

# Complementary texture and intensity gradient estimation and fusion for watershed segmentation

Padraig Corcoran · Adam Winstanley · Peter Mooney

Received: 15 August 2009 / Revised: 1 July 2010 / Accepted: 2 December 2010 / Published online: 17 December 2010  
© Springer-Verlag 2010

**Abstract** In this paper, we identify two current challenges associated with watershed segmentation algorithms which attempt to fuse the visual cues of texture and intensity. The first challenge is that most existing techniques use a competing gradient set which does not allow boundaries to be defined in terms of both visual cues. The second challenge is that these techniques fail to account for the spatial uncertainty inherent in texture gradients. We present a watershed segmentation algorithm which provides a suitable solution to both these challenges and minimises the spatial uncertainty in boundary localisation. This is achieved by a novel fusion algorithm which uses morphological dilation to integrate intensity and texture gradients. A quantitative and qualitative evaluation of results is provided demonstrating that our algorithm outperforms three existing watershed algorithms.

**Keywords** Feature fusion · Spatial uncertainty · Texture · Watershed segmentation

## 1 Introduction

In many object-recognition systems image segmentation is typically used as a pre-processing step to define objects upon which a classification method operates [1,2]. Image segmentation is a very active research area with thousands of papers published on the subject [3]. Despite many promising approaches it is often difficult to produce segmentation of sufficiently high quality that could allow reliable calculation

of object specific information such as object shape. Consequently, many of the current best approaches to object-recognition do not employ segmentation [4,5]. The goal of this work is to implement a segmentation algorithm which is capable of providing a set of meaningful objects for an object-recognition system.

Most segmentation strategies can be classified as region- or boundary-based methods. In region-based methods a grouping of homogenous areas is performed to produce segmentation. In contrast boundary-based methods attempt to extract the boundaries between these homogenous areas. The watershed transform is a technique which combines both region- and boundary-based methods [6]. Pixels are grouped around the regional minima of a gradient image and boundaries are located along the crest lines of this image. The watershed transform is a function of an input gradient image. Therefore in order to obtain accurate segmentation a suitable gradient image must first be estimated. Many strategies exist for calculating a gradient image in terms of the visual cues of texture and intensity but all are subject to two challenges. We now describe each of these under the headings of *description*, *current solutions* and *proposed solution*.

### 1.1 Challenge 1: non-harmonising gradients

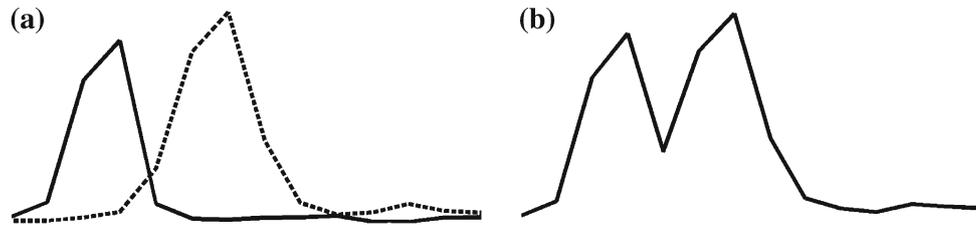
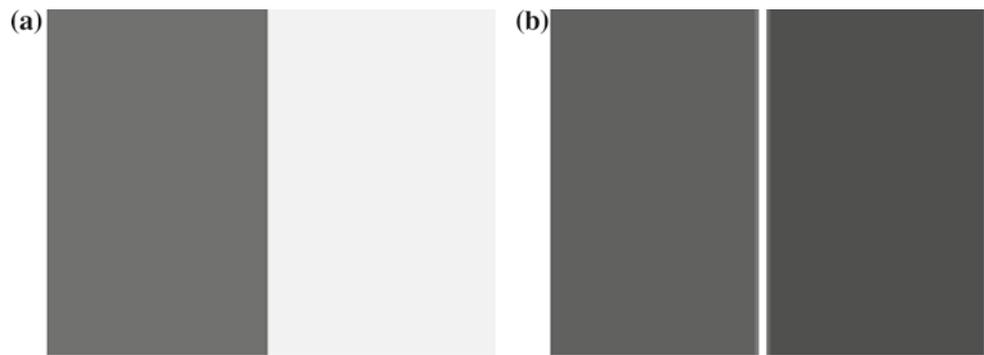
#### 1.1.1 Description

Raw intensity and texture gradients do not harmonise. A gradient operator applied to a raw intensity image will not only respond to intensity boundaries between objects but it will also respond to intensity variation due to object texture which results in a large number of false-positives. Texture is a spatial property and therefore any features used to describe it must be calculated within a neighbourhood. At object boundaries this neighbourhood can contain a large

---

P. Corcoran (✉) · A. Winstanley · P. Mooney  
Department of Computer Science,  
National Centre for Geocomputation,  
National University of Ireland Maynooth, Kildare, Ireland  
e-mail: padraigc@cs.nuim.ie

**Fig. 1** An image containing a pure intensity edge and corresponding texture feature image are shown in **a** and **b**, respectively. Larger response values in **b** are represented by the colour white



**Fig. 2** Corresponding intensity and texture gradients for an individual edge are plotted by *solid* and *dashed lines*, respectively, in **a**. Although both gradients correspond to the same edge they are located in different

spatial locations. The result of performing a summation of the individual gradient in **a** is displayed in **b**

intensity edge and multiple objects textures. This can result in unwanted unique responses, known as texture boundary-responses, along object boundaries [7]. For example, consider the image in Fig. 1a of a pure intensity edge. Despite the absence of texture the corresponding texture feature image in Fig. 1b contains a relatively large response along the intensity edge. A gradient operator applied to such a raw texture feature image will not give the desired measure of object to object gradient at the object boundary. Instead a double response will occur where each response represents an object to boundary-response gradient. These double responses represent an inaccurate estimate of object boundary gradient and can be falsely interpreted as an object lying along the boundary.

### 1.1.2 Current solutions

Many authors [8–10] propose to use texture and intensity gradients in their raw form and perform fusion using the principle of a *competing* gradient set. This strategy is based on the assumption that each object boundary is predominantly a texture or intensity boundary and only the corresponding gradient should be used to define it. The authors in [8, 10] use a measure of *texturedness* to modulate both texture and intensity gradients. Intensity gradients in the presence of texture are inhibited reducing the number of false-positives due to texture intensity variation. On the other hand texture gradients in the absence of texture are suppressed removing false-positives due to boundary-responses at pure intensity boundaries. In [9] the authors only modulate the intensity gradients and remove texture boundary-responses

by performing separable median filtering. Many boundaries in natural images are defined by both the visual cues of texture and intensity. Since all the above competing gradient set strategies estimate each individual gradient value in terms a single visual cue they result in inaccurate boundary gradients and reduced discrimination strength.

### 1.1.3 Proposed solution

We propose to extract a complementary set of texture and intensity gradients. This set can then be fused to give a resulting gradient image which represents the summation of both texture and intensity gradients. Since each gradient value is estimated using both visual cues this strategy offers increased discrimination strength over competing gradient approaches.

## 1.2 Challenge 2: spatial uncertainty in texture gradients

### 1.2.1 Description

Texture, unlike intensity, is a spatial phenomenon and therefore the locations of its boundaries are subject to a degree of spatial uncertainty. This phenomenon is known as the uncertainty principle and asserts that we cannot simultaneously know ‘what happens when’ [11]. Consequently, texture and intensity gradients corresponding to the same boundary may be located in different spatial locations. This point is illustrated using a one-dimensional example extracted from a real image in Fig. 2. The solid and dashed lines in Fig. 2a represent intensity and texture gradients respectively for a single edge.

The intensity gradient does not suffer from spatial uncertainty and is therefore spatially located at the actual boundary which is left of centre. On the other hand the texture gradient suffers from spatial uncertainty and as a result is not spatially located at the actual boundary location. This uncertainty in the spatial location of boundary gradients must be taken into consideration when fusion is performed.

### 1.2.2 Current solutions

The authors are unaware of any existing works which attempt to provide a solution to this challenge. In all cases the gradient magnitude at a particular spatial location is determined by summing the gradient values at the same spatial location in each of the corresponding gradient images [8,9]. This approach is based on the invalid assumption that all gradient values corresponding to the same edge will be situated at the same spatial location in each gradient image. For a particular edge the result of such a summation will not be the desired sum of all gradients corresponding to the edge in question. This point is illustrated by Fig. 2b which corresponds to the summation of the intensity and texture gradients of Fig. 2a. The two peaks in this image correspond to the individual intensity and texture gradients. Neither peak represents the desired sum of the two gradients in question.

### 1.2.3 Proposed solution

To address this challenge we propose to post-process the texture gradient images such the spatial location of each edge and its corresponding gradient value intersect. Since all gradient values for a particular edge intersect at their corresponding edge location the sum at this location will be the desired sum of all gradient values corresponding to that edge.

This paper presents a strategy for computing a gradient image which represents a solution to the two challenges presented above. The layout of this paper is as follows. In Sect. 2, we describe the proposed complementary texture and intensity gradient set. Section 3 describes a novel algorithm used to fuse these complementary gradients and the watershed algorithm. In Sect. 4, we present an evaluation of the proposed segmentation method. Finally in Sect. 5, we draw conclusions from this work and propose future research directions.

## 2 Complementary gradient set extraction

To overcome the challenge of non-harmonising gradients, we propose to extract a complementary set of texture and intensity gradients which will be later fused using a strategy presented in Sect. 3. The individual gradients in this set are not novel when used in isolation. The novelty presented in

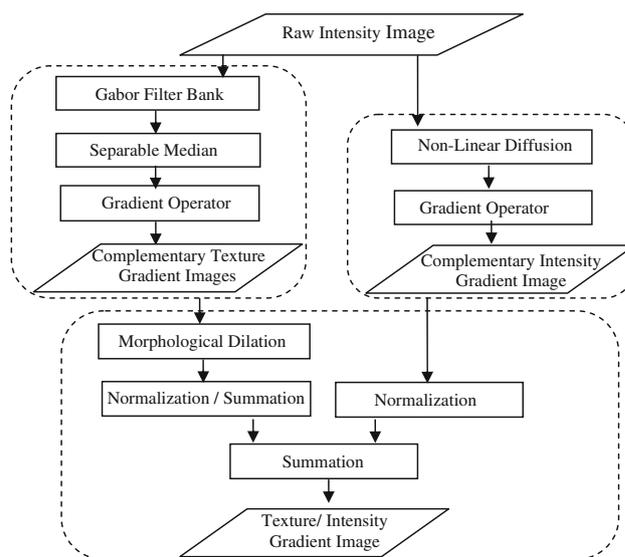


Fig. 3 Flowchart of the overall gradient extraction process

this work is the proposal that if fused correctly, their fusion represents an estimate of total boundary gradient in terms of both visual cues.

