

© Springer Verlag. The copyright for this contribution is held by Springer Verlag. The original publication is available at www.springerlink.com.

Classification of Endoscopic Images Using Delaunay Triangulation-based Edge Features

M. Häfner¹, A. Gangl², M. Liedlgruber³, A. Uhl³, A. Vécsei⁴, and F. Wrba⁵

¹ Department for Internal Medicine, St. Elisabeth Hospital, Vienna

² Department of Gastroenterology and Hepatology, Medical University of Vienna, Austria

³ Department of Computer Sciences, Salzburg University, Austria

⁴ St. Anna Children's Hospital, Vienna, Austria

⁵ Department of Clinical Pathology, Medical University of Vienna, Austria

{mliedl, uhl}@cosy.sbg.ac.at

Abstract. In this work we present a method for an automated classification of endoscopic images according to the pit pattern classification scheme. Images taken during colonoscopy are transformed using an extended and rotation invariant version of the Local Binary Patterns operator (LBP). The result of the transforms is then used to extract polygons from the images. Based on these polygons we compute the regularity of the polygon positions by using the Delaunay triangulation and constructing histograms from the edge lengths of the Delaunay triangles. Using these histograms, the classification is carried out by employing the k-nearest-neighbors (k-NN) classifier in conjunction with the histogram intersection distance metric.

While, compared to previously published results, the performance of the proposed approach is lower, the results achieved are yet promising and show that a pit pattern classification is feasible by using the proposed system.

1 Introduction

Today, the third most common malignant disease in western countries is colon cancer. Therefore a regular colon examination is recommended, especially for people at an age of 50 years and older. Currently the gold standard for colon examination is colonoscopy, which is performed by using a colonoscope. Modern colonoscopes are able to take pictures from inside the colon which allows to obtain images for a computer-assisted analysis with the goal of detecting tumorous lesions. To get highly detailed images a magnifying endoscope is used [1]. Such an endoscope represents a significant advance in colonoscopy as it provides images which are up to 150-fold magnified, thus uncovering the fine surface structure of the mucosa as well as small lesions.

In Sect. 2 we review the classification of pit patterns of the colonic mucosa. Section 3 describes the feature extraction process, including image transformation using a LBP extension, polygon extraction, Delaunay-based feature computation, histogram creation, and the classification. Experimental results and configuration details of the classification system proposed are given in Sect. 4. Section 5 concludes the paper.

2 Pit Pattern Classification

Polyps of the colon are a frequent finding and are usually divided into metaplastic, adenomatous, and malignant. As resection of all polyps is time-consuming, it is imperative that those polyps which warrant endoscopic resection can be distinguished: polypectomy of metaplastic lesions is unnecessary and removal of invasive cancer may be hazardous. For these reasons, assessing the malignant potential of lesions at the time of colonoscopy is important.

The most commonly used classification system to distinguish between non-neoplastic and neoplastic lesions in the colon is the pit pattern classification, originally reported by Kudo et al. [2]. This system allows to differentiate between normal mucosa, hyperplastic lesions (non-neoplastic), adenomas (a pre-malignant condition), and malignant cancer based on the visual pattern of the mucosal surface. Thus this classification scheme is a convenient tool to decide which lesions need not, which should, and which most likely can not be removed endoscopically. The mucosal pattern as seen after dye staining and by using magnification endoscopy shows a high agreement with the histopathologic diagnosis. Due to the visual nature of this classification it is also a convenient choice for an automated image classification.

As illustrated in Fig. 1(a)-(f) in this classification scheme exist five main types according to the mucosal surface of the colon. Type III is divided into types III-S and III-L, designating the size of the pit structure. It has been suggested that type I and II pattern are characteristic of non-neoplastic lesions (benign and non-tumorous), type III and IV are found on adenomatous polyps, and type V are strongly suggestive of invasive carcinoma, thus highly indicative for cancer.

