CURRENT CHALLENGES ON REAL-TIME POLYP DETECTION IN COLONOSCOPY VIDEOS

From Image Segmentation to Classification. A Pattern Recognition-based Approach

Keywords: Medical Imaging, Colonoscopy, Pattern Recognition, Segmentation, Polyp Detection, Region Description, Machine Learning, Real-time.

Abstract: In this paper we present our approach on real-time polyp detection in colonoscopy videos. Our method consists of three stages: Image Segmentation, Region Description and Image Classification. Taking into account the constraints of our project, we introduce our segmentation system that is based on the model of appearance of the polyp that we have defined after observing real videos from colonoscopy processes. The output of this stage will ideally be a low number of regions of which one of them should cover the whole polyp region (if there is one in the image). This regions will be described in terms of features and, as a result of a machine learning schema, classified based on the values that they have for the several features that we will use on their description. Although we are still on the early stages of the project, we present some preliminary segmentation results that indicates that we are going in a good direction.

1 INTRODUCTION

The main lesions associated to the intestine are: bleeding, lump, Crohn disease and cancer. In this paper we present our approach to cancer polyp detection in colonoscopy videos but, before entering into implementation details, we will introduce the disease and the current methods to detect its appearance.

1.1 Colon Cancer

Colorectal cancer (also called colon cancer), with an approximate number of 655.000 deaths worldwide per year, has become the fourth leading cause of death by cancer in the United States and the third leading cause in the Western world (National Cancer Institute, 2010). Colon cancer includes cancerous growths in the colon, rectum and appendix. Colon cancers arise from adenomatous polyps in the colon which can be identified by its prominent (flat or peduncular) shape.

Invasive cancers that are confined within the wall of the colon can be curable with surgery. If they are not treated, they spread to regional lymph nodes and, finally, in the cases where cancer metastasizes to distant sites, they may not be curable. Colorectal cancer can take many years to develop and early detection of colorectal cancer increases greatly the chances of being cured. Even taking this account their efficiency, colorectal cancer screening rates remain low (Allen et al., 2010) and it is often recommended for individuals who are at increased risk.

1.2 Screening Techniques

There are several different tests available for this purpose, and we will divide them into several groups according to their main principles.

The first group is composed by techniques that imply the on-line intervention of ah physicist, such as Digital Rectal Exam (DRE) (Baumgart et al., 2010) or Endoscopy, where a lighted probe is inserted into the rectum and colon to check for polyps and other abnormalities. There is another group of techniques that can aid in cancer polyps visualization and detection, consisting of introducing some particle into the patient's body and observe the reaction of the patient to this substance. Belonging to this group we have Virtual Colonoscopy (National Digestive Diseases In-

formation Clearinghouse, 2010b), where a 3-D model of the inside of the large intestine is created from CT or MRI images, or Positron Emission Tomography (PET) (Weston et al., 2010) where radioactive sugar, which collects in tissues with high metabolic activity, is injected into the patient. Finally we have Wireless Capsule Video Endoscopy (WCVE) where a capsule attached with a camera, a battery and a set of lamps for illumination is swallowed by the patient. It emits a radio frequency signal which is received and stored in an external device (Eliakim, 2010) (giving a video movie as a result that can be analyzed later).

The final group is composed by techniques that explore the status of the patient without any intervention such as Computed Axial Tomography (Vries et al., 2010) or Blood tests. Although all the techniques are being used currently, there is a clear tendency of using colonoscopy (and virtual colonoscopy). As we will have to deal with colonoscopy videos in our project, we will explain this method more in depth.

1.3 Colonoscopy

Colonoscopy is a procedure used to see inside the colon and rectum and it can detect inflamed tissue, ulcers, and abnormal growths (National Digestive Diseases Information Clearinghouse, 2010a). During colonoscopy, patients lie on their left side on an examination table. The doctor inserts a long and flexible tube called colonoscope into the anus and guides it slowly through the rectum and into the colon. The scope inflates the large intestine with carbon dioxide gas in order to give the doctor a better view. A small camera is mounted on the scope and transmits a video image from inside the large intestine to a computer screen, allowing the doctor to examine carefully the intestinal lining. During this process the doctor can remove polyps and later test them in a laboratory to look for signs of cancer while the physicist can also take some samples from tissues during colonoscopy (process known as biopsy) to do a later analysis.

Colonoscopy is widely accepted as the definitive method for diagnosing colon cancer because it allows a direct visualization of the intestinal surface but it has its drawbacks, such as (Winawer et al., 1993) the risk of perforation, the intervention cost, visualization difficulties, bad patient-doctor ratio, need of preparation, etc. Although this methods performs well, there can be still some errors related to the visibility of some zones and to the skills of the physician.

