

DETECTING SHADOWS IN QUICKBIRD SATELLITE IMAGES

V. Arévalo^{a,*}, J. González^a, J. Valdes^b, G. Ambrosio^b

^a Dept. of System Engineering and Automation, University of Málaga, Campus Teatinos–Complejo Tecnológico s/n, 29071 Málaga (Spain) - {varevalo,jgonzalez}@ctima.uma.es

^b DECASAT S.L., Parque Tecnológico de Andalucía, Severo Ochoa 4, 29590 Campanillas, Málaga (Spain) - {jmvaldes,gambrosio}@decasat.com

KEY WORDS: Shadow detection, invariant colour indices, high resolution imagery, QuickBird.

ABSTRACT:

High resolution images provided by latest missions such as QuickBird, Ikonos or OrbView have opened a new and important age in the remote sensing field, mainly in those applications where is crucial to extract detailed information such as buildings, streets, vehicles, etc. The improvement in resolution makes also that something inherent to an image as shadows take on special significance. This paper reviews some previous published methods to detect shadows and analyzes their suitability for being applied to colour satellite images. Based on this study, this work presents a procedure which exploits colour invariants and image edges to accurately identify shadows in QuickBird images. The presented method has been successfully tested with images acquired under different lighting conditions and covering both urban and rural areas.

1. INTRODUCTION

High resolution images provided by latest missions such as QuickBird, Ikonos or OrbView have opened a new range of applications in the remote sensing field because of the possibility of extracting detailed information from the images. These applications include some that are becoming common in the last years, such as natural disaster monitoring (i.e. flooding, earthquake, etc.) or urban change detection, and others that will be a real possibility in the next years, such as urban scene reconstruction, cartography update, urban inventory, etc.



Figure 1: Shadows in QuickBird imagery.

The improvement in spatial resolution of satellite imagery makes also that something inherent to images as shadows take on special significance for different reasons. On the one hand, they cause the partial or total loss of radiometric information in the affected areas, and consequently they make more difficult or even fail image analysis processes like object detection and recognition, temporal change detection, 3D scene reconstruction, etc. (see figure 1). On the other hand, we may take advantage of the presence of shadows as a valuable cue for inferring 3D scene information based on the position and shape of cast shadow, for example, for building detection, delineation, and height estimation (Irvin & McKeown, 1989). It is clear, then, the convenience of accurately detecting shadowed areas in

satellite images in order to restore the affected areas and/or to infer 3D information.

Shadow detection has received great attention within the computer vision field in the last years, though not much work has been done towards the application of these results to high resolution satellite images recently available. This paper describes the main approaches to detect shadows in images, and analyzes their suitability for being applied to colour satellite imagery. Based on this study, we propose a procedure which exploits two invariant colour indices of complementary spaces as well as edge information to effectively and accurately detect shadows in QuickBird images (0.6 m./pixel). Additional information such as sun azimuth or height information (DEM) is not considered here.

Very briefly, the proposed method deals with shadow detection through a region growing procedure which basically consists of two stages:

1. Small groups of pixels with very high confidence of being shadow are selected as seeds of shadow regions. They are obtained from a neighbourhood of local maxima in the c_3 component of the $c_1c_2c_3$ colour space (Gevers & Smeulders, 1999). Each shadow area is then characterized by a Gaussian distribution of the c_3 values of the pixels within this region.
2. From these seeds, the shape of the shadow region is recursively extended by adding adjacent pixels which are consistent with the above distribution. To detect as precise as possible the shape of the shadowed area, this process takes into account region boundary information provided by an edge detector.

One of the problems when using the c_3 component is their instability for certain colour values which leads to the misclassification of non-shadow pixels as shadow (false positives). As reported in (Salvador et. al. 2001; Gevers & Smeulders, 1999), this occurs for both pixels with low values of saturation and for pixels with extreme intensity values (low and

* Corresponding author.

high). To avoid this, some components of the *HSV* and *RGB* spaces are checked during the above two stages. The presented method has been successfully tested with images acquired under different lighting conditions taken in diverse seasons and with different sun elevation angles and covering both urban and rural areas.

