Extraction and characterization of regions of interest in biomedical images

Yakov Keselman
Department of Computer Science
Rutgers University
Piscataway, NJ 08855, USA
keselman@paul.rutgers.edu

Evangelia Micheli-Tzanakou
Department of Biomedical Engineering
Rutgers University
Piscataway, NJ 08855, USA
etzanako@biomed.rutgers.edu

Abstract

A system that could serve as a first stage of a two–stage automated biomedical image classification system is presented. The underlying image segmentation algorithm and several region representations are discussed. Provided examples of extracted (robustly and quickly) regions show flexibility and potential of the approach.

1 Introduction

The problem of automated classification of various biomedical images obtained as a result of screening (e.g. X–ray or mammography) can be approached in two stages. During the first stage "regions of interest", i.e. regions whose presence in the image to a large degree affects how the image is later classified, are extracted. During the second stage the extracted regions are analyzed in more detail in terms of their shape, internal properties (e.g. intensity, texture), as well as locations with respect to the image and to other such regions. It is believed that provided both the stages are performed sufficiently accurately, the resulting system could be of great use in many different biomedical applications, including cell identification, mammography, etc.

The previous efforts in the field mostly concentrated on the second stage where various low–level image processing tools (including wavelet and Fourier analysis, neural networks, pattern matching, etc. – see e.g. [6, 7, 10]) were successfully applied. The first stage, related to problems occuring in low–level image understanding, also received some attention, e.g. in [3, 2, 9]. In general such region extracting systems tend to be less automated because the notion of "region of interest" is in many cases problem–dependent and a successful extraction algorithm should incorporate some field–specific knowledge directly. At present such systems mostly offer some means of interaction to improve the often far from perfect initial segmentation of images.

The approach taken by the author of the system to be described also falls into this category. Namely, after an initial segmentation of an image, the user is offered a set of operations on the resulting regions by means of which it is possible to obtain a desirable segmentation in a few steps. Motivations for various components of the system and their descriptions are given below.

2 Image Segmentation

The system is capable of working with both gray level (e.g. mammography) and color (e.g. cellular and retinal) images. The ability results from the features of the underlying simple, robust, and diverse *mean shift* clustering algorithm [1](first applied to image segmentation by D. Comaniciu and P. Meer [4]). For a detailed account of the algorithm as well as its properties the reader is referred to the mentioned paper.

The *mean shift* algorithm itself is designed to find *Modes* (or centers of the regions of high concentration) of data represented as arbitrary–dimensional vectors. The algorithm proceeds as follows ([4]):

1. Choose the radius of the search window.
2. Choose the initial location (center) of the window.
3. Compute the mean (average) of the data points over the window and translate the center of the window into this point.
4. Repeat step above 3 until translation distance of the center becomes less than a preset threshold.

The algorithm could then be applied to gray level/color segmentation in the following fashion (simplified version of [4]):

1. Map the image domain into the feature space (1–dimensional for gray level images and generally 3–dimensional for color images).
2. Define an adequate number of search windows at random locations in the feature space.
3. Find centers of high density regions by applying the *mean shift* algorithm to each window.
4. Find regions in the image domain corresponding to high density regions in the feature space.
5. Do some postprocessing on the image domain regions based on constraints and prior assumptions.

Since high density regions in the feature space correspond to sufficiently big numbers of pixels in a narrow range of intensities/colors in the image domain, provided the pixels form connected regions (as is often the case for relatively smooth images), the algorithm essentially finds relatively big connected regions which have sufficiently small variations in intensity/color and thus perceived as well–defined regions by humans as well. In practice, the algorithm proceeds by placing randomly one search window at a time, finding the corresponding mode, and removing all the feature vectors in the final window from the feature space. Thus one would expect to find bigger regions first. Given some prior information on the image (e.g. lower intensities correspond to background as it is often the case in mammograms), the strategy could be modified by placing the first window at a preferred position. Provided that additional information about the regions of interest is known (e.g. the regions are characterized by colors rather than their intensities), it could easily be used by the system by applying the same algorithm in a lower–dimensional space.

One of the shortcomings of the segmentation algorithm is that it does not take into account spatial information at the processing step. This could result in regions with many holes and irregular boundaries due to relatively large variation over a small range of pixels (see e.g. Figure 1). However depending on prior information about the image (e.g. that the regions of interest should not have such holes as it is the case for nuclei of cells), such artifacts could be eliminated at the postprocessing stage. The postprocessing stage usually includes such simple morphological operations as removing isolated pixels and small (as defined by the user) regions and smoothing boundaries of regions.