In the following sections we will present the objective of our project and then explain the our approach for each stage of our method.

2 OUR PROJECT

The objective of our project is to develop a tool that can indicate the doctor, in real-time, which areas of the colon are more likely to contain a cancer polyp by means of computer vision techniques. In order to achieve this goal we base our approach, which general scheme that can be seen in Figure 1, on a common Pattern Recognition (Devijver and Kittler, 1982).

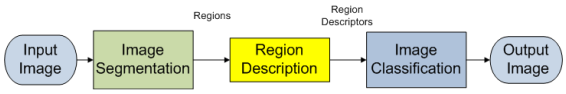

Figure 1: General Scheme of Our Approach.

Our approach consists of three consecutive process (Image Segmentation, Region Description and Image Classification). The choice of the concrete stages of our approach is done having in mind the requirements of our project.

So, the first stage will consist of segmenting automatically the image in order to end up with a reduced number of regions that may contain relevant information. So the objectives are twofold: reduce the number of regions that should be analyzed and eliminate those that are no relevant for our application. With this approach, which acts as the Feature (in this case Regions of Interest) Detection step, we can reduce the dimensionality of our problem and therefore ease the following steps.

The second and third stages are closely related one to the other. Once we have a few regions from the segmentation stage, we need to find some characteristics in these regions that can denote the presence or not of a cancer polyp. The idea here is to use a combination of Feature Descriptors that can define together what is a polyp region and what is not by observing the values of the regions for a series of parameters. But first our method will have to learn, from examples, what is a polyp region and what is not. This will be done in the learning step of our Image Classification stage.

Image Classification is divided into two processes: learning and testing, based on usual machine learning approaches. In the learning stage we will train our method with a large number of examples of both polyp-containing and non-polyp-containing regions, described using the Feature Descriptors from the second stage. Once a new example arrives it will be incorporated into the testing stage and, by using machine learning algorithms, classified.

As it can be seen all the stages are connected and the success of the whole approach depends on the individual success of each of them. We believe that a

good classification system, that is the underlying aim of our project, will work better if its inputs are good. And to get good inputs we need to describe the best that we can the data that we have. And considering our real-time requirements, it is better that amount of data that we have is small and also relevant.

In order to develop our approach we rely on a database of thousands of images extracted from 15 different videos of colonoscopy interventions. As of now we have only implemented the segmentation stage of our approach but in the following sections we will present our ideas of all the stages.

3 IMAGE SEGMENTATION

Our method for image segmentation has to take into account the special characteristics of the images that we are working with so, before describing our method we will present a summary of what we should take into account in our method after what we have learnt by observing the videos and images.

3.1 Observing Our Data

As we have mentioned before, we are dealing with colonoscopy images obtained from real interventions videos. While observing the videos, we have found out that the lighting of the probe can give us hints about what is a polyp in an image. As the light falls perpendicularly to the walls of the colon, it creates shadows around the surfaces at which it is. More precisely, when the light falls into a prominent surface, it creates a bright spot (with high grey-scale value) surrounded by darker areas, which are the shadows, generating edges and valleys in the intensity image. This can be better understood by looking at Figure 2, where we show an indication of the effect of light in the intensity profiles, that depend on how the polyp appears (cenital or lateral view).

Even considering this evidences of prominent surface appearance, there are some challenges that need to be overcome:

- Non-uniform Polyp Appearance: first, the appearance of the polyp by itself is not uniform, as it can be seen in Figure 3, going from peduncular to flat shapes. Second, in most of the images we will not have a clear vision of the polyp, viewing them from cenital or lateral views, which makes difficult a shape-based recognition scheme work.
- Uniform Colour Pattern: As it also can be seen in Figure 3, a segmentation scheme based purely on color has difficulties to segment correctly the image, because all the tissues in the image present a very similar color distribution.
- Effect of the Reflections: As the camera lights the image, reflections appears. And the majority of Feature Detectors (such as SIFT (Lowe, 1999), SURF (Bay et al., 2006), MSER (Matas et al., 2004) or Harris (Harris and Stephens, 1988)). As it can be seen in Figure 4, these detectors mark interest points around reflections instead of pointing the parts of our images that separates structures (in our case, polyps).
- Over and under Segmentation: We can have two problems related to the number of segmented regions. Oversegmentation is related to having a large number of very small regions, which implies a high computation cost to analyze them all) and under-segmentation to the fact of having a smaller number of bigger regions, but still higher than the number of structures that a human could identify

