour Images Using a Multivariate Percentage Occupancy Hit-or-Miss Transform

https://doi.org/10.1515/mathm-2020-0111
Received February 23, 2021; accepted November 17, 2021.

Abstract: The extension of Mathematical Morphology to colour and multivariate images is challenging due to the need to define a total ordering in the colour space. No one general way of ordering multivariate data exists and, therefore, there is no single, definitive way of performing morphological operations on colour images. In this paper, we propose an extension to mathematical morphology, based on reduced ordering, specifically the morphological Hit-or-Miss Transform which is used for object detection. The reduced ordering employed transforms multivariate observations to scalar comparisons allowing for an order to be derived and for both flat and non-flat structuring elements to be used. We also compare other definitions of the Hit-or-Miss Transform and test alternative colour ordering schemes presented in the literature. Our proposed method is shown to be intuitive and outperforms other approaches to multivariate Hit-or-Miss Transforms. Furthermore, methods of setting the parameters of the proposed Hit-or-Miss Transform are introduced in order to make the transform robust to noise and partial occlusion of objects and, finally, a set of design tools are presented in order to obtain optimal values for setting these parameters accordingly.

Keywords: Image processing; Mathematical morphology; Hit-or-Miss Transform; Template matching; Object detection; Pattern recognition

1 Introduction

Mathematical Morphology (MM), first introduced and formalised by Matheron [25] and Serra [37, 38] and later extended by Heijmans [13], is a fundamental set of non-linear image processing techniques. MM applies the mathematical concepts of set theory, specifically lattice theory, in order to study the shape, or morphology, of objects in various image analysis tasks such as object detection, edge detection, segmentation and image de-noising [32, 41]. Although, the use of MM was initially limited to binary images, there has since been much effort to extend MM techniques for application to greyscale images [30, 40] and beyond for multivariate images [2, 20, 50].

The Hit-or-Miss Transform (HMT) is a useful tool in MM and is often used for detecting objects based on their size and shape. In common with many other MM operators, the definition of the HMT was originally

***Corresponding Author: Fraser Macfarlane,** Department of Electronic and Electrical Engineering, University of Strathclyde, Glasgow, UK, E-mail: fraser.macfarlane@strath.ac.uk
Paul Murray, Department of Electronic and Electrical Engineering, University of Strathclyde, Glasgow, UK
Stephen Marshall, Department of Electronic and Electrical Engineering, University of Strathclyde, Glasgow, UK
Benjamin Perret, Université Paris Est, Laboratoire d'Informatique Gaspard-Monge, ESIEE Paris, France
Adrian Evans, Department of Electronic and Electrical Engineering, University of Bath, Bath, UK
Henry White, BAE Systems, Air Sector, Filton, Bristol, UK

Open Access. © 2021 Fraser Macfarlane et al., published by De Gruyter. This work is licensed under the Creative Commons Attribution alone 4.0 License.

restricted to binary images [37]. Multiple proposals for greyscale extensions followed, which are reviewed in [28], and a unified formulation now exists [30]. A further extension to colour and higher dimensionality domains is non-trivial. The fundamental issue is that MM requires a complete lattice [34] in order for its operators to be accurately defined. Defining a complete lattice, i.e. an *infimum* and a *supremum*, on multivariate data is a challenging task as there is no generally accepted or unambiguous way to order such data [4]. As a result, no unified morphological framework exists for use in colour or multivariate images. One way of implementing MM in the colour domain is to introduce a scheme for ordering multivariate data into a morphological framework [2, 36]. Other methods include redefining morphological frameworks using techniques such as fuzzy logic [35] and supervised ordering [44].

Another increasingly common method for object detection is the use of deep learning techniques. Deep learning has been employed to great effect in various image processing applications such as segmentation, edge detection as well as for object tracking [46] and object detection and subsequent classification [16]. Deep-learning, especially convolutional neural network (CNN), based methods offer object detection typically based on some bounding box regression and classification [33]. However, morphology provides a rapid, pixel-wise decision for object detection and pattern recognition based on size and shape without the need for training data. Instance segmentation based networks such as Mask R-CNN [12] offer similar pixel-wise classification however are comparatively complex and computationally expensive. Recently, various investigations into ways of combining morphology into state of the art deep learning methods have been introduced. All morphological operations are performed using a structuring element, or SE, and as such the design of these SEs is of great importance. Traditionally, SEs are designed by hand with some prior knowledge of the operation however, more automatic methods to optimise the design of SEs using supervised [22] and deep learning [24, 31, 39] have been explored. Morphology can also be deployed within deep learning frameworks, forming the non-linear filters required in feature pooling stage [9] or in other layers within a CNN [51]. Deep learning can also be used to learn and mimic morphological operations such as erosion and dilation [9] as well as the HMT [14] or used in morphological operations for various image analysis tasks [24, 26]. All of these proposed methods are solely restricted to single channel images, or can be applied marginally to colour images to provide a naive multi-channel approach. In the future, we may investigate how to couple our techniques with deep learning however the immediate aims of this work are to improve upon the field of colour morphology.

In this paper we propose a novel multivariate extension of the HMT, the HMT was chosen due to its simplicity, speed, and interpretability. Pairs of SEs can be designed with little to no prior knowledge and it can easily be altered for use in various scenarios. Our proposed method can be applied directly to colour images in any colour space as well as to other high dimensionality multivariate images such as hyperspectral images or volumetric images. This is achieved by computing the dissimilarity, using reduced orderings, between an SE and a region in an image bounded by the support of that SE. These dissimilarity measures can then be ordered and morphological operations can be applied to this ordered set. By translating the SE to all points in an image, the morphological analysis of that image can take place. We show how our techniques can be made robust to noise by relaxing the fit of SEs to the image under test allowing for noise to "puncture" the SE or to allow for objects of interest with slight disparities in size, shape or texture to be detected with a generalised SE. Additionally, we present a novel approach for the optimal setting of parameters for use with our technique based on the Percentage Occupancy (PO) Plots, presented in [27], which informs the decision of how much relaxation to apply to each SE for an object to be detected correctly. We compare our method with other recently proposed extensions of the HMT to multivariate data with our approach performing favourably.

Preliminary work is presented in [23], while this paper presents a more detailed analysis of the transform, the method for setting the optimal noise robustness parameters as well as comparisons with other relevant state of the art techniques. The rest of this paper is structured as follows: Section 2 presents an overview of related work on morphology and its extension to colour and multivariate images. Sections 3 and 4 contain our proposal for an extension of the HMT as well as the results and comparisons between our method and other colour HMTs and related state of the art object methods for object detection. Section 5 details the summary and conclusions of this work.

2 Related Work

2.1 Mathematical Morphology

The basis of any structural morphological operator is the SE which defines a region for some morphological operation to take place [41]. In the binary case a SE is simply a window defining the region for the operation, when extended to greyscale images, the SE can be "flat" or "non-flat", where the window has weights which can affect the outcome of MM operations.

2.1.1 Erosion and Dilation

The fundamental operators of MM are Erosion and Dilation. The Erosion on a binary image, X, 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,

$$\varepsilon_S(X) = \{x \mid S_x \subseteq X\}. \tag{1}$$

This is equivalent to all points in the image where an SE, S, is translated to a point x, S_x, and matches, or fits inside, the foreground. The dilation by a reflected SE, $\hat{S}$, (δ_S), known as the dual of the erosion, is the locus of all points where the intersection of $\hat{S}$ and X is non-empty,

$$\delta_S(X) = \{x \mid \hat{S}_x \cap X \neq \emptyset\}. \tag{2}$$

This is equivalent to all points in the image where a translated SE touches, or intersects with, the foreground. Equations (1) and (2) can be extended for application to a greyscale image G, with support $g \in G$. The greyscale erosion is defined as $[\varepsilon_S(G)](g) = \min(G(g+s))$, and greyscale dilation as $[\delta_S(G)](g) = \max(G(g+s))$, where S denotes a greyscale SE with support $s \in S$.

2.1.2 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 (SE_{FG}) is designed to fit within the object under study in an image, X, and the background SE (SE_{BG}) is designed to fit the background of the object or the complement (X^c) of image X. These two SEs can be collated into one composite SE S which is defined as $S = [SE_{FG}, SE_{BG}]$ where $SE_{FG} \cap SE_{BG} = \varnothing$. The HMT of X using a composite SE, S, can be implemented as the intersection of two separate erosions:

$$\mathrm{HMT}_S(X) = \varepsilon_{SE_{FG}}(X) \cap \varepsilon_{SE_{BG}}(X^c) \tag{3}$$

An example of the HMT operating on a synthetic binary image is shown in Figure 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 white squares on a black background.

Equation (3) could be extended to greyscale images but this requires taking the complement of a greyscale image, the result of which would ultimately be dependent on the value used for complementing the image. In addition, the HMT is not an increasing transform, this is to say that given $x \leq y$, this does not imply that $\mathrm{HMT}(x) \nleq \mathrm{HMT}(y)$. This means that its extension to greyscale is not a straightforward endeavour [32, 41]. The general equation for the HMT used in this paper is a form of Soille's Unconstrained HMT (UHMT) [41] with this and various other greyscale HMT extensions reviewed in detail in [28]. The UHMT is based on the binary HMT from (3) and exploits the duality between erosion and dilation, removing the need to erode the complement of the image. Using the greyscale erosion and dilation defined previously and following the definition of the UHMT [41] we can define the greyscale HMT as:

$$[\mathrm{HMT}_S(X)](\mathbf{x}) = \begin{cases} 1 & \text{if } [\delta_{S_{\mathrm{BG}}}](\mathbf{x}) < [\varepsilon_{S_{\mathrm{FG}}}](\mathbf{x}) \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

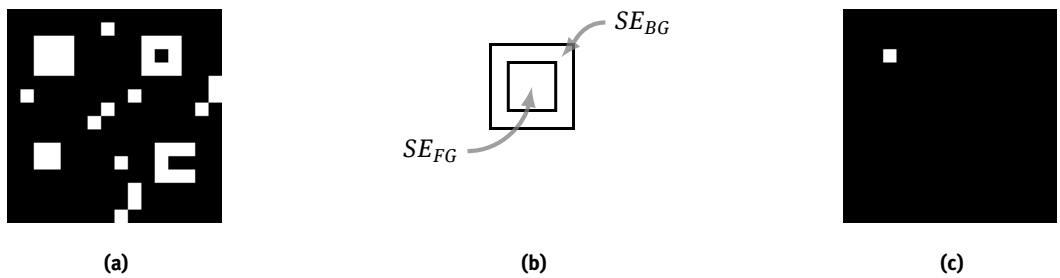

Figure 1: Example of a binary HMT. **a)** 16×16 binary test image made up of multiple sizes of shapes. **b)** Binary composite SE where SE_{FG} is 3×3 and SE_{BG} is 5×5. **c)** Output of the HMT of (a) using (b).

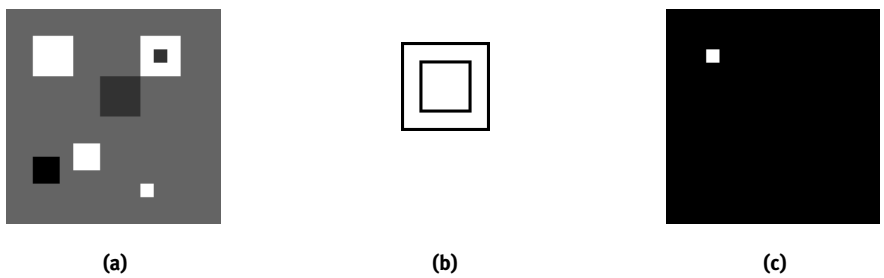

Figure 2: Example of a greyscale HMT. **a)** 16×16 greyscale test image made up of of shapes with varying intensity and size. **b)** Greyscale composite SE where SE_{FG} is 3×3 and SE_{BG} is 5×5. **c)** Output of the HMT of (a) using (b).

An example of object detection using this greyscale HMT on a synthetic greyscale image can be seen in Figure 2. Similarly to the binary image test shown in Figure 1, the SE was designed to detect 3×3 squares brighter than their background.

While it performs well in many applications, the HMT is very susceptible to errors when 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. This is apparent in Figures 1 and 2 where, in several cases, there are white square shapes with a single non-white pixel (representing noise) that are not detected using the standard binary and greyscale HMTs.