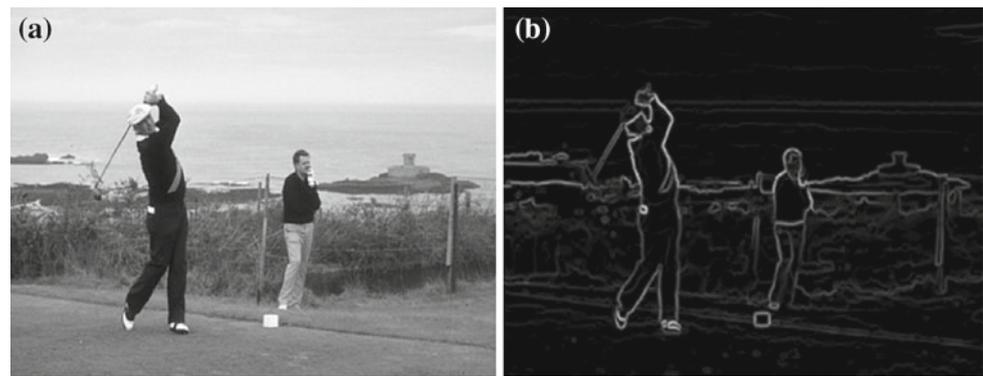
### 2.1 Complementary intensity gradient extraction

In this section, we describe how the complementary intensity gradient image, denoted  $IG$  is extracted. A flowchart of the overall this process is shown in the top right of Fig. 3. In summary, to generate an intensity image containing only edges due to object boundaries while removing those due to texture intensity variation we employ a non-linear diffusion process. Accurate gradients are then extracted from this diffused image using a multi-scale gradient operator.

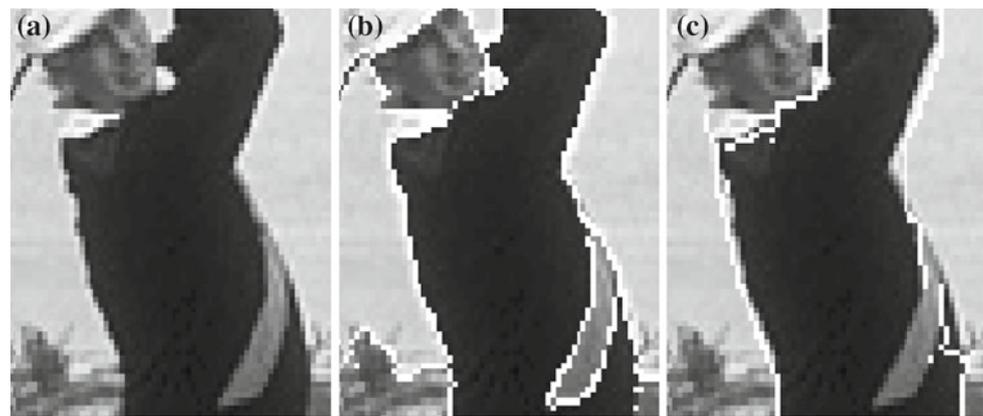
In non-linear diffusion, first introduced in [12], the amount of diffusion performed at each location is controlled by an edge-stopping function of gradient magnitude. The textures contained in natural image will be of different scales and therefore it is recommended to use a spatially varying diffusion scale. This is implemented using the strategy proposed in [13]. This strategy will smooth the texture intensity variation of each object and will replace it with an area of uniform intensity with values equal to the mean grey value for that object. This effect is known as the principle of *grey level invariance* [14].

Following generation of the diffused image a corresponding gradient image is extracted using the multi-scale gradient operator proposed in [15]. Using a process of edge tracking this operator combines large scale gradient estimation with small scale localisation. The resulting gradient image contains accurate intensity gradient values for each object-to-object boundary and few false-positives due to texture intensity variation. As a consequence of the principle of *grey*

**Fig. 4** A complementary intensity gradient image  $IG$  extracted from **a** is displayed in **b**



**Fig. 5** A section of Fig. 4a is displayed in **a**. Following non-maximum suppression those intensity and texture gradients extracted above a certain threshold are represented by white and superimposed on the original image in **b** and **c**, respectively



*level invariance* the gradients located at the boundary between two objects specify the contrast between the mean grey values of the objects in question.

An example of a complementary intensity gradient image is displayed in Fig. 4. Figure 5b displays a section of the image in Fig. 4a with those corresponding complementary intensity gradient magnitudes, following non-maximum suppression, which are above a certain threshold represented by white overlay. Due to the edge preserving properties of the non-linear diffusion these gradients are localised to the pixel scale.

## 2.2 Complementary texture gradient extraction

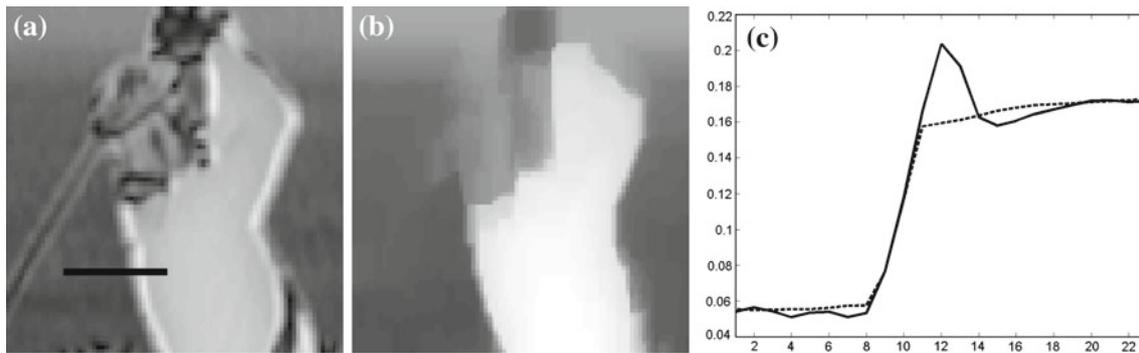
In this section, we describe how a set of  $n$  complementary texture gradient images are extracted. Each individual image is denoted  $TG_i$  for  $i = 1$  to  $n$ . A flowchart of the overall process is shown in the top left of Fig. 3. In summary, a set of raw texture features are extracted using a Gabor filter bank. These features are then filtered with separable median filters to remove texture boundary-responses and reduce noise. Finally gradients are extracted. The remainder of this section describes each of these steps.

### 2.2.1 Gabor filter bank

Gabor filters have been increasingly considered and applied to image analysis [16]. Spatially a Gabor function is a Gaussian of aspect ratio  $\sigma_x/\sigma_y$  modulated by a sinusoid of frequency  $F$  and orientation  $\theta$ . In this work, we implemented the Gabor filter bank design strategy proposed by Claudi and Jernigan [16]. The Empirical Rule informs us that for a Gaussian function 99.7% of the area under this function falls within plus and minus  $3\sigma$  of the mean [17]. Therefore a Gaussian or Gabor function has a spatial extent of approximately  $6\sigma$ . This fact will be utilised in Sect. 3. To extract features from the raw Gabor filter outputs we calculate the magnitude response as recommended by [16].

### 2.2.2 Median filtering

As stated earlier any texture features will exhibit texture boundary responses resulting in inaccurate texture boundary gradients. To overcome this problem we propose that any strategy which attempts to remove texture boundary-responses should exhibit the property that it *returns a single edge at each boundary from which it is possible to extract*



**Fig. 6** A section of a texture feature image extracted from Fig. 4a is displayed in **a**, and the result following median filtering in **b**. A 1-D plot of a cross-section of **a**, represented by the *black line* in **a**, is plotted in

**c** using a *solid line*. The result following median filtering is represented by a *dashed line* in **c**. An edge is returned within the area of the original boundary-response

an accurate object to object gradient. Most existing solutions to the texture boundary-response problem do not exhibit this property. Shao and Forstner [18] proposed to smoothed the gradient image with a Gaussian function converting the double peaks at boundaries due to boundary-responses into single peaks. Although the generation of false objects along boundaries is prevented, this strategy does not return accurate boundary gradient values. Kruijzinga and Petkov [19] modelled grating cells in the primary visual cortex (V1). Grating cells respond strongly to a grating of bars but respond weakly to single bars which do not form part of such a grating. The problem with this approach is that natural images do not contain gratings and this model will not respond in any way when applied to such images. In [20] the authors choose a set of texture features which give a ramp-edge response at boundaries in remotely sensed datasets. If segmentation is carried out at a sufficiently large scale the boundary-response values will be assigned to classes either side of the boundary. It is difficult to prove that a given feature extraction algorithm will always result in a ramp-edge at object boundaries for a given data type. This approach is not data and feature extraction algorithm independent and only those features which give a ramp edge boundary-response for a given data type may be used.

An approach which tackles the issues resulting from boundary-responses and exhibits the above property which we require is separable 2-D median filtering [9]. This technique will remove noise having a scale less than half the scale of the median filters applied. Consequently, any separable median filters used to remove boundary-responses must be at least twice the scale of the boundary-responses if they are to be removed. O'Callaghan and Bull [9] propose to remove texture boundary-responses using a separable median filtering approach and we have adapted this strategy. The maximum possible scale of a texture boundary-response is equal to the scale of the corresponding Gabor filter. Therefore to remove boundary responses a median filter

of twice the scale of the corresponding Gabor function is applied.