Furthermore lesions of type I and II can be grouped into non-neoplastic lesions and types III to V can be grouped into neoplastic lesions. This allows a grouping of lesions into two classes, which is more relevant in clinical practice as indicated in a study by Kato et al. [3].

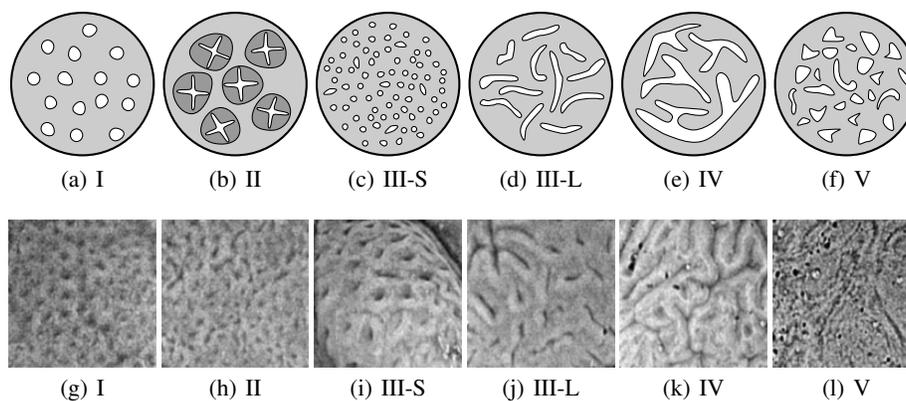


Fig. 1. Pit pattern classification according to Kudo et al. (a)-(f) Schematically and (g)-(l) example images for the respective classes taken from the available image database

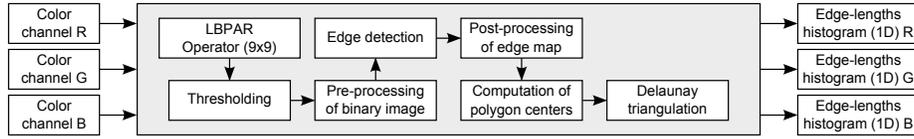


Fig. 2. This figure illustrates the different parts of the feature extraction process. The parts within the gray box are carried out for each color channel under consideration separately.

Using a magnifying colonoscope together with indigo carmine dye spraying, the mucosal crypt pattern on the surface of colonic lesions can be observed [4]. Several studies found a good correlation between the mucosal pit pattern and the histological findings, where especially techniques using magnifying colonoscopes led to excellent results [3].

From Fig. 1 we notice that pit pattern types I to IV can be characterized fairly well, whereas type V is a composition of unstructured pits. At a first glance this classification scheme seems to be straightforward and easy to be applied. But it needs some experience and exercising to achieve fairly good results [5].

As evident from Fig. 1(g)-(l), pit pattern types I and II are regular to some extent and the pits are distributed more tightly. Types III to V in contrast are more irregular in terms of the pit distribution, showing a lower pit density or even a complete absence of pits. These observations are the basis for the method presented in the following sections.

3 Proposed Approach

In the past we have already shown that an automated classification of endoscopic images based on the pit pattern scheme is feasible. But in our previous work we mainly focused on general purpose features describing texture properties (e.g. [6–8]), dealing with the two-classes case as well as with the six-classes case. By contrast, the method proposed in this work aims at distinguishing between non-neoplastic and neoplastic images only. It is furthermore based on high level features obtained by measuring the density of pits visible within the images. This is inspired by the fact that the pit distributions in non-neoplastic images are more dense than in case of the neoplastic ones, as already pointed out above. An overview of the feature extraction process is shown in Fig. 2.

3.1 Local Binary Patterns

Prior to any further processing all color channels of the input images are transformed separately using a modified Local Binary Patterns operator (LBP) based on block averaging which we already used successfully to classify endoscopic images [8]. In contrast to the standard operator, which is described in more detail in [9], we compute the average over neighboring blocks and compare the average of the center block against the averages of the neighboring blocks to obtain the LBP number. By adjusting the block size used it is possible to find a trade-off between noise-suppression (higher block widths) and detail preservation (smaller block widths). Throughout this work we used a rather

high block width of 9 pixels to suppress noise which otherwise would have had a negative influence on the subsequent edge detection.

