Pattern Recognition Letters 33 (2012) 615–622

Contents lists available at SciVerse ScienceDirect

Pattern Recognition Letters

journal homepage: www.elsevier.com/locate/patrec

Image segmentation using local spectral histograms and linear regression

Jiangye Yuan [a,*], DeLiang Wang [b], Rongxing Li [a]

[a] Mapping and GIS Laboratory, Department of Civil and Environmental Engineering and Geodetic Science, Columbus, OH 43210, United States
[b] Department of Computer Science and Engineering, Center for Cognitive Science, The Ohio State University, Columbus, OH 43210, United States

ARTICLE INFO

Article history:
Received 2 June 2011
Available online 17 December 2011
Communicated by S. Wang

Keywords:
Texture segmentation
Spectral histogram
Linear regression

ABSTRACT

We present a novel method for segmenting images with texture and nontexture regions. Local spectral histograms are feature vectors consisting of histograms of chosen filter responses, which capture both texture and nontexture information. Based on the observation that the local spectral histogram of a pixel location can be approximated through a linear combination of the representative features weighted by the area coverage of each feature, we formulate the segmentation problem as a multivariate linear regression, where the solution is obtained by least squares estimation. Moreover, we propose an algorithm to automatically identify representative features corresponding to different homogeneous regions, and show that the number of representative features can be determined by examining the effective rank of a feature matrix. We present segmentation results on different types of images, and our comparison with other methods shows that the proposed method gives more accurate results.

1. Introduction

The goal of image segmentation is to partition an image into a number of regions so that each region should be as homogeneous as possible and neighboring ones should be as different as possible. It is a critical task for a wide range of image analysis problems. Although substantial progress has been made in this area (Shi and Malik, 2000; Chan and Vese, 2001; Comaniciu and Meer, 2002; Felzenszwalb and Huttenlocher, 2004; Sharon et al., 2006; Li et al., 2008; Wang et al., 2009), image segmentation remains an unsolved problem in computer vision.

An image segmentation method should work on different kinds of imagery including texture and nontexture images. It is widely recognized that a visual texture is very difficult to characterize. A large number of methods have been proposed to deal with texture images. These methods address two issues: underlying texture model that defines region homogeneity and a framework or strategy for producing segmentation (Paragios and Deriche, 2002; Deng and Clausi, 2004; Kim and Kang, 2007; Lehmann, 2011). In general the two issues are not treated independently. A successful segmentation methodology should couple an accurate texture model and an effective segmentation strategy.

Numerous texture models have been proposed. The most successful ones focus on the following two aspects. One is on filtering, which typically uses filterbanks to decompose an image into a set of sub-bands. Filtering methods have received experimental supports on human texture perception and have shown impressive performance for texture segmentation and classification (Jain and Farrokhsia, 1991; Dunn and Higgins, 1991; Randen and Hakon, 1999; Clausi and Jernigan, 2000). The other is on statistical modeling, which characterizes texture regions as resulting from some underlying stochastic processes. For example, autoregressive models (Mao and Jain, 1992) and Markov random fields (Cross and Jain, 1983; Cesmeli and Wang, 2001; Benboudjema and Pieczynski, 2007) emphasize global appearance and are robust to noise.

Building on the above themes, a local spectral histogram has been proposed as a texture model, which consists of marginal distributions of chosen filter responses in an image window (Liu and Wang, 2002). Local spectral histograms provide a generic statistic model for texture as well as nontexture regions. Using local spectral histograms as features, the segmentation problem can be approached by measuring the distance among the features (Liu and Wang, 2006). However, since the local spectral histograms computed over the windows straddling boundaries do not give distinctive features, such methods have difficulty in accurately localizing region boundaries. As Malik et al. (2001) point out, this problem widely exists in the approaches based on measuring texture descriptors over local windows. To address this problem, quadrant filters or similar strategies are often employed, which compute features from shifted local windows around a pixel and make a best choice by examining certain criteria (Kim and Kang, 2007; Liu and Wang, 2006). Another popular technique is to use local windows of different sizes, also referred to as scales

* Corresponding author. Tel.: +1 614 214 0686.
E-mail address: yuan.75@osu.edu (J. Yuan).