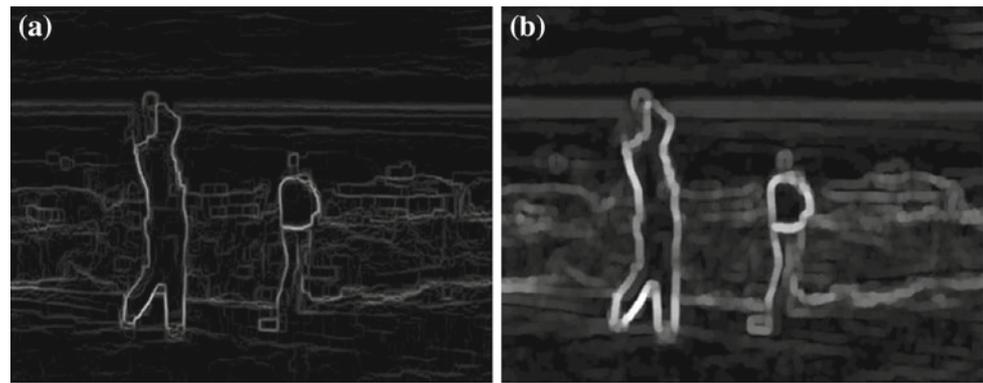
### 2.2.3 Gradient extraction

Smoothing is known to improve the texture discrimination strength of Gabor features [16]. The separable median filtering presented earlier for removing boundary responses also performs the dual task of smoothing the feature images. Now that texture boundary responses have been removed and the feature images smoothed we extract gradients using the same operator as Sect. 2.1. Figure 5c displays gradients extracted from a texture feature image overlaid on the original image and represented by white. The texture gradients, for example those corresponding to the boundaries of the man, are not entirely localised accurately.

This spatial uncertainty in the texture gradients is due to the fact that texture is a spatial property and is subject to the uncertainty principle which states that there is a trade off between spatial and spatial-frequency resolutions [11]. In order to discriminate between textures we require the spatial-frequency resolution of our features to be sufficiently high which in turn results in a lower spatial resolution and introduces a loss in boundary localisation. The Gabor function minimises this trade-off between spatial and spatial-frequency resolution [11]. When a texture boundary-response is filtered with a median filter an edge will result within the location of the original boundary-response. Since the maximum scale of a boundary-response is equal to the scale of the corresponding Gabor filter applied, this edge will be localised correctly to the scale of the corresponding Gabor filter. This fact is illustrated using an example in Fig. 6.

Comparing Fig. 5b with c, we can see that the complementary texture gradients are localised to a broader scale than the corresponding complementary intensity gradients. When fusing such gradient images this property must be taken into consideration.

**Fig. 7** A texture gradient image extracted from Fig. 4a is displayed in **a**. The result following the dilation of **a** is displayed in **b**



### 3 Gradient fusion and segmentation

The goal of the work presented in this paper is provide a segmentation technique which integrates the visual features of texture and intensity. In Sect. 2, we presented a strategy for calculating a set of complementary intensity and texture gradients denoted  $IG$  and  $TG_i$ , respectively. In order for a segmentation to accurately exploit this gradient set a suitable fusion strategy must be used. A flowchart of the overall fusion strategy we propose is in the lower half of Fig. 3. First, a morphological dilation operation is applied to all texture gradient images. Next, a series of normalisation and summation steps are then performed to produce a final gradient image. This gradient image is then used as input to the watershed algorithm to generate a segmentation result. The following subsections describe each of these steps in turn.

#### 3.1 Morphological dilation

This section describes the first step in the proposed fusion strategy which is morphological dilation. The standard approach to fusing multi-spectral and texture gradient images is to perform a weighted summation [8,9]. This strategy is based on the invalid assumption that different gradients corresponding to the same boundary will be located in the same spatial location. Texture unlike intensity is a spatial phenomenon and therefore is subject to the uncertainty principle as discussed in Sect. 1. This uncertainty must be taken into consideration when performing fusion. If corresponding texture and intensity gradients are located in different positions, a sum of their values will result in a double peak effect, with neither peak being a true sum of the individual gradients. These effects are illustrated in Fig. 2b.

To overcome this issue we use grey scale morphological dilation. The dilation of the image  $f$  at the location  $(x, y)$  by a structuring element  $b$  is given by Eq. 1 [6].

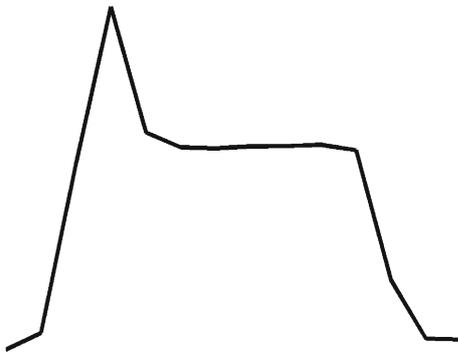
$$[f \oplus b](x, y) = \max_{(s,t) \in b} \{f(x-s, y-t)\} \quad (1)$$

For a Gabor filter to respond to an edge the edge must be located within the region over which the filter integrates. Therefore the distance between the location of an edge detected by a Gabor filter and its true location will be less than or equal to the size or scale of this region. The dilation of the gradient image by a structuring element of equal scale will therefore result in the gradient value corresponding to the edge in question, assuming it is a local maximum, being assigned to the true location of the edge. We dilate the texture gradient images using a flat disk-shaped structuring element so that the scale of dilation is equal in all directions. An example of a texture gradient image and the result following dilation is displayed in Fig. 7.

The location of an intensity gradient corresponding to an edge will equal the true location of the edge in question and that location only. Following dilation the location of a texture gradient corresponding to an edge will equal the true location of the edge and the neighbourhood of this location. A logical AND of these locations gives a single location corresponding to the true location of the edge. Following the summation of gradients with equal location, this location will be the only one which contains a sum of both corresponding gradient values. Hence the gradient will be maximised at this location.

If the watershed transform locates a boundary resulting from these gradient values it will locate the boundary at this maximum location, which is the location of the intensity boundary. Since the intensity gradient minimises spatial uncertainty the resulting boundary will in turn minimise this uncertainty. This property is illustrated in 1-D example in Fig. 8 where the peak of in this example is an accurate sum of individual gradients and the location of the intensity boundary. In Sect. 4, we will illustrate this property in the context of 2-D images. Chen et al. [21] proposed an alternative approach to minimise boundary localisation error using intensity features by applying an iterative post-processing of an initial crude segmentation using intensity features.

A side-effect of dilating the texture gradient images is that if a boundary has a texture gradient in a single texture



**Fig. 8** The result of dilating the texture gradient in Fig. 2a before summing it with the intensity gradient in Fig. 2a is displayed

feature image and no intensity gradient, this boundary could be located most accurately if no dilation is performed. Performing dilation in this case would actually increase uncertainty. For example, if the texture gradient image of Fig. 7a was the only gradient image which contained non-zero values it would not need to be fused with other gradient images. Therefore dilation would not be required and the watershed transform could be applied directly to this image. Textures will generally not have a sufficiently narrow frequency bandwidth such that they only respond to a single filter. Clausi [16] refers to this effect as “leakage”. It is therefore likely that an object boundary will have a gradient in more than a single feature image and dilation must then be performed in order to generate an accurate sum of gradient values.

### 3.2 Normalisation and summation

Following dilation of texture gradients, all intensity and texture gradient must be summed to form a final single gradient image. Before this summation can be performed, each must be normalised to correctly weigh its contribution. Shao and Forstner [18] proposed to normalise by variance or maximum filter response. In this section, we describe the normalisation and summation steps used in this work. These steps are represented graphically by the flowchart displayed in Fig. 3.

O’Callaghan and Bull [9] normalised by maximum filter response and then normalised by the sum of individual filter responses. In [8] the authors chose to perform no weighting operation and instead summed the gradient values in their original forms. There is no strong evidence to suggest any one normalisation strategy is superior and usually a trial and error approach is required to define a suitable one. The normalisation approach used in this work contains the following steps. We normalise each gradient image, denoted  $TG_i$ , by its maximum response. We then sum these images and normalise again by the maximum value to give a single gradient image denoted  $TG$ . These steps are described by Eq. 2. An

example of a gradient image  $TG$  is displayed in Fig. 9a.

$$TG(x, y) = \frac{\sum_i \widehat{TG}_i(x, y)}{\max_{x, y} (\sum_i \widehat{TG}_i(x, y))} \quad (2)$$

where

$$\widehat{TG}_i(x, y) = \frac{TG_i(x, y)}{\max_{x, y} (TG_i(x, y))}$$

Given the gradient image  $IG$  and  $TG$  we fuse these into a single gradient image as follows. We divide  $IG$  by four times its median value and  $TG$  by its median value. This step aligns the noise floor of each image and was originally proposed in [9]. The additional factor of four in the intensity gradient normalisation reflects the fact that the intensity gradient image has sharp peaks while the texture gradient image has broader peaks. The latter must be amplified to avoid being dominated by the former. Finally,  $IG$  and  $TG$  are summed to form a final single gradient image  $CG$  and again this is normalised to the range [0 1]. This fusion process is represented by Eq. 3.

$$CG(x, y) = \frac{\sum_i \widehat{CG}(x, y)}{\max_{x, y} (\sum_i \widehat{CG}(x, y))} \quad (3)$$

where

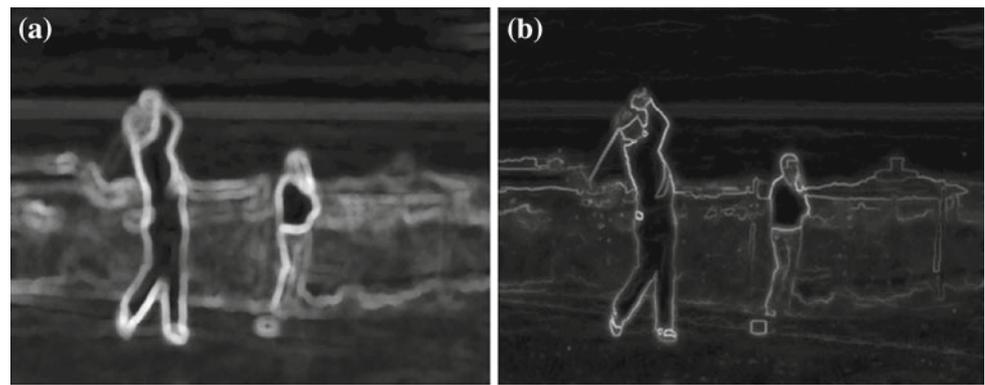
$$\widehat{CG}(x, y) = \frac{IG(x, y)}{4w_I} + \frac{TG(x, y)}{w_T}$$

In Eq. 3  $w_I$  and  $w_T$  represent the median values of  $IG$  and  $TG$ , respectively. All further segmentation analysis is performed on this image; an example of which is displayed in Fig. 9b.