As implemented in the system, after inputting a few parameters, the user obtains an initial segmentation of the image. If the resulting regions are "too small" (as it it the case for the two regions in Figure 3) or "too big" an adjustment of the parameters and resegmentation is required. At present the implementation uses the same window size over the whole spectrum (of values of each color space coordinates), but specification of different window sizes for different parts of spectra could easily be added. In many instances the incoming images are quite far from being smooth (due to low resolution of scanning devices) and a few simple smoothing algorithms (some being plain image filters, and some are wavelet–based) are optionally applied to achieve better segmentation results. In addition, to extract relatively small regions of interest (and to reduce segmentation time) the provided image cropping feature could be used. As another implementation detail, before segmenting color images, pixels, usually represented in the RGB color space, are mapped into the $L^*u^*v^*$ color space (see e.g. [5]) which has a "brightness" component represented by L^* and two "chromatic" components represented by u^* and v^*. It is argued that the latter color space is more isotropic and thus is better suitable for the used Mode finding algorithm. Finally, for efficiency reasons "rectangular" rather than "circular" windows are used. The difference between the two is non–existent for 1–dimensional case and is sufficiently small even for 3–dimensional one.

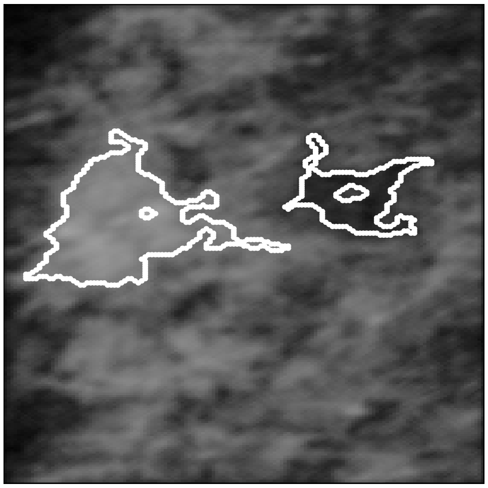

Figure 1: Two possible regions of interest in a mammogram (partial view). The region on the right does not necessarily correspond to a "suspicious" region.

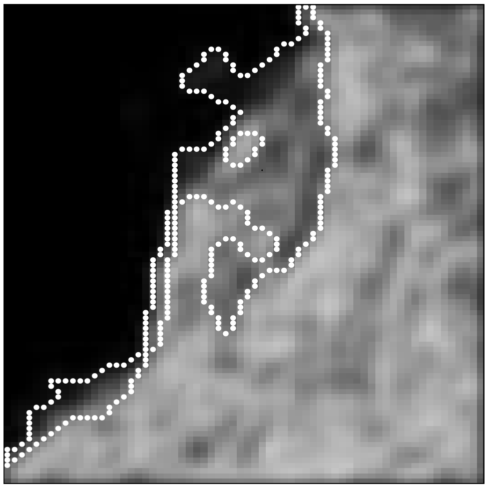

Figure 2: A region of interest in a retinal image.

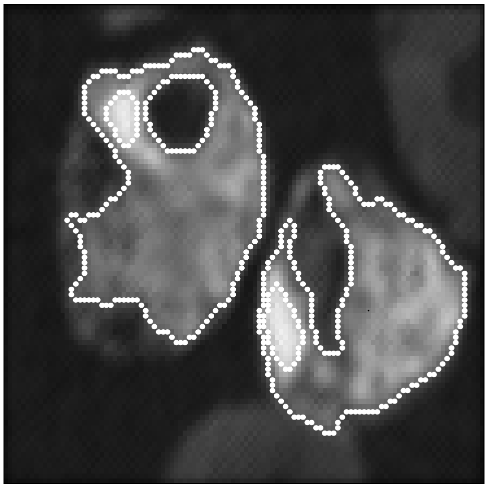

Figure 3: Two extracted contours of cells with not quite satisfactory results. See text for details.

3 Regions representations and processing

Provided the initial segmentation is satisfactory, the system allows to work with regions resulting from segmentation. At the moment, regions are assumed to be connected (taking the view that two different regions of about the same intensity/color may be present in the image independently of each other), although the operation of merging allows to form regions consisting of connected subcomponents if desired. In general it is hard to make such decisions at the system level due to the absence of any prior knowledge about the regions nature. If the resulting region encompasses too big a portion of the image and a finer segmentation does not provide satisfactory results (as it might be the case for the region in Figure 2), it is possible to use the included spline–based "Cutting" and "Merging" features to select the region of interest manually.

Each region possesses several (computed on request) representations. For the purpose of rough region shape matching (which will likely precede more fine region matching at the second stage) they are represented by singular values of the matrix formed by placing ones at the positions of pixels belonging to the region and zeros everywhere else (the size of the matrix is defined by the extreme values of X and Y coordinates of pixels of the region and not of the image itself). The advantage of such a representation lies in its (near) invariance to rotations by multiples of 90° and scaling of the region. Indeed, assume that the region R_B was obtained from region R_A by replicating each of its pixels over a rectangle of size m by n. Since the rank of matrix B representing region R_B is the same as that of matrix A representing region R_A, the number of non–zero eigenvalues will be the same for both matrices. Also if $\lambda_i(B)$ is the i-th largest eigenvalue of B and $\lambda_i(A)$ is the i-th largest eigenvalue of A, then $\lambda_i(B) = m \cdot n \cdot \lambda_i(A)$. Since the eigenvalues of a matrix, of its reflection about the central row, and of its transpose coincide, the invariance to rotations by the above angles follows. To accomodate rotations by other angles at the matching step the region is rotated with a fixed step size (e.g. 5° or 10°) and the corresponding eigenvalues are computed. Then the best match is taken. The performed experiments confirmed the applicability of the representation for such rough matching.