There are multiple ways of relaxing the hard constraints of the HMT to detect these imperfect objects including the notion of a Percentage Occupancy HMT (POHMT) [27] which allows for a percentage error in the matching of an SE to an image. These Percentage Occupancy measures can be implemented by replacing the hard limits of erosions (minimum filters) and dilations (maximum filters) with rank k filters [27]. 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:

$$[\zeta_{S,k}(I)](x) = k\text{th} \operatorname*{rank}_{s \in S}\{I(x+s)\} \tag{5}$$

A rank order filter with rank k will return the kth 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, or Cardinality (Card), 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 [27]:

$$k_P = \left\lfloor \left(\frac{P}{100} \times \text{Card}(S) \right) \right\rfloor \tag{6}$$

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,

(a) (b) (c) (d)

Figure 3: a) Binary test image from Figure 1a. **b)** POHMT of (a) using Figure 1b and P = 70%. **c)** Greyscale test image from Figure 2a. **d)** POHMT of (c) using Figure 2b and P = 80%.

$[\zeta_{S_{FG},k_{100-P}}(I)](x)$, and the second is a relaxed background dilation, $[\zeta_{S_{BG},k_P}(I)](x)$, such that:

$$[\mathrm{POHMT_S}(X)](x) = \begin{cases} 1 & \text{if } [\zeta_{S_{FG},k_{100-P}}(I)](x) > [\zeta_{S_{BG},k_P}(I)](x) \\ 0 & \text{otherwise.} \end{cases} \tag{7}$$

In (7) 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 Figure 3 where the images from Figure 1 and 2 were input into a POHMT with P = 90% and the result is a more accurate object detection in the presence of the simulated "noisy" pixels. This can be seen in the results shown in Figure 3 as the squares with non-white centre pixels are detected despite the non-ideal shape or intensity profile.

2.2 Ordering relations

Having explained and reviewed the extensions of the HMT to greyscale and the incorporation of noise robustness, we now examine methods to extend the HMT to multivariate images. According to Barnett [5], there are four main ways of ordering multivariate, vectorial data. The first, marginal or *m-ordering* is concerned with treating each vector as a set of scalars, in the case of image processing techniques - processing each channel in an image separately, which can introduce false colour as shown in Figure 4. The second method is partial or *p-ordering* which separates the vectors into groups which can then be distinguished by rank or extremeness. The third method is conditional or *c-ordering* which involves ordering one of the marginal sets based on some condition. This can be repeated on subsequent sets if the previous result is inconclusive - i.e. first ordering the Red channel in an RGB image followed by the Green and Blue if Red intensity values are repeated and a true order cannot be determined. C-ordering, or lexicographical ordering as it is sometimes referred to in literature, has the disadvantage that the infimum and supremum defined are wholly dependant on the order in which the channels are processed [15] which may introduce biases when detecting various objects of interest. The final method, and the one employed in our proposed extension, is reduced or *r-ordering* which involves measuring the vector against some reference vector or by projecting it into a reduced space. This creates a scalar value based on the disparity of each pixel vector to the reference, which can be sorted easily.

Both *r-ordering* and *p-ordering* rely on a pre-ordering of data, meaning they lack the anti-symmetry characteristic. As well as this, two distinct vectors can be equidistant from a reference vector and no longer be distinct when r-ordering takes place [2]. As such, our approach incorporates reference colour values locally into the transform, by encoding them in the SE, as opposed to being a global reference.

An example of *m-ordering* and *r-ordering* is shown in Figure 4 for ten synthetic colour observations. The effect of false colour introduction can be seen when comparing the marginal (Figure 4a) with the reduced (Figure 4b) ordering schemes, where with the reduced scheme the colours are preserved. In the *m-ordering* scheme, each colour channel is ordered independently leading to the introduction of false colours. In the *r-ordering* scheme, a reference colour, on the right, is used to determine the distance from each colour to the reference which is a scalar value suitable for ordering. Once the colours have been ordered in this way they are able to be re-represented as ordered colour values based on similarity. This produces a colour order where

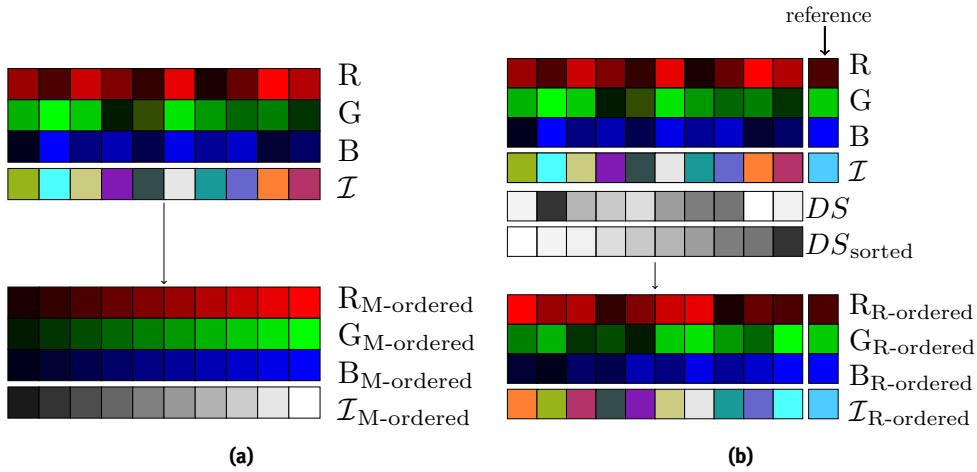

Figure 4: a) *M-ordering scheme* and **b)** *R-ordering scheme* on synthetic colour data samples. The m-ordering scheme orders each band independently meaning the original colours of each pixel are lost. The r-ordering scheme retains the original colour information as each pixel is ordered based on some reference colour.

colours close to the desired colour are high in the order whereas any complimentary colours, the furthest from the reference, are lowest.

2.3 Alternate Proposals for a Colour HMT

There exist multiple proposals for colour morphology frameworks and colour HMTs in particular each with varying ways of defining morphological operators in colour and multivariate domains. Many of these general methods are reviewed in [4]. The most common ways morphology has been applied to multivariate spaces is through the ordering relations such as those described in Section 2.2.

In [7] and [18] the authors compare the marginal, channel-wise, processing with various vectorial ordering strategies. Both come to the conclusion that a marginal strategy can work in some cases but produces false colour artefacts which are potentially detrimental to further processing steps and the definition of a colour ordering is needed.

In [1] the author investigates the combination of r-ordering and c-ordering using the Mahalanobis colour distance in order to apply morphological operations to colour images. Similarly in [2] the authors also use reduced and conditional ordering but extend the theory to multiple alternative colour spaces. Neither defines a colour HMT.

In [48], the authors use the HSV colour space along with lexicographical (conditional) ordering, where the intensity values are considered most important over Saturation and in turn Hue and colours are ordered in relation to this scheme. Other recent approaches involve using the Loewner ordering method [6] or other alternative supervised ordering methods [44] or by attempting to learn and infer a complete lattice for colour image data [21].

In addition to colour morphology frameworks, there have been multiple colour and multivariate HMTs defined in the literature. An extension of Barat's HMT for colour morphology is proposed in [19] and uses the notion of convergence and colour distances within the CIELAB colour space in order to define colour erosion and dilation. In [45], the authors define an multivariate HMT where each channel is processed independently as in the marginal strategy but with a individual SE designed for each band. In [3], a Vectorial HMT (VHMT) based on a vector ordering, lexicographical in this case, is defined using two colour templates or structuring functions (SF), a lower SF and an upper SF, where objects that fit between the two are detected. The VHMT is seen to be sensitive to noise and therefore a robust VHMT was developed, also utilising rank-order filters. The rank parameters are obtained by defining a maximum number of non-matching pixels in both the foreground and background SEs that can be allowed for the object to still be detected. A downside of using such an

approach is that this measure will be highly specific to an application whereas in this work, alongside that presented in [29], a percentage is used to generalise more appropriately to this problem. In [43] the authors define operations based on h-supervised ordering. This supervised approach requires training data in order to define the foreground and background operators necessary. This may lead to problems in scenarios with few objects or conversely with many objects with varied appearance where multiple training sets may be required. In [42], the authors use a Percentage Occupancy HMT [27] to detect buildings in multispectral imagery. They define a local order within the SE based on extremeness where each element is compared to every other and the pair of infimum and supremum can be found.

In [50] the authors define a lexicographical ordering scheme based on colour similarity. This is closely related to our approach where the similarity in this case is calculated based on the likeness between the image values and the corresponding values in the SE - this allows for non-flat SEs to be utilised. For an image, I, and a composite SE, S, the Similarity HMT (SHMT) is defined as $\mathrm{SHMT}(\mathbf{I};\mathbf{S}) = \left\{ h^s_{\mathrm{similarity}} \left(\mathcal{E}_{h^s_{\mathrm{similarity}}} \mathbf{I}(x) \right) = 1, x \in \mathbf{E} \right\}$ for all points x in I where $\mathbf{E}$ is the 2D support of the image in $\mathbb{Z}^2$ and $h^s_{\mathrm{similarity}}$ and $\mathcal{E}_{h^s_{\mathrm{similarity}}}$ are the colour similarity and colour similarity based erosion respectively. This SHMT approach has the drawback of being restricted to only colour images with three channels due to the definition of similarity functions based on Finlayson's colour image formation model. The authors also present a noise robust variant, the SHMT_η, where η is some threshold value for similarity. Along with the limitation in spectral dimension, this threshold value can cause errors in detection. The SHMT_η is defined as $\mathrm{SHMT}_\eta(\mathbf{I};\mathbf{S}) = \left\{ h^s_{\mathrm{similarity}} \left(\mathcal{E}_{h^s_{\mathrm{similarity}}} \mathbf{I}(x) \right) \geq \eta, x \in \mathbf{E} \right\}$. A single noisy pixel can cause the similarity to decrease and consequently the threshold must also decrease which can introduce false positive detections if it is set too low. In our approach we combine a similar dissimilarity threshold with relaxation in the erosion and dilation using rank-order filters as opposed to hard minimum and maximum filters in order to attain improved results. This reduces the susceptibility of our method to noise and reduces the number of false positives introduced when compared with the SHMT.

3 Proposed Reduced Ordering-based HMT

Our proposed algorithm, the Multiple Dimensional Percentage Occupancy Hit-or-Miss Transform (MDPOHMT), provides a robust, controllable approach using a reduced ordering scheme based on point-wise distances between two multivariate SEs and an image. Each colour or multivariate observation can be thought of as a single point in a multivariate space of the same domain. By using a reduced ordering, these multivariate observations, which are themselves difficult to order and quantify, are reduced into an easily sorted scalar quantity based on some reference – in this case the corresponding pixel of the coincident SE. By using reduced ordering, the correlation of the input vectors are also preserved unlike in some methods. SEs of the same dimensionality as an image can be designed in order to detect objects of interest through translation and comparison at each pixel in a set of query images. For a more varied set of objects, or ones which vary in scale, multiple SEs can be generated and tested rapidly. By using multivariate SEs, both flat and non-flat, *i.e.*, multi-coloured, SEs can be created in order to fit to more complex or textured objects in images. This also constrains the colours considered as a match since the desired colour information is encoded into the SE. For simplicity, in this paper, the distance between an image pixel and the reference is calculated using the Euclidean distance in the RGB colour space, however the choice of colour space and distance measure is an important consideration and is often highly dependent on the application. Our proposed MDPOHMT is intended as a general technique which can be readily applied to any image regardless of the colour space and distance measure used.

3.1 Definition

We consider the set of images $\mathcal{I}$ from $\mathbb{Z}^n$ the discrete domain of dimension $n > 0$ into the interval $[0, 1]^n$: *i.e.*, $\mathcal{I} = \mathbb{Z}^n \rightarrow [0, 1]^n$. Given an image $I \in \mathcal{I}$, the support of I, denoted by $\mathrm{supp}(I)$, is the set of points where I

is non-zero: *i.e.*, $\text{supp}(I) = \{x \in \mathbb{Z}^n, I(x) \neq 0\}$. In the following, we will always assume that the considered images have a finite support. Let $\hat{d}$ be a normalised positive dissimilarity measure on $[0, 1]^n$, in the case of the Euclidean distance function utilised as a dissimilarity function, the normalisation is achieved by dividing the dissimilarity values, d, by the number of channels:

$$\hat{d}(a, b) = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(a_i - b_i)^2} \tag{8}$$

Given two images A and B in $\mathcal{I}$ and a point p in $\mathbb{Z}^n$, we define the indexed family of normalised dissimilarities, DS, between A and B at point x by:

$$DS(A, B, x) = \left\{\hat{d}(A(x + b), B(b))\right\}_{b \in \text{supp}(B)} \tag{9}$$

Given a finite indexed family $F = \{f_j\}_{j \in J}$ of elements in $[0, 1]$ and an integer k between 1 and $\text{Card}(J)$, we denote by $rank_k(F)$ the k-th smallest element of F (thus we have $rank_1(F) = \min(F)$ and $rank_{\text{Card}(F)}(F) = \max(F)$).