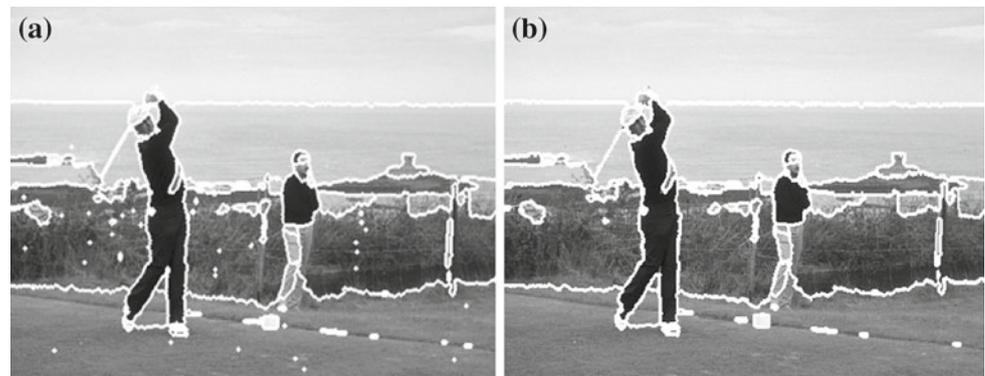
### 3.3 Segmentation

To generate segmentation from the gradient image  $CG$  we use the watershed transform [6]. The most intuitive formulation of the watershed transform is based on flooding simulation where the gradient image is considered to be a topographic surface. The watershed transform determines the watershed lines on this surface. To achieve this holes are punched in each regional minima of the image. The image is then flooded by letting water rise from each at a uniform rate. When the water from two distinct minima merge a segmentation boundary is created. The gradient image may contain noise in the form of unwanted extra regional minima and flooding from these will result in a greater number of segments than desired. To overcome this problem the gradient image is filtered to remove all regional-minima except those of a depth greater than a specified parameter  $h$ . This is achieved using a technique known as the  $h$ -minima transform [6] which can determine all regional minimum in a given gradient having a depth less than a specified parameter  $h$ . Varying  $h$  controls the scale of

**Fig. 9** The complementary texture gradient image extracted from Fig. 4a is displayed in **a**. This image is fused with its corresponding complementary intensity gradient image, which is displayed in Fig. 4b, to form the final gradient image displayed in **b**



**Fig. 10** The segmentation result for Fig. 4a achieved by fusing the proposed complementary gradients is shown in **a**. A post-processing of this result is performed to remove small segments with the result shown in **b**



regional-minima which are suppressed and in turn the scale of segmentation.

Figure 10a shows a segmentation result achieved using this fusion method. From this example it can be seen that this method tends to over-segment in some textured regions returning a large number of very small segments. For an example of this see the bottom of Fig. 10a. To remove this inaccuracy we implemented a post-processing step where segments of small size, that is less than 10 pixels in size, are removed and their corresponding regions assigned to the largest neighbouring segment. This size was determined by trial and error. Figure 10b shows the improvement in segmentation this step brings.

## 4 Results

To perform a quantitative evaluation of the proposed watershed segmentation method a state of the art dataset and performance metric were used. The dataset in question is the Berkeley segmentation dataset [22] which contains a large collection of images of natural scenes and is very suitable for determining the performance of any segmentation method. Segmentation is an ill-posed problem with most images having many possible segmentation solutions. To capture this variability in possible solution, each image in the dataset has between five and ten corresponding ground-truths. The data-

set contains 300 images which are divided by its creators into 200 training and 100 test images. Using all images in the dataset performance was determined using a raining-followed-by-testing evaluation methodology.

In order to determine segmentation performance a state of the art technique known as the Normalised Probabilistic Rand (NPR) index [23] was used. The NPR index quantifies the agreement of segmentation with its corresponding set of ground-truths with greater agreement resulting in a higher index value. The range of the NPR index is  $[-1, 1]$ . It is normalised such that a random segmentation will result in an index close to 0. Any segmentation which results in an index significant greater than 0 can be regarded as useful [23]. The index is sensitive to errors in boundary localisation and the degree of such errors is reflected in corresponding index values. From extensively working with the NPR index we observed that an increase of around 0.05 in the index generally corresponds to good improvement in segmentation quality. An increase of around 0.1 or above generally corresponds to a very good improvement. These facts may be confirmed by viewing segmentation results and corresponding index values presented later.

The remainder of this section is divided into the following three parts. Firstly, we present a comparison of segmentation results achieved by applying the marker-controlled watershed to the complementary texture and intensity gradients in isolation and to their fusion. Next, we evaluate the

contribution made by each major component in the proposed segmentation algorithm by determining performance in its absence. Finally, we present a comparison of segmentation results achieved by the proposed segmentation method relative to three other existing methods.

#### 4.1 Individual complementary gradients

In this section we present a comparison of segmentation results achieved by applying the marker-controlled watershed to the complementary texture and intensity gradients in isolation and their fusion. We perform this evaluation to demonstrate that when such gradients are fused correctly they complement each other to generate superior results. For all results in this sub-section we choose not to perform post-processing to remove small segments because we wanted to directly evaluate the strengths and weaknesses of our fusion strategy.

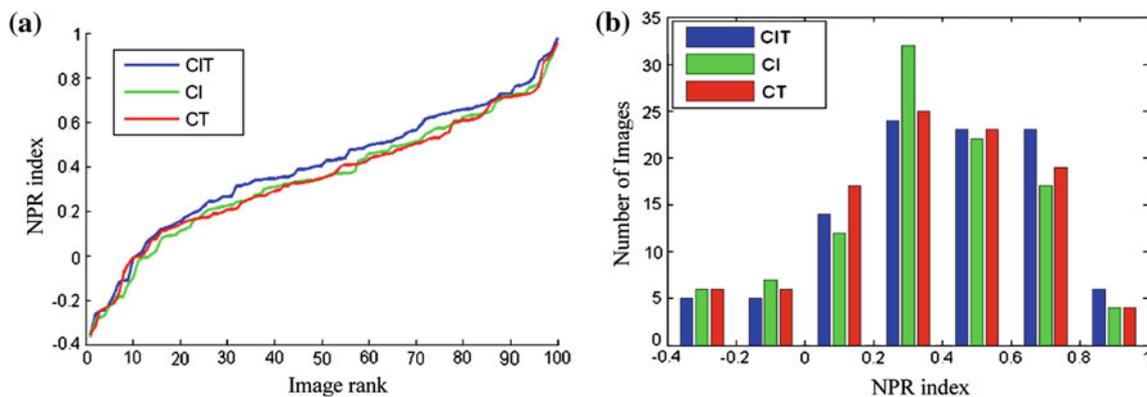
Thereafter we will refer to the segmentation derived by applying the marker-controlled watershed transform to the complementary intensity gradients, the complementary texture gradients and the fused complementary gradients as the Complementary Intensity (CI) watershed, the Complementary Texture (CT) watershed and the Complementary Intensity Texture (CIT) watershed respectively. All gradient images were normalised to the range  $[0, 1]$  before appli-

cation of the h-minima transform. Each of these watershed algorithms requires the specification of the h-minima transform scale parameter  $h$ . For each algorithm this parameter was optimised on the training set using all parameters from the following set:  $h = \{0.0100, 0.0311, 0.0522, 0.0733, 0.0944, 0.1156, 0.1367, 0.1578, 0.17889, 0.2\}$ . These  $h$  values were chosen because they result in segmentations ranging from over- to under-segmented. The degree of over- or under-segmentation was determined by visual inspection of corresponding segmentations. The parameters of each algorithm were optimised on the training set and these were used when evaluating the corresponding performance on the test set. These optimal parameters and the mean and variance of NPR values on the test are presented Table 1. On the test set the CIT watershed achieved a mean NPR index of 0.397. This represents a good improvement in segmentation quality relative to the CI and CT watersheds by which accomplished mean NPR index values of 0.350 and 0.352 respectively on the same data. The fact that each method produced a mean significantly above zero informs us that they generally produce a segmentation which is much better than random.

Using the same approach as [23] the individual NPR values are presented in plot and histogram form in Fig. 11. Figure 11a shows the NPR index on each image for each algorithm in the test set. The indices are plotted in increasing order for each algorithm; hence image 10 refers to the

**Table 1** Table showing for each algorithm the optimal parameters on train set with corresponding mean and variance of NPR index values on test set

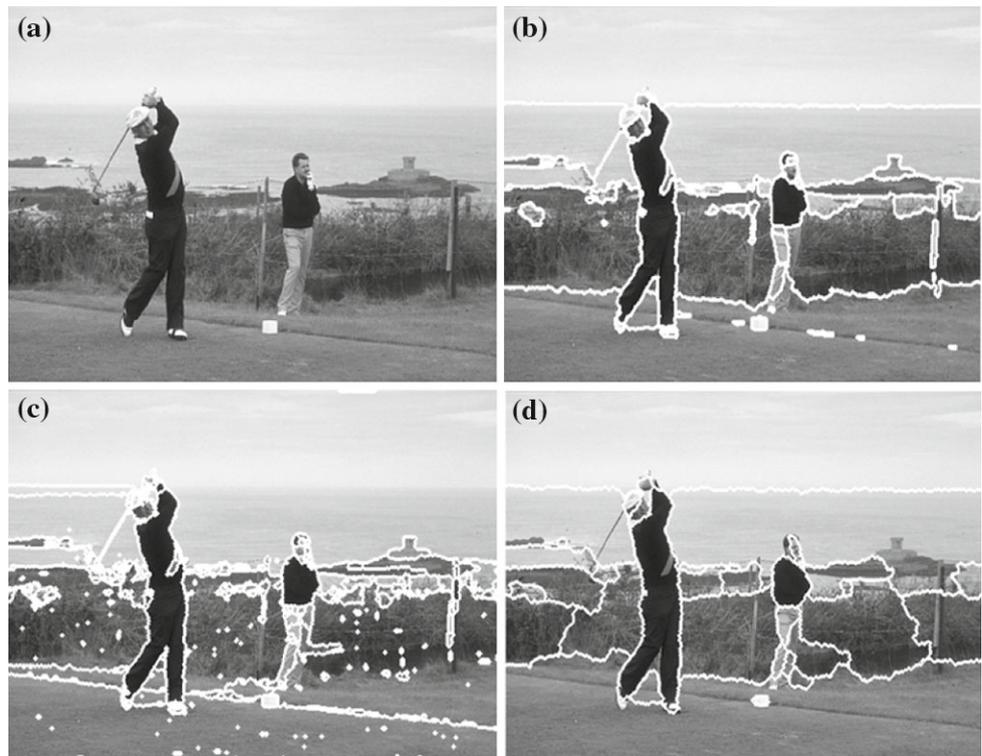
|                 | CI      | CT      | CIT     | CIT-PP  | JSEG   | Canny   | BG+TG   |
|-----------------|---------|---------|---------|---------|--------|---------|---------|
| $\sigma$ -value | –       | –       | –       | –       | –      | 3.5     | –       |
| h-value         | 0.03111 | 0.07333 | 0.05222 | 0.05222 | –      | 0.09444 | 0.01673 |
| Mean NPR        | 0.350   | 0.352   | 0.397   | 0.4176  | 0.322  | 0.363   | 0.376   |
| Variance NPR    | 0.085   | 0.080   | 0.086   | 0.0862  | 0.0770 | 0.1056  | 0.0854  |