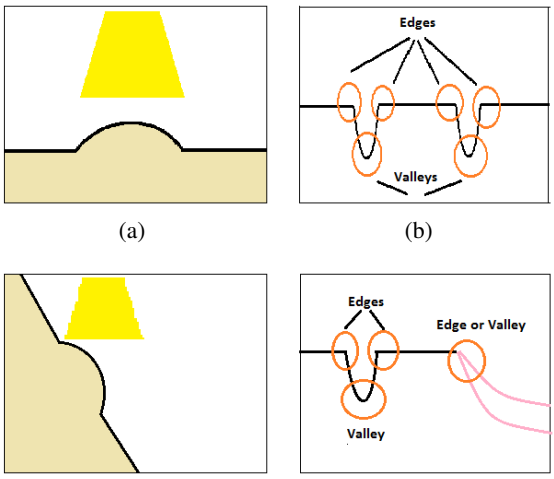

Figure 2: (a) Simulation of an illuminated prominent surface (cenital) (b) Simulation of grey-scale profile (c) Simulation of an illuminated prominent surface (lateral) (d) Simulation of grey-scale profile.

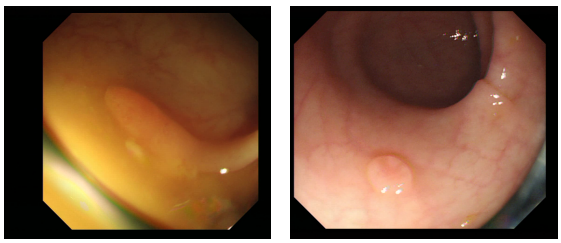

Figure 3: (a) Peduncular Polyp (b) Flat Polyp.

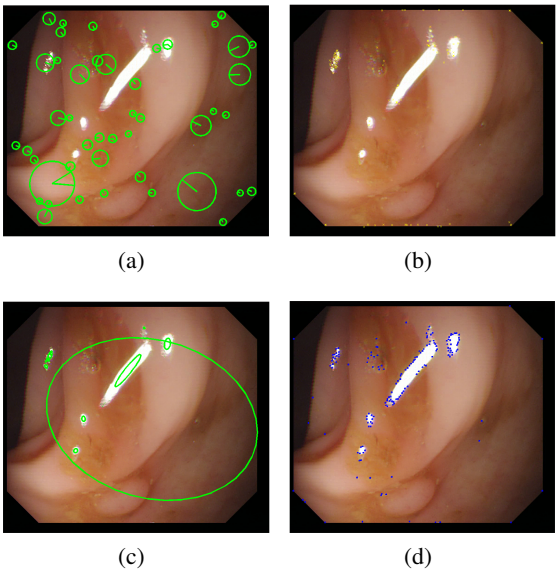

(a) (b)

(c) (d)

Figure 4: Detection by (a) SIFT (b) SURF (c) MSER (d) Harris.

on the image). An example of both can be seen in Figure 5.

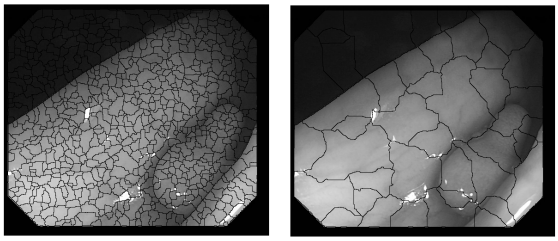

(a) (b)

Figure 5: (a) Oversegmentation (648 regions) (b) Undersegmentation (44 regions).

So, taking these considerations into account, we base our segmentation method on a model of appearance of a polyp that we can roughly define as a prominent shape enclosed in a region that can be identified by the presence of edges and valleys. But we have to take this as an indication, and we also should try to overcome the challenges we have presented.

3.2 Our Segmentation Approach

In this subsection first we present the basics of each step of our segmentation approach and at the end we show in Figure 7, step by step, a complete graphical example.

1. **Image Preprocessing:** Before applying any segmentation algorithm there are some operations that should be done to the original input image (Figure 7 a) in order to overcome some of the challenges that were presented before. These preprocessing operations include: converting to grayscale ((Figure 7 b), image deinterleaving (as our images come from a high definition interleaved video), correction of the reflections (Figure 7 c) and obtaining the complemented version of the image (Figure 7 d).
2. **Segmentation:** In this step we have several alternatives. We can use either simpler (in terms of computation cost) methods such as watersheds (Vincent and Soille, 1991) or go with algorithms that are more powerful in terms of segmentation such as Mean-shift (Wang et al., 2004), or Normalized Cuts (Shi and Malik, 2000). We have chosen to use watersheds in behalf of reducing the computation cost as much as possible and also because of the more complex approaches are generally color-based and, as we have mentioned before, this is not useful in our case. Another point of our approach is that, instead of using the preprocessed image obtained in the first step, we use gradient information. As it can be seen in Figure 6 by using gradient information we get a first segmentation that encloses better the structure of the shapes that appear on the image.