The remainder of this paper is organized as follows. In section 2 we review some of the most representative methods to detect shadows. In section 3, we describe some colour spaces of special interest for shadow detection. In section 4, the proposed method is described. In section 5, we present some experimental results. And finally, some conclusions and future work are outlined.

2. AN OVERVIEW OF SHADOW DETECTION METHODS

In this section we first give a brief introduction to the nature of shadows (types and structure) and then review some of approaches followed in the literature to detect shadows in digital images.

Shadows occur when objects totally or partially occlude direct light from a source of illumination. Shadows can be divided into two classes: *cast* and *self* (see figure 2). A cast shadow is projected by the object in the direction of the light source; a self shadow is the part of the object which is not illuminated by direct light. The part of a cast shadow where direct light is completely blocked by its object is called *umbra*, while the part where direct light is partially blocked is called *penumbra*. See (Funka-Lea & Bajscy, 1995) for a deeper analysis of the physics of shadows.




Figure 2: Different types of shadows.

When dealing with aerial or satellite images neither self shadows nor the distinction between umbra and penumbra are of significance. Thus, most of the proposed methods presented in the remote sensing field only deal with cast shadows, and follow two general approaches: model-based methods or image-based methods (based on shadow properties).

2.1 Model-based methods

In this approach, the 3D geometry and illumination of the scene are assumed to be known. This includes the sensor/camera localization, the light source direction, and the geometry of observed objects, from which *a priori* knowledge of shadow areas is derived. For example, we can consider polygonal regions to approximate the shadows of buildings or urban elements in some simple urban scenes. However, in complex scenes with a great diversity of geometric structures, as it is usually the case of QuickBird images, these models are too restrictive to provide a good approximation. In addition, in most applications the geometry of scene and/or the light sources are unknown. These facts limit the applicability of this approach to

very constrained scenes, as the ones reported in (Koller, et. al., 1993) and (Irvin & McKeown, 1989).

2.2 Image-based methods

This approach makes use of certain image shadow properties such as colour (or intensity), shadow structure (umbra and penumbra hypotheses), boundaries, etc., without any assumption about the scene structure. Nevertheless, if any of that information is available, it could also be used to improve the detection process performance. Some common ways of exploiting image shadow characteristics are:

- The value of shadow pixels must be low in all the RGB bands. Shadows are, in general, darker than their surrounding, thus they have to be delimited by noticeable borders (shadow boundaries) (Salvador, et. al., 2001; Jiang & Ward, 1994).
- Shadows do not change the surface texture. Surface markings tend to continue across a shadow boundary under general viewing conditions (Jiang & Ward, 1994).
- In some colour components (or combination of them) no change is observed whether the region is shadowed or not, that is, they are invariant to shadows (Salvador, et. al., 2001; Etemadnia & Reza-Alsharif, 2003) (see section 3 for further information).

When some scene knowledge is also available it can be used in combination with the above procedures. Examples of how to apply such knowledge are:

- One or more edges of the cast shadow are oriented exactly in the light direction (Massalabi, et. al., 2004).
- Shadow size depends on the light source direction and the object height (Stevens, et. al., 1995).

Our technique for shadow detection is a pure image-based method, since we exclusively apply image properties. Although the sun azimuth and sensor/camera localization are typically available for satellite images (i.e. in QuickBird), we do not make use of this information since, in general, it is not possible to derive precise model of the acquired area (in most cases, the 3D geometry of scene is unknown or of a great complexity). In particular, we exploit both shadow invariant components and edge information to segment shadowed regions. Next, we describe different invariant colour spaces and analyse their suitability for being applied to shadow detection in colour high resolution satellite images.

3. INVARIANT COLOUR SPACES

Colour can be represented in a variety of three dimensional spaces, such as *RGB*, *HSV*, *XYZ*, $c_1c_2c_3$, $l_1l_2l_3$, *YCrCb*, *Lab*, *Luv*, etc. (Ford & Roberts, 1998). Each colour space is characterized by interesting properties which make it especially appropriated for a specific application. Among these properties, we can stand out the invariant features. For example, some colour spaces are invariant to changes in the imaging conditions including viewing direction, object surface orientation, lighting conditions, and shadows. Traditional colour spaces such as normalized *RGB* (*rgb*), hue and saturation (from *HSV*) and more recently, $c_1c_2c_3$ (Gevers & Smeulders, 1999) are colour representations that have revealed some kind of shadow invariant property. Of remarkable performance is the latest one, $c_1c_2c_3$, which has been successfully used by (Salvador, et. al., 2001) to extract shadows in simple images with few single-colour objects and flat (non-texture) background. Obviously, these premises can not be assumed in high resolution colour