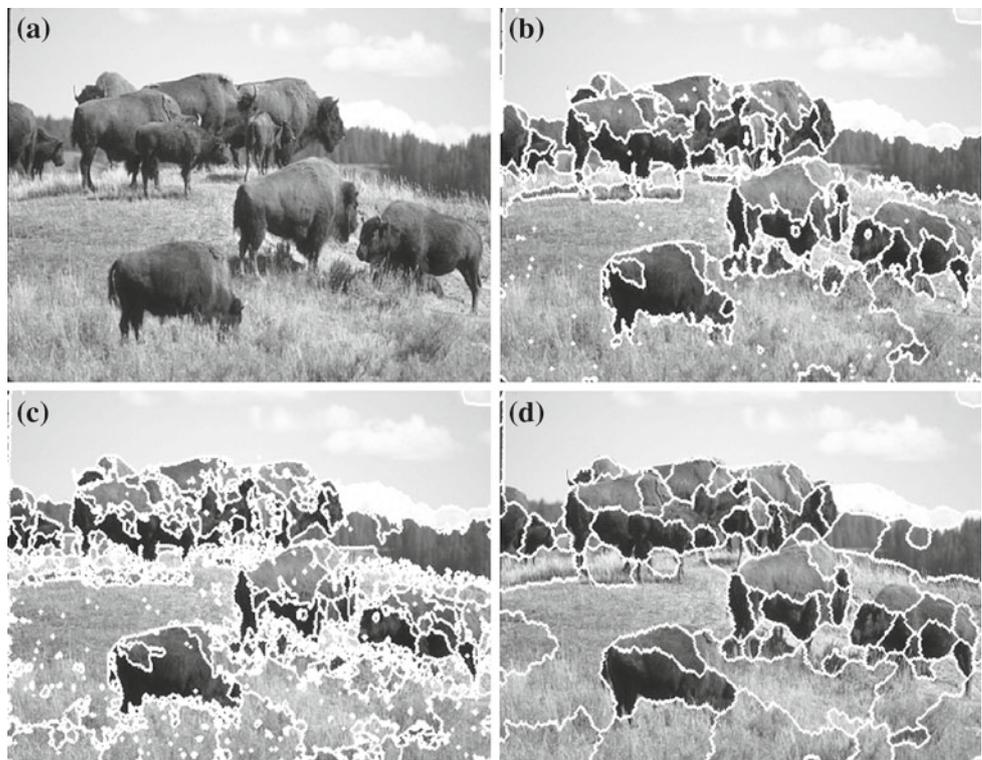
**Fig. 11** NPR index values on test set. Plot **a** shows the indices achieved on each image individually, ordered by increasing values. The  $x$ -axis represents the image index while corresponding NPR value is represented by the  $y$ -axis. Plot **b** shows the same information in the form of a group histogram. The  $x$ -axis represents the range of NPR index

values for each group while  $y$ -axis represents the number of images falling within the corresponding range. For example the first group of bars represents the number of images with index values between the range  $-0.4$  and  $-0.2$

**Fig. 12** An image taken from the test dataset is displayed in **a**. The corresponding segmentation results achieved using the CIT, CI and CT watersheds are displayed in **b**, **c** and **d**, respectively, and have NPR index values of 0.7951, 0.7217 and 0.7615, respectively



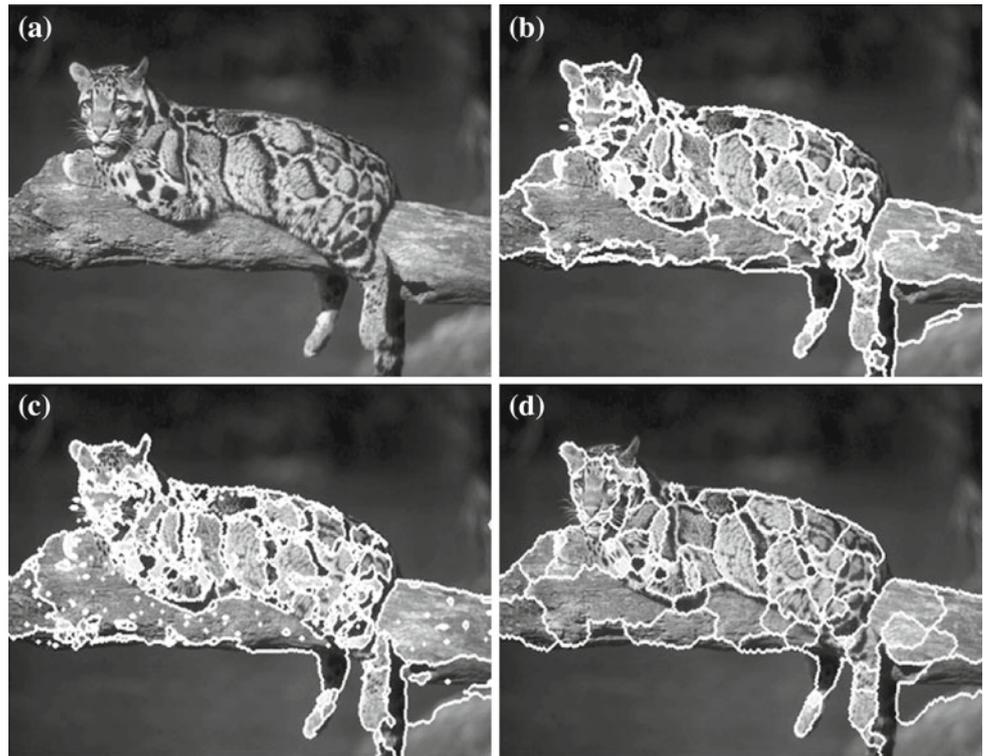
**Fig. 13** An image taken from the test dataset is displayed in **a**. The corresponding segmentation results achieved using the CIT, CI and CT watersheds are displayed in **b**, **c** and **d**, respectively, and have NPR index values of 0.658, 0.491 and 0.485, respectively



images with the tenth lowest index value for each algorithm, and may not represent the same image across algorithms. From this plot we see that the CIT watershed curve is generally above those corresponding to the CI and CT watersheds,

indicating its superior performance. The CI and CT watershed curves are close together indicating a similar level of performance. Figure 11b shows the individual NPR values in histogram form. An optimal histogram would be centred to

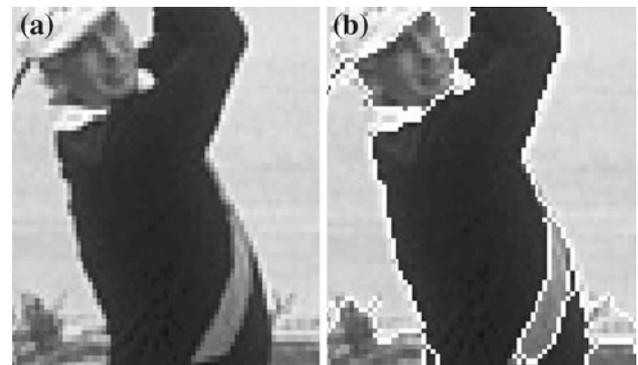
**Fig. 14** An image taken from the test dataset is displayed in **a**. The corresponding segmentation results achieved using the CIT, CI and CT watersheds are displayed in **b**, **c** and **d**, respectively, and have NPR index values of 0.7712, 0.7498 and 0.7601, respectively



the right with minimal spread. The CIT histogram is centred to the right more so than the corresponding CI and CT histograms. The CIT histogram spread is similar to that of the CI and CT histograms but it is less skewed to the left. Both CI and CT histograms are centred in the roughly the same region and have similar variance. The corresponding mean and variance values in Table 1 verify these facts. The above histogram analysis again reflects the better performance of the CIT watershed relative to the CI and CT watersheds. This superior performance is not surprising. It has been shown several times in the past [24–27] that a fusion of texture and multi-spectral features results in superior segmentation performance compared to using these features in isolation.

Some segmentation results achieved by each algorithm on the test dataset with corresponding NPR index values are displayed in Figs. 12, 13 and 14. In all results segmentation boundaries are represented by the colour white. These segmentation examples illustrate a number of properties of the CI, CT and CIT watersheds. In all examples it was clear by visual inspection and reference to the corresponding NPR index values that the CIT watershed outperformed both the CI and CT watersheds.