0167-8655/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.patrec.2011.12.003

(Liang and Tjahjadi, 2006; Martin et al., 2004). Boundaries are determined by analyzing information across different scales. Although reasonable results are achieved, these methods inevitably involve more computation and hand chosen parameters.

In this paper, we propose a new segmentation method that leverages local spectral histograms to discriminate region appearances and at the same time accurately localizes boundaries. Using local spectral histograms as features, we regard a pixel location as a linear combination of representative features, which encodes a natural criterion to identify boundaries. The relationship between features of all the pixels and representative features is modeled by a multivariate linear regression, and the segmentation problem can be directly solved by least square estimation. The rest of the paper is organized as follows. The local spectral histogram representation is introduced in Section 2. Section 3 presents our segmentation algorithm in detail. In Section 4, we show experimental results and comparisons. Section 5 concludes the paper.

2. Local spectral histograms

Motivated by perceptual observations, the spectral histogram model has been proposed to characterize texture appearance (Liu and Wang, 2002). For a window $\boldsymbol{W}$ in an input image, a set of filter responses is computed through convolving with a chosen bank of filters $\{F^{(\alpha)}, \alpha = 1, 2, \ldots, K\}$. For a sub-band image $\boldsymbol{W}^{(\alpha)}$, a bin of its histogram can be written as

$$H_W^{(\alpha)}(z_1, z_2) = \sum_{\vec{v} \in \boldsymbol{W}} \int_{z_1}^{z_2} \delta(z - \boldsymbol{W}^{(\alpha)}(\vec{v}))dz. \quad (1)$$

Here z_1 and z_2 specify the range of the bin. $\vec{v}$ represents a pixel location, and δ denotes the Dirac delta function. In this paper, we use 11 equal-width bins for each filter response. The spectral histogram with respect to the chosen filters is then defined as (Liu and Wang, 2002)

$$H_W = \frac{1}{|\boldsymbol{W}|}\left(H_W^{(1)}, H_W^{(2)}, \ldots, H_W^{(K)}\right) \quad (2)$$

where $|\,|$ denotes cardinality. The spectral histogram is a normalized feature statistic, which can compare image windows of different sizes. For each pixel location, the local spectral histogram is computed over the window centered at the pixel. The size of the window is called integration scale. When the filters are selected properly, the spectral histogram is sufficient to capture texture appearance (Liu and Wang, 2002).

In this paper, we use seven filters: the intensity filter, two LoG (Laplacian of Gaussian) filters with the scale values of 0.2 and 1.0, and four Gabor filter with the orientations of 0°, 45°, 90°, and 135° and the scale value of 3. The parameters of filters are not adjusted for individual images, but for a type of images a set of filters is chosen based on general image attributes.

In order to extract meaningful texture features, the integration scale is set to be relatively large, which makes computing local spectral histograms computationally expensive. A fast implementation method is therefore introduced in (Liu and Wang, 2002). For an input image, an integral histogram image is defined as follows: at location (x,y) the integral histogram is calculated using the pixel values above and to the left of (x,y) (see Fig. 1(a)). The integral histogram image can be efficiently computed in one pass over the image. Given the integral histogram image, the histograms of arbitrary rectangular regions can be obtained with four references. As illustrated in Fig. 1(b), we can compute the histogram of region R using the following four references: $L_4 + L_1 - L_2 - L_3$. Hence, once the integral histogram image is computed, we only need three vector arithmetic operations to obtain any local spectral histogram regardless of window size. A detailed description of the fast implementation can be found in (Liu and Wang, 2002).