satellite images where the observed scenes are highly textured, objects may have many different colours and the scene, in general, is very complex (see figure 3). Nevertheless, to evaluate the limitations of this space for shadow detection in QuickBird imagery we have performed tests over a broad set of images, acquired under different lighting conditions and covering both urban and rural areas. The result of our tests verifies the suitability of the c_3 component to identify shadowed regions which produce much higher response than those non-shadowed (see figure 4). The c_3 band is computed from the RGB representation through the following non-linear transformation:

$$c_3 = \arctan\left(\frac{B}{\max\{R, G\}}\right) \quad (1)$$

where R , G , and B = the red, green, and blue components of each pixel in the image, and MAX and MIN = the maximum and minimum of these values.



Figure 3: Typical urban scenes which illustrate the complexity of high resolution colour satellite images (e.g. highly textured, objects may have many different colours, etc.).

Despite of its promising possibilities our tests have revealed also the following problems:

- The c_3 band is quite noisy, which causes the misclassification of shadow pixels as non-shadow (true negatives) as well as inaccuracies in the shadow boundaries.
- Equation (1) becomes unstable for low saturation (S) values, which causes the misclassification of non-shadow pixels as shadow (false positives). This behaviour has been also reported by (Gevers & Smeulders, 1999).
- Colours close to the white (high values of V) and the blue (high values of B) are wrongly detected as shadows (false positives). In our test images, typically this turns to happen in saturated areas and for the water of swimming pools (as illustrated in figure 4).

To overcome the above problems, in our approach we incorporate the following two actuations:

- To minimize the noise effect, the c_3 image is smoothed and we make use of the image gradient to delimitate more precisely the shadow areas.
- We check the components S , V (from HSV) and B (from RGB) and do not classify the pixel as shadow if it is one of the above cases. This may give rise to small gaps in a shadow region, but it is a minor cost to pay for avoiding false positives, especially since these small gaps may be effectively filled in by a morphological filter.

Next, we describe the developed method, which is based on the segmentation of a smoothed c_3 image while taking into account the above considerations.




Figure 4: c_3 components corresponding to the test images shown in figure 3. Local maxima produced by colours close to the white (left) and the blue (right) are marked in both images.




Figure 5: Horizontal scan-lines extracted from (a) the c_3 component and (b) the smoothed c_3 component.

4. DESCRIPTION OF THE PROPOSED METHOD

Figure 6 shows the block diagram of the proposed system which comprises the two stages described next.

4.1 Pre-processing stage

The system input is a RGB image from which the following components are computed: c_3 , saturation (S) and intensity (V). The c_3 image is convolved with an 5×5 average kernel to minimize the effects of the noise (see figure 5) and the magnitude of gradient of the intensity image (V) is computed from a 5×5 Sobel detector (see figure 7.c).

4.2 Shadow detection stage

A generic region growing process starts with a small group of pixels (called seed region) and recursively adds neighbour pixels as long as they verify certain region membership rules.

This technique is particularly suitable for our problem here because many shadow regions may exist throughout the image and each one may have different radiometric profiles (shadow strength may vary from one zone of the image to another). Obviously, one of the pivotal points when applying this

technique is that of reliably placing the seeds in the image: at least one per shadow region (no matter if they are redundant) and no seed at non-shadow pixels. Next we describe our implementation of this technique in depth.