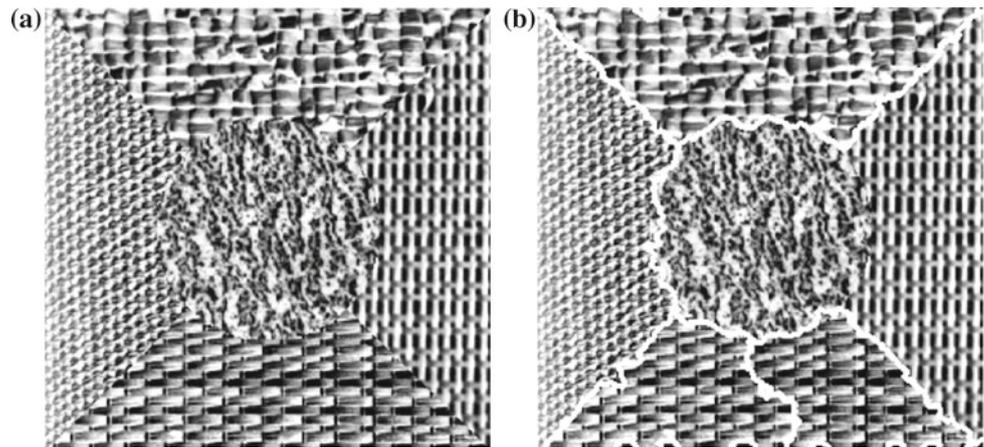
Boundaries which are represented in only a single complementary watershed segmentation, either CI or CT watershed, will still likely be represented in the segmentation derived from their fusion in the CIT watershed. For example, the boundary of the island in the background on the right hand side of Fig. 12a is represented by the CI watershed but is



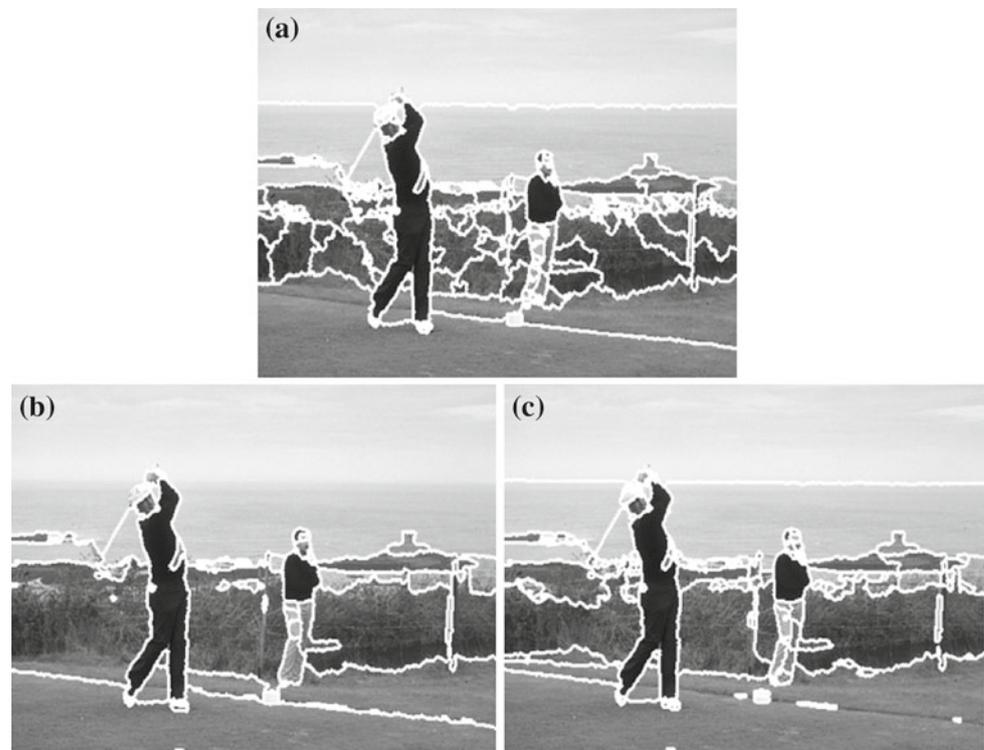
**Fig. 15** A crop of the image in Fig. 4a is displayed in **a**. The corresponding segmentation result achieved using the CIT watershed is displayed in **b**. The boundaries of the man, which have both an intensity and texture boundary, are localised to the spatial scale of the intensity boundary

not represented by the CT watershed. Despite this fact this boundary is represented by the CIT watershed. This illustrates that the CIT watershed fuses boundary information contained in both complementary gradient images to achieve superior segmentation performance. Figure 13 illustrates that the individual CI and CT watersheds may individually over-segment an image region but this effect will be reduced following fusion. In this example the grass region is over-segmented by both the CI and CT watersheds but this is significantly reduced by the CIT watershed.

**Fig. 16** A Brodatz texture and its corresponding and the corresponding segmentation achieved using the CIT watershed algorithm are displayed in **a** and **b**, respectively



**Fig. 17** Segmentation results of Fig. 12a without intensity smoothing, texture smoothing and texture gradient dilation are displayed in **a**, **b** and **c**, respectively



The image of the tiger in Fig. 14a represents a very challenging image to segment. This is due to the fact that the tiger is covered in patches which have both strong intensity and texture boundaries. As a result the CI and CT watershed both over-segment the tiger into a set of smaller regions which generally correspond to individual patches. Consequently, such regions are also represented in the CIT segmentation. This example highlights an inability of the proposed technique in some cases to distinguish object boundaries and strong object texture. This inaccuracy could be introduced by integrating prior knowledge in the form of specific object shape information or more generic rules regarding the shape of genuine object contours [31].

It is evident from segmentation results that the CI watershed tends to over-segmented in some textured regions returning a large number of very small segments in such regions. For example see the bottom of Fig. 12c. Much of these very segments are removed in CIT watershed segmentation results but some still remain; again see Fig. 12. These may be removed using the post-processing step we proposed earlier.

As discussed earlier the CIT watershed minimises the boundary localisation error when it fuses the individual complementary gradient images and this point is illustrated in Fig. 15. It was shown in Fig. 5 that the boundary of the man in this figure has both an intensity and texture gradient

respectively. In the CIT watershed segmentation these boundaries are localised to the spatial scale of the intensity boundary therefore minimising spatial uncertainty.

To demonstrate the ability of the proposed algorithm to segment very textured images we applied it the pure texture image in Fig. 16a. The corresponding segmentation result is displayed in Fig. 16b and is very accurate despite the presence of some over and under segmentation. Although other techniques exist which offer superior performance for pure texture images these have been optimised for this task and are not as generic as the proposed technique [16].

#### 4.2 Individual algorithm components

In order to determine the contribution of each major component in CIT watershed we evaluated segmentation performance when the component in question was removed. Using this strategy we evaluated the contribution of smoothing the intensity features, smoothing the texture features and dilating the texture gradients. Through analysing results on multiple images the following conclusion were drawn and are illustrated using the sample results in Fig. 17. It was generally found that failure to smooth intensity features resulted in an over-segmented image in textured regions; for example see Fig. 17a. Failure to smooth texture features did not have a significant impact on results. For example the segmenta-

tion result of Fig. 17b is very similar the corresponding CIT segmentation of Fig. 12b. We attribute this to the fact that texture features are calculated over an area which has the inherent effect of smoothing. Removal of the dilation step in general resulted in segmentations which were significantly worse both in terms of over- and under-segmentation. Again referring to Fig. 12b, the CIT segmentation is superior. To determine segmentation performance in a quantitative manner we used the training and testing methodology of Sect. 4.1. The results of this analysis are shown in Table 2. Comparing these results to that of the CIT algorithm in Table 1 we see that the removal of each component resulted in reduced performance. In particular the removal of intensity smoothing and gradient dilation returned a significant reduction in performance.

#### 4.3 Fused complementary gradients

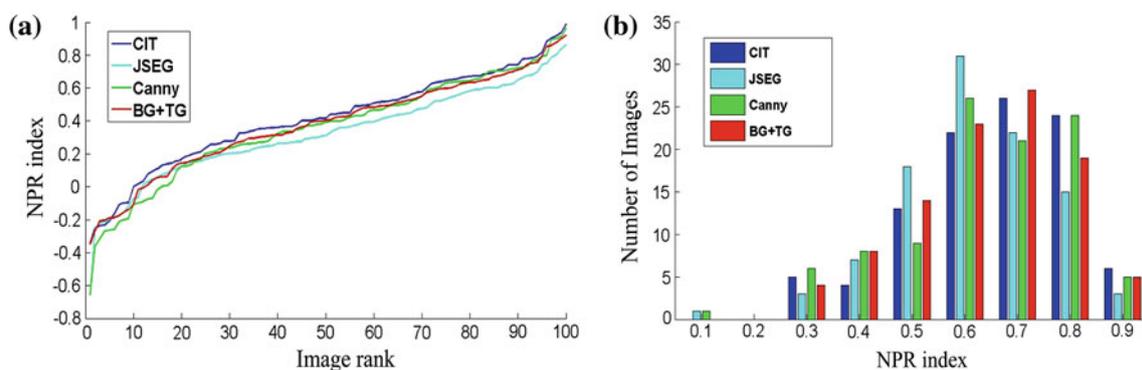
In this section, we evaluate the CIT watershed which has been post-processed to remove small segments. We refer to this segmentation algorithm as the CIT-PP watershed and evaluate it against segmentation results produced by the following benchmark methods:

1. BG + TG watershed: watershed transform applied to the Brightness Gradient (BG) + Texture Gradient (TG) operator of Martin et al. [7].
2. Canny watershed: watershed transform applied gradients extracted from a Gaussian smoothed image.
3. JSEG: segmentation method proposed by Deng and Manjunath [28].

The BG + TG operator is an algorithm which given ground-truth will learn to determine gradient values using texture and intensity features. Since it returns gradient values this method can be used as input for the watershed transform to

**Table 2** Table showing for each algorithm the optimal parameters on train set with corresponding mean and variance of NPR index values on test set

|                 | CIT                    |                      |             |
|-----------------|------------------------|----------------------|-------------|
|                 | No intensity smoothing | No texture smoothing | No dilation |
| $\sigma$ -value | –                      | –                    | –           |
| h-value         | 0.0522                 | 0.0311               | 0.0522      |
| Mean NPR        | 0.282                  | 0.383                | 0.307       |
| Variance NPR    | 0.092                  | 0.088                | 0.094       |



**Fig. 18** NPR index values on test set. Plot a shows the indices achieved on each image individually, ordered by increasing values. Plot b shows the same information in the form of a group histogram

**Fig. 19** The segmentation results achieved for **a** using the Canny, BG + TG, CIT and JSEG algorithms are displayed in **b**, **c**, **d** and **e**, respectively, and have corresponding NPR index values of 0.4940, 0.5383, 0.5970 and 0.5916 respectively