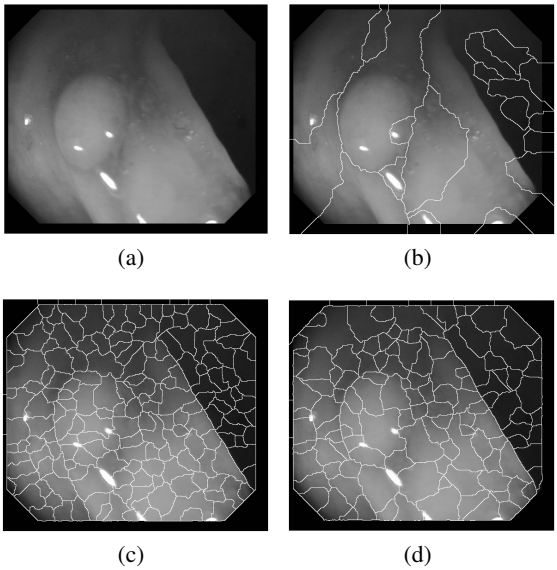

(a) (b)

(c) (d)

Figure 6: (a) Original Image Preprocessed (b) Original Image Segmented (c) Gradient Modulus Image Segmented (d) Morphological Gradient Image Segmented.

After this step our image will be divided in a large number of regions (Figure 7 e). This number will be reduced by merging neighbor regions.

3. **Frontier-based Region Merging:** As a first region merging approach, we focus on merging small neighbor regions that are separated by weak frontiers (Figure 7 f). We denote as weak those frontiers which present a low degree of certain measures, measured as the percentage of pixels that fall under the dark side of a series of masks that we have defined. These masks take into account the appearance model of the polyp that we defined before and therefore include an edge mask and a valley mask (López et al., 1999) (apart from some others such anisotropic filtering mask).

4. **Region-based Region Merging:** In this step, as it did not happen in the previous one, we consider when merging not only the weakness of the frontiers (by using different measures such as if the frontier is kept after smoothing the image) that separate them but also the grey-level content of the regions. First we categorize the regions and frontiers, in terms of amount of information that they contain (i.e., low information means a region with a very high or very dark mean grey level and very low standard deviation) and then merge compatible regions (with the same degree of information) that have weak frontiers with the same degree of information. The objective here is to end up with a reduced number of large regions which content is as uniform as possible and with clear frontiers of separation.

The objective of our segmentation and region merging sub-steps is twofold: first obtaining a good segmentation of the image that fits the structure of the several objects that appear in it, such polyps and second, join and label those parts of the image where we know we will not find a polyp inside and therefore, we should not process them in order to save resources. We show in Figure 7 one complete segmentation process, showing the output at the end of each step.

As it can be seen in the example, we start with a color image with reflections, that we correct and then pass this image, complemented, as input to the watersheds transform. The first segmentation that this method provides divides the image in 184 regions. By merging small regions with weak frontiers we reduce the number of regions to 136 regions. Adding region characterization to the region merging process lets us reduce the number of regions up to 9.

