

**A COMPLETE SYSTEM FOR POLYPS FLAGGING IN
VIRTUAL COLONOSCOPY**

By

Marcelo Fiori


Pablo Musé

and

Guillermo Sapiro

IMA Preprint Series # 2367

(April 2011)



INSTITUTE FOR MATHEMATICS AND ITS APPLICATIONS

UNIVERSITY OF MINNESOTA
400 Lind Hall
207 Church Street S.E.
Minneapolis, Minnesota 55455-0436

Phone: 612/624-6066 Fax: 612/626-7370

URL: <http://www.ima.umn.edu>

A Complete System for Polyps Flagging in Virtual Colonoscopy

Marcelo Fiori, Pablo Musé, and Guillermo Sapiro

Abstract—Computer tomographic colonography, combined with computer-aided detection, is a promising emerging technique for colonic polyp analysis. We present a complete pipeline for polyp flagging based on a simple segmentation technique that enhances polyps, a multi-scale candidate polyp delineation, and new texture and geometric features that consider both the information in the candidate polyp location and its immediate surrounding area. The proposed algorithm is tested with ground truth data, including flat and small polyps which are hard to detect even with optical colonoscopy, obtaining very promising results. For polyps greater than 6mm in size we achieve 100% sensitivity with just 0.8 false positives per study, and for polyps greater than 3mm in size we achieve 100% sensitivity with 2.2 false positives per study.

Index Terms—Virtual colonoscopy, CT colonography, computer-aided detection, colonic polyp detection, colon segmentation, curvature motion, geometric features, texture features, differential features, classification.

I. INTRODUCTION

COLORECTAL cancer is the second leading cause of cancer-related deaths in the United States (only surpassed by lung cancer), and the third cause worldwide [1]. The early detection of polyps is fundamental, allowing to reduce mortality rates up to 90% [2]. Nowadays, optical colonoscopy (OC) is the most used detection method due in part to its relative high performance. However, this technique is invasive, very expensive, and still prone to miss polyps (in particular small or flat ones), making it hard to use in large screening campaigns.

Computer Tomographic Colonography (CTC), or Virtual Colonoscopy (VC), is a promising alternative technique that emerged in the 90's [3]. It uses volumetric Computed Tomographic data of the cleansed and air-distended colon. It is less invasive and less expensive than optical colonoscopy, and as a consequence, much more suitable for screening campaigns once its performance is demonstrated. CTC also has the potential to outperform OC, in particular for small or flat polyps, or those polyps in regions of the colon where OC has been shown to perform poorly. In addition, regarding optical colonoscopy, only around 70% of the colon is explored visually. Incomplete studies due to obstructing colorectal lesions, colon twists, or anatomical variations are not rare (5% to 15% [4]), and there is the additional risk of perforating the colon. In a large study by Kim *et. al* [5], where about 3000 patients went through OC

and another 3000 through VC, seven perforations occurred in the OC while none were recorded in the VC.

Nevertheless, it takes more than 15 minutes for a trained radiologist to complete a CTC study, and the performance of the overall optical colonoscopy is still considered better. In this regard, Computer-Aided Detection (CAD) algorithms can play a key role, assisting the expert to both reduce the procedure time and improve its accuracy.

Colon lesions can be classified according to their size, measured in diameter, and according to their morphology, into pedunculated, sessile, or flat. Flat lesions are of special interest because these are an important source of false negatives in CTC, and they are around 10 times more likely to contain high-grade epithelial dysplasia¹ [6], [7], [8].

There are numerous discussions regarding the potential risks of the polyps according to their size. Even though some authors consider that “small” polyps may not represent risk, the majority of gastroenterologists disagree [9]. Summers claims that one of the major challenges in the field is in increasing sensitivity for smaller polyps [10]. At the same time, Bond [11] declares that the major disadvantage of VC is its current low performance for flat lesions.

The goal of the work presented in this paper is to exploit VC precisely to automatically flag colon regions with high probability of being polyps, with special attention to results in challenging small and flat polyps.² Toward this aim, we propose a complete pipeline that starts with a novel and simple segmentation step. We then introduce geometrical and textural features that take into account not only the candidate polyp region, but the surrounding area at multiple scales as well. This way, our proposed CAD algorithm is able to accurately detect candidate polyps by measuring local variations of these features. The whole algorithm is completely automatic and produces state-of-the-art results.

The rest of this paper is organized as follows. In Section I-A we briefly review prior related work and in Section I-B we present an overview of the whole proposed pipeline. We address the colon segmentation problem in Section II and the feature extraction and classification in Section III. In Section IV we describe the evaluation method and results. The discussion is presented in Section V and we conclude in Section VI.

Marcelo Fiori and Pablo Musé are with the Instituto de Ingeniería Eléctrica, Universidad de la República, Uruguay.

Guillermo Sapiro is with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, Minnesota, 55414, USA.

¹An abnormality of development in cells that may become *cancer in situ* or *invasive cancer*.

²This paper extends our previous conference publication [12].

A. Virtual Colonoscopy CAD Review

Automatic polyp detection is a very challenging problem, not only because the polyps can have different shapes and sizes, but also because they can be located in very different surroundings. Most of the previous work on CAD of colonic polyps is based on geometric features, some using additional CT image density information, but none of them takes into account the (geometric and texture) information of the tissues *surrounding* the polyp. This is a crucial issue since it is well known that the tissue properties of the colon may vary with location. This local and adaptive differential analysis is part of the contributions of this work.