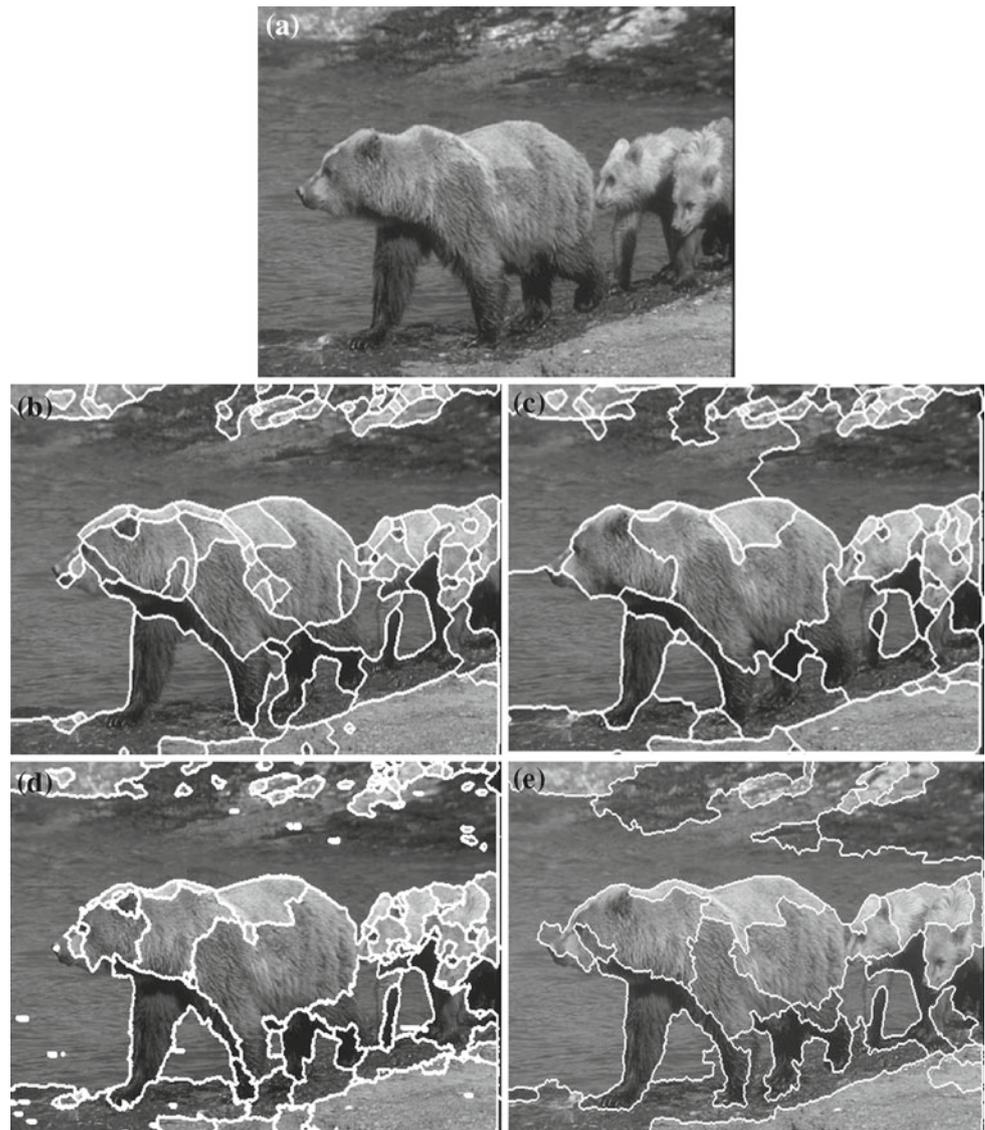
generate segmentation. To operate this algorithm requires the specification of the  $h$ -minima transform scale parameter  $h$ .

The Canny watershed algorithm which we have selected for use in this study has strong similarities to the edge detection technique proposed by Canny [29]. In both algorithms an initial gradient image is calculated by performing Gaussian smoothing followed by the application of a gradient operator. Canny then applied non-maximum suppression to suppress all gradient values which are not local maximums and remove multiple responses at each boundary. If the watershed transform locates a boundary due to a particular gradient response it will locate it once along the crest of that response. It therefore performs the same function as non-maximum suppression. This algorithm requires the specification of the  $h$ -minima transform scale parameter  $h$  and the scale of the Gaussian used for smoothing.

The JSEG algorithm is a multi-scale region based segmentation algorithm which fuses texture and intensity information to derive segmentation. It has the ability to parameterise itself based on the content of the image in question and therefore does not need require any input from the user. This should not be seen as a disadvantage as this method has been shown to produce useful segmentation results [28]. A C++ implementation of this algorithm, which is available from Deng's website [30], was used.

The CIT-PP watershed was optimised on the training set using the parameter set:  $h = \{0.0100, 0.0311, 0.0522, 0.0733, 0.0944, 0.1156, 0.1367, 0.1578, 0.17889, 0.2\}$ . The BG + TG watershed was optimised on the training set using the parameter set:  $h = \{0.0001, 0.0056, 0.0112, 0.0167, 0.0223, 0.0278, 0.0334, 0.0389, 0.0445, 0.0500\}$ . Finally the Canny watershed was optimised on the training set using

**Fig. 20** The segmentation results achieved for **a** using the Canny, BG + TG, CIT and JSEG algorithms are displayed in **b**, **c**, **d** and **e**, respectively, and have corresponding NPR index values of 0.4076, 0.3727, 0.4469 and 0.4274 respectively



the parameter set  $h = \{0.0100, 0.0311, 0.0522, 0.0733, 0.0944, 0.1156, 0.1367, 0.1578, 0.17889, 0.2\}$  and  $\sigma = \{0.5, 1.5, 2.5, 3.5, 4.5, 5\}$ . In the above parameter sets the  $\sigma$  parameters were chosen because they result in images ranging from slightly to significantly smoothed; the  $h$  parameters were chosen because they result in segmentations ranging from over- to under-segmented. All gradient images were normalised to the range  $[0 \ 1]$  before application of the  $h$ -minima transform.

The parameters of each algorithm were optimised on the training set and these were used when evaluating the performance of each on the test set. These optimal parameters used, along with the mean and variance of NPR values on the test are presented in Table 1. On the test set the CIT-PP watershed achieved a mean NPR index of 0.4176. The improvement in mean performance compared to the

CIT watershed is impressive and we believe this justifies that use of such a post-processing step. This result represents a good improvement in performance relative to the Canny watershed, the BG + TG watershed and the JSEG algorithm which accomplished mean NPR index values of 0.3589, 0.3767 and 0.322 respectively on the same data.

The individual NPR values for each algorithm are presented in plot and histogram form in Fig. 18. In the plot the CIT-PP watershed curve is generally above those corresponding to the other algorithms showing that it generally performs better. The curves of the BG + TG and Canny watersheds are similar with the BG + TG curve being slightly higher. Of all algorithms the JSEG algorithm curve was the worst and is significantly below the CIT-PP curve. The histograms for the CIT-PP and BG + TG have a similar level of spread, while the

**Fig. 21** The segmentation results achieved for **a** using the Canny, BG + TG, CIT and JSEG algorithms are displayed in **b**, **c**, **d** and **e**, respectively, and have corresponding NPR index values of 0.2941, 0.76184, 0.84947 and 0.5254, respectively



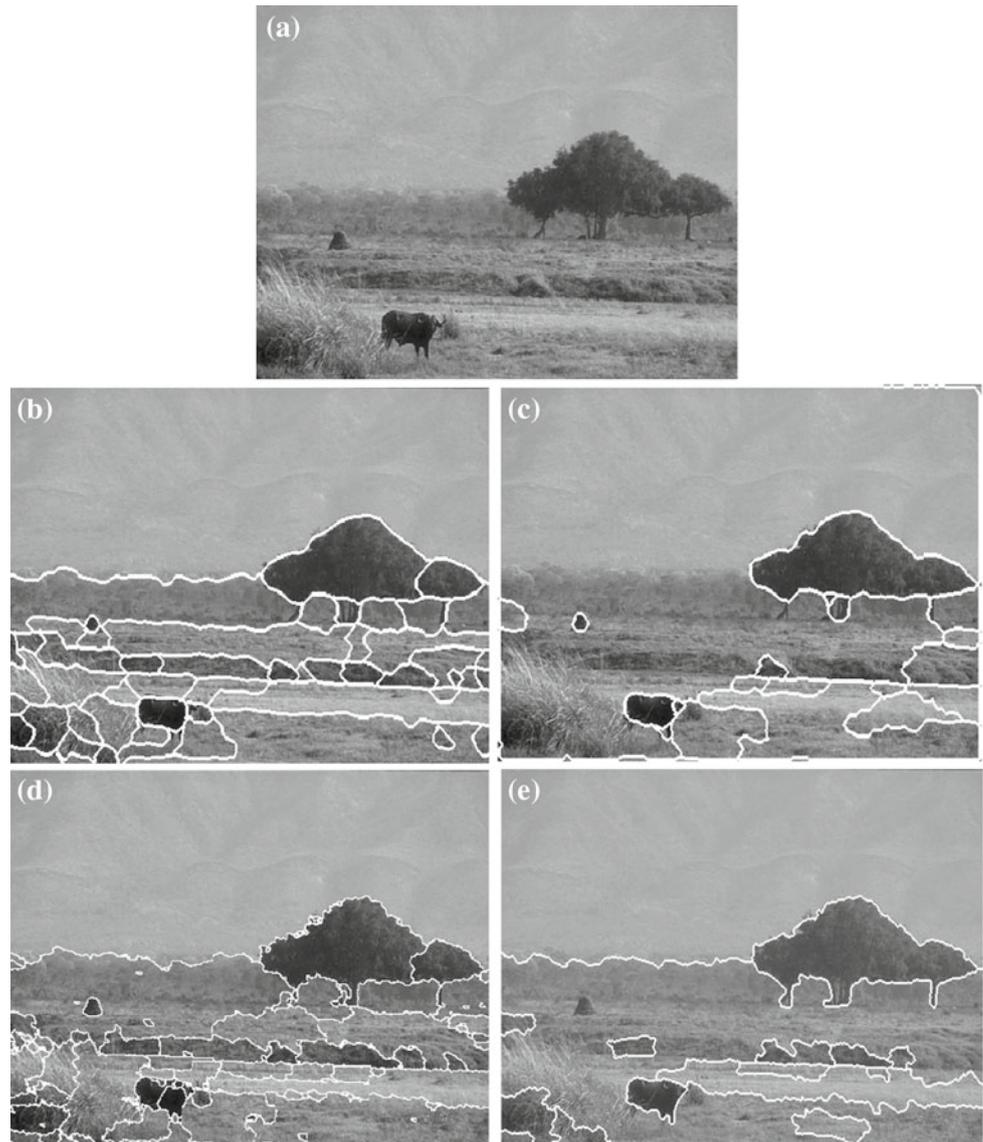
Canny histogram has slightly more spread and the JSEG histogram slightly less. The CIT-PP histogram is centred more to the right than any other. The corresponding mean and variance values in Table 1 verify these facts. From this histogram analysis we can say that of all algorithms tested the CIT-PP watershed does not have the least variance in segmentation quality but it generally performs better.