The motivation behind extracting edges from LBP transformed channels is that the pit structures we try to locate can be identified more easily since pits usually are surrounded by brighter areas. In terms of LBP searching for pits thus corresponds to locating LBP numbers above some certain threshold. Another advantage of the LBP operator is that it is known to be invariant against global illumination changes in images.

Furthermore, we achieve rotation invariance by circularly rotating each LBP number obtained until the minimum is reached [10]. This way we are able to cope with changes in the direction of illumination across different images. In the remaining part of this work this combination of averaged LBP blocks and rotation invariance is abbreviated with LBP_{AR}.

3.2 Polygon Extraction

In order to extract polygons from a LBP channel we first apply a global thresholding. The choice for the threshold used throughout this work is motivated by the appearance of an ideal pit ($t = 127$). It is chosen such that at least 7 of the 8 neighbor block averages must be higher than the center block average for a pixel to be assumed to be part of a pit.

Prior to edge detection we pre-process the binary image by using a set of six different morphological operators (O_C , O_B , O_I , O_H , O_M , and O_R). First, we apply a closing (O_C) - using a disk of radius 1 as structuring element - to remove small "holes". The small radius has been chosen to not disturb the shape of the pits too much but to only fill small holes and cancel out small notches eventually present along the borders of pit areas. Then we bridge (O_B) unconnected pixels by setting pixels to white which lie between two unconnected, white neighbors (using a 3×3 -neighborhood). Furthermore we remove isolated pixels (O_I) followed by setting black pixels surrounded by white ones to white (O_H). Finally, we cancel out pixels which have less than five white neighbors if only half of the pixels in the 3×3 -neighborhood or less are set to white (O_M). This step helps to minimize the number of small spurs which might eventually have endured the previous steps.

To extract edges from the resulting binary image we use the Canny edge detector without multi-resolution feature synthesis [11] which may produce polygons having a boundary with gaps. Thus we post-process the edges by using the morphological operators from above, except for the closing (in the same order). To obtain the final edge map we remove all interior pixels of the white areas (O_R).

This processing of the binary map and edges ensures that we end up with closed polygons only, which are smooth and free of unwanted artifacts. We get only closed polygons since when applying the filling of closed areas (O_H) polygons previously not closed are not affected (not filled) as can be seen in Fig. 3(e). By subsequently applying O_M these polygons are removed, which can be noticed from Fig. 3(f).

Apart from that we see from Fig. 3(a) that some images contain ridges which can be considered to be artifacts. By applying the post-processing steps to the edges ridges touching the image border are removed, thus reducing the number of these artifacts.