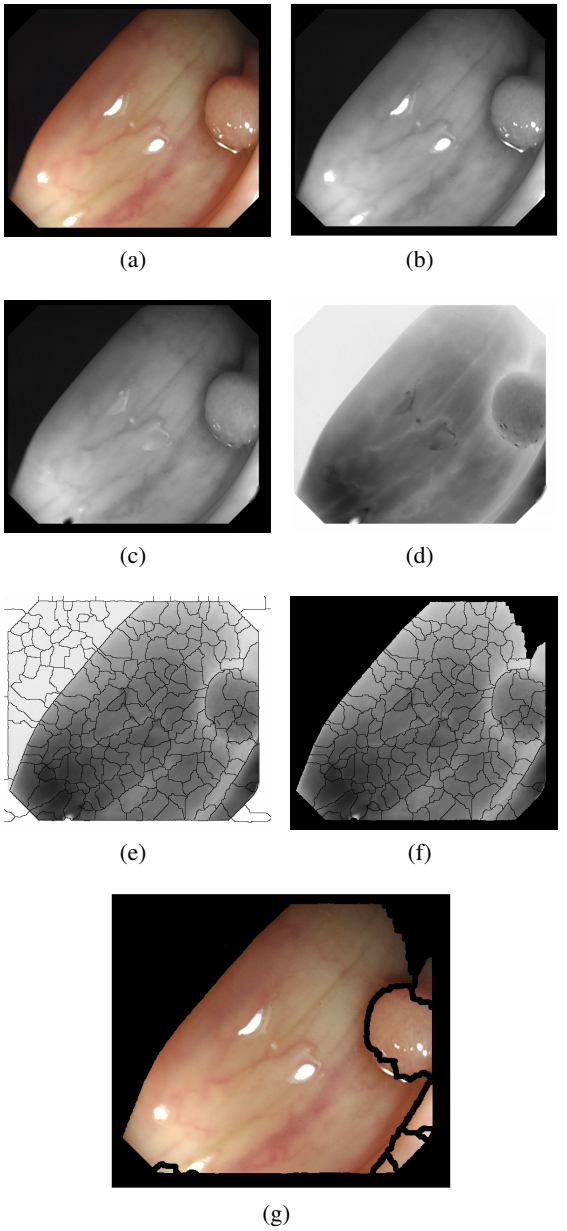

Figure 7: (a) Original Image (b) Grey-scale (c) Reflection Corrected (d) Preprocessed Image (e) After segmentation (184 regions) (f) After Region-based Merging (136 regions) (f) Final Segmented Image (9 regions).

4 REGION DESCRIPTION

Although this stage is not implemented, we will present in this section our ideas about how to do a profitable Region Description, taking into account the nature of our problem. If we study the bibliography of Feature Descriptors, we can separate them into four groups: Shape Descriptors, Color Descriptors, Texture Descriptors and Motion Descriptors. Our approach to this stage is not to rely on one only type of descriptors (as it can be seen in Figure 8) and, as possible, try to use really informative descriptors.

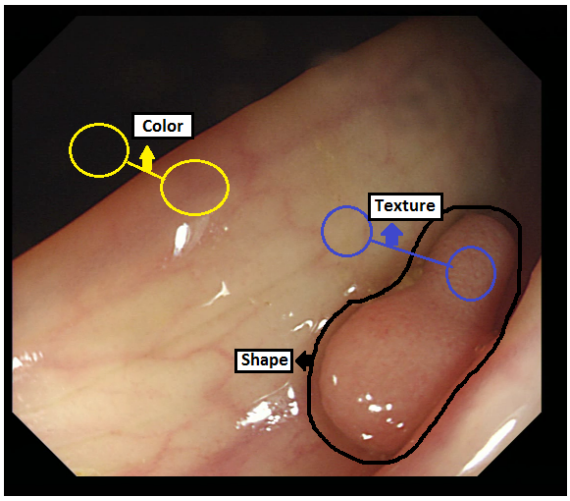

Figure 8: Examples of Shape, Color and Texture cues.

If we take a look at the example, we can see that each of the types of descriptors may have a role in our system. For example, we can see that the polyp is enclosed by a closed contour and could be approximated by an ellipsoidal shape, so here we could use Shape Descriptors in order to eliminate some regions. Color as it can be seen can be an important cue when defining what is clearly not a polyp and we can also observe the difference of texture between the polyp region (more granular) and non-polyp region (more plain). The problem here, as it was shown in Figure 3, is that some of the assumptions that we have made for this example may not be extrapolated to all the images. In the rest of this section we will present each group of Feature Descriptors and indicate in Table 1 some of the most important ones of each group.

- **Shape Descriptors:** Shape can be a key aspect in our description stage, but it needs that a good segmentation has been done. Depending on the pose of the camera, which affects how they appear, and the type of polyp (peduncular or flat) that we are aiming to find, we can have several methods, which are divided into the ones that consider only the contour of the shape (Contour-based) and the ones that consider all the region within (Region-based). A further division is done taking into account if they use all the available information (Global-based) or divide it into small parts (Structural).
- **Color Descriptors:** In our case, color cannot be used as a principal descriptor because on many images the interior of the polyp region has the same color appearance than the outside, but it can be used to assess the detection of polyp frontiers (in the cases where polyps are surrounded by blood vessels).
- **Texture Descriptors:** Although there is a great diversity of Texture Descriptors (and the most well-known methods are into this group) their use in this type of images presents problems, many related to the non-uniform appearance of the polyp. We cannot state that polyps by themselves have a concrete texture pattern but we do not discard at all their use because, for example, veins and blood vessels do have a texture pattern and we can use this information to discard regions.
- **Motion Descriptors:** As we are dealing with video images, we can think of using Motion Descriptors as a way to track polyps, that is, once we have found a region which is very likely to contain a polyp inside we can predict the position of this same region in the next frame and give an extra help in this next frame region classification.