Some segmentation examples achieved by each algorithm on the test dataset with corresponding NPR index values are displayed in Figs. 19, 20, 21 and 22. In all examples, boundaries are represented by the colour white. These segmentation examples reveal a number of properties of the all segmentation algorithms. On all images it is clear by visual inspection and reference to the corresponding NPR index values that the CIT-PP watershed generally outperforms the other methods. Although all methods over-segment to a degree, the bench-

mark methods tend to over-segment more than the CIT-PP watershed. Examples where this is evident are the central bear in Fig. 20 and the trees in the background of Fig. 21. Due to the fact that it does not integrate texture information, the Canny watershed regularly fails to detect a number of key boundaries. Example of this would be the boundary between the two animals in Fig. 19 and the boundary between the trees and grass in left of Fig. 21. The JSEG method tends to over-segment and fails to detect key boundaries at the same time. This point is illustrated in Fig. 21 where JSEG over-segmented the grass area at the bottom of the image while it failed to detect the elephant head boundaries in the centre of the image.

Due to image segmentation being an ill-posed problem it would be difficult to eliminate all over-segmentation from the CIT-PP watershed even through the incorporating top-down

**Fig. 22** The segmentation results achieved for (a) using the Canny, BG + TG, CIT and JSEG algorithms are displayed in b, c, d and e, respectively, and have corresponding NPR index values of 0.6420,  $-0.1420$ , 0.6555 and 0.6193, respectively



knowledge into the segmentation process. One individual's perception of an accurate segmentation may be perceived as inaccurate by another individual. It is because of this fact that the Berkeley dataset contains multiple ground-truths for each image. It is also because of this fact that the NPR index measures performance by comparison to such a set of ground truths. Top-down knowledge could improve the segmentation in some cases. For example the central bear in Fig. 20 has some very strong internal texture and intensity edges; therefore top-down knowledge would be required to prevent these from being represented as genuine boundaries. Despite this fact the proposed method has been shown in both a quantitative and qualitative manner to outperform existing state of the art techniques which also attempt to derive segmentation in a bottom-up manner.

From the above segmentation examples, particularly the trees in the background of Fig. 21, we see that object boundaries which have both an intensity and texture boundary are localised very accurately by the CIT-PP watershed. The other methods do not minimise this spatial uncertainty and therefore their corresponding segmentation boundaries are not localised so accurately.

## 5 Conclusions

This paper presents a new watershed based segmentation algorithm which overcomes the challenges of a competing gradient set and the spatial uncertainty inherent in texture gradients. A complementary texture and intensity gradient

set is described which allows boundary gradients to be measured in terms of both visual cues. It therefore offers a more accurate measure of boundary gradient compared to existing competing gradient sets, where each boundary gradient can only be described in terms of a single visual cue. Texture, unlike intensity, is a spatial property and therefore its gradients suffer from the uncertainty principle resulting in poorer boundary localisation compared to intensity gradients. A new fusion strategy, namely the CIT-PP watershed, is proposed which accounts for this uncertainty and minimises the error in boundary localisation.

An in-depth qualitative and quantitative evaluation of the CIT-PP watershed against a total of seven other segmentation methods is presented. The accuracy of the proposed fusion strategy is demonstrated by an evaluation against segmentation achieved by the individual features in isolation. To evaluate the usefulness of the CIT-PP watershed against other previously published techniques in the domain of computer vision, three benchmark methods were selected. In all cases the CIT-PP watershed produced impressive results and outperformed all other methods. Although segmentation results achieved by the CIT-PP watershed are of a good quality some over-segmentation was evident. Reducing such false-positives may be achieved by the integration of top-down knowledge into the segmentation process. Colour is another very important cue used by the visual system to define boundaries. In this work we have ignored its presence and concentrated on the visual cues of intensity and texture. We believe that incorporating colour information would increase segmentation performance and this shall be the focus of future research work.

## References

- Blaschke, T.: Object-based contextual image classification built on image segmentation. In: IEEE Workshop on Advances in Techniques for Analysis of Remotely Sensed Data, Washington DC, USA, pp. 113–119 (2003)
- Hofmann, T., Puzicha, J., Buhmann, J.M.: Unsupervised texture segmentation in a deterministic annealing framework. *IEEE Trans. Pattern Anal. Mach. Intell.* **20**(8), 803–818 (1998)
- Zhang, Y.-J.: An overview of image and video segmentation in the last 40 years. In: Zhang, Y.-J. (ed.) *Advances in Image and Video Segmentation.*, pp. 1–15. IRM Press, Pennsylvania (2006)
- Roth, V., Ommer, B.: Exploiting low-level image segmentation for object recognition. *Lecture Notes in Computer Science*, vol 4174, pp. 11–20. Springer, Heidelberg (2006)
- Agarwal, S., Awan, A., Roth, D.: Learning to detect objects in images via a sparse, part-based representation. *IEEE Trans. Pattern Anal. Mach. Intell.* **26**(11), 1475–1490 (2004)
- Soille, P.: *Morphological image analysis: principles and applications.* Springer, Berlin (2002)
- Martin, D.R., Fowlkes, C.C., Malik, J.: Learning to detect natural image boundaries using local brightness, color, and texture cues. *IEEE Trans. Pattern Anal. Mach. Intell.* **26**(5), 530–549 (2004)
- Chaji, N., Ghassemian, H.: Texture-gradient-based contour detection. *European Association for Signal Processing (EURASIP) J Appl. Signal Process.* **2006**, 1–8 (2006)
- O'Callaghan, R.J., Bull, D.R.: Combined morphological-spectral unsupervised image segmentation. *IEEE Trans. Image Process.* **14**(1), 49–62 (2005)
- Malik, J., Belongie, S., Leung, T., Shi, J.: Contour and texture analysis for image segmentation. *Int. J. Comput. Vis.* **43**(1), 7–27 (2001)
- Petrou, M., Sevilla, P.G.: *Image processing: dealing with texture.* Wiley, London (2006)
- Perona, P., Malik, J.: Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Mach. Intell.* **12**(7), 629–639 (1990)
- Black, M., Sapiro, G.: Edges as outliers: anisotropic smoothing using local image statistics. In: *International Conference on Scale-space Theories in Computer Vision*, Corfu, Greece, pp. 259–270 (1999)
- Deng, H., Liu, J.: Development of anisotropic diffusion to segment texture images. *J. Electron. Imag.* **12**(2), 307–316 (2003)
- Corcoran, P., Winstanley, A.: Watershed segmentation using a multiscale ramp edge merging strategy. In: *Proceedings of International Conference on Machine Vision and Image Processing*, NUI Maynooth, Ireland, pp. 158–168 (2007)
- Clausi, D.A., Jernigan, M.E.: Designing Gabor filters for optimal texture separability. *Pattern Recognit.* **33**(11), 1835–1849 (2000)
- DeVeaux, R., Velleman, P.F., Bock, D.E.: *Stats: Data and Models.* Addison Wesley, Reading (2004)
- Shao, J., Forstner, W.: Gabor wavelets for texture edge extraction. In: *ISPRS Commission III Symposium on Spatial Information from Digital Photogrammetry and Computer Vision*, Munich, pp. 745–752 (1994)
- Kruizinga, P., Petkov, N.: Nonlinear operator for oriented texture. *IEEE Trans. Image Process.* **8**(10), 1395–1407 (1999)
- Jobanputra, R., Clausi, D.A.: Preserving boundaries for image texture segmentation using grey level co-occurring probabilities. *Pattern Recognit.* **39**(2), 234–245 (2006)
- Chen, J., Pappas, T.N., Mojsilovic, A., Rogowitz, B.E.: Adaptive perceptual color-texture image segmentation. *IEEE Trans. Image Process.* **14**(10), 1524–1536 (2005)
- Martin, D., Fowlkes, C., Tal, D., Malik, J.: A database of human segmented natural images and its applications to evaluating segmentation algorithms and measuring ecological statistics. In: *International Conference on Computer Vision*, Vancouver, Canada, pp. 416–423 (2001)
- Unnikrishnan, R., Pantofaru, C., Hebert, M.: Toward objective evaluation of image segmentation algorithms. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**(6), 929–944 (2007)
- Zhang, Y.: Optimization of building detection in satellite images by combining multispectral classification and texture filtering. *ISPRS J Photogramm Remote Sens.* **54**(1), 50–60 (1999)
- Herold, L.X.M., Clarke, K.C.: Spatial metrics and image texture for mapping urban land use. *Photogramm. Eng. Remote Sens.* **69**(9):991–1001 (2003)
- Puissant, A., Hirsch, J., Weber, C.: The utility of texture analysis to improve per-pixel classification for high to very high spatial resolution imagery. *Int. J. Remote Sens.* **26**(4), 733–745 (2005)
- Ma, W.-Y., Manjunath, B.S.: EdgeFlow: a technique for boundary detection and image segmentation. *IEEE Trans. Image Proc.* **9**(8), 1375–1388 (2000)
- Deng, Y., Manjunath, B.S.: Unsupervised segmentation of color-texture regions in images and video. *IEEE Trans. Pattern Anal. Mach. Intell.* **22**(8), 800–810 (2001)
- Canny, J.F.: A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.* **8**(6), 679–698 (1986)

30. JSEG C++ implementation. <http://vision.ece.ucsb.edu/segmentation/jseg/>. Accessed 30-5-2009 (Accessed 25-6-2010)
31. Corcoran, P., Mooney, P., Winstanley, A., Tilton, J.: Convexity grouping of salient part contours. In: First Intentional Workshop on Parts and Attributes, European Conference on Computer Vision (Under Review) (2010)

### Author Biographies

**Padraig Corcoran** received a BSc in Computer Science and Software Engineering from the NUIM in 2004. He completed his PhD in computer science at the National Centre for Geocomputation (NCG) and Department of Computer Science NUIM 2008.

**Adam Winstanley** received a BA in Archaeology from the University of Cambridge in 1978. He gained MSc and PhD degrees in Computer Science from Queen's University, Belfast in 1987 and 1991 respectively. Currently, he is Senior Lecturer and Head of Dept. of Computer Science and Senior Research Associate of the NCG at NUIM.

**Peter Mooney** holds a BSc and PhD in Computer Science from the National University of Ireland, Maynooth (NUIM). He is currently Environmental Research Data Manager with the Irish Environmental Protection Agency and a Postdoctoral Research Fellow in the Department of Computer Science at NUIM.