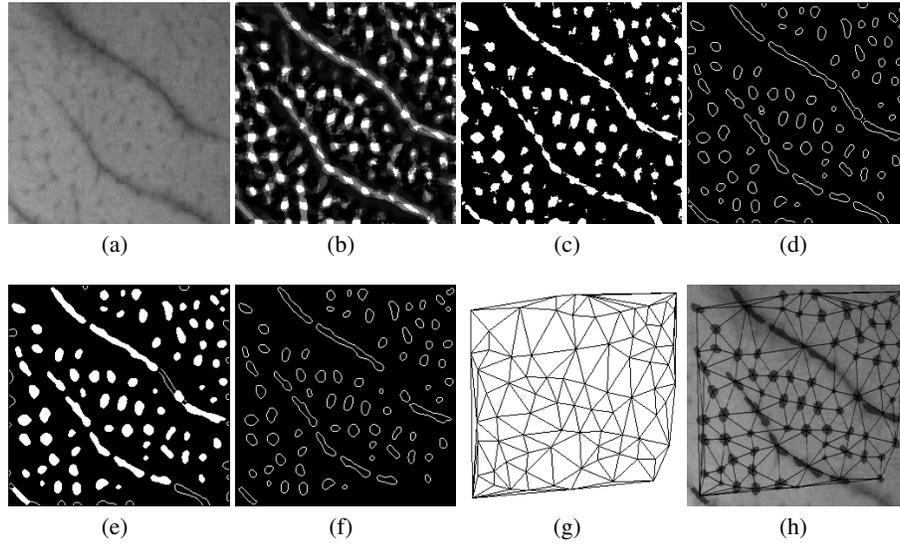


Fig. 3. Some steps from the process of obtaining polygons from a color channel of a pit pattern type I image (a) the red color channel of the input image, (b) the according LBP transformed image, (c) the result of thresholding, (d) result of the Canny edge detector, (e) the pits (white and filled) and the parts which get discarded by the edge map post-processing (not filled), (f) the final edge map, (g) the respective Delaunay triangulation, and (h) the triangulation with the pits overlaid to the original color channel.

After tracing the edges of the connected components we determine the polygon center for each polygon as the mean position of all edge pixels belonging to the polygon. Some of these steps are illustrated in Fig. 3(a)-(f).

3.3 Delaunay Triangulation

To measure the density of pits within an image we aim at constructing a mesh from the previously extracted polygon centers. Then we deduce the density of the pits from the edge lengths within the mesh. For this purpose we employ the Delaunay triangulation based on the Quickhull algorithm [12].

This algorithm basically transforms the 2D points to 3D (lifted to a paraboloid), computes the convex hull in 3D, and projects the lower part of the hull back to 2D to obtain the triangulation. This way we get a set of non-overlapping triangles with the minimum of the inner angles maximized. An example triangulation for a pit type I image is shown in Fig. 3(g).

Figure 4 shows sample images from our image database along with the respective Delaunay triangulations and the detected pits. From this figure we notice that non-neoplastic images exhibit a higher density with respect to the arrangement of the detected pits. But we also notice that in case of the non-neoplastic images ridges have a negative influence on this density in some parts of the images.