Given an image I and our two structuring element images SE_{FG} and SE_{BG}, and using a similar process to the greyscale POHMT from (7) the fitting of the foreground SE and the image or the background SE and the image at a point x can be defined using rank order filters, identical to those from (5), in place of erosion and dilation. The rank of these filters, k, is based on the ranked dissimilarity using percentage occupancy values of $50 < P \leq 100$, with the value of k being obtained from (6) and can be expressed as

$$[\xi_{SE,k}(I)](x) = k\text{th rank}\left\{DS(I, SE, x)\right\} \tag{10}$$

The range of PO values must lie between 50 and 100% as 50% occupancy would simply be a median filter and anything below would cause an erosion to tend towards a dilation and vice versa. Combining two of these operations, one which emulates an erosion and another for a dilation, along with a variable similarity threshold, $T \in [0, 1]$, allows us to define a new multivariate HMT-like operator as seen in (11), where k_P is the appropriate rank based on some desired PO value given in (6).

$$[\text{MDPOHMT}_S(I)](x) = \begin{cases} 1 & \text{if } [\xi_{SE_{FG},\text{Card}(SE_{FG})-k_P+1}(I)](x) < [\xi_{SE_{BG},k_P}(I)](x) \\ & \text{and } [\xi_{SE_{FG},\text{Card}(SE_{FG})-k_P+1}(I)](x) \leq T \\ 0 & \text{otherwise} \end{cases} \tag{11}$$

If P = 100%, (10) reduces to $\min(DS(I, S_{FG}, x)) - \max(DS(I, S_{BG}, x))$, which means that we compare the largest dissimilarity between the image and the foreground SE to the smallest dissimilarity between the image and the background SE. If the image is close to the foreground SE and far from the background SE the fitting value should be large. Relaxing the precision by reducing P leads to an increased true positive rate but can also result in an increased false positive detection. As well as the rank parameter, the threshold T can be adjusted to specify a maximum dissimilarity between the image and the corresponding SE that can register a hit, this can account for variability of colour in various objects. For example, if we are interested in red objects, the threshold can be set in such a way as to rule out other colour possibilities. Figure 5 displays the operation of the MDPOHMT in pictorial form, showing how adding rank order filter equivalents in the place of erosion and dilation can aid in noise robustness.

3.2 Rank estimation and noise robustness

It has been shown previously how rank order filters can be used in order to add robustness to both binary and greyscale HMTs. The PO plot method defined in [27] can be used to estimate the most appropriate rank from example objects and indeed can be extended for use with our MDPOHMT. When centred at a pixel in a greyscale image, the intensities of the image pixels are compared with their corresponding SE values. The occupancy can then be calculated for foreground and background by determining the percentage of pixels that

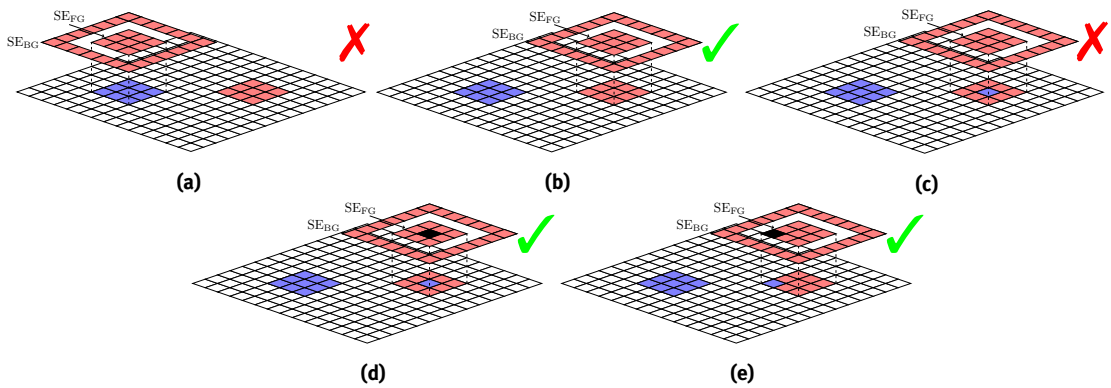

Figure 5: MDPOHMT in pictorial form depicting an image plane as well as the composite SE, $S = [SE_{FG}, SE_{BG}]$. In **a)** the algorithm will return a relatively large dissimilarity between the blue image and the red SE_{FG}. **b)** SE_{FG} matches the image producing a small to zero dissimilarity relative to that produced by comparing SE_{BG} with the white background and thus detects the object. **c)** The object has been corrupted with noise in the form of a flipped pixel and therefore cannot be detected without the use of percentage occupancy or another relaxation method on the foreground erosion. **d)** By relaxing the transform using percentage occupancy, the pixel containing the noise can be overlooked as the foreground erosion no longer takes the maximum distance instead taking a lower ranked distance. **e)** The relaxation is designed to be invariant to the location of the noisy pixels.

lie below some grey-level, t, for the foreground SE, essentially the percentage that "fit" within the foreground of the object and conversely the percentage that are above t in the background SE. This produces two occupancy vectors which contain the foreground and background occupancies at each grey-level respectively. When plotted against one another, these two vectors produce what is known as a PO plot at the particular pixel under test.

In our case, when applied to multivariate images, these estimated occupancies need to relate to colour and multivariate dissimilarity as opposed to greyscale intensity. This, along with the lack of grey-level analogue in a colour image, means a slightly different approach is required in order to extend the notion of PO plots for use with colour images. This extension of the PO plot rank parameter setting method to the multivariate domain is one of the main contributions of this work and provides a simple and effective approach to improve the recognition rate of objects and patterns in the presence of noise.

In order to obtain the occupancy vectors using multivariate images and SEs the probability distribution function (PDF) of the ordered normalised dissimilarity for each SE is estimated. This is achieved by binning the dissimilarities in a histogram for both the foreground, H_{FG}, and background, H_{BG} before determining the occupancy vectors PO_{FG} and PO_{BG} using cumulative histograms where $PO_{FG}(n) = \sum_{m=1}^{n} H_{FG}(m)$ and $PO_{BG}(n) = 1 - \sum_{m=1}^{n} H_{BG}(m)$.

Here, the occupancy vectors indicate whether the foreground SE or background SE is more similar to the image at any given rank. Following the notion that the foreground SE must fit the image better than the background SE, a critical point corresponding to the minimum occupancy required to detect an object can be found as $P = \max[\min\{PO_{FG}, PO_{BG}\}]$. This critical point can also be found by plotting both PO_{FG} and PO_{BG} together and finding their point of intersection. Plotting PO_{FG} against PO_{BG} producing a PO plot also highlights this critical point as the location of the graphs intersection with the straight line $PO_{FG} = PO_{BG}$.

An example of this can be seen in Figures 6b and 6e for clean and noisy test images respectively. The noise was added using the methods described in [8]. By finding the maximum of the minimum of all points in the cumulative distribution functions (CDF) of the BG and FG, or the point of intersection between the two, (indicated by the horizontal black lines in Figure 6b and 6e) an estimate of PO can be found. In both the clean and noisy cases, the object of interest is the red square in the upper left of each image. In the case of the uncorrupted object (Figure 6b) it was, as expected, 100% occupancy whereas for the object corrupted by Gaussian noise (Figure 6e) it fell to 84% in order for the object to be detected. The second method involves plotting the estimated CDF of the foreground against that of the background and by analysing where this curve crosses the diagonal the PO can be estimated. In the example shown in Figure 6 analysing the PO plots gave the same values of 100% and 84% occupancy to detect the clean and noise corrupted objects as was

the case when analysing the CDF. Once the PO estimation has been made, the rank parameter, k_P, can be found for each SE using (7). Figure 7 shows the PO plots for each of the red, green and blue SEs when centred

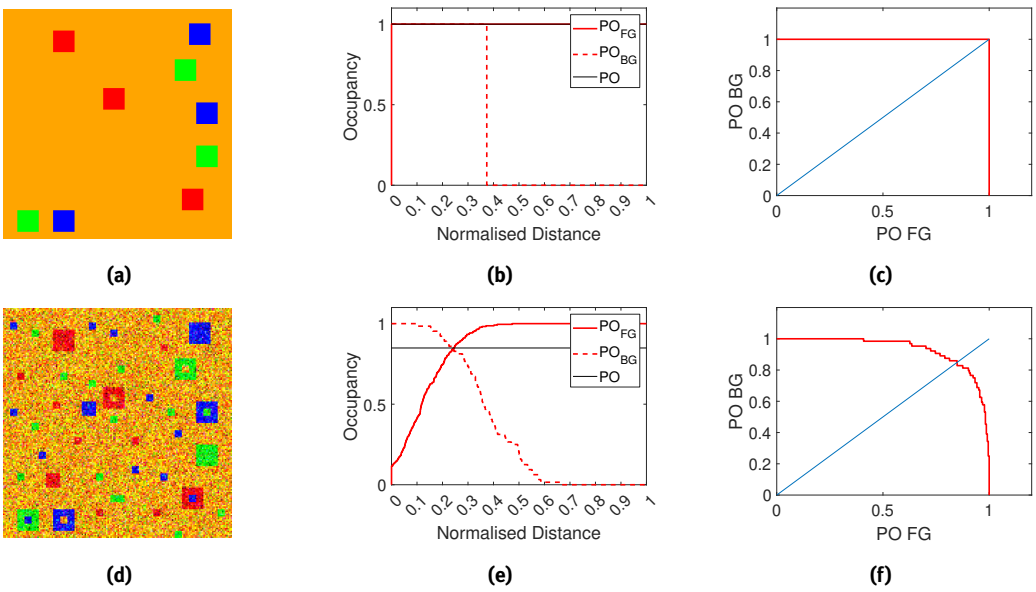

Figure 6: PO plots and setting of the appropriate PO and rank parameters. **a)** Clean 160 × 160 test image. **b)** Occupancy of the foreground and background SEs. **c)** PO plot of the foreground and background SEs. **d)** Noisy 160 × 160 test image. **e)** Occupancy of the foreground and background SEs. **f)** PO plot of the foreground and background SEs.

on their respective objects in the two test images. In the case of the clean image (Figure 7a) all nine objects

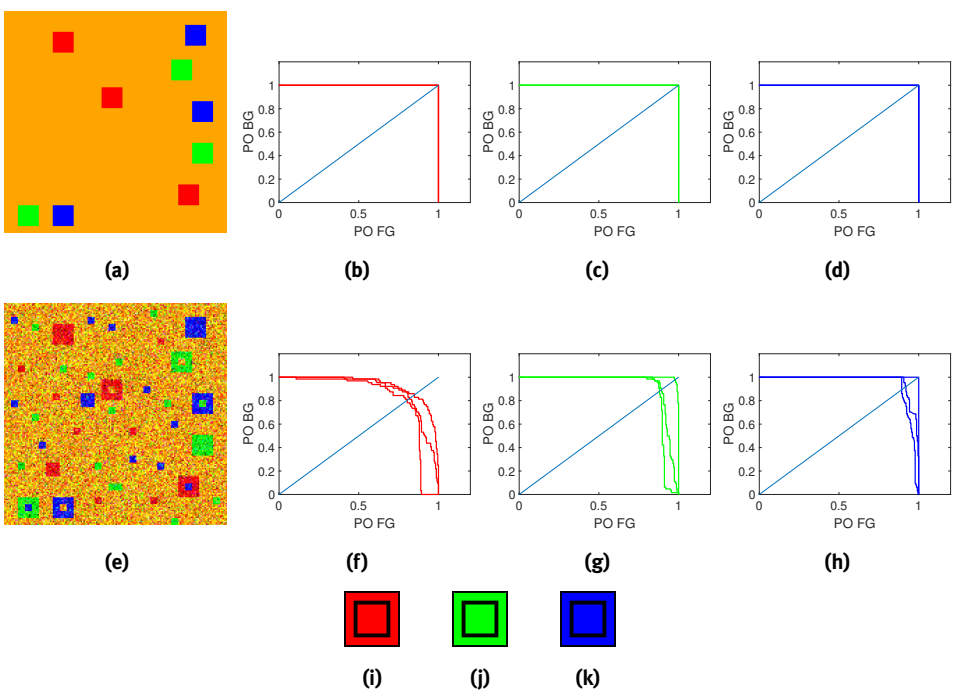

Figure 7: Percentage occupancy plots of coloured square in test images. **a)** Clean 160 × 160 test image. **b)** PO plots for the red objects. **c)** PO plots for the green objects. **d)** PO plots for the blue objects. **e)** Noisy 160 × 160 test image. **f)** PO plots for the red objects. **g)** PO plots for the green objects. **h)** PO plots for the blue objects. **i)** Red SE. **j)** Green SE. **k)** Blue SE.