Type	Methods
Shape	• Contour-based – Global: Wavelets (Chuang and Kuo, 1996), Fourier (Kauppinen et al., 1995), Shape Signature (Zhang and G.Lu, 2004). – Structural: Chain Code (Sun and Wu, 2007), Blurred Shape Model (Radeva, 2007), Shape Context (Bohg and Kragic, 2009). • Region-based – Global: Zernike Moments (Khotanzad and Hong, 1990), Shape Matrix (Zhang and Lim, 2007), Angular Radial Partitioning (Chalechale et al., 2004). – Structural: Skeletons.
Color	Scalable Color Descriptor (Borghesani et al., 2009), Color Structure Descriptor (Kundu et al., 2009), Color Constant Color Indexing (Funt and Finlayson, 1995).
Texture	SIFT (Lowe, 1999), SURF (Bay et al., 2006), Texture Browsing Descriptor (Lee and Chen, 2005), Local Binary Patterns (Ojala et al., 2000), Co-ocurrence Matrices (Carr and de Miranda, 1998).
Motion	Optical Flow (Horn and Schunck, 1981), Angular Circular Motion (Erol and Kossentini, 2001).

Table 1: Feature Descriptors.

5 IMAGE CLASSIFICATION

As happens with the Region Description step, this stage has not been implemented yet in our method

because of two reasons: first, time constraints (since we are on the sixth month since the project started its development) and secondly, and more important, because we have a strong belief in that the classification system is good as long as its inputs (outputs from the previous stages) are good. As in many Pattern Recognition-based methods, we will use a machine learning approach (Bishop et al., 2006), which is a scientific discipline that is concerned with the design and development of algorithms that allow computers to evolve behaviors based on empirical data, such as from sensor data or databases. A learner can take advantage of examples (in our case, descriptions of polyp and non-polyp containing regions) in order to capture characteristics of interest of their unknown underlying probability distribution.

We will provide the system examples of regions that contain polyps and examples of regions that no contain polyps. These example regions will be described and incorporated into the machine learning algorithm of choice in order to learn a polyp and non-polyp pattern. New input images will be automatically segmented, described and incorporated into the testing step with the objective of finding out if they are more near to be a polyp candidate region or non-polyp candidate region. In our case our success will be measured not only on terms of how many true positives we get but also on the number of false negatives. It seems clear that it is harmless to identify one non-polyp region as a polyp region that the opposite.

Machine learning algorithms are commonly organized into a taxonomy, based on the desired outcome of the algorithm. This classification goes from supervised learning (where we give the system supervised training data, as we pretend to do on our case) to unsupervised learning (where, without previous knowledge, the system seeks to determine how the data are organized), presenting also several degrees of supervision in the learning (semi-supervised learning). There are several machine learning algorithms that can be used to fulfill our objectives, such as Support Vector Machines (Hearst et al., 1998), Neural Networks (Dayhoff and DeLeo, 2001), Decision Trees (Mitchell, 1997) or Bayesian Networks (Friedman et al., 1997) and our intention is, after a deep study of their characteristics, choose the one that suits better our problem.

6 RESULTS SO FAR: DISCUSSION AND FUTURE WORK

Detecting polyps in colonoscopy video is not an easy task but our intention has never been to provide an algorithm that segments perfectly the polyp but to help the physicist which zones in the image are more likely to contain polyps, and do it on real-time. In previous sections we have explained our approach and, until now, we have only developed part of the image segmentation stage and studied the state of the art on Feature Detection and Feature Description. As the requirements of our project include real-time constraints, we are strongly focused on reduce as possible the dimensionality of the problem by means of a correct image segmentation. In Figure 9 we show some preliminary segmentation results.

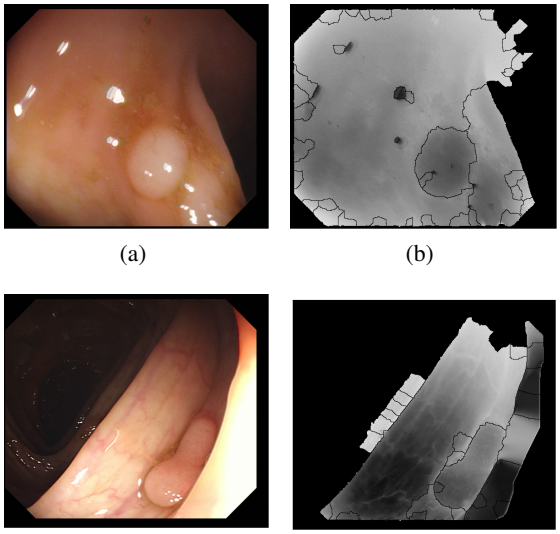

(a) (b)

(c) (d)

Figure 9: (a-c) Original Images (b-d) Segmented Images.