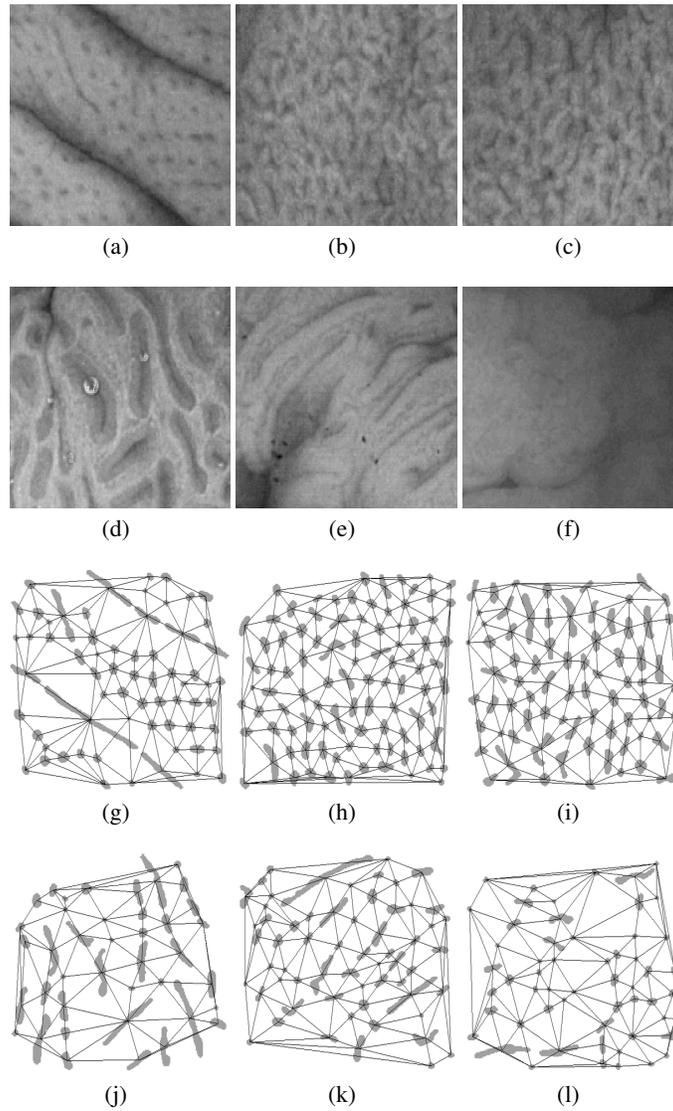


Fig. 4. Results of the Delaunay triangulation along with the detected pits. (a)-(c) example images from the non-neoplastic class (red channel), (d)-(f) neoplastic images, and (g)-(l) the according Delaunay triangulations along with the detected pits.

3.4 Histogram Creation and Classification

Based on the triangulations we create 1-dimensional histograms from the edge lengths of all triangles for each color channel of an image separately. To concentrate on triangles not located on the border of the triangulation we iterate over all triangles and use each edge of each triangle to update the histogram. This way edges shared by two triangles

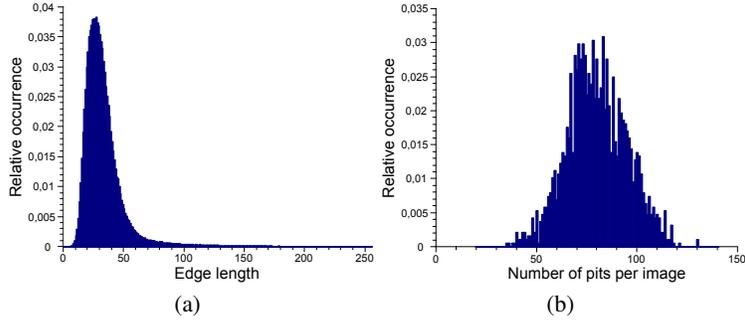


Fig. 5. (a) Relative occurrences of the different edge lengths in our image database and (b) the relative occurrences of the different number of detected pits across all images.

contribute to the histogram twice, while edges located on the border of the triangulation result are used only once.

If only a few pits are detected within an image the respective edge lengths will be rather high. For a high number of detected pits (higher density) the distances between them will get smaller, hence lowering the respective Delaunay edge lengths too.

Since the number of edges between images most likely will vary we normalize each histogram such that the histogram bins sum up to 1. This makes the histograms comparable during the classification process. Moreover, since all our images have a dimension of 256×256 pixels the upper limit for an edge length is $\sqrt{256^2 + 256^2} \approx 362$ (corresponding diagonal). But it is very unlikely that pits are only detected in the image corners. This implies that it is also unlikely that the maximum possible edge length occurs. Apart from that, the more pits we detect the more likely it is that the distances between neighboring pits get smaller.

The images used throughout our experiments show a maximum edge length of approximately 249, but most edge lengths lie between 10 and 100, as can be observed from Fig. 5(a). We also detect a rather high number of pits in each of our images (between 35 and 130), as can be seen from Fig. 5(b).

Based on these observations we consider the range for the edge lengths between 1 and 256 as a reasonable choice and therefore use this range throughout our experiments.

For the classification of unknown images we employ the k-NN classifier along with the histogram intersection distance metric, defined as

$$d(H_i, H_j) = \sum_{k=1}^B \min(H_{i,k}, H_{j,k}), \quad (1)$$

where H_i and H_j are two normalized histograms, B denotes the number of bins used in our histograms, and $H_{i,k}$ and $H_{j,k}$ represent the value of the k -th bin of histogram H_i and H_j , respectively. We also carried out experiments using the Euclidean distance metric and the Bhattacharyya distance metric but the histogram intersection always yielded a slightly better classification performance.