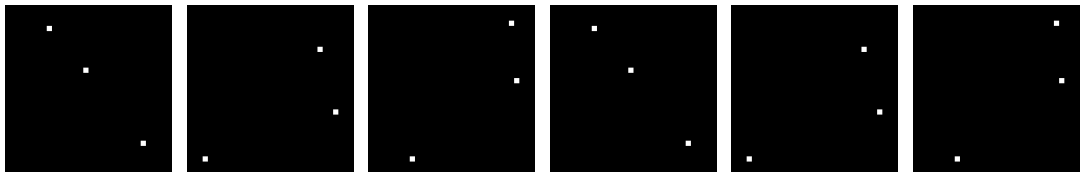

(a) (b) (c) (d) (e) (f)

Figure 8: Results of the MDPOHMT on the two test images from Figures 6a and 6d. **a)** Red squares detected in the clean image using the MDPOHMT with $P = 100\%$. **b)** Green squares detected in the clean image using the MDPOHMT with $P = 100\%$. **c)** Blue squares detected in the clean image using the MDPOHMT with $P = 100\%$. **d)** Red squares detected in the noisy image using the MDPOHMT with $P = 70\%$. **e)** Green squares detected in the noisy image using the MDPOHMT with $P = 70\%$. **f)** Blue squares detected in the noisy image using the MDPOHMT with $P = 70\%$.

are fully fitted with their respective SEs, as shown in Figures 7i-7k. As each of the SEs match the objects in question, this produces the square PO plots, which as the plot intersects with the top right corner, indicates a perfect fit. The clean image was then corrupted with random Gaussian noise as well as foreground removal to create non-uniform shapes. The percentage occupancy was again calculated at the same positions resulting in the PO plots shown in Figures 7f-7h where the PO plot begins to curve, indicating a non exact fit. The PO and resulting rank estimation can be carried out through investigation of these plots, namely where each of the occupancies intersect with the diagonal.

Applying the MDPOHMT to the test image shown in Figure 7a with 100% occupancy results in an ideal fitting as expected through analysis of the PO plots and verified by the HMT result shown in Figures 8a-8c for the red, green and blue objects respectively. However, applying the MDPOHMT with the same parameters to the corrupted image (Figure 7e) results in none of the objects being detected. This is because the SE no longer fits the image as it is being "punctured" by noise. By estimating and setting the appropriate rank using PO plots and applying this to the MDPOHMT the objects are again detected correctly, Figures 8d-8f. Despite each of the objects having varying noise and flipped pixel colours, through analysis of the PO plots and appropriate rank selection they are all successfully detected.

The rank parameter and selection is intended to be intuitive as possible and represents a percentage fit between an image under test and a probing SE. A percentage is used as there is no need to define which pixels must fit inside the SE. While the PO plot method provides a way of setting the optimal parameter for detecting an object, there may exist some cases where few or no samples are available to train such methods. In these cases, the rank parameter can be estimated easily simply by trialling various percentages, under the assumption that a high percentage fit will restrict the detected objects and a lower percentage will relax the transform [29]. The transform is computationally inexpensive enough to be performed multiple times with different test rank parameters until a positive outcome is reached.

4 Experimental Results

In this section, we apply our algorithm to various synthetic and natural colour and hyperspectral images. We first compare our MDPOHMT to an existing greyscale HMT in order to highlight the need for a generalisation to colour and multivariate imagery in order to better distinguish between objects of interest. We then compare our approach with other colour HMTs and investigate the benefits of the noise robustness offered by using PO. We also show how the MDPOHMT can be used to detect objects with varying size, shape, colour and textural information from high resolution images. Finally we show how the reduced ordering can be performed on higher dimensional data and apply the MDPOHMT to hyperspectral object detection using a synthetic dataset.

4.1 Image Acquisition

The primary source of images used when testing our MDPOHMT is the DOTA dataset [47] which contains thousands of images taken from Google Earth with 15 object classes to be detected. For the purposes of testing the MDPOHMT, a subset of these images which contain swimming pools are used as they have varying size, shape, orientation and colour information as well as having varying amounts of pixels on target, providing a very varied target set. The dataset contains a ground truth for its training and validation subsets, ground truths for further images are created in order to increase the volume of data available to test on. In total, 7 high resolution colour images were selected for testing with a total of 1257 instances of the swimming pool object included. Other natural and synthetic images were used for testing our MDPOHMT against other greyscale and colour HMTs.

4.2 Comparison with Soille's Unconstrained HMT (UHMT)

The greyscale UHMT [41] and other greyscale HMTs have been shown to be useful for various object detection applications [11, 32]. However, there are situations where colour information being taken into account can be advantageous and provide an extra dimension to a detection decision.

Figure 9 shows an example use of the MDPOHMT in detecting traffic lights of various colours. In Figure 9a a colour image of a traffic light with its red and amber lights lit is shown. From this image, two flat SEs were created (Figures 9b and 9c respectively) in order to detect the red and amber lights. The SEs were originally made from cropped regions of the image as to match exactly with a single pixel but were flattened using the average value in each region in order to be generalised. Figure 9d shows the locations where the red light is detected successfully using the red SE from Figure 9b. Similarly, the amber light is detected as shown in Figure 9e using the amber SE shown in Figure 9c.

This same task was carried out using the greyscale unconstrained HMT (UHMT). Figure 10a shows the image from 9a converted to greyscale using the weights used to calculate the luminance, Y, in the NTSC colour space. The red and amber SEs from Figures 9b and 9c were also converted to greyscale and are shown in Figures 10b and 10c. Applying these greyscale SEs, or indeed a set of binary SEs, to the image yields the results shown in Figures 10d and 10e where both the red and amber lights are being detected by both SEs. This is

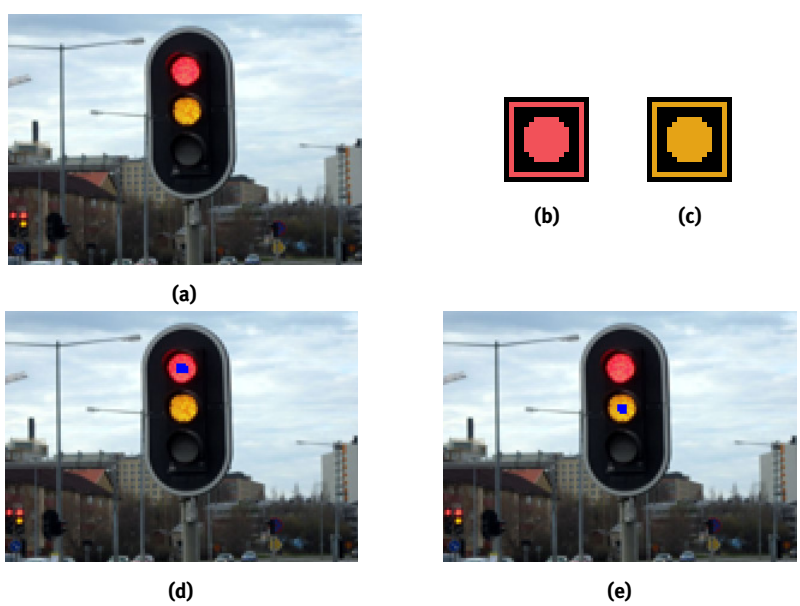

Figure 9: a) Original colour image of a traffic light. **b)** Circular SE used to find the red light. **c)** Circular SE used to find the amber light. **d)** Results of taking the MDPOHMT of (a) using (b) with P = 90%. **e)** Results of taking the MDPOHMT of (a) using (c) with P = 90%.

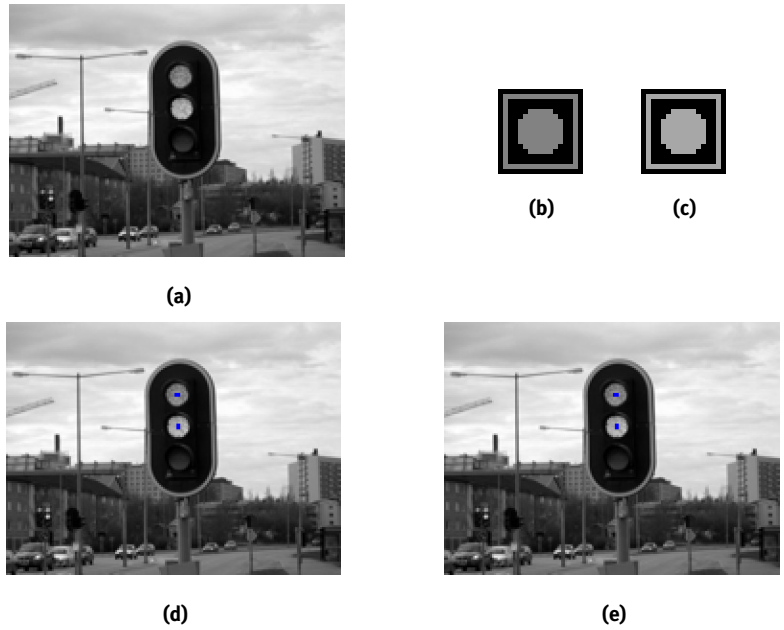

Figure 10: a) Original image of a traffic light converted to greyscale. **b)** Circular SE used to find the red light. **c)** Circular SE used to find the amber light. **d)** Results of taking the greyscale UHMT of (a) using (b). **e)** Results of taking the greyscale UHMT of (a) using (c).

due to both SEs fitting both objects based solely on similar intensities and with no colour information available to discriminate between the two there is confusion. This highlights the benefits of incorporating colour information into the HMT in order to better discriminate between colours with similar greyscale intensities.

4.3 Comparison with the colour Similarity HMT (SHMT)

The MDPOHMT is based on a reduced ordering between a multivariate image and a multivariate SE, turning this multivariate data into a scalar for ordering. Another similar approach is the SHMT [50] which calculates the colour similarity between a multivariate image and SE pair. These similarity values along with the image's individual colour channels are ordered lexicographically, where if there are two identical similarity values they are ordered by the corresponding green values, if their green values are identical then they are ordered by red and then, if their red values are identical, by blue. The infimum, or erosion, of the now ordered similarity values is taken and the SHMT is defined as all points in the image where this value of similarity is equal to one. The SHMT_η further constrains this where the minimum similarity in the SE window must be greater than some threshold η. A comparison between our MDPOHMT and the SHMT on a synthetic test image is shown in Figure 11.