Alternatively, regions can be represented by singular values of matrices formed by placing ones only at the boundary pixels of regions. There are advantages and disadvantages to both representations. For example, cutting out a hole in a region will most likely not lead to big changes in the latter representation, provided the boundary of the hole in not very long compared to the boundary of the region (as in Figure 3 in the region on the left), but will likely lead to noticeable changes in the former representation making it more suitable for a more precise shape matching.

An even finer matching could be done by placing not ones, but actual pixel values in the corresponding matrix (or matrices in case of color images) and computing its singular values. An advantage of such representations based on singular values over ones based on representing boundaries as parameterized functions is that it allows to treat regions even in the case when their boundary is not connected (i.e. the region consists of several pieces or has holes). Also small local changes in boundaries (e.g. protrusions and intrusions) will likely not lead to any significant changes in eigenvalues, thus offering greater stability under small perturbations. A more fine eigenvalues–based approach is described in [8].

Another representation of regions, mostly suitable for their visualization, is a 3–dimensional intensity–based one, where the regions are represented by graphs of functions defined on the regions pixels and whose values at the pixels are taken to be intensities of the pixels. A few simple attributes (e.g. area on which the function is defined and volume of the resulting "body") are computed to easily distinguish non–similar regions.

4 Working experience

The system provides a nice graphical user interface to all its functionality and could be easily extended in several directions, including using different image segmentation algorithms, color spaces, and region representations. It processes images in a variety of formats (including "bmp", "jpeg", "tiff"). To increase segmentation speed of color images only one or two most significant intensity/color components (as determined by singular value decomposition) could be used. From the working experience, the system has shown to be an easy to use and valuable tool in several applications, which can greatly reduce the effort spent on extraction and classification of biomedical images and which could serve as a first stage of very helpful diagnostic systems.

5 Directions for future work

At present all the advantages and disadvantages of initial image partitioning arise from the segmentation algorithm used. It will be desirable to incorporate some (prior) shape/intensity/color/texture information in its work, which should lead to a more automated system. Although the latter 3 characteristics are somewhat easier to use (they also could be "learned" through experience with images from some particular domain), the shape information will likely require a very different segmentation algorithm, probably using the approaches from [2, 11]. The matching component of the system is also quite rudimentary but can be linked to existing finer matching algorithms if such need arises.

References

[1] Y.Z. Cheng. Mean shift, mode seeking, and clustering. *PAMI*, 17(8):790–799, August 1995.

[2] L. Cohen and I. Cohen. Finite element methods for active contour models and balloons for 2-d and 3-d images. *IEEE Trans. Pattern Analysis and Machine Intelligence*, PAMI-15:1133–1147, 1993.

[3] D. Comaniciu, D. Foran, and P. Meer. Shape-based image indexing and retrieval for diagnostic pathology. In *2nd International Workshop on Statistical Techniques in Pattern Recognition*, 1998.

[4] D. Comaniciu and P. Meer. Robust analysis of feature spaces: Color image segmentation. In *IEEE Conf. Comp. Vision Patt. Recogn.*, pages 750–755, 1997.

[5] G. Gunter and W. S. Stiles. *Color Science – Concepts and Methods, Quantitative Data and Formulae*, chapter 13. John Wiley & Sons, Inc., 2 edition, 1982.

[6] Stephane G. Mallat. A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Trans. PAMI*, 11(7):674–693, July 1989.

[7] E. Micheli-Tzanakou, E. Uyeda, R. Sharma, R. Ramanujan, and J. Doug. Comparison of neural network algorithms for face recognition. *Simulation*, 64(1):15–27, July 1995.

[8] S. Sclaroff and A. Pentland. Modal matching for correspondence and recognition. *PAMI*, 17(6):545–561, June 1995.

[9] J. Shi and J. Malik. Normalized cuts and image segmentation. In *IEEE Conf. Comp. Vision Patt. Recogn.*, pages 731–737, 1997.

[10] Y. Wu, K. Doi, M. Giger, and R. Nishikawa. Computerized detection of clustered microcalcifications in digital mammograms: Applications of artificial neural networks. *Medical Physics*, 19(3):555–560, May/June 1992.

[11] S.C. Zhu, T.S. Lee, and A.L. Yuille. Region competition: Unifying snakes, region growing, energy/Bayes/MDL for mutil-band image segmentation. In *Int. Conf. Comp. Vision*, pages 416–423, 1995.