To combine multiple color channels we compute the distances for each color channel separately and multiply them to obtain the final distance D . This can be formulated

as

$$D(I_a, I_b) = \prod_{i=1}^C d(H_i^{(a)}, H_i^{(b)}), \quad (2)$$

where I_a and I_b denote two images, C is the number of color channels considered for combination, and $H_i^{(a)}$ and $H_i^{(b)}$ represent the histograms for the i -th color channel considered of image I_a and I_b , respectively. There are also other possibilities for a combination, for example summing up the distances instead of multiplying them by replacing the product in (2) by a sum. But since the product is more tolerant against outliers - one similar color channel in terms of histogram distance leads to a very small total distance between two images already - we favor the product instead of a sum.

4 Experiments

4.1 Settings

The image database used throughout our experiments consists of 627 images acquired between the years 2005 and 2008 at the Department of Gastroenterology and Hepatology (Medical University of Vienna) using a zoom-colonoscope (Olympus Evis Exera CF-Q160ZI/L) with a magnification factor set to 150.

Lesions found during colonoscopy have been examined after application of dye-spraying with indigocarmine as routinely performed in colonoscopy. Biopsies or mucosal resection have been performed in order to get a histopathological diagnosis. Biopsies have been taken from type I, II, and type V lesions, as those lesions need not to be removed or cannot be removed endoscopically. Type III and IV lesions have been removed endoscopically. Out of all acquired images, histopathological classification resulted in 178 non-neoplastic and 449 neoplastic cases which is used as ground truth for our experiments.

Using leave-one-out cross-validation, 626 out of 627 images are used as training set. The remaining image is then classified. This process is repeated for each image.

To find the optimal values for B (number of histogram bins used) and k for the k -NN classifier we carry out a naive search testing all possible combinations for $k = 1, \dots, 25$ and $B = 16, \dots, 256$ (for different color channel combinations).

4.2 Results

From the results shown in Table 1 we see that the proposed method achieves very promising results – in particular when combining two or more color channels. The best result has been obtained by combining the red and the blue channel, resulting in an overall classification accuracy of 93,3%. But also combining all color channels available yielded a high result of 93%.

From the results we also see that in case of the single channel results the green channel yielded the worst results. Also in case of combined channels the results always drop as soon as the green channel is taken into consideration.

Despite the high overall classification results we also notice that there is an imbalance between the two classes. While the results for the neoplastic images are always

Table 1. Overall classification rates obtained by different color channel combinations along with the respective choices for k and B (compared to the results published in [8]).

	Non-neoplastic	Neoplastic	Total	k	B
R	61,2	98,0	87,6	5	249
G	57,9	94,9	84,4	11	243
B	74,2	91,3	86,4	8	179
R+B	83,2	97,3	93,3	15	202
R+G	68,0	98,2	89,6	7	249
G+B	78,7	92,9	88,8	10	220
R+G+B	77,5	99,1	93,0	11	53
[8]	98,3	99,5	99,2	–	–

above 90% the results for the first class vary between approximately 58% and 83% only. This effect is especially apparent in case of the single channel results. When considering the ground truth we notice that the number of neoplastic images is about 1.6 times higher compared to the other class, which is one reason for this behavior.

Compared to the results we published in [8] we see that especially in case of the non-neoplastic images the results of the proposed approach are still very low. This is most possibly due to ridges, which – although not characteristic for non-neoplastic images – sometimes appear in these images too (see Fig. 3). As we also notice from Fig. 3(g) these ridges have a noticeable influence on the triangulation result.

Additional problems arise from image artifacts and noise which are quite frequently misinterpreted as being pits. As a consequence neoplastic images get more similar to non-neoplastic ones in terms of the Delaunay edge length histograms which makes misclassification of such images more likely. Although this problem exists, especially in case of neoplastic images, this is not evident from Table 1 due to the imbalance between the two classes.

5 Conclusion and Future Research

In this work we presented a method for an automated pit pattern classification system which - in contrast to all our previously published methods - is strongly linked to the visual appearance of the pits on the colonic mucosa. Although, compared to previously obtained results, this method still delivers lower recognition rates, the results we currently achieve are very promising already - especially when combining different color channels for the classification.