Figures 10f and 10k show a colour test image and the same colour image corrupted with Gaussian noise as well as flipped pixels. These images are used to verify and compare the SHMT and the MDPOHMT. The composite SE used is the same as that shown in Figure 7i. The SEs are designed to detect the three 3×3 red objects in the test images. Figure 10j shows the results of the SHMT on the clean test image from Figure 10f with the corresponding similarity at each pixel shown in Figure 10i. At all pixels where the similarity between the SE and the image is one, *i.e.*, they match, the output of the SHMT registers a hit. Figure 10h shows the results of the MDPOHMT with the corresponding foreground dissimilarity shown in 10g. The MDPOHMT registers a hit where the foreground SE matches the image better than the background SE, or where there is lower dissimilarity in the foreground than the background. It can be observed that the two methods of ordering the data are somewhat inverted, with the SHMT being based on colour similarity and the MDPOHMT being based on a colour distance, or dissimilarity. This is apparent when comparing Figures 10i and 10g where, at

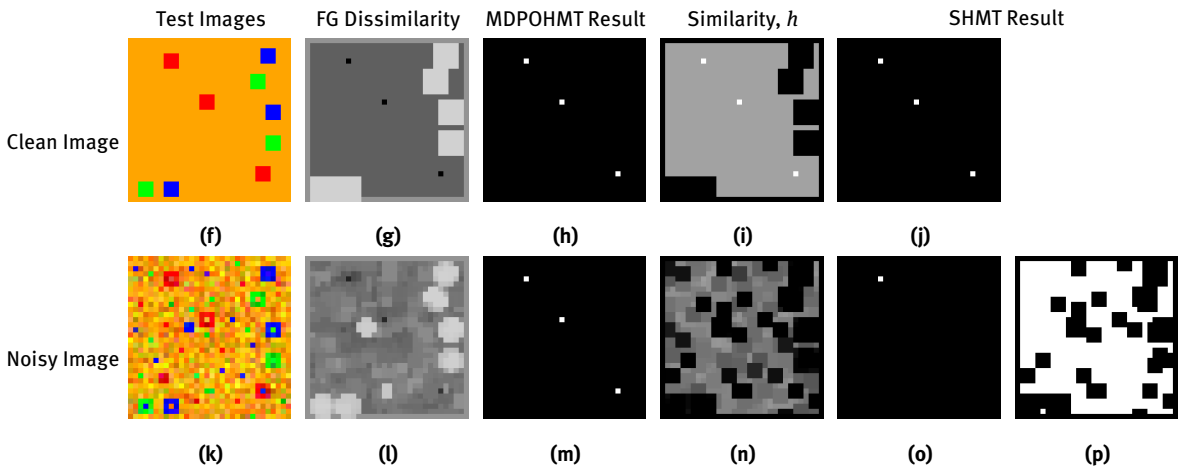

Figure 11: Comparisons between the MDPOHMT and the SHMT. **a)** Dissimilarity between the image from c) and the SE with P = 100%. **b)** Results of the MDPOHMT on the image from c) where P = 100%. **c)** Clean test image (32×32). **d)** Similarity of each pixel in the test image in c) to the SE. **e)** Result of the SHMT on the image from c). **f)** Dissimilarity between the image from h) and the SE with P = 80%. **g)** Results of the MDPOHMT of h) where P = 80%. **h)** Noisy test image (32×32). **i)** Similarity of each pixel in the test image in h) to the SE. **j)** Results of taking the $\text{SHMT}_{0.6}$ of h). **k)** Results of taking the $\text{SHMT}_{0.05}$ of h).

the pixels where the HMT registers a hit, the similarity values of the SHMT are one and the dissimilarity of the MDPOHMT values are zero.

The noise robustness of each method was then assessed by applying both the SHMT and MDPOHMT to the Gaussian noise corrupted image in Figure 10k. When applying the standard SHMT where $\eta = 1$, *i.e.*, where no thresholding is applied, the transform fails to detect any of the three red objects as the SE no longer matches the image perfectly and the colour similarity, h, will be in the range $0 \leq h < 1$. By investigating the similarity at each of the appropriate pixels (Figure 10n) a valid threshold value can be obtained to ensure all objects of interest can be detected.

The SHMT_η applies a single threshold η across the similarity image in order to relax the similarity required for the detection of an object in the presence of noise. The result of applying a threshold of $\eta = 0.6$ on the similarity image shown in Figure 10n can be seen in Figure 10o where only one of the three objects of interest is detected successfully. The other objects, containing flipped blue and orange pixels respectively, have a low minimum colour similarity under the probing SE, with pure red and pure blue having zero chromatic similarity. Due to this low colour similarity, the threshold must be lowered below the minimum similarity value present when the SE is centred over an object of interest. In order for all three of the red objects to be detected in the noisy case, the threshold needs to be set less than $\eta = 0.05$ to compensate for the flipped blue pixel in the lower right red square. However, lowering the threshold to this level introduces a large number of false positives detections as seen in Figure 10p.

The SHMT_η is essentially operating as an erosion and as such only considers the minimal value of similarity between the image pixels and corresponding reference values within the support of an SE. This could be improved by using a rank order filter to relax this erosion in a similar manner to the proposed MDPOHMT. If the hard minimum similarity is relaxed situations, like that presented in Figure 10k, where there are low similarity pixels present in the support of an SE causing the SHMT to miss an object can be avoided. By incorporating such rank order filters into the proposed method, the MDPOHMT can be made more robust to these effects of noise. The results of taking the MDPOHMT with a P = 80% are shown in Figure 10m with the foreground dissimilarity shown in 10l, with the locations of the three objects visibly darker than the rest of the dissimilarity image.

Further comparisons with the SHMT using more natural images were then carried out. Figure 12c shows an image of white and yellow flowers in a grassy background. This image was used in [49] to validate the performance of the SHMT and SHMT_η. Figures 12a and 12d show the SE from [49] used in the SHMT and the similarity of each pixel respectively. The majority of the image is black meaning there is low chromatic

similarity between the image and SE at those pixels. The SHMT_η is essentially a thresholding operation on this similarity image where the threshold is set to the value of η. Setting the threshold at lower values of η allows for the detection of decreasingly similar objects as shown in Figures 12e-12g which show the results of carrying out the SHMT_η where η is reduced in each subsequent test: $\eta = \{0.7, 0.5, 0.3\}$. The SHMT_η operator does not allow for any relaxation of the shape of an SE and therefore requires much more specific SE design and can be more sensitive to changes in the object as well as noise or occlusion.

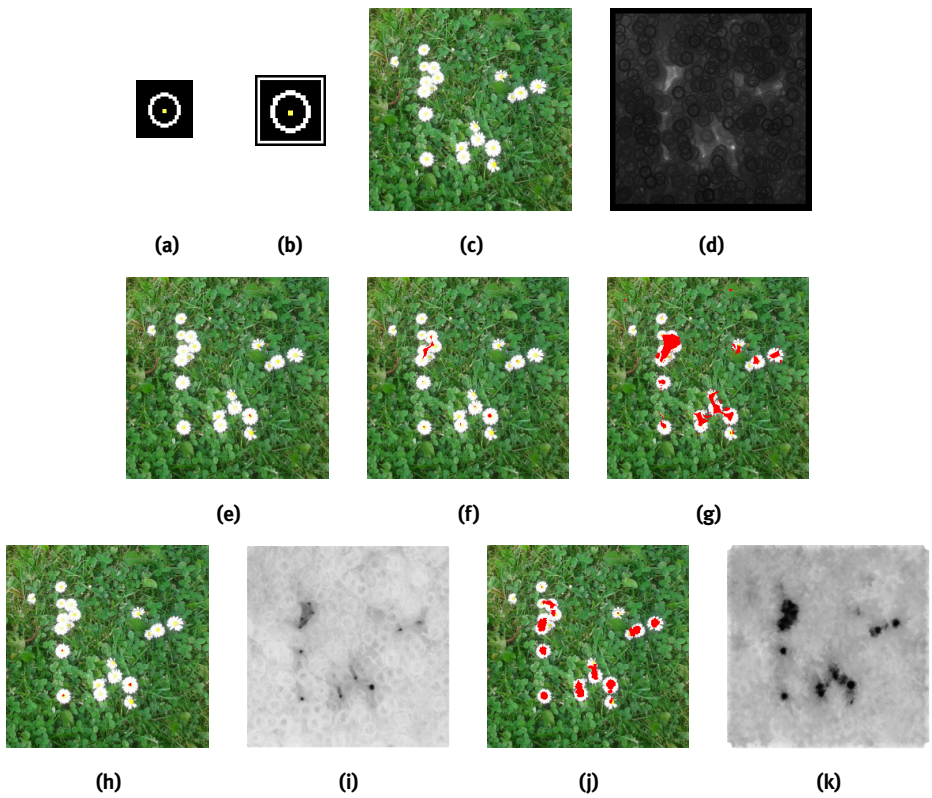

Figure 12: Comparison between the SHMT and the MDPOHMT on a real image. **a)** SE for SHMT_η. **b)** Composite SE for MDPOHMT. **c)** Test image "flowers". **d)** SHMT chromatic similarity. **e)** Results of the $\text{SHMT}_{0.7}$. **f)** Results of the $\text{SHMT}_{0.5}$. **g)** Results of the $\text{SHMT}_{0.3}$. **h)** MDPOHMT with P= 90%. **i)** Dissimilarity with P = 90%. **j)** MDPOHMT with P = 60%. **k)** Dissimilarity with P = 60%.

When $\eta = 0.7$ (Figure 12e) only a couple of pixels in the centre of one flower are detected, this was the point most similar in shape to the designed SE. Reducing this further to $\eta = 0.5$ (Figure 12f), half of the flowers present in the scene have at least one positively detected pixel, however a patch of false positive detections spreads across a number of closely grouped flowers. At $\eta = 0.3$ (Figure 12g) this patch has grown to cover multiple flowers, this is due to the similarity being set so low as to introduce pixels which no longer exhibit high similarity. Referring to Figure 12d, which shows the similarity of each pixel to the probing SE, it is possible to see that there are some flowers which can never be detected with this SE - highlighted by zero or low similarity. Attempting to lower η further in order to detect them would yield results similar to those in Figure 10p, where multiple false positive pixels are introduced. In [49], the author explains that some flowers cannot be detected as they are smaller than the generalised SE or that they are damaged to the point they are too dissimilar.

We then compared the performance of the SHMT with our MDPOHMT using the same test image with the results shown in Figure 12. The foreground SE used is identical to that used in the SHMT (Figure 12a) and the background SE is a hollow white square which can be seen in the composite SE shown in Figure 12b. The MDPOHMT outperforms the SHMT even with high values of PO, at P = 90% it detects the same flowers as the

SHMT with $\eta = 0.5$ however much more precisely with no false positives (Figure 12h). In an attempt to detect the other damaged or smaller flowers, the PO value was reduced to P = 60% as shown in Figure 12j. This improves the number of flower centres detected but there is also some patches of the white petal detected, this is because as the PO decreases, lower percentages of pixels are allowed to fit and in this case as the white petals are a feature of the foreground SE (Figure 12b). Areas where white petals are prevalent are also detected as they are not dissimilar to the SE. However, by lowering PO we are not decreasing the similarity required for a fit, merely allowing dissimilar pixels within the support of the SE and ignoring them in order to detect areas that have a lower percentage match. As discussed in Section 2.3 and shown in Figure 10p, the SHMT is susceptible to a single noisy pixel with low chromatic similarity to the target requiring such a low threshold to be set that large numbers of false positives are introduced. The h-similarity measure may be incorporated as a distance measure within the MDPOHMT and by relaxing the percentage fit using rank order filters, as well as varying the thresholded value η, the detection results can be improved when compared to using the SHMT_η.

4.4 Application of the MDPOHMT to high resolution Aerial Imagery

The MDPOHMT is also designed to generalise to objects of interest over multiple images. In order to validate this, the MDPOHMT was applied to a set of aerial images from the DOTA dataset to test its performance for detecting various sizes and shapes of colour objects, in this case swimming pools. This class of objects was chosen for the challenge that it presented with varying size, shape, colour and level of occlusion. In order to detect these objects first a generalised SE must be designed from examples or from prior knowledge of the objects of interest. After viewing some examples, the average colour size and shape can be used to create an appropriate SE.

In order to generalise well to objects, including those which are noisy or partially occluded, the fit often needs to be relaxed. This is preferable over creating multiple bespoke SEs as this not only takes time but increases the time required to probe an image with these multiple SEs. The optimal relaxation parameters can be found in one of two ways, either utilising the PO plots from [27] in order to train the parameters on similar objects of interest or to perform a greedy search optimisation and use the parameters which maximise some desired measure such as accuracy or F1 score. In order to select the optimal parameters using PO plots, multiple pixels in a training image which contained positive examples of swimming pools were selected. The MDPOHMT was then taken at each of these points with the designed SE and the foreground and background occupancy for each point was plotted in a PO plot, and an appropriate value can be selected.

Figure 13 shows the approach for estimating the appropriate PO parameters, where multiple points in a representative image (Figure 13a) are selected and the occupancy of both the foreground and background SEs are taken and used to generate a PO plot as in Figure 13b. By observing the plot of foreground against background occupancy appropriate values for the PO parameter are set. In the case of Figure 13, a PO value of P = 80% would be sufficient to detect each of the three swimming pools in the scene when the SE is centred at that particular pixel. This was validated by performing a grid search with multiple PO values as inputs in order to maximise the F1 score.

As well as the PO parameter, an additional colour distance threshold, T, where, $0 \leq T \leq 1$, was introduced to ensure that the dissimilarity between the image and the foreground SE for any prospective match was limited. In order to set this parameter, a number of cropped regions from each of the test images were extracted and the MDPOHMT was applied with the PO set to P = 80%. A number of colour distance thresholds were tested on the basis of the resultant F1 score in order to optimise this parameter. It was found that the optimal threshold value for a set of training images is $T = 0.15$.

With this we now have the set of parameters to test on the full set of images. Figure 14 shows the results of taking the MDPOHMT on a full image from DOTA, where detected swimming pools are outlined in green and missed pools in red. Examples of the swimming pools that are detected and missed are shown in Figure 15. The full results from applying the MDPOHMT to the images from DOTA are shown in Table 1, where the image shown in Figure 15 is "*Image 7*" in Table 1.