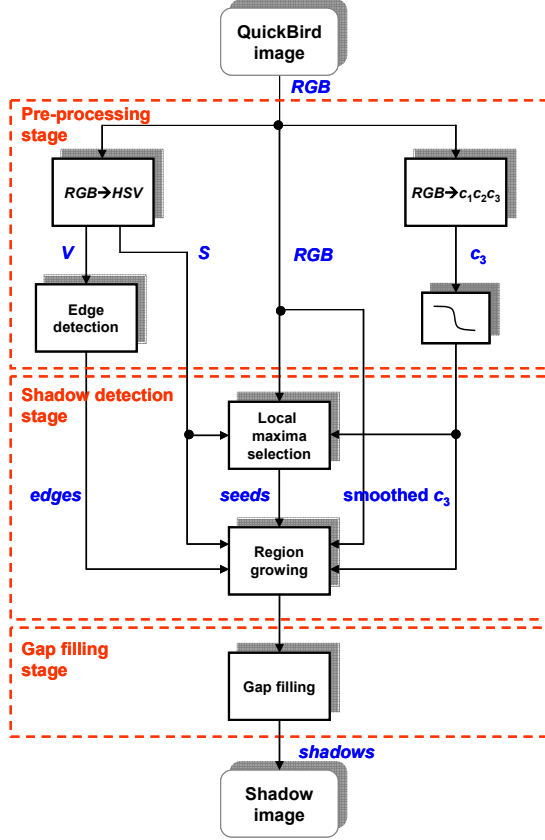


Figure 6: The structure of the proposed shadow detection system.

4.2.1 Seed selection: A seed is a square region (window) of 9×9 pixels. It is placed at pixels of the smoothed c_3 image that verify the following conditions:

1. The centre is a local maxima and its 9×9 neighbourhood must have values higher than the mean of the whole c_3 image (see figure 8). The size of this window gives the minimum size allowed for the shadows.
2. The means of the B (blue) and intensity (V) components of the window pixels must be lower than certain thresholds T_B and T_V , respectively, to avoid the problems mentioned in the previous section. In the same way, the mean of the saturation (S) component of the window pixels must be higher than a threshold T_S .
3. None of the window pixels are of another previous seed.

In our implementation, these values have been empirically set to $T_B=0.65$, $T_V=0.85$ and $T_S=0.02$.

These conditions make highly likely that the seed window corresponds to a shadow. Since in a typical shadow region several seeds are usually identified, we do not matter very much if some of them are ruled out by these demanding premises: the important point is not to generate false positive seeds. When several seeds are placed within the same shadow they eventually meet each other during the growing process and end up being a single region. The only concern here could be the higher computational requirements involved in comparison with having just one seed per shadow region, though it is not a

critical issue and therefore it has been obviated in our implementation.

Each of the seeds is now taken as a shadow prototype which is characterized by means of Gaussian distribution $N(c_3)$ of their c_3 values (see figure 7). The process described below relies on this information for growing the seeds over the shadow region.

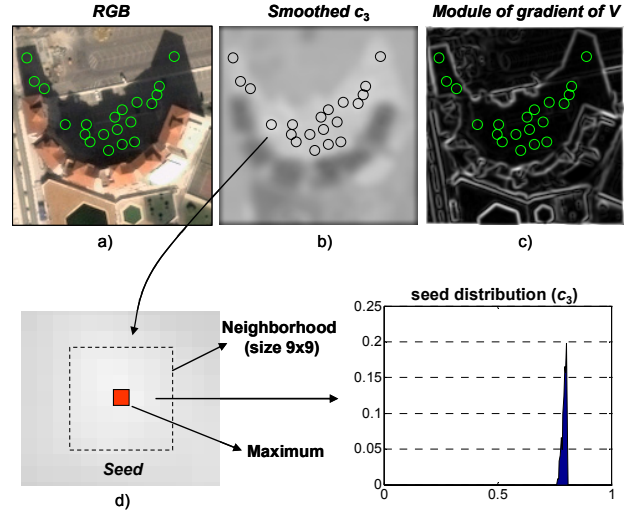


Figure 7: Shadow seeds identified in the smoothed c_3 component (first row). (d) Typical shadow seed and distribution estimated from the c_3 values of the pixels contained inside of the maximum neighbourhood.

4.2.2 Growing process: This is procedure recursively executed for all the 8-neighbours of the boundary pixels of the shadow region. Starting with the seed window, a pixel is classified as shadow and added to the region if it satisfies the following conditions:

- It has not previously added to any shadow region.
- It has to be below a certain Mahalanobis distance d_0 from the mean \hat{c}_3 of the region, that is, its c_3 value follows the Gaussian distribution $N(c_3)$:

$$\frac{(c_3 - \hat{c}_3)^2}{\sigma^2} < d_0 \quad (2)$$

where σ is the estimate standard deviation of the region.