The segmentation results shows that first, we are reducing the search are of the image by eliminating some areas (that are show in black in the segmented image) and second, that we end up with a reduced number of regions (considering that we start with more than 600 regions, we end up with less than 40). The segmentation is not perfect yet, being room for improvements in both preprocessing and region merging stages. Also we need to tune the region merging methods in order to obtain more stable results for a wider variety of images.

Once we have a stable segmentation algorithm, our next step will be to study in depth the kind of images that we have (and the final segmented regions that we will generate) in order to decide which of the available Feature Descriptors can be useful for us or, if necessary, develop some that may fulfill our expectations. As we said before, we strongly believe that choosing a descriptor for only one of the four categories that we have presented is not the way to do

things. As it is shown in Figure 8, we state that a combination of several types of Feature Descriptors will be needed to define what is a polyp-containing region and what is not.

As our objective is to classify the regions of the image into polyp-containing candidates, we will need to learn what a polyp is in order to classify new input images. In the classification stage we plan to use a machine learning procedure. We will train our system off-line so we will characterize polyp and non-polyp regions. Once a new image arrives, we will segment it, describe the regions and put them into our classification algorithm. The idea here is to measure the level of 'polypness' of a region but also taking into account that detecting what is not a polyp can also be useful while detecting what is a polyp.

As we have said throughout the paper, our approach is guided by what we have learnt by observing the images from our databases. We have explored the images in order to find which cues can be relevant when defining what is a polyp, in order to adapt our segmentation method to it. We do not have to forget the requirements of our project, which includes real-time objectives, which is the reason of our strong focus on image segmentation in order to reduce the dimensionality of the problem. If we get to deal with a reduced number of regions we can maybe make a bigger effort in the Feature Description stage which could lead to a stronger definition of a polyp region which is crucial in the final classification stage.

7 CONCLUSIONS

The objective of our project is to detect in real-time cancer polyps in colonoscopy videos. To do so, we propose a Pattern Recognition scheme divided in three main stages. The first one, Image Segmentation, is done with two objectives: reduce the dimensionality of the problem as much as possible and to provide as result to the later stages of the process chain a set of regions of interest. In order to segment the image we have studied in depth the structure of several polyp-containing images from our video database, with the aim of finding cues that can let us discern which regions have some evidence of containing polyps and which not (in this case, we will not analyze again these regions). As of now, our method reduces greatly the number of regions, offering as output a small number of them that, in some cases, cover the whole shape of the polyp in just one region.

Once we have this subset of regions of interest, the next step, Region Description, consists of describing them in terms of features in order to characterize them. This step will be the seed to the final stage, Image Classification, where our current idea is to implement a machine learning approach. The ideal output will be a mask superimposed to the image that will enhance the parts where the physicist should pay more attention.

Currently we have only implemented the Image Segmentation stage but the results obtained gives us hope that in a near future we will be able to start with the description step. The task is not easy, but our objectives are clear and we have a path to follow.

REFERENCES

Allen, J., Barlow, W., Duncan, R., et al. (2010). NIH State-of-the-Science Conference Statement: Enhancing Use and Quality of Colorectal Cancer Screening. *NIH consensus and state-of-the-science statements*, 27(1).

Baumgart, L., Gerling, G., and Bass, E. (2010). Characterizing the range of simulated prostate abnormalities palpable by digital rectal examination. *Cancer epidemiology*, 34(1):79.

Bay, H., Tuytelaars, T., and Gool, L. V. (2006). SURF: Speeded Up Robust Features. *Computer Vision–ECCV 2006*, pages 404–417.

Bishop, C. et al. (2006). *Pattern recognition and machine learning*. Springer New York:.

Bohg, J. and Kragic, D. (2009). Grasping familiar objects using shape context. In *International Conference on Advanced Robotics*.

Borghesani, D., Grana, C., and Cucchiara, R. (2009). Color Features Performance Comparison for Image Retrieval. *Image Analysis and Processing–ICIAP 2009*, pages 902–910.

Carr, J. R. and de Miranda, F. P. (1998). The semivariogram in comparison to the co-occurrence matrix for classification of image texture. *Geoscience and Remote Sensing, IEEE Transactions on*, 36(6):1945–1952.

Chalechale, A., Mertins, A., and Naghdy, G. (2004). Edge image description using angular radial partitioning. *IEEE Proceedings-Vision Image and Signal Processing*, 151(2):93–101.

Chuang, C. H. and Kuo, C. C. J. (1996). Wavelet Descriptor of Planar Curves: Theory and Applications. In *IEEE Transactions on Image Processing*.

Dayhoff, J. and DeLeo, J. (2001). Artificial neural networks. *Cancer*, 91(S8):1615–1635.

Devijver, P. and Kittler, J. (1982). *Pattern recognition: A statistical approach*. Prentice Hall.