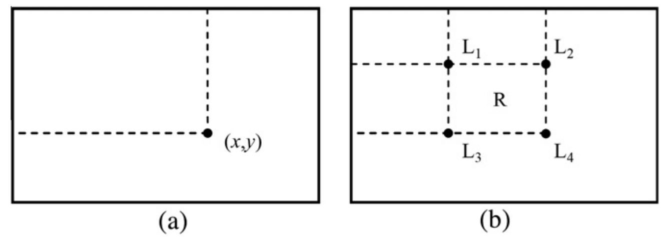

Fig. 1. Illustration of fast implementation for computing local spectral histograms. (a) The integral histogram value at location (x,y) is the histogram of the image window above and to the left of (x,y). (b) The histogram of region R can be computed using four references: $L_4 + L_1 - L_2 - L_3$.

3. Segmentation algorithm

As discussed in Section 1, although local spectral histograms provide an effective feature, segmentation methods using feature distance to measure region homogeneity tend to produce inaccurate boundaries caused by features extracted in image windows that cross multiple regions. Here, we describe a new segmentation algorithm based on linear regression, which can produce segmentation with high accuracy and great efficiency.

3.1. Segmentation using linear regression

Fig. 2(a) illustrates the difficulty of extracting features over a large window. Here, an image contains five homogenous regions. Given a pixel location A, the corresponding local spectral histogram is computed using a square window. Since this window straddles two different regions, the extracted feature is not discriminative. As a result, it is difficult to correctly classify the corresponding pixel by measuring feature similarity.

Let us define a window $\boldsymbol{W}$ consisting of disjoint connected subregions $\{\boldsymbol{W}_1, \boldsymbol{W}_2, \ldots, \boldsymbol{W}_s\}$, and $H_{W_i}^{(\alpha)}$ is the histogram computed from region $\boldsymbol{W}_i$ and filter α. Since $H_W^{(\alpha)} = H_{W_1}^{(\alpha)} + H_{W_2}^{(\alpha)} + \cdots + H_{W_s}^{(\alpha)}$, we can rewrite the spectral histogram H_W as

$$\begin{aligned} H_W &= \frac{1}{|\boldsymbol{W}|}\left(H_W^{(1)}, H_W^{(2)}, \ldots, H_W^{(K)}\right) \\ &= \frac{1}{|\boldsymbol{W}|}\left(\sum_{i=1}^{s} H_{W_i}^{(1)}, \sum_{i=1}^{s} H_{W_i}^{(2)}, \ldots, \sum_{i=1}^{s} H_{W_i}^{(K)}\right) \\ &= \frac{|\boldsymbol{W}_1|}{|\boldsymbol{W}|}\left(\frac{1}{|\boldsymbol{W}_1|}\left(H_{W_1}^{(1)}, H_{W_1}^{(2)}, \ldots, H_{W_1}^{(K)}\right)\right) \\ &\quad + \frac{|\boldsymbol{W}_2|}{|\boldsymbol{W}|}\left(\frac{1}{|\boldsymbol{W}_2|}\left(H_{W_2}^{(1)}, H_{W_2}^{(2)}, \ldots, H_{W_2}^{(K)}\right)\right) + \cdots \\ &\quad + \frac{|\boldsymbol{W}_s|}{|\boldsymbol{W}|}\left(\frac{1}{|\boldsymbol{W}_s|}\left(H_{W_s}^{(1)}, H_{W_s}^{(2)}, \ldots, H_{W_s}^{(K)}\right)\right). \end{aligned}$$

With the definition in (2), we have

$$H_W = \frac{|\boldsymbol{W}_1|}{|\boldsymbol{W}|}H_{W_1} + \frac{|\boldsymbol{W}_2|}{|\boldsymbol{W}|}H_{W_2} + \cdots + \frac{|\boldsymbol{W}_s|}{|\boldsymbol{W}|}H_{W_s}. \quad (3)$$

Therefore, a spectral histogram of an image window can be linearly decomposed into spectral histograms of its subregions, where weights are proportional to region areas.

Because spectral histograms can characterize image appearance, we assume that spectral histograms within a homogeneous region are approximately constant. This assumption along with Eq. (3) implies that the spectral histogram of a local window can be approximated by a weighted sum of the spectral histograms