We also identified ridges as a potential problem being very likely one cause for a lowered classification performance. In future work we will therefore focus on minimizing the effect of ridges to a maximum possible extent. Besides that we will also have to investigate other features in order make the system work in the six-classes case as well. This case has been neglected completely in this work due to the nature of the features used, since these rely on differences in the density of pit distributions across different image classes. In the six-classes case this is unfortunately not sufficient since these dif-

ferences are not that distinct between all of the six classes. We will also have to make the pit detection more robust to improve the discrimination between the image classes.

Another interesting possibility will be to use the method proposed as part of an ensemble classifier, since this method works completely different compared to our previous approaches.

Acknowledgements

This work is partially funded by the Austrian Science Fund (FWF) under Project No. L366-N15 and by the Austrian National Bank “Jubiläumsfonds” Project No. 12514

References

1. Bruno, M.J.: Magnification endoscopy, high resolution endoscopy, and chromoscopy; towards a better optical diagnosis. *Gut* **52**(4) (June 2003) 7–11
2. Kudo, S., Hirota, S., Nakajima, T., Hosobe, S., Kusaka, H., Kobayashi, T., Himori, M., Yagyuu, A.: Colorectal tumours and pit pattern. *Journal of Clinical Pathology* **47** (1994) 880–885
3. Kato, S., Fu, K., Sano, Y., Fujii, T., Saito, Y., Matsuda, T., Koba, I., Yoshida, S., Fujimori, T.: Magnifying colonoscopy as a non-biopsy technique for differential diagnosis of non-neoplastic and neoplastic lesions. *World Journal of Gastroenterology: WJG* **12**(9) (March 2006) 1416–1420
4. Kudo, S., Tamura, S., Nakajima, T., Yamano, H., Kusaka, H., Watanabe, H.: Diagnosis of colorectal tumorous lesions by magnifying endoscopy. *Gastrointestinal Endoscopy* **44**(1) (July 1996) 8–14
5. Hurlstone, D.: High-resolution magnification chromoendoscopy: Common problems encountered in “pit pattern” interpretation and correct classification of flat colorectal lesions. *American Journal of Gastroenterology* **97** (2002) 1069–1070
6. Häfner, M., Gangl, A., Liedlgruber, M., Uhl, A., Vécsei, A., Wrba, F.: Pit pattern classification using multichannel features and multiclassification. In T.P. Exarchos, A. Papadopoulos, D.F., ed.: *Handbook of Research on Advanced Techniques in Diagnostic Imaging and Biomedical Applications*. IGI Global, Hershey, PA, USA (2009) 335–350
7. Häfner, M., Gangl, A., Liedlgruber, M., Uhl, A., Vecsei, A., Wrba, F.: Combining Gaussian Markov random fields with the discrete wavelet transform for endoscopic image classification. In: *Proceedings of the 17th International Conference on Digital Signal Processing (DSP’09)*, Santorini, Greece (2009)
8. Häfner, M., Gangl, A., Liedlgruber, M., Uhl, A., Vécsei, A., Wrba, F.: Pit pattern classification using extended local binary patterns. In: *Proceedings of the 9th International Conference on Information Technology and Applications in Biomedicine (ITAB’09)*, Larnaca, Cyprus (November 2009)
9. Ojala, T., Pietikäinen, M., Harwood, D.: A comparative study of texture measures with classification based on feature distributions. *Pattern Recognition* **29**(1) (January 1996) 51–59
10. Mäenpää, T., Pietikäinen, M.: Texture analysis with local binary patterns. In: *Handbook of Pattern Recognition and Computer Vision*. 3rd edn. World Scientific (2005) 197–216
11. Canny, J.: A computational approach to edge detection. *IEEE Transactions on Pattern Recognition and Machine Intelligence* **8**(6) (September 1986) 679–698
12. Barber, C., Dobkin, D., Huhdanpaa, H.: The Quickhull algorithm for convex hulls. *ACM Transactions on Mathematical Software* **2**(4) (December 1996) 469–483