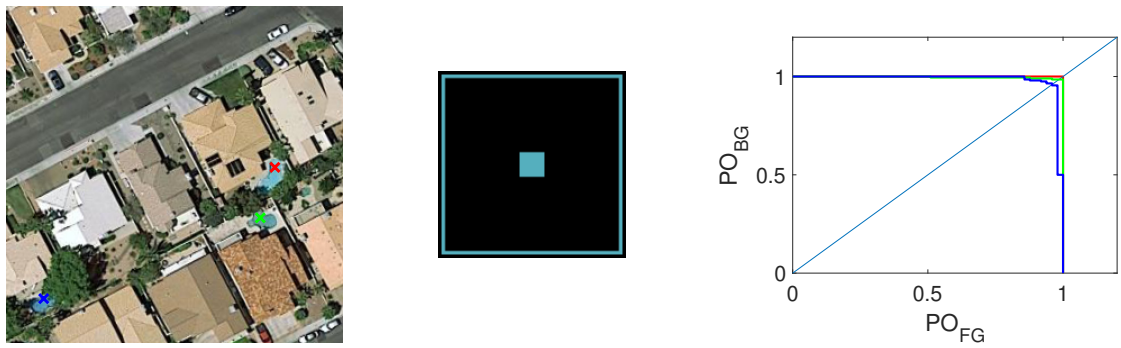

(a) (b) (c)

Figure 13: a) Cropped image with multiple pixels of interest overlaid. **b)** Composite SE used to detect the swimming pools in a). **c)** PO Plots at the three selected pixels of interest.

On average the method performs well, achieving an accuracy of 84% with high F1 scores as well as high precision and recall. The number of false positive detections is very low, with only eleven false positives detected over all of the seven images. The majority of objects are successfully detected, as indicated by the high number of true positives as well as the high accuracy. Those that are not detected are often missed due to significant occlusion, shadow, or other changes in colour. The occlusions can be overcome by reducing the PO value further, however, with this comes the risk of also increasing the number of false positives if the transform is relaxed in excess. Alternative SEs can be designed by looking at examples that a single SE has missed in order to improve the detection rate and recall of the MDPOHMT. In this particular example, another three SEs were used to detect objects the first had missed, using information from the examples of false negatives to design appropriate SEs. This process, however, comes at the expense of execution time, highlighted by the results in Table 1. This has the desired effect of increasing the recall of the MDPOHMT at the expense of some precision due to false positives introduced by each SE.

In order to compare the MDPOHMT with other related techniques, the SHMT was also applied to the same subset of the DOTA dataset with identical SEs to those used in the MDPOHMT. The results of the SHMT with a single SE and with multiple SEs are shown in Table 1. From these results we can see that the SHMT fails to generalise due to the need for the SE to fit the query image more exactly. However, similarly to the results of the MDPOHMT, the addition of other SEs improves the results but the average F1 score is still 0.211 below that of the MDPOHMT. This can be seen further in Figure 16 where the precision-recall curves of the MDPOHMT and SHMT with both a single and multiple SEs are plotted.

Figure 16 shows the results of the MDPOHMT with both a PO value of $P = 100\%$ as well as a relaxed transform with $P = 80\%$. Both the MDPOHMT and SHMT perform similarly with a single SE both achieving high precision before decreasing. Using multiple SEs with the MDPOHMT improves the recall slightly as indicated by the curve shifting to the right before tailing off. Using multiple SEs with the SHMT has a similar effect, with the recall being slightly increased, however, this comes at the expense of precision with false positives introduced in both methods. Examples of the false positive detections introduced can be seen in Figure 17 and are mostly grey roofs with other blue objects such as umbrellas, shipping containers and other clutter within the image, these can ultimately be filtered out through further optimisation of the SEs used. By reducing the PO value of the MDPOHMT to 80% (Figure 16), the recall can be improved without a significant decrease in precision. This is true for both the single SE test and multiple SE test. From the results in both Table 1 and Figure 16 we can show that our MDPOHMT offers an advantage over similar techniques. From Table 1 we can also see that the MDPOHMT has a faster execution time compared to the SHMT as the h-similarity measure is more computationally expensive to determine.

Figure 14: Image from the DOTA test set used to validate the detection of swimming pools where objects detected using the MDPOHMT are highlighted in green and those missed are shown in red.

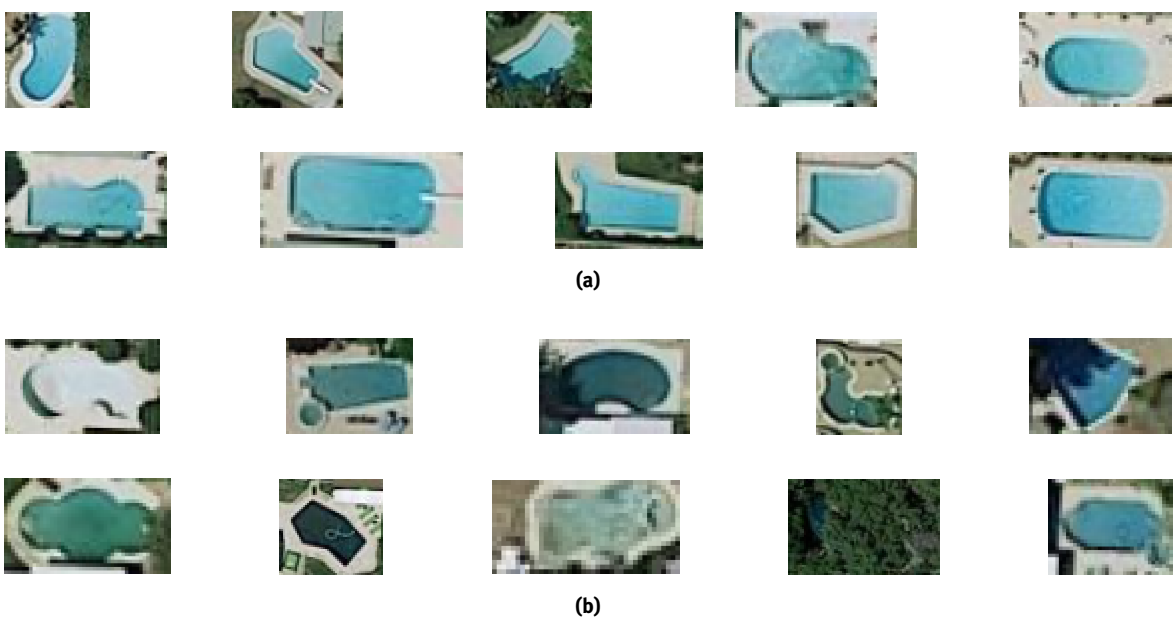

Figure 15: Examples of: **a)** 10 positively detected swimming pools, **b)** 10 missed swimming pools from Figure 14

4.5 Application of the MDPOHMT to object detection in hyperspectral imagery

Another limitation of the SHMT, as discussed in Section 2.3, is that the model for similarity can only consider 3 channel colour images in either the RGB or L*a*b colour spaces [50]. A general method for images in higher dimensionalities such as those present in multispectral or hyperspectral imaging applications would be of benefit in object and target detection tasks. The reduced ordering method employed in our MDPOHMT is not

Table 1: MDPOHMT and SHMT results on a subset of the DOTA Dataset using both single and multiple SEs.

MDPOHMT - $P = 80\%$, $T = 0.15$								
Image Number	1	2	3	4	5	6	7	Average
Processing Time (s)	4.2	2.5	1.1	0.9	2.8	2.1	2.6	-
No.True Positives	236	125	95	91	197	164	157	-
No. False Negatives	59	15	23	20	28	22	37	-
No. False Positives	0	1	0	3	2	0	5	-
Accuracy	80.00	89.29	80.51	81.98	87.56	88.17	80.93	84.06
F1 Score	0.8889	0.9398	0.8920	0.8878	0.9290	0.9371	0.8820	0.9081
Precision	1.0000	0.9921	1.0000	0.9681	0.9899	1.0000	0.9691	0.9885
Recall	0.8000	0.8929	0.8051	0.8198	0.8756	0.8817	0.8093	0.8406
MDPOHMT Multiple SEs - $P = 80\%$, $T = 0.15$								
Image Number	1	2	3	4	5	6	7	Average
Processing Time (s)	16.5	10.1	4.5	3.7	11.4	8.3	10.6	-
No.True Positives	264	132	102	96	202	173	168	-
No. False Negatives	31	8	16	15	23	13	26	-
No. False Positives	9	5	21	22	12	20	26	-
Accuracy	89.49	94.29	86.44	86.49	89.78	93.01	86.60	89.44
F1 Score	0.9296	0.9521	0.8465	0.8384	0.9203	0.9129	0.8660	0.8952
Precision	0.9670	0.9635	0.8293	0.8136	0.9439	0.8964	0.8660	0.8971
Recall	0.8949	0.9429	0.8644	0.8649	0.8978	0.9301	0.8660	0.8944
SHMT$_{0.9}$								
Image Number	1	2	3	4	5	6	7	Average
Processing Time (s)	22	14.2	6.4	5.3	15.9	11.7	14.6	-
No.True Positives	76	54	47	45	94	102	94	-
No. False Negatives	219	86	71	66	131	84	100	-
No. False Positives	0	0	0	0	0	0	0	-
Accuracy	25.76	38.57	39.83	40.54	41.78	54.84	48.45	41.40
F1 Score	0.4097	0.5567	0.5697	0.5769	0.5893	0.7083	0.6528	0.5805
Precision	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
Recall	0.2576	0.3857	0.3983	0.4054	0.4178	0.5484	0.4845	0.4140
SHMT$_{0.9}$ Multiple SEs								
Image Number	1	2	3	4	5	6	7	Average
Processing Time (s)	88.6	56.4	25.5	20.6	62.7	46.8	58.3	-
No.True Positives	110	76	58	59	118	121	115	-
No. False Negatives	185	64	60	52	107	65	79	-
No. False Positives	2	2	0	1	2	1	5	-
Accuracy	37.29	54.29	49.15	53.15	52.44	65.05	59.28	52.95
F1 Score	0.5405	0.6972	0.6591	0.6901	0.6841	0.7857	0.7325	0.6842
Precision	0.9821	0.9744	1.0000	0.9833	0.9833	0.9918	0.9583	0.9819
Recall	0.3729	0.5429	0.4915	0.5315	0.5244	0.6505	0.5928	0.5295

restricted to $\mathbb{R}^3$, such is the case for RGB or other colour images, and can be readily applied to $\mathbb{R}^N$ images with no additional complexity, other than the additional dimensionality of the data. While we use the Euclidean distance measure here in order to present the MDPOHMT, other distance measures, or indeed target detection methods such as the Spectral Angle Mapper (SAM), which are explicitly suited to hyperspectral imaging applications may be incorporated as the reduced ordering method within the MDPOHMT. Figure 18 shows the set of synthetic images created in order to test the MDPOHMT on hyperspectral data. Each of the images created have a spatial dimension of 215×215 with 80 spectral bands and contains two endmember spectra from the United States Geological Survey (USGS) Spectral Library [17], shown in Figure 18a, one comprised of a mixture of antigorite and grass which create the background and another which forms the foreground at varying abundancies throughout the image. The abundance measure was varied from 100% to 1%, as seen in Figure 18b in order to verify whether the MDPOHMT could be used in scenarios where there was no pure target spectra present in a pixel. In order to also validate the noise robustness of the MDPOHMT when applied to

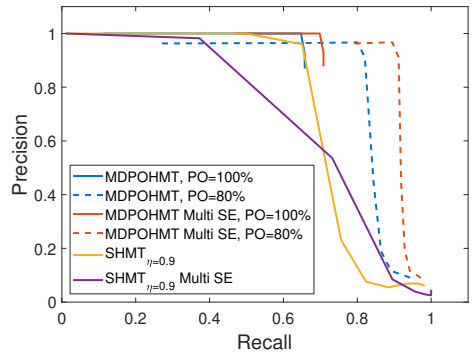

Figure 16: Precision-Recall curves of the SHMT with $\eta = 0.9$ and MDPOHMT where $P = 100\%, 80\%$ with both single and multiple SEs.