Eliakim, R. (2010). Video capsule endoscopy of the small bowel. *Current opinion in gastroenterology*, 26(2):129.

Erol, B. and Kossentini, F. (2001). Local motion descriptors. In *IEEE Workshop on Multimedia Signal Processing*, pages 467–472. Citeseer.

Friedman, N., Geiger, D., and Goldszmidt, M. (1997). Bayesian network classifiers. *Machine learning*, 29(2):131–163.

Funt, B. V. and Finlayson, G. D. (1995). Color Constant Color Indexing. *IEEE transactions on Pattern analysis and Machine Intelligence*, 17(5):522–529.

Harris, C. and Stephens, M. (1988). A Combined Corner and Edge Detection. In *In Proceedings of The Fourth Alvey Vision Conference*, pages 147–151.

Hearst, M., Dumais, S., et al. (1998). Support vector machines. *IEEE Intelligent Systems and Their Applications*, 13(4):18–28.

Horn, B. K. P. and Schunck, B. G. (1981). Determining Optical Flow. *Artificial intelligence*, 17(1-3):185–203.

Kauppinen, H., Seppanen, T., and Pietikainen, M. (1995). An experimental comparison of autoregressive and Fourier-based descriptors in 2D shape classification. *IEEE transactions on pattern analysis and machine intelligence*, page 201.

Khotanzad, A. and Hong, Y. H. (1990). Invariant image recognition by Zernike moments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 489–497.

Kundu, M., Banerjee, M., and Bagrecha, P. (2009). Interactive Image Retrieval in a Fuzzy Framework. *Fuzzy Logic and Applications*, pages 246–253.

Lee, K. L. and Chen, L. H. (2005). An efficient computation method for the texture browsing descriptor of MPEG-7. *Image and Vision Computing*, 23(5):479–489.

López, A., Lumbreras, F., et al. (1999). Evaluation of methods for ridge and valley detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(4):327–335.

Lowe, D. G. (1999). Object Recognition from Local Scale-Invariant Features. In *International Conference on Computer Vision, Corfu, Greece*, pages 1150–1157. Published by the IEEE Computer Society.

Matas, J., Chum, O., Urbana, M., and Pajdla, T. (2004). Robust wide-baseline stereo from maximally stable extremal regions. *Image and Vision Computing*, 22(10):761–767.

Mitchell, T. (1997). Decision tree learning. *Machine Learning*, 414.

National Cancer Institute (2010). What You Need To Know About Cancer of the Colon and Rectum.

National Digestive Diseases Information Clearinghouse (2010a). Colonoscopy.

National Digestive Diseases Information Clearinghouse (2010b). Virtual Colonoscopy.

Ojala, T., Pietikainen, M., and Maenpa, T. (2000). Gray scale and rotation invariant texture classification with local binary patterns. *Computer Vision-ECCV 2000*, pages 404–420.

Radeva, P. (2007). Multi-class binary object categorization using blurred shape models. In *Progress in pattern recognition, image analysis and applications: 12th Iberoamerican Congress on Pattern Recognition, CIARP 2007*, page 142.

Shi, J. and Malik, J. (2000). Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence*, 22(8):888–905.

Sun, J. and Wu, X. (2007). Shape Retrieval Based on the Relativity of Chain Codes. *Lecture Notes in Computer Science, Multimedia Content Analysis and Mining*, pages 76–84.

Vincent, L. and Soille, P. (1991). Watersheds in digital spaces: an efficient algorithm based on immersion simulations. *IEEE transactions on pattern analysis and machine intelligence*, 13(6):583–598.

Vries, A. D., Bipat, S., et al. (2010). Polyp measurement based on CT colonography and colonoscopy: variability and systematic differences. *European radiology*, 20(6):1404–1413.

Wang, J., Thiesson, B., et al. (2004). Image and video segmentation by anisotropic kernel mean shift. *Computer Vision-ECCV 2004*, pages 238–249.

Weston, B., Iyer, R., et al. (2010). Ability of integrated positron emission and computed tomography to detect significant colonic pathology. *Cancer*, 116(6):1454–1461.

Winawer, S., Zauber, A., et al. (1993). Prevention of colorectal cancer by colonoscopic polypectomy. *The New England journal of medicine*, 329(27):1977.

Zhang, D. and G.Lu (2004). Review of Shape Representation and Description Techniques. *Pattern Recognition*, 37(1):1–19.

Zhang, D. and Lim, M. C. Y. (2007). An efficient and robust technique for region based shape representation and retrieval. In *6th IEEE/ACIS International Conference on Computer and Information Science, 2007. ICIS 2007*, pages 801–806.