- It satisfies the above condition 2 imposed to the seed pixels.
- The magnitude of the gradient of V is below a given threshold $T_E = 0.25^\dagger$.

If the pixel is incorporated to the region, the Gaussian distribution $N(c_3)$ is updated (that is, \hat{c}_3 and σ). The process ends when none of the neighbour pixels have been added to the region.

Figure 8 illustrates the growing process in one-dimension. It displays several scan-lines: the module of gradient of the V image, the c_3 and the $Blue$ component. We can observe three detected seed windows and the shadow boundaries (indicated by vertical lines) where the region growing process ends because of the high gradient magnitude.

[†] We have tried also with the gradient of the c_3 smoothed image but because of the noise (still there) the results are not reliable enough.

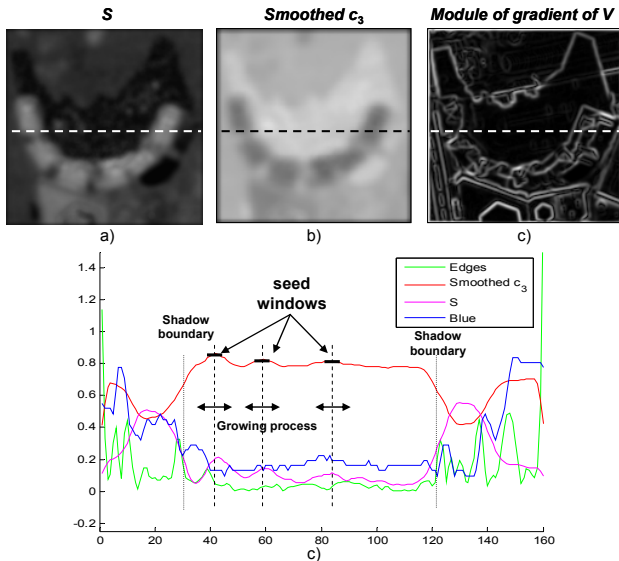


Figure 8: The seed selection and growing stages are illustrated on a horizontal scan-line of the images: (a) the S component, (b) the smoothed c_3 component and (c) the module of gradient of V . (d) All scan-lines plotted together.

4.3 Gap-filling stage

After the region growing small gaps in the shadow region may appear because of the noisy nature of the c_3 image. Thus, we apply a typical pair of morphological operations over the binary shadow image consisting of a dilatation with a 2×2 filter, followed by a 2×2 erosion (Sonka et al., 1996). This fills in gaps up to 2 pixels wide.

5. EXPERIMENTAL RESULTS

The proposed method has been tested with a variety of QuickBird images acquired under different lighting conditions (diverse seasons and sun elevation angles) and covering both urban and rural areas

For showing results of the proposed method, in this paper we have selected two image pieces of the city of Málaga (Spain): one covering a residential area with buildings and some forest (figure 9.a), which is particular interesting for checking the method behaviour on high-textured areas and with small and irregular shadows; and the other one is a typical urban area (i.e. high buildings with a variety of shapes, colours, and dark regions) (figure 10.a) that allows us to illustrate the method performance in the presence of saturations and dark regions.

Figure 9.b and figure 10.b show the shadows detected in the above images for the set of parameters (thresholds) mentioned in section 4. In reference to these figures, we would like to catch your attention in the following concrete points of interest:

- In figure 9.b, although some small shadows have not been detected, the method performs quite acceptably given the difficulty, even for a human, to distinguish between the tree crown and shadows. We have checked that most of the undetected shadows are because no seed was placed in it. This could be corrected to a certain extend, but at the expenses of some false positives. Please notice that, with

the current set of parameters, practically no false positives appear in the images.

- Figure 10.b shows less and bigger shadow regions than the earlier image, which are pretty well delimited. Shadows in this image are not easy to detect because shadows are cast over a highly texture surface (see the bottom-left building) and also because some elements could be wrongly classified as shadows (i.e. the dark hexagons, the saturated side of the buildings, etc.).