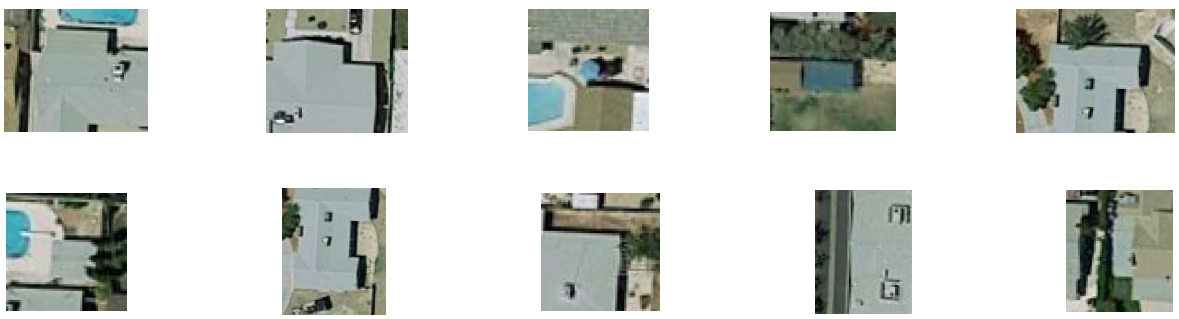

(a)

Figure 17: Examples of 10 False positive swimming pool detections when using multiple SEs with the MDPOHMT on the image shown in Figure 14.

hyperspectral imagery, both impulsive noise with 10% probability as well as zero-mean Gaussian noise with $\sigma^2 = 0.1$ were added to the image (Figure 18c).

The MDPOHMT was first applied to the image that has not been perforated by noise (Figure 18b) with SEs designed using the desired target spectra from Figure 18a, the results of this operation are shown in Figure 19a. All 100 of the desired target objects are correctly detected with no false positive detections. Even those with little (≤ 10%) abundancy are detected as the foreground SE will have marginally lower dissimilarity with these mixed regions than the background SE has with the background.

The MDPOHMT was then applied with P = 100% to the noise compromised image (Figure 18c) with the number of correctly identified objects dropping from 100 to 22, as seen in Figure 19b where white pixels indicating a True Positive (TP) detection and magenta pixels indicating a False Negative (FN) detection. This drop in detections can mostly be attributed to the impulsive noise added to the image as this can "puncture" the SE causing the MDPOHMT to fail. In order to rectify this, a number of example points were selected from the image with varying target abundancies being represented. PO plots for each of these points were generated and are shown in Figure 19c. Using the combination of the PO plots at each pixel it was found that the 1% abundancy target could not be detected at all due to the required PO falling below 50%. The other test objects could be detected with a PO value set to P = 70% as represented by the magenta cross in Figure 19c. The results of the HMT in Figures 19a, 19b and 19d have been enlarged in order to increase their legibility, each detection is a single pixel.

With this new occupancy value of P = 70%, the MDPOHMT could be reapplied to the noisy image. Upon relaxing the PO value the number of positively detected objects rose from 22 to 92 largely overcoming the impulsive noise present in the image. Only 8 objects were not detected when the noise was accounted for, each of those with spectral abundancy ≤ 10%. This can be attributed primarily to the Gaussian noise overpowering

the target spectra and blending it with the background spectra due to the low abundancies of target spectra and subsequent high sensitivity to changes.

The spectra present in Figure 18a are near-uniform over the portion of the spectrum considered here. Thus, as seen in Figure 19a, they are relatively easy to delineate and detect with even the 1% abundance target being detected in the noise free image. Despite the impulsive and Gaussian noise applied to the image, this uniformity is unrealistic and is intended to be an example of the extension of the MDPOHMT for application to spectral data. In order to assess how reliably the MDPOHMT can detect targets in a more complex domain, an additional set of images was created again using the USGS [17] Spectral Library along with the Hyperspectral Imagery Synthesis (EIAs) MATLAB toolbox [10] which allowed for the creation of randomised abundance maps for hyperspectral image synthesis.

Four spectra from the USGS Spectral Library were selected, three natural materials: dry grass, spruce needles, and limestone which would form the background of the synthetic scene and concrete which again forms the in-painted targets at abundancies decreasing from 100% to 1%. Each of these endmember spectra present in the image are displayed in Figure 20a with both the grass and limestone spectra being similar to that of the target concrete. The background scene was generated using the Gaussian Fields method [10] and the targets were subsequently inserted, with an abundance of 100%, i.e., pure spectra, in the bottom right corner of Figure 20b and an abundance of 1% in the top left corner. The generated image is shown in Figure 20b and Figure 20c shows the image with additional Gaussian and impulsive noise generated using the same parameters as those in Figure 18c.

The MDPOHMT was first applied to the noise free image in Figure 20b, the results of which can be seen in Figure 21a where 95 of the 100 targets are detected successfully despite the likeness between the targets and their background in places. When compared with the results on the first, simpler, synthetic image (Figure 18b), there are False Positive (FP) detections introduced due to areas of spectrally similar background being detected incorrectly. These FPs are represented as green pixels in the MDPOHMT results images. In total, 33 objects are detected incorrectly. One method for removing such FP detections is to lower the foreground distance threshold, T however, this may also lead to a drop in TP detections due to some targets having low abundancies and less similarity to the probing SE.

Applying the MDPOHMT with $\mathrm{P} = 100\%$ to the noise degraded image (Figure 20c) yields similar results to that of Figure 19b with only 20 of the 100 targets being detected, as shown in Figure 21b. This can again be improved upon by lowering the PO, with the appropriate value set using a PO Plot (Figure 21c). Through investigation, a PO value of $\mathrm{P} = 80\%$ (represented by the magenta cross in Figure 21c) was sufficient to detect each of the targets tested, excluding the target with 1% abundance which required a PO value of $\mathrm{P} < 50\%$. Applying the MDPOHMT to the noisy image with $\mathrm{P} = 80\%$ results in the image shown in Figure 21d where 88 of the 100 targets are detected with only 8 false positive detections. By relaxing the percentage fit through reducing the PO measure based on the PO Plot analysis the impulsive noise better accounted for.

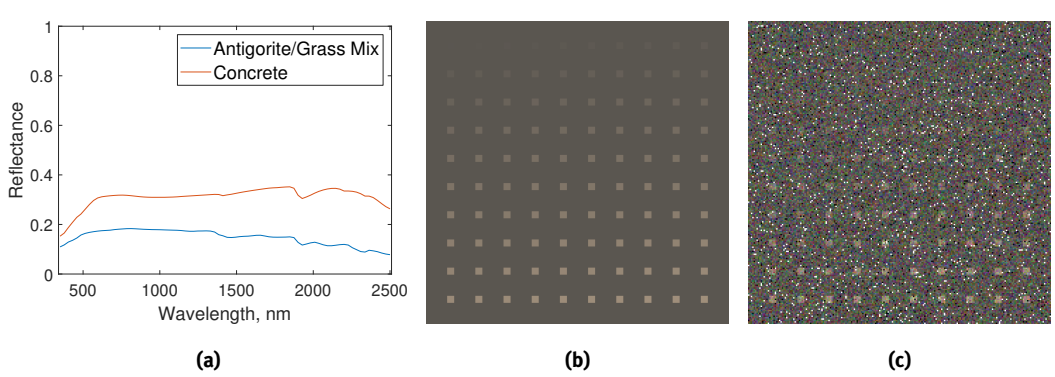

Figure 18: Synthetic Images generated to assess the MDPOHMT for use on hyperspectral images. **a)** The two endmember spectra used to create the image - the antigorite/grass mixed spectra forming the background and the concrete forming the in-painted targets. **b)** False-colour synthetic image created with 100 targets ranging from 100% abundance to 1%. **c)** False-colour synthetic image from b) with 10% impulsive noise and Gaussian noise with $\sigma^2 = 0.1$ and $\mu = 0$.

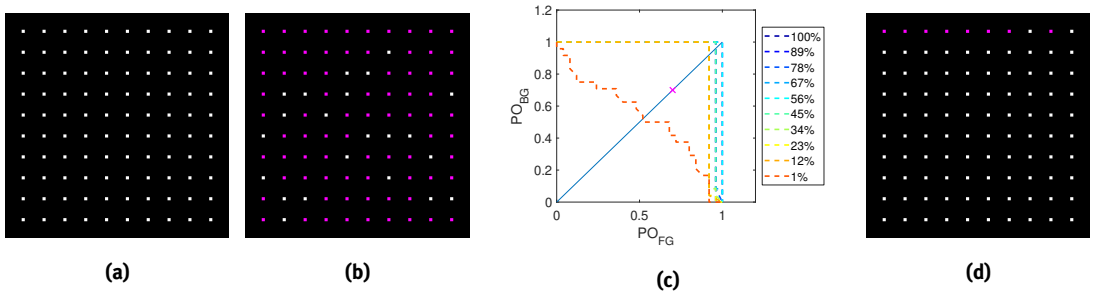

Figure 19: MDPOHMT results on the generated synthetic hyperspectral images **a)** MDPOHMT results on the image from Figure 18b with P = 100%. **b)** MDPOHMT results on the image from Figure 18c with P = 100%. **c)** PO plots at selected test pixels with decreasing target spectra abundancy. **d)** MDPOHMT with P = 70% selected using the PO plot in c).

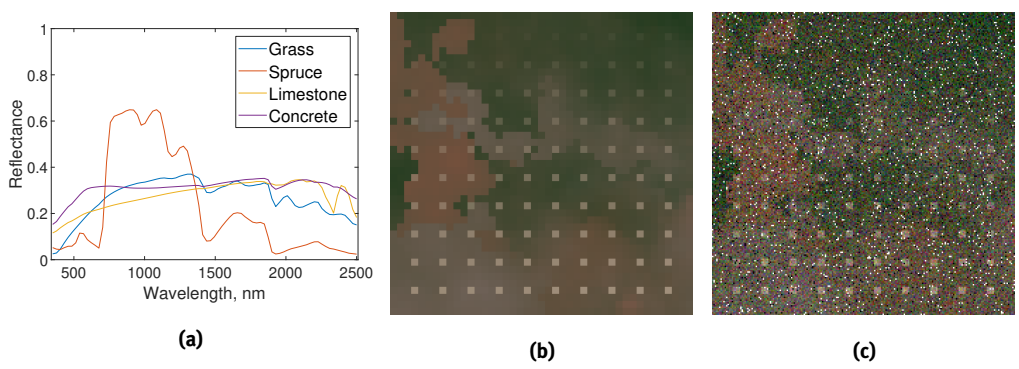

Figure 20: Hyperspectral image synthesis with a more complex, non-uniform, scene. **a)** The four endmember spectra used to create the synthetic image. **b)** False-colour synthetic image created with 100 targets ranging from 100% abundance to 1%. **c)** False-colour synthetic image from b) with 10% impulsive noise and Gaussian noise with $\sigma^2 = 0.1$ and $\mu = 0$.

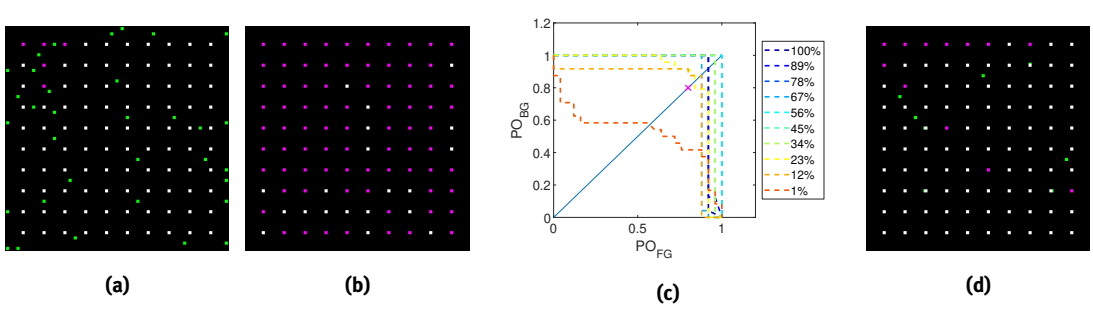

Figure 21: a) MDPOHMT results on the image from Figure 20b with P = 100%. **b)** MDPOHMT results on the image from Figure 20c with P = 100%. **c)** PO plots at selected test pixels with decreasing target spectra abundancy. **d)** MDPOHMT with P = 80% selected using the PO plot in c).

