Shadow Detection based on Colour Segmentation and Estimated Illumination

Xiaoyue Jiang[1]
x.y.jiang@bham.ac.uk

Andrew J. Schofield[1]
a.j.schofield@bham.ac.uk

Jeremy L. Wyatt[2]
jlw@cs.bham.ac.uk

[1] School of Psychology, University of Birmingham
Birmingham, UK

[2] School of Computer Science, University of Birmingham
Birmingham, UK

Shadows are unavoidable in natural images. They are useful because they can provide information about both the light source and the shape of objects but they also degrade the performance of algorithms for tasks such as segmentation, object detection, tracking, and shape reconstruction. In all of these cases it is beneficial to detect shadows in images. In this paper we improve on existing supervised learning methods for shadow detection by incorporating an illumination map – extracted using a correlation-based approach we developed recently [1] – alongside colour features. These are given as input to a decision tree trained using AdaBoost. Our method outperforms the approach of Lalonde *et al*. [2] for both shadow edge detection and ground shadow edge detection on a database of images from Zhu *et al*. [3].

The illumination map of an image provides information not available to most of the shadow detection methods. It describes all the changes in illumination within a scene, including shading caused by changes in surface orientation, and incident illumination changes such as attached and cast shadows. However, extracting illumination maps from images is itself a challenging task, and thus so far such maps have not been used widely in shadow detection. Jiang *et al*. [1] proposed a steerable filter framework to derive illumination and reflectance maps of a scene from a single image based on the relationship between luminance and texture, colour, and local contrast. We extend this framework for cast shadow detection in two ways. First, local correlation strategy is proposed to improve Jiang *et al*. [1] for illumination map extraction in high-frequency components. For low frequency bands, the global correlation is still applied. Second, we extend the method of Lalonde *et al*. for shadow detection by incorporating the features extracted from the illumination map as additional inputs to a decision tree that has been trained using supervised learning.

The structure of our approach is captured in Fig. 1(a). The upper path corresponds broadly to the approach of Lalonde *et al*. without texture features. When trained their decision tree (Decision Tree 1) produces an initial shadow edge map which is an input for illumination map estimation as shown in the lower path in Fig. 1(a). The details of the improved illumination map estimation process are shown in Fig. 1(b). Jiang *et al*. look for global correlations between AM (*Amplitude Modulation* or local luminance contrast) and LM (*Luminance Modulation* or average local luminance) to spot the difference between reflectance changes (where AM and LM are uncorrelated) and illumination changes (where AM and LM are correlated). Their global strategy works well at low frequencies, but not for high frequency changes. We therefore calculate local correlations between LM and AM, and separately between LM and colour features. Both these help us to separate high-frequency filter responses, at a local level, for later reconstruction into either illumination or reflectance maps. For the lower-frequency bands, the global correlation between luminance and contrast in each band still determines the corresponding weights for reconstruction [1].

Having improved the separation of intrinsic images in the lower path, we then extract features from the illumination map. Using these illumination ratio features and the colour features extracted in the upper path, a second decision tree (also trained using AdaBoost) produces an improved estimate of the shadow edge probability map. This method discriminates more effectively between shadow- and reflectance-edges. This revised shadow edge probability map can then be improved further using a Conditional Random Field model with constraints that edges should be continuous and that neighbouring pixels should have the same state.

The ROC curves (Fig. 1(c)) shows the improved performance of the proposed algorithm. From resulting images (Fig. 1(c)), we can see our algorithm offers improved sensitivity for shadow detection and reduced false alarm rates.

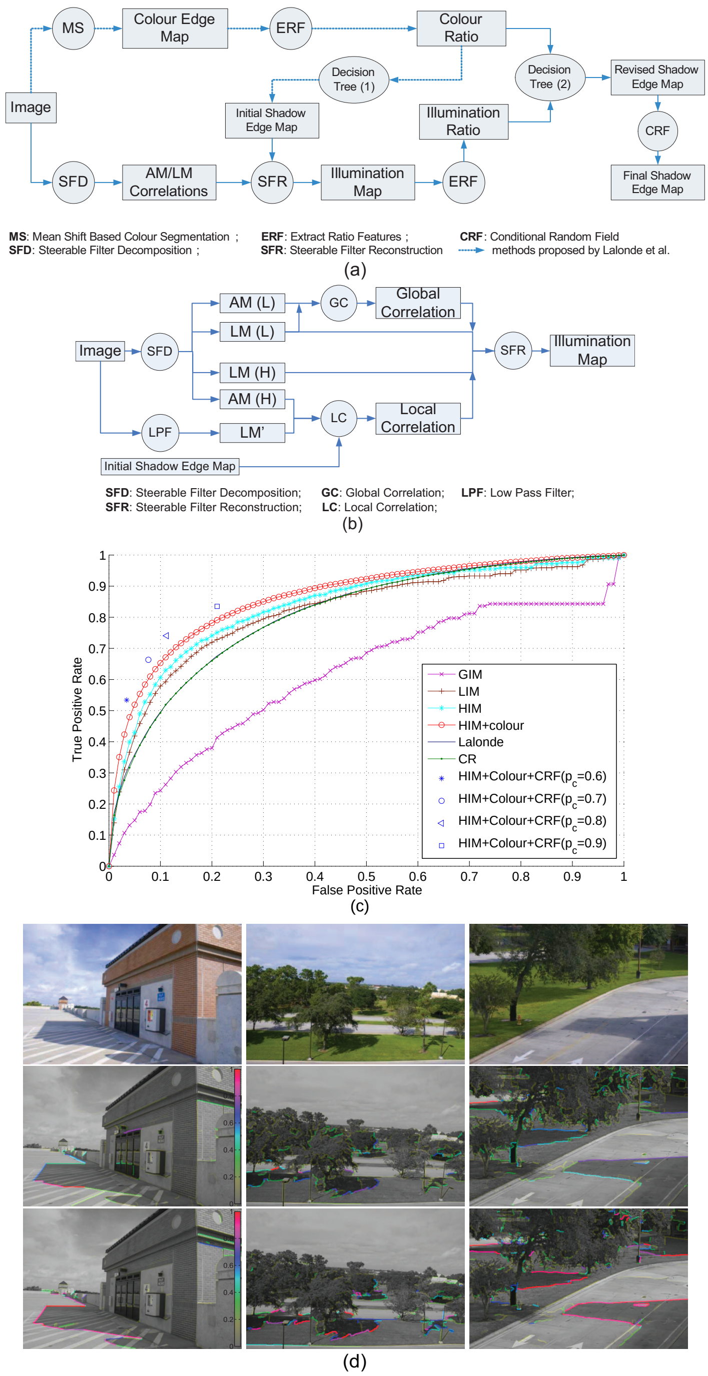

Figure 1: (a) A flow chart of our proposed algorithm. (b) the flow chart of estimating illumination map, which corresponds to the first two stages of the bottom row in(a). (c) ROC curves of shadow edge detection for the dataset from Zhu *et al*. [3]. (d) shadow detection results: for each column, top, original image; middle, shadow detection results based on colour features; bottom, results of proposed algorithm.

[1] X. Jiang, A. J. Schofield, and J. L. Wyatt. Correlation-based intrinsic image extraction from a single image. *ECCV*, 4:58–71, 2010.

[2] J.-F. Lalonde, A. A. Efros, and S. G. Narasimhan. Detecting ground shadow in outdoor consumer photographs. *ECCV*, 2:322–335, 2010.

[3] J. Zhu, K. G. G. Samuel, S. Z. Masood, and M. F. Tappen. Learning to recognize shadow in monochromatic natural images. *CVPR*, 2010.