To quantify the effectiveness of the method we have also compared two of the detected shadow images (pixel-by-pixel) with those manually identified by an operator (who has drawn the contour with the mouse). By taking these images as “ground truth” we have computed the rates of failures and successes shown in table 1.

| | Figure 9.b | Figure 10.b |
|------------------------|------------|-------------|
| False positives | 4% | 1% |
| True negatives | 23% | 7% |
| True positives | 77% | 93% |

Table 1: Percentage of false positives, true negatives and true positives of two images of size 2076×1024 pixels from where the pieces of figure 9 and figure 10 have been cut out.

6. CONCLUSIONS AND FUTURE WORK

In this paper we have presented a procedure for automatically detecting shadows in QuickBird satellite images. The proposed approach is based on other works reported in the computer vision that where developed and tested for quite different images. In particular, we have exploited the sensitivity to shadowed pixels of one of the components (c_3) of a colour space named $c_1c_2c_3$. To overcome the limitations that the c_3 band presents, we have developed a region-growing procedure that takes also into account some components of other colour spaces (S , V and B).

Although our approach needs the manual selection of certain parameters (thresholds) they do not need to be tuned very precisely to end up with satisfactory results, though this is something that we would like to avoid in future realizations. Experimentally, we have verified the reliability and effectiveness of the method for a wide variety of image conditions, though the rate of true negatives (non-detected true shadows) is still high for some applications as urban change detection.

As future work, we would like to exploit shadow areas for both restoring the radiometric information and for inferring 3D information from them (e.g. height of buildings).

REFERENCES

- Etemadnia, H. and Reza-Alsharif, M. (2003). Automatic image shadow identification using LPF in homomorphic processing system. *In Proc. VII Digital Image Computing: Techniques and Applications*, pp. 429-438, Sydney.
- Ford, A. and Roberts, A. (1998). Colour space conversions. *Technical report*, Westminster University, London.

Funka-Lea, G. and Bajscy, R. (1995). Combining colour and geometry for the active, visual recognition of shadows. *In IEEE International Conference on Computer Vision*, pp. 203-209.

Gevers, T. and Smeulders, A.W.M. (1999). Colour-based object recognition, *Pattern Recognition*, 32, pp. 453-464.

Irvin B. and McKeown J.R. (1989). Methods for exploiting the relationship between buildings and their shadows in aerial imagery. *IEEE Transactions on System, Man and Cybernetics*, 19(6), pp. 1564-1575.

Jiang, C. and Ward, M.O. (1994). Shadow segmentation and classification in a constrained environment. *In CVGIP: Image Understanding*, 59(2), pp. 213-225.

Koller, D., Daniilidis, K., and Nagel, H. (1993). Model-based object tracking in monocular image sequences of road traffic scenes. *International Journal of Computer Vision*, 10(3), pp. 257-281.

Massalabi, A., He, D.C., Bénié, G.B. and Beaudry, É. (2004). Restitution of information under shadow in remote sensing high space resolution images: Application to IKONOS data of Sherbrooke city. *In Proc. XX ISPRS Congress*, Istanbul.

Salvador, E., Cavallaro, A. and Ebrahimi, T. (2001). Shadow identification and classification using invariant colour models. *In IEEE International Conference on Acoustic, Speech, and Signal Processing*, 3, pp. 1545-1548, Salt Lake City, Utah.

Sonka, M., Hlavac, V., and Boyle, R. (1996). *Image Processing, Analysis and Machine Vision*, International Thomson Computer Press, London.

Stevens, M.R., Pyeatt, L.D., Houlton, D.J. and Goss, M. (1995). Locating shadows in aerial photographs using imprecise elevation data. *Computer Science Technical Report CS-95-105*, Colorado State University, USA.

ACKNOWLEDGMENTS

The ©DigitalGlobe QuickBird imagery used in this study is distributed by Eurimage, SpA. (www.eurimage.com) and provided by Decasat S.L., Málaga, Spain. (www.decasat.com).



Figure 9: QuickBird image from a high-textured residential area which contains regular buildings and forest. The number of shadow seeds identified in this image is 189.




Figure 10: QuickBird image from a high-textured urban area which contains cast shadows, intensity saturations and dark regions. The number of shadow seeds identified in this image is 92.