5 Conclusion

In this paper we presented and described our proposed novel method based on the Hit-or-Miss Transform, the Multi-Dimensional Percentage Occupancy Hit-or-Miss Transform. The MDPOHMT is intended as a general technique which extends the morphological HMT for use in object detection and pattern recognition applications in colour and multivariate images. This is achieved using a reduced ordering scheme which reduces multivariate data into a scalar value via some distance or similarity measure. We have validated our approach, first using synthetic colour images in order to display the correct operation of our technique followed by object detection and discrimination experiments on natural images. We tested the MDPOHMT on the DOTA aerial image dataset in order to display how the technique performs in a large dataset with objects of interest varying in size, shape, colour, texture, and occlusion levels. We also compared our technique with other similar HMT operators with our method performing favourably.

In order to make our MDPOHMT more robust to noise, rank-order filters are used in place of the hard maximum and minimum filters commonly employed in morphological techniques. This allows for the transform to be flexible in cases where the image contains noise or objects of interest are occluded. By relaxing the constraints of traditional erosion and dilation operations, objects with a high enough percentage fit to the corresponding probing SEs can be detected. By using a percentage fit, the transform remains unbiased to changes in size or shape of the objects of interest or the SEs used in detecting them. The parameters used in this relaxation are intuitive and can can be determined through trial and error using only a few training samples or indeed estimated and refined empirically, by using the extended PO Plots introduced in this work. Through the use of percentage occupancy and rank order filters combined with the use of PO plots for setting the appropriate parameters, the MDPOHMT has been shown to perform favourably when compared with related single channel and multivariate HMT definitions on the considered pattern recognition tasks.

Through experimentation using the DOTA dataset we found that our MDPOHMT operator, in conjunction with the Euclidean distance-based reduced ordering, can be very sensitive to changes in colour. The Euclidean distance measure was used as an example of the reduced ordering necessary for the operation of our technique. Other distance measures may be more appropriate for specific applications, especially in the case of hyperspectral imaging where dedicated spectral measures may be implemented. Similarly, alternative colour spaces, such as L*a*b or HSV, may be more suited to colour morphology as discussed in the literature. A general way of overcoming this sensitivity to changes in colour or spectral information in objects of interest is to simply design multiple composite SEs, one for each image or object however this would be a rather involved and bespoke process. By using multivariate SEs and encoding the SE with the desired colour information, as well as size and shape, the issue of colour perception is limited - as it simply needs to match the SE - while also allowing for non-flat objects to be detected with a single composite SE. As with the vast majority of morphological operators the SE choice and design is still an active area of research and optimal design of SE for particular applications is of interest in future work. An alternative way to relax the sensitivity to changes in object colour could be to have a multivariate threshold or decision space rather than a scalar distance threshold. In our testing we found that, like other HMT definitions, the MDPOHMT itself is not scale invariant, this is something we will investigate further but can again be overcome by simply creating multiple SEs based on the scale and orientation of the objects to be detected, or by designing SEs which generalise between objects.

While in this paper we focus on extending the Hit-or-Miss Transform in particular, the methods used here may be able to be extended and applied to other morphological operators such as erosions and dilations and subsequently openings and closings. Similarly to the techniques described in Section 1 our MDPOHMT may be integrated into a machine learning framework. This may have the effect of augmenting feature extraction either as an individual HMT layer or alternatively using the notion of percentage occupancy and rank-order filters within the pooling stage.

Acknowledgements: This work was supported by BAE Systems, the Engineering and Physical Sciences Research Council of the UK (EPSRC) Grant numbers EP/S000631/1 and EP/R512114/1 (ref. no. 1969791), and the UK MOD University Defence Research Collaboration (UDRC) in Signal Processing.

References

[1] H. M. Al-Otum. Morphological operators for color image processing based on Mahalanobis distance measure. volume 42, page 12. SPIE, 2003.

[2] J. Angulo. Morphological colour operators in totally ordered lattices based on distances: Application to image filtering, enhancement and analysis. *Computer Vision and Image Understanding*, 107(1):56–73, 2007.

[3] E. Aptoula, S. Lefevre, and C. Ronse. A hit-or-miss transform for multivariate images. *Pattern Recognition Letters*, 30(8):760–764, 2009.

[4] E. Aptoula and S. Lefèvre. A comparative study on multivariate mathematical morphology. *Pattern Recognition*, 40(11):2914–2929, 2007.

[5] V. Barnett. The ordering of multivariate data. *Journal of the Royal Statistical Society*, 139(3):318–355, 1976.

[6] B. Burgeth and A. Kleefeld. An approach to color-morphology based on Einstein addition and Loewner order. *Pattern Recognition Letters*, 47(Supplement C):29–39, 2014.

[7] M. L. Comer and E. J. Delp. Morphological operations for color image processing. *J. Electronic Imaging*, 8(3):279–289, 1999.

[8] A. N. Evans and X. U. Liu. A morphological gradient approach to color edge detection. *IEEE Transactions on Image Processing*, 15(6):1454–1463, 2006.

[9] G. Franchi, A. Fehri, and A. Yao. Deep morphological networks. *Pattern Recognition*, page 107246, 2020.

[10] Grupo de Inteligencia Computacional. Hyperspectral imagery synthesis (EIAs) toolbox. *Universidad del País Vasco/Euskal Herriko Unibertsitatea (UPV/EHU), Spain*.

[11] N. R. Harvey, R. B. Porter, and J. Theiler. Ship detection in satellite imagery using rank-order greyscale hit-or-miss transforms. Report, 2010.

[12] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask R-CNN. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2980–2988, 2017.

[13] H. Heijmans. *Morphological Image Operators*. Academic Press, 1994.

[14] M. A. Islam, B. Murray, A. Buck, D. T. Anderson, G. Scott, M. Popescu, and J. Keller. Deep morphological hit-or-miss transform neural network. *arXiv preprint arXiv:1912.02259*, 2019.

[15] M. Ivanovici, A. Căliman, N. Richard, and C. Fernandez-Maloigne. Towards a multivariate probabilistic morphology for colour images. In *Conference on Colour in Graphics, Imaging, and Vision*, volume 2012, pages 189–193. Society for Imaging Science and Technology, 2012.

[16] L. Jiao, F. Zhang, F. Liu, S. Yang, L. Li, Z. Feng, and R. Qu. A survey of deep learning-based object detection. *IEEE Access*, 7:128837–128868, 2019.

[17] R. F. Kokaly, R. N. Clark, G. A. Swayze, K. E. Livo, T. M. Hoefen, N. C. Pearson, R. A. Wise, W. M. Benzel, H. A. Lowers, and R. L. Driscoll. USGS spectral library version 7. Report 2327-638X, 2017.

[18] P. Lambert and J. Chanussot. Extending mathematical morphology to color image processing. *Proc. CGIP*, pages 158–163, 2000.

[19] A. Ledoux, N. Richard, and A.-S. Capelle-Laizé. Color hit-or-miss transform (cmomp). In *Signal Processing Conference (EUSIPCO), 2012 Proceedings of the 20th European*, pages 2248–2252. IEEE, 2012.

[20] S. Lefèvre, E. Aptoula, B. Perret, and J. Weber. *Morphological template matching in color images*, pages 241–277. Springer, 2014.

[21] O. Lezoray, C. Charrier, and A. Elmoataz. Learning complete lattices for manifold mathematical morphology. In *International Symposium on Mathematical Morphology*, pages 1–4, 2009. Computer Science [cs]/Image ProcessingConference papers.

[22] R. P. Loce and E. R. Dougherty. Facilitation of optimal binary morphological filter design via structuring element libraries and design constraints. *Optical Engineering*, 31(5):1008–1025, 18, 1992.

[23] F. Macfarlane, P. Murray, S. Marshall, B. Perret, A. Evans, and H. White. A colour hit-or-miss transform based on a rank ordered distance measure. In *2018 26th European Signal Processing Conference (EUSIPCO)*, pages 588–592, 2018.

[24] J. Masci, J. Angulo, and J. Schmidhuber. A learning framework for morphological operators using counter–harmonic mean. In *International Symposium on Mathematical Morphology and Its Applications to Signal and Image Processing*, pages 329–340. Springer, 2013.

[25] G. Matheron. *Random Sets and Integral Geometry*. Wiley, New York, NY, 1975.

[26] R. Mondal, M. S. Dey, and B. Chanda. Image restoration by learning morphological opening-closing network. *Mathematical Morphology - Theory and Applications*, 4(1):87–107, 2020.

[27] P. Murray and S. Marshall. A new design tool for feature extraction in noisy images based on grayscale hit-or-miss transforms. *IEEE Transactions on Image Processing*, 20(7):1938–1948, 2011.

[28] P. Murray and S. Marshall. *A Review of Recent Advances in the Hit-or-Miss Transform*, volume 175, book section 5, pages 221–282. Elsevier, 2013.

[29] P. Murray, S. Marshall, and E. Bullinger. The percentage occupancy hit or miss transform. In *Signal Processing Conference, 2009 17th European*, pages 253–257. IEEE, 2009.

[30] B. Naegel, N. Passat, and C. Ronse. Grey-level hit-or-miss transforms—part I: unified theory. *Pattern Recognition*, 40(2):635–647, 2007.

[31] K. Nogueira, J. Chanussot, M. D. Mura, W. R. Schwartz, and J. A. d. Santos. An introduction to deep morphological networks. *arXiv preprint arXiv:1906.01751*, 2019.

[32] B. Perret, S. Lefevre, and C. Collet. A robust hit-or-miss transform for template matching applied to very noisy astronomical images. *Pattern Recognition*, 42(11):2470–2480, 2009.

[33] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.

[34] C. Ronse. Why mathematical morphology needs complete lattices. *Signal Processing*, 21(2):129–154, 1990.

[35] M. Sangalli and M. E. Valle. Color mathematical morphology using a fuzzy color-based supervised ordering. Fuzzy Information Processing, pages 278–289. Springer International Publishing, 2018.

[36] M. Sangalli and M. E. Valle. Approaches to multivalued mathematical morphology based on uncertain reduced orderings. In *International Symposium on Mathematical Morphology and Its Applications to Signal and Image Processing*, pages 228–240. Springer, 2019.

[37] J. Serra. *Image Analysis and Mathematical Morphology*, volume 1. Academic press, New York, 1982.

[38] J. Serra. *Image Analysis and Mathematical Morphology, Vol. 2: Theoretical Advances*, volume 2. Academic Press, New York, NY, 1988.

[39] X. Shen and W. Bao. Hyperspectral endmember extraction using spatially weighted simplex strategy. *Remote Sensing*, 11(18):2147, 2019.

[40] P. Soille. Advances in the analysis of topographic features on discrete images. In *Discrete Geometry for Computer Imagery*, pages 271–296. Springer, 2002.

[41] P. Soille. *Morphological Image Analysis: Principles and Applications*. Springer Science & Business Media, 2 edition, 2013.

[42] K. Stankov and D.-C. He. Detection of buildings in multispectral very high spatial resolution images using the percentage occupancy hit-or-miss transform. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 7(10):4069–4080, 2014.

[43] S. Velasco-Forero and J. Angulo. Hit-or-miss transform in multivariate images. In *International Conference on Advanced Concepts for Intelligent Vision Systems*, pages 452–463. Springer, 2010.

[44] S. Velasco-Forero and J. Angulo. Supervised ordering in $\mathbb{R}^p$: Application to morphological processing of hyperspectral images. *IEEE Transactions on Image Processing*, 20(11):3301–3308, 2011.

[45] J. Weber and S. Lefèvre. Spatial and spectral morphological template matching. *Image and Vision Computing*, 30(12):934–945, 2012.

[46] S. Wu and Y. Xu. DSN: A new deformable subnetwork for object detection. *IEEE Transactions on Circuits and Systems for Video Technology*, pages 1–1, 2019.

[47] G.-S. Xia, X. Bai, J. Ding, Z. Zhu, S. Belongie, J. Luo, M. Datcu, M. Pelillo, and L. Zhang. DOTA: A large-scale dataset for object detection in aerial images. In *Proc. CVPR*, 2018.

[48] H. Xu, Y. Zhang, and H. Zhao. Edge detection of color image using mathematical morphology in HSV color space. In *Proceedings of 2012 2nd International Conference on Computer Science and Network Technology*, pages 2112–2116, 2012.

[49] C.-W. Yeh. *Colour morphology and its approaches*. Thesis, 2015.

[50] C. W. Yeh and D. Pycock. Similarity-based colour morphology. In *2013 5th Computer Science and Electronic Engineering Conference (CEEC)*, pages 71–76, 2013.

[51] Y. Zhang, S. Blusseau, S. Velasco-Forero, I. Bloch, and J. Angulo. Max-plus operators applied to filter selection and model pruning in neural networks. In *International Symposium on Mathematical Morphology and Its Applications to Signal and Image Processing*, pages 310–322. Springer, 2019.