Early work on CAD methods by Vining *et al.*, [13], is based on the detection of abnormal wall thickness. Summers *et al.*, [14], detect polyps greater than 10mm by computing mean curvatures and sphericity ratio, and present results over a large screening patient population. Yoshida *et al.*, [15], use the *shape index* (defined later in this paper) and *curvedness* as geometric features, applying fuzzy clustering and then using directional gradient concentration to reduce false positives. Paik *et al.*, [16], also use geometrical features, computing the Surface Normal Overlap (SNO) instead of calculating curvatures. Wang *et al.*, [17], compute a global curvature, extract an ellipsoid, and analyze morphological and texture features on this ellipsoid. They reach 100% sensitivity with a relative low false positives (FP) rate, using heuristic thresholds and texture features. Hong *et al.*, [18], map the 3D surface to a rectangle, use 2D clustering, and reduce false positives with shape and texture features. Sundaram *et al.*, [19], compute curvatures via the Smoothed Shape Operators method, and use principal curvatures and Gaussian curvatures to detect polyps. Götürk *et al.*, [20], propose a technique to reduce the false positives based on features calculated from three random orthogonal sections, and then classifying with SVM. All these described techniques based on local geometric computations suffer from a high dependence on the regularity of the polyp shape itself, ignoring how pronounced it is with respect to the surrounding area. Using geometry alone is also very sensitive to the colon segmentation.

More recently, van Wijk *et al.* [21] proposed a PDE motion that flattens only the polyp-like shapes, and then they consider the difference between the original and the processed images. The main drawback of this approach is that the resulting flattened polyps look like a flat lesion, therefore the algorithm does not detect flat lesions. Suzuki *et al.* [22] use artificial neural networks to reduce the false positive number of the previous algorithm by Yoshida *et al.* described above [15]. Although there were no polyps submerged in fluid in the database, the results are very promising, achieving 96.4% sensitivity (over 28 polyps) with 1.1 FP per case. However, the sensitivity is not perfect, an evaluation with small lesions and polyps submerged in liquid should be performed, and the FP number can be further improved. Konukoglu *et al.* [23] introduced a preprocessing stage that enhances the polyps via a PDE evolution based on the heat equation, and showed how it improves the CAD performance. Proprietary algorithms, [24], [25], have been reported as well, but with no better results

than the methods mentioned above. Although the comparison of the experimental results is delicate since different databases were used, all these approaches can be improved as here demonstrated, either detecting a more general class of lesions or directly on the classification performance.

The results reported by the algorithms presented above were reported on using databases containing polyps greater than 6mm in size (or greater than 10mm in some databases). To the best of our knowledge, no algorithm reported in the literature can detect small polyps properly. On the other hand, for polyps bigger than 6mm in size, no algorithm can achieve 100% sensitivity with less than one false positive per study. Therefore, it is important to keep improving these techniques.

B. Overview of the Proposed System

The main goal we are addressing in this work is to highlight/flag all the candidate polyps, so the radiologist can quickly check them. It is crucial to minimize the false negatives, keeping a reasonable false positives number. We achieve this by a four-steps process, Figure 1, which is completely automatic and constitutes the entire end-to-end algorithm, from data to candidate polyps flagging.

In the first step, the colon segmentation, the input is the CT volume data, and the output is a 3D mesh representing the colon surface. In the second step, from the segmented mesh we perform a multi-scale search of candidates in order to capture the appropriate polyp size, obtaining a set of candidate patches. The CT volume data, the segmented colon surface, and the candidate patches are the input to the third step, where we compute geometrical and textural features. The final step consists of a machine learning algorithm that classifies polyps and non-polyps patches from the computed features. In the following sections we describe each of these steps in detail.

II. COLON SEGMENTATION

The segmentation of the colon surface, which is critical in particular to compute geometric features, is divided into two parts: a pre-processing stage for dealing with the air-liquid composition of the colon volume, and a second stage that consists on smoothing the pre-processed image and obtaining the final colon surface by thresholding the smoothed volume. The overall procedure here presented is very simple and computationally efficient, leading to the state-of-the-art classification results later reported.

A. Classifying CT regions

One of the strongest difficulties concerning the segmentation of the colon from abdominal volumes in CT is the presence of liquid and its interfaces with air and tissue. Figure 2 shows a CT slice and its pixel values over the highlighted vertical segment. At first sight there are three clearly distinguishable classes: the lowest gray levels correspond to air, the highest levels correspond to liquid, and the middle gray values correspond to tissue. Nevertheless, there are around 6 interface voxels between air and liquid whose gray values, due to continuity, lie within the normal tissue range. Therefore,




Fig. 1. Basic pipeline of the proposed polyp flagging system.

a naïve approach based on gray values only, ignoring the physical nature of the tissue and its environment, is not suitable for proper tissue classification and segmentation.




Fig. 2. CT slice and typical values for air, liquid and normal tissue.

The proposed approach addressing this issue consist on generating a *probability map* \mathbf{P} , that is, a 3D image with the same size of the CT volume (around $512 \times 512 \times 400$ for our data), whose values are the (estimated) probability of the voxel of being inside the colon. We will then process this map to estimate the interior colon wall.

We obtain \mathbf{P} by computing for each voxel the three probabilities of belonging to each of the three classes of interest (namely, air, liquid and interface, the union of these three classes forms the inside of the colon), and then taking \mathbf{P} as the maximum of these three probabilities.

Since the gray value distributions of air (class w_1) and liquid (class w_2) are very distinguishable and stable among the different studies, they can be empirically learned by manually segmenting these two classes in a given CT study, and then constructing the probability distribution functions $p(x|w_1)$ and $p(x|w_2)$ by standard kernel density estimation [26]. From this, the distributions $p(w_1|x)$ and $p(w_2|x)$ can be computed, assuming uniform class priors. The air and liquid probabilities for a voxel in the testing CT are then easily computed by simply evaluating these estimated distributions on the gray value of the CT voxel.

The challenging component of computing \mathbf{P} is the computation of the interface probability. Here we take advantage of the physics of the problem, and in particular of the gravity and the position of the patient: the person is laid horizontally so the interface between the liquid and the air is a plane parallel to the floor. The voxels situated on the interface then have a large gradient in the z (vertical) direction, since the values of

air and liquid are on opposite ends. However, the transition is about 6 voxels wide for the standard data resolution used in this work, and the computations should be done taking this into account. Additionally, if a given voxel belongs to the interface layer, it is expected that at least half of the neighbor voxels at the same horizontal plane also belong to the interface layer.

The implementation of these criteria is as follows. A cubic neighborhood around each voxel \mathbf{x} is considered, and for each one of the “columns” that result of fixing the x and y coordinates, the probabilities of the upper voxels of being air and the lower voxels of being liquid are accumulated. If the tested voxel belongs indeed to the interface layer, then all these air and liquid probabilities will be high. The interface probability is then an increasing function of this accumulated measures. The algorithm in Figure 3 provides a pseudo-code that represents this procedure.

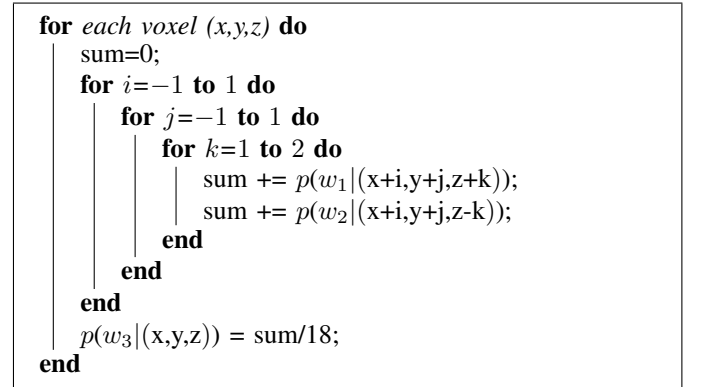


Fig. 3. Interface probability computation

After the computation of the probability map some spurious (isolated) voxels may have high probability of being liquid (bones for example, or simply noise), and we know that the colon volume should be connected, so we clean the probability map by keeping the connected component³ containing a chosen voxel used as seed. The seed is automatically detected choosing the voxel with greater value of P_i , since these high probabilities occur only at the true interface between air and

³ Actually, since the probability map is not binary, a (conservative) threshold of 0.6 is considered to separate the connected components.

liquid. Alternatively, the seed could also be provided by the user for a more interactive system if desired.

Figure 4 shows a slice of the original volume data and the same slice of the computed probability map \mathbf{P} .




Fig. 4. CT slice and its corresponding slice on the probability map.

B. Smoothing and colon surface computation

In order to eliminate noise and to obtain a smoother colon surface after the segmentation, the 3D probability map \mathbf{P} should be first smoothed.

One of the fundamental aspects that we look for in this smoothing process is the preservation of the shape of the polyps, while obtaining a smooth enough surface to reliably compute geometric features such as curvature. Of course, the ultimate goal is to derive a method to process the surface that simplifies and improves the polyp/non-polyp classification system.

At this point we have the previously computed probability map \mathbf{P} , which will be processed by a Partial Differential Equation (PDE) of the form

$$\frac{\partial u(\mathbf{x}, t)}{\partial t} = \beta |\nabla u| \quad , \quad u(\mathbf{x}, 0) = \mathbf{P}(\mathbf{x}) \quad , \quad (1)$$

where \mathbf{P} is the initial condition. After making a few iterations of this PDE evolution, an iso-level surface of the resulting 3D image $u(\mathbf{x}, T)$ will be extracted, and that will be the resulting colon surface. Since we are interested in a certain iso-level surface of the volume, it will be useful to point out the relation that states how the iso-levels evolve when the 3D image evolves according to (1). Namely, the basic equation of the Level Set Method [27] states that if $u(\mathbf{x}, t)$ evolves according to (1), then its iso-levels (level sets) satisfy

$$\frac{\partial \mathcal{S}}{\partial t} = \beta \vec{\mathcal{N}} \quad , \quad (2)$$

where \mathcal{S} is any iso-level surface and $\vec{\mathcal{N}}$ its unit normal.

Therefore, in what follows we study surface motions, in particular when the deformation term β is a function of the principal curvatures κ_{max} and κ_{min} . In the actual implementation, this term (and therefore the principal curvatures) have to be computed from the function $u(\mathbf{x}, t)$, a standard procedure in the implementation of curvature-based surface motions [28]. We first review the classical motions by curvature and then propose some variations used as part of our developed pipeline.

The classical mean curvature motion

$$\frac{\partial \mathcal{S}}{\partial t} = \mathcal{H} \vec{\mathcal{N}}$$

has problems regarding topological changes, consider for example the deformation of a dumbbell [29], which is a surface homeomorphic to a sphere that evolves into two separate topological spheres. Nevertheless, making only a few iterations this problem does not affect our surface. With very few iterations a very smooth surface is obtained, but the polyps are flattened fast too, as it can be observed in Figure 5. However, if used carefully, it is a very good alternative to more standard Gaussian filtering.




Fig. 5. Mean Curvature Motion: original surface and the result after 2, 8, 15, 30 and 50 iterations. Note how both the surface (as desired) and potential polyps (undesired) are smoothed and flattened.

The motion by Gaussian curvature K on the other hand, has several problems with surfaces containing nonconvex parts, which is our case. Indeed, only a few iterations are needed for the evolving surface to explode (decompose into disconnected parts).

A suitable version of the motion by Gaussian curvature, namely the affine motion

$$\frac{\partial \mathcal{S}}{\partial t} = (K^+)^{1/4} \vec{\mathcal{N}} \quad \text{where} \quad K^+ = \max(K, 0) \quad ,$$

has a better behavior in general, in terms of topological preservation. However, when used with only a few iterations, the results are comparable with the mean curvature.

The last classical motion here discussed is the motion by minimal curvature κ_{min} . Caselles [30] showed with experimental results that this evolution is the only among those just described that preserves the topology of a bent dumbbell (homeomorphic to a sphere). In our particular application, this evolution yields very good results in terms of both surface smoothing and polyp enhancement.

In what follows we try to take advantage of the known properties of our shapes of interest to propose motions that behave better in our application. Inspired in the exponent $1/4$ of the affine motions in dimension 3, and the best (in terms of polyp-shape preservation) of the motions presented above, we first tested the evolution that obeys the equation

$$\frac{\partial \mathcal{S}}{\partial t} = \kappa_{min}^{1/4} \vec{\mathcal{N}} \quad .$$

It turns out that the motion governed by this equation yields better results than the ones presented. Figure 6 shows the result after a few iterations, and Figure 7 evidences the difference with a comparative image: the result of the motion by κ_{min} is in gray while the result of the motion by $\kappa_{min}^{1/4}$ is shown in orange. It can be seen that the polyp surface in the latter is above the polyp surface of the former, and the surrounding zone is the other way around, showing that the evolution




Fig. 6. Evolution by $\kappa_{min}^{1/4}$: original surface and the result after 2, 8, 15, 30 and 50 iterations.




Fig. 7. Comparison between evolutions. Motion by κ_{min} in gray vs. motion by $\kappa_{min}^{1/4}$ in orange.

by $\kappa_{min}^{1/4}$ better enhances the polyp when compared with the motion by minimal curvature.

In order to further improve the smoothing process, we will characterize the polyps shape and distinguish even more the evolution for those zones. Towards this aim, we present a surface property that will be used in the feature extraction stage as well.

A measure of the local shape of a surface is the so-called *shape index* [31],

$$SI := -\frac{2}{\pi} \arctan \left(\frac{\kappa_{max} + \kappa_{min}}{\kappa_{max} - \kappa_{min}} \right),$$

where κ_{max} and κ_{min} are the principal curvatures. A complementary measure called *curvedness* C , is defined as

$$R := \sqrt{\frac{\kappa_{max}^2 + \kappa_{min}^2}{2}}, \quad C := \frac{2}{\pi} \ln R.$$

The $(\kappa_{max}, \kappa_{min})$ plane is then transformed into the (SI, C) plane. While the value of SI is scale-invariant and measures the local shape of the surface, the value of C indicates how pronounced it is. Figure 8 shows different shapes and their corresponding shape index. Due to the chosen orientation, shape index values close to -1 (protrusions) are of special interest for polyp detection.

Back to the PDE motion, the next step is to include this information concerning the shape of the surface in order to make potential polyps evolve differently than the rest of the colon surface. More precisely, we modify the best motion so far ($\kappa_{min}^{1/4}$), in such a way that the resulting motion further enhances the potential polyps. In order to achieve this, it is needed to characterize the potentially polyp points and modify the deformation function according to that. The shape index characterization is very precise, but the computation of the




Fig. 8. Some shapes and their corresponding shape index values.

shape index of the level surface from the embedding function u is delicate. Indeed, the typical scheme is to compute \mathcal{H} and K by direct formulae and then the principal curvatures (which are necessary to compute the shape index) are obtained by solving a second order equation, leading to

$$\kappa_{max,min} = \mathcal{H} \pm \sqrt{\mathcal{H}^2 - K},$$

so very little errors in the estimation of \mathcal{H} and K may become important errors in the estimation of the principal curvatures. In order to overcome this difficulty we directly compute the principal curvatures [32]. Let P_v be the operator that represents the projection onto the orthogonal complement of v , that is, $P_v = I - \frac{vv^T}{|v|^2}$, let H_u be the Hessian matrix of the function u , and consider the symmetric matrix

$$J = \frac{1}{|\nabla u|} P_{\nabla u} H_u P_{\nabla u}.$$

The orthogonal complement of ∇u is the tangent plane, where the second fundamental form is defined. It is easy to see that ∇u is an eigenvector of J corresponding to the zero eigenvalue, due to the projection. It can be shown that the other two eigenvalues are the principal curvatures and the corresponding eigenvectors are the principal directions.

This computation scheme leads to much more accurate estimation of the curvatures, and therefore better estimations of the shape index. Both techniques (the direct formulae and this latter one) yield very similar results for the mean curvature \mathcal{H} . The important difference is on the computation of the Gaussian curvature K , where the latter technique leads to much more stable results. When the absolute value of the Gaussian curvature K is small, sometimes the different estimations have different signs, which leads to very significant errors in the shape index estimation, since SI is very sensitive to the signs.

We now define a function that acts as a multiplying factor of the term $\kappa_{min}^{1/4}$, making the surface evolve slower at the points of interest. One option is to choose this function to depend on the shape index only, assigning low values to shape index near -1 , and values close to unity to other points. A smooth function $g(SI)$ verifying these constraints is shown in Figure 9.




Fig. 9. Function $g(SI) = \frac{1}{\pi} \arctan((SI - 0.75) \cdot 10) + \frac{1}{2}$, multiplying factor for PDE curvature evolution.

The final motion then becomes

$$\frac{\partial \mathcal{S}}{\partial t} = g(SI) \kappa_{min}^{1/4} \vec{N}.$$

This proposed evolution keeps all the advantages of the motion by $\kappa_{min}^{1/4}$ and in addition, the polyps are flattened slower, so at the end the obtained surface is smooth and the polyps are still outstanding.

At this point, after choosing the appropriate diffusion, we have a smoothed probability map $u(\mathbf{x}, T)$ indicating the volume inside of the colon. We then extract the surface of the colon, using the marching cubes algorithm [33] to obtain the iso-value surface of level $\alpha \in [0, 1]$. The choice of the value α can be made by maximizing some criteria, in order to obtain the best surface in a given sense [34]. This optimization-oriented method was tested, but in our particular application all the consistent surfaces are very close to each other, and all of them are reasonable segmentations of the colon. Therefore, the computational effort is not justified and we simply use a fixed value, extracting the iso-level surface for $\alpha = 0.7$. Note that this choice can be safely made once for all, since all the studies are performed under very similar conditions. The result of this stage is then a triangulated surface \mathcal{S} representing the colon wall.

III. FEATURE EXTRACTION

For each case study, we consider as input both the raw CT images and the segmentation of the colon volume as obtained following the procedure described in the previous section. The first stage of the proposed polyps flagging algorithm consists in detecting surface patches that are candidates of being polyps. The complete set of connected points that constitutes the candidate patch is found by progressively growing the patch and keeping the one that maximizes the geometric dissimilarity with the surrounding area, in the sense of the features presented below (the starting point for this growth is also detailed below). The detection system is therefore based on differential (non-absolute) features, thereby better adapting to intrinsic variabilities both of the colon region and of the potential polyps, as further explained next.

All the polyp detection methods reported in the literature try to detect or classify the polyps from properties defined only within the candidate region, without considering the data

surrounding the region. However, it is important to analyze the context in which the candidate patch is located, not only because different sections of the colon present different characteristics, but also because polyps can be situated over different structures such as folds or plain colonic wall. A good feature including the shape of the neighborhood for example, can help in the discrimination between irregular folds and polyps over folds. In addition, looking for significant differences in the gray level imitates the human-based inspection, which highlights zones that contrast with their vicinity.

In this regard, most of the features described in this section take into account the local information of the area surrounding the candidate patch. Polyps (actually all the candidate patches) are then characterized not only by their intrinsic geometry and structure, but also by their relationship with the surrounding area. This makes the features more robust to the particular local phenomena, in a context where the natural variations of the properties of the colon tissue impact the measures and make absolute thresholds or decision rules impractical.

A. Candidate detection and geometrical features

The starting point for the geometric features described in this section is the segmented surface \mathcal{S} . Let us consider the shape index as a function $SI : \mathcal{S} \rightarrow [-1, 1]$, and recall that the polyps have shape index values close to -1 . Therefore, it is expected that a region (patch) of the surface that corresponds to a polyp contains at least one local minimum of the shape index function. The detection of the candidate patches starts from this observation, and follows a multi-scale search. For each local minimum $x_0 \in \mathcal{S}$ of the function SI , several level sets of SI ($\mathcal{P}_1 \dots \mathcal{P}_n$) around x_0 are tested, and the level set \mathcal{P}_i that maximizes the distances between the histograms described below, is the final considered candidate patch, which we simply denote by \mathcal{P} . A total of $n = 7$ level sets are tested, corresponding to the shape index values from -0.8 to -0.5 with a -0.05 step. The following description is given for the final chosen patch \mathcal{P} , but the ring and histogram computations are made for all the level sets \mathcal{P}_i in order to select the most appropriate of them (the most appropriate scale).

Given a candidate patch \mathcal{P} , a ring \mathcal{R} around \mathcal{P} is computed, in order to consider geometrical measurements with respect to the area surrounding the patch. The ring is calculated by dilating the patch \mathcal{P} a certain geodesic distance, such that the areas of \mathcal{P} and \mathcal{R} are equal. The geodesic distance computation is performed using the algorithm in [35]. Figure 10 shows a candidate patch (actually a true polyp), and its corresponding ring.

Histograms of the shape index values are then computed for the patch \mathcal{P} and the ring \mathcal{R} , and two different distances between them are computed: the L_1 distance and the symmetric Kullback-Leibler divergence. If the patch corresponds to a polyp-like shape then the values of the histogram \mathcal{P} will be concentrated around the -1 extrema, on the other hand, the histogram \mathcal{R} will be inclined to the other extreme in case of a polyp on a normal colon wall (concave), or with tendency to values near -0.5 if the polyp is on a fold. These two features give a measure of the geometric local variation of the candidate

patch \mathcal{P} . We assume that there are no other polyps in \mathcal{R} or that they do not significantly affect the statistics on the ring.




Fig. 10. Ring (in blue) surrounding a candidate polyp (in orange).

Although these two are the most discriminative features, we also consider the following additional ones since they still help to discriminate some typical false positives:

- The mean value of the shape index over the patch \mathcal{P} , which describes the shape of the selected patch.
- The area of the patch, since we want to detect polyps in a certain range of size.
- The growth rate of the areas at the multi-size stage, meaning the ratio between the area of the chosen patch $\mathcal{P} = \mathcal{P}_i$ and the area of the immediately smaller tested patch \mathcal{P}_{i-1} ; this feature measures how fast the shape of the patch is changing, in a context where it is difficult to quantize the variation of the shape.
- And finally the *shape factor*

$$SF = \frac{4\pi \cdot Area}{Perimeter^2},$$

which measures the shape of the patch border, how efficiently the perimeter is used in order to gain area,⁴ and it favors circle-like patches (like the polyp patch in Figure 10), avoiding elongated patches (like the false positives in folds).

We then end-up with a total of 6 geometric features for detecting candidate polyps, namely: L_1 and Kullback-Leibler distance between shape index histograms of patches and corresponding rings, the mean shape index over \mathcal{P} , the area of the candidate patch, the growth rate of the areas and the shape factor.

B. Texture features

Due to the differences in biological activity of polyp cells, the gray-level of the CT image and its texture can be very helpful for detecting polyps. This is in particular useful for flat or small polyps, where the geometric information is limited. Some work has been done on the inclusion of texture features (inside the candidate polyps only), in order to reduce false positives [17]. According to the reported results, there is a lot of room for improvement in texture features. We propose

both the use of new texture features and the inclusion of the information on the candidate's surrounding area.

First, for each polyp candidate $\mathcal{P} \subset \mathcal{S}$, a volume V_1 is calculated, containing the patch \mathcal{P} and a portion of the inner tissue next to the patch. A second volume V_2 , surrounding V_1 is calculated, containing normal tissue, in order to compare it with the polyp candidate tissue.

The features chosen are a subset of the classical Haralick texture features [36], namely, entropy, energy, contrast sumMean, and homogeneity. Seven co-occurrence matrices (considering seven directions in \mathbb{R}^3 , $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$, $(1, 0, 1)$, $(1, 1, 0)$, $(0, 1, 1)$ and $(1, 1, 1)$) are calculated with the voxels of V_1 , and all the five features are averaged over the seven directions. The analogous computation is made for V_2 , and the differences between the two volumes, for each texture feature, is considered. Additionally, the mean gray levels of the voxels in both volumes is computed, and their difference is considered as a feature. In this way, six texture features are considered. This approach for computing the texture features, measuring differences with the surrounding area, leads to better discrimination than the features computed just for V_1 , as demonstrated next.

IV. CLASSIFICATION AND RESULTS

A total of 31 cases of the WRAMC database were used to test the proposed CAD algorithm,⁵ with 49 polyps detected by optical colonoscopy, including two flat polyps. Among these 49 polyps, 34 are greater than $6mm$ in size, and the other 15 are between $3mm$ and $6mm$ in size.

After the candidates detection with the multi-scale approach, the number of true polyps was much lower than the number of non-polyps patches, a relation on the order of 500:1, which is a significant problem for the learning stage of the classifier. Three techniques were considered to overcome this shortcoming.

The approach of Domingos [37] is called MetaCost, and it consists of combining several instances of the classifier instead of stratification (modify the proportion of classes in the training data according to the costs). This method does not work with "stable" classifiers (those that produce similar models with slightly different training sets) like SVM or Naive Bayes.

In the work by Elkan [38] the class imbalance problem is tackled and a scheme called Cost Sensitive Learning is proposed. Unlike the MetaCost approach, it tries to balance the classes before the learning stage. The author proves a theorem giving a formula for how to rebalance the data, given the target threshold probability. The implementation used (from *Weka*) reweights the training instances according to the cost matrix, or predicts the class with minimum expected misclassification cost.

Finally, the Synthetic Minority Over-sampling TEchnique (SMOTE) is a method to generate artificial instances of the minority class, in order to get a balanced data to learn from.

⁴The maximum value for the shape factor is 1 and its achieved only by the circle.

⁵Data provided courtesy of Dr. Richard Choi, Virtual Colonoscopy Center, Walter Reed Army Medical Center.

The new artificial instances are created as a linear combination of the existing instances of the minority class. Therefore, there is an underlying assumption that the optimal partition in the feature space gives convex sets, which may not be the case in several applications.

We tested all these options and the best results were obtained using Cost Sensitive Learning, as expected from the comments above.

The numerical results listed below were obtained by classifying with SVM using Cost Sensitive Learning, after normalizing the data; Naïve Bayes performed similarly.

Using the leave-one-out strategy (i.e., testing with one case and training with the rest) all the 49 polyps were detected with an average of 2.2 FP per patient case (1.3 FP per polyp). These values are comparable with the state-of-the-art results [17], [22], but include very small and therefore challenging polyps in our database. A more precise comparison of results is not necessarily useful, since in general each work considers its own database.

On the other hand, when testing our CAD pipeline with the polyps greater than 6mm in size (34 in total, including two flat polyps), which is the typical framework in the VC CAD literature, the results are further improved, a 100% sensitivity is achieved with just 0.8 FPs per study.

	Texture features	
	Absolute	Differential
Sensitivity	96%	100%
FP per case	3.1	2.2

TABLE I
COMPARISON OF ABSOLUTE AND DIFFERENTIAL TEXTURE FEATURES,
WITH POLYPS GREATER THAN 3mm IN SIZE.

	Polyp sizes	
	> 3mm	> 6mm
Sensitivity	100%	100%
FP per case	2.2	0.8

TABLE II
COMPARISON OF PERFORMANCE BY POLYP SIZE.

Table I shows the comparison between absolute and differential texture features. The classification was performed using all the geometric features and either the absolute texture features (computed just for V_1), or the differential texture features, using the standard leave-one-out strategy. The results show that, when combined with the differential geometric features, differential texture features are significantly more discriminative than the absolute ones. Table II compares the classification results according to the polyps size. Again, the work with such small, as well as flat polyps, is unique to the framework here presented; see next.

V. DISCUSSION

A. Small polyps, big polyps, flat lesions

It is clear that both the small polyps and flat lesions are much more difficult to detect than the other polyps. What is not clear is if the same kind of algorithm and features are suitable for detecting all the range of polyp types and sizes. We

showed that the proposed combination of features, although it may not be optimal for every specific type of lesion, is able to correctly detect all of them.

All the stages in the pipeline, specially the segmentation and the features, contribute to the good classification results for the whole database. However, it would be interesting to study which pre-processing techniques and features are better for each type and size of polyps, and eventually propose different CAD systems for each class of polyp. Nevertheless, the 100% sensitivity together with the 2.2 FP rate for polyps greater than 3mm in size is as remarkable as the 0.8 FP rate for polyps beyond 6mm in size.

B. Geometric and texture importance

Although the geometrical features are the most discriminative ones, the texture ones still play a fundamental role in the classification. Adding the texture features to the geometric ones, the sensitivity reaches 100%, and at the same time the false positives rate decreases by 30%.

Figure 11 shows a detected polyp, where the geometry is crucial, because the gray-level does not present considerable local variations. This is specially true in polyps located over tagged material. On the other hand, in the flat polyp of Figure 12, the geometry is weakly discriminative (although the measure considering the ring enhances the detectability), and the texture features lead to a correct classification.

Texture information is very important also because it is more robust to segmentation errors, as the texture features are computed integrating from the volumetric data itself (once the local volumes have been considered). Moreover, the differential texture features (the differences between V_1 and V_2), outperform the absolute texture features (just computed in V_1), as shown in Table I.




Fig. 11. Polyp with no texture information.




Fig. 12. Polyp with texture information, but weak geometric information.

C. Qualitative analysis of false positives

In addition to the number of false positives, it is very important to study how these FP patches look like, since some of them can be quickly ruled out by the expert and some can be avoided by improving some aspects of the segmentation step.

About half of the false positives are quite reasonable, in the sense that they are (usually small) sections of the colon that are polyp-like shaped, Figure 13, specially taking into account that we designed the system to also detect small and flat polyps. On the other hand, about 20% of the FPs were in fold sections of the wall, Figure 13, or parts of the insufflation tube. All these patches are easily ruled out by inspection. Another 20% of the FPs were caused by segmentation errors, this number significantly reduced thanks to the texture features as discussed above.




Fig. 13. False positives: fold and patch similar to polyp.

VI. CONCLUSION

We introduced a complete pipeline for a Computer Aided Detection algorithm that flags candidate polyp regions. The segmentation stage is very simple and fast, and its main novelty is the smoothing PDE which enhances the polyps, leading to a better detection. In addition to the incorporation of the Haralick texture features, the main novelties of the classification stage are in the consideration of the surrounding area for each candidate polyp (we compute differential features instead of absolute ones), and the strategy for testing regions of multiple sizes. Differential features are significantly more discriminative than the absolute ones, as they emphasize local deviations of the geometry and texture over the colon. Testing regions of different sizes allows to precisely delineate polyps. The obtained quantitative results are very promising, detecting 100% of the true-polyps, including flat and small ones, with a low rate of false positives. Additional improvement of the segmentation and, in collaborations with radiologists, finding features that are tailored to polyp-like geometries, can further improve these results.

ACKNOWLEDGMENTS

Work supported in part by NSF, NGA, ARO, AFOSR and ANII-Uruguay. We thank Sergio Aguirre from Echopixel for important feedback regarding VC.

REFERENCES

- [1] World Health Organization. (2011) Cancer. [Online]. Available: <http://www.who.int/mediacentre/factsheets/fs297>
- [2] American Cancer Society. (2010) Colorectal cancer facts & figures 2008-2010. [Online]. Available: <http://www.cancer.org/Research/CancerFactsFigures>
- [3] D. Vining, D. Gelfand, R. Bechtold, E. Scharling, E. Grishaw, and R. Shifrin, "Technical feasibility of colon imaging with helical CT and virtual reality," *Am J Roentgenol*, vol. 162, no. Suppl, 1994.
- [4] A. Blachar, M. Graif, A. Kessler, and J. Sosna, "State-of-the-art CT colonography: Update on technique and performance," *Current Colorectal Cancer Reports*, vol. 3, no. 1, pp. 49–54, May 2007.
- [5] D. Kim, P. Pickhardt, A. Taylor, and W. Leung, "CT colonography versus colonoscopy for the detection of advanced neoplasia," *New England Journal of Medicine*, vol. 358, no. 1, p. 89; author reply 90, Jan. 2007.
- [6] J. L. Fidler, C. D. Johnson, R. L. MacCarty, T. J. Welch, A. K. Hara, and W. S. Harmsen, "Detection of flat lesions in the colon with CT colonography," *Abdom Imaging*, vol. 27, no. 3, pp. 292–300, 2002.
- [7] J. Fidler and C. Johnson, "Flat polyps of the colon: accuracy of detection by CT colonography and histologic significance," *Abdom Imaging*, vol. 34, no. 2, pp. 157–171, 2009.
- [8] R. A. Wolber and D. A. Owen, "Flat adenomas of the colon," *Hum Pathol*, vol. 22, no. 1, pp. 70–4, 1991.
- [9] A. J. Aschoff, A. S. Ernst, H.-J. Brambs, and M. S. Juchems, "CT colonography: an update," *European Radiology*, vol. 18, no. 3, pp. 429–37, Mar. 2008.
- [10] R. Summers, "CT colonography computer-aided detection: Effect on radiologist observers," in *MICCAI*, 2010.
- [11] J. H. Bond, "Colorectal cancer screening: the potential role of virtual colonoscopy," *Journal of Gastroenterology*, pp. 92–96, 2002.
- [12] M. Fiori, P. Musé, S. Aguirre, and G. Sapiro, "Automatic colon polyp flagging via geometric and texture features," in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, vol. 1. IEEE, Jan. 2010, pp. 3170–3173.
- [13] D. Vining, Y. Ge, D. Ahn, and D. Stelts, "Virtual colonoscopy with computer-assisted polyps detection," *Computer-Aided Diagnosis in Medical Imaging Elsevier Science B. V.*, pp. 445–452, 1999.
- [14] R. M. Summers, J. Yao, P. J. Pickhardt, M. Franaszek, I. Bitter, D. Brickman, V. Krishna, and J. R. Choi, "Computed tomographic virtual colonoscopy computer-aided polyp detection in a screening population," *Gastroenterology*, vol. 129, no. 6, pp. 1832–44, 2005.
- [15] H. Yoshida and J. Näppi, "Three-dimensional computer-aided diagnosis scheme for detection of colonic polyps," *IEEE Trans Med Imaging*, vol. 20, no. 12, pp. 1261–74, 2001.
- [16] D. S. Paik, C. F. Beaulieu, G. D. Rubin, B. Acar, R. B. Jeffrey, J. Yee, J. Dey, and S. Napel, "Surface normal overlap: a computer-aided detection algorithm with application to colonic polyps and lung nodules in helical CT," *IEEE Trans Med Imaging*, vol. 23, no. 6, pp. 661–75, 2004.
- [17] Z. Wang, Z. Liang, L. Li, X. Li, B. Li, J. Anderson, and D. Harrington, "Reduction of false positives by internal features for polyp detection in CT-based virtual colonoscopy," *Med Phys*, vol. 32, no. 12, pp. 3602–16, 2005.
- [18] W. Hong, F. Qiu, and A. Kaufman, "A pipeline for Computer Aided Polyp Detection," *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 5, pp. 861–868, 2006.
- [19] P. Sundaram, A. Zomorodian, C. Beaulieu, and S. Napel, "Colon polyp detection using smoothed shape operators: preliminary results," *Med Image Anal*, vol. 12, no. 2, pp. 99–119, 2008.
- [20] S. B. Gökürk, C. Tomasi, B. Acar, C. F. Beaulieu, D. S. Paik, R. B. Jeffrey, J. Yee, and S. Napel, "A statistical 3-D pattern processing method for computer-aided detection of polyps in CT colonography," *IEEE Trans Med Imaging*, vol. 20, no. 12, pp. 1251–60, 2001.
- [21] C. van Wijk, V. F. van Ravesteijn, F. M. Vos, and L. J. van Vliet, "Detection and segmentation of colonic polyps on implicit isosurfaces by second principal curvature flow," *IEEE Trans Med Imaging*, vol. 29, no. 3, pp. 688–698, 2010.
- [22] K. Suzuki, H. Yoshida, J. Nappi, S. G. Armato, and A. H. Dachman, "Mixture of expert 3D massive-training ANNs for reduction of multiple types of false positives in CAD for detection of polyps in CT colonography," *Medical Physics*, vol. 35, no. 2, p. 694, 2008.
- [23] E. Konukoglu, B. Acar, D. S. Paik, C. F. Beaulieu, J. Rosenberg, and S. Napel, "Polyp enhancing level set evolution of colon wall: method and pilot study," *IEEE Trans Med Imaging*, vol. 26, no. 12, pp. 1649–56, Dec. 2007.

- [24] L. Bogoni, P. Cathier, M. Dundar, A. Jerebko, S. Lakare, J. Liang, S. Periaswamy, M. E. Baker, and M. Macari, "Computer-aided detection (CAD) for CT colonography: A tool to address a growing need," *Br J Radiol*, vol. 78, no. suppl_1, pp. S57–62, 2005.
- [25] S. A. Taylor, S. Halligan, D. Burling, M. E. Roddie, L. Honeyfield, J. McQuillan, H. Amin, and J. Dehmeshki, "Computer-assisted reader software versus expert reviewers for polyp detection on CT colonography," *Am. J. Roentgenol.*, vol. 186, no. 3, pp. 696–702, 2006.
- [26] M. P. Wand and M. C. Jones, *Kernel Smoothing (Monographs on Statistics & Applied Probability)*. Chapman and Hall/CRC, Dec. 1994.
- [27] S. Osher and J. Sethian, "Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations," *Journal of Computational Physics*, vol. 79, pp. 12–49, 1988.
- [28] S. Osher and R. Fedkiw, *Level Set Methods and Dynamic Implicit Surfaces*. Springer, Oct. 2002.
- [29] J. Sethian, "A review of recent numerical algorithms for hypersurfaces moving with curvature-dependent speed," *Journal of Differential Geometry*, vol. 31, pp. 131–161, 1989.
- [30] V. Caselles and C. Sbert, "What is the best causal scale space for three-dimensional images?" *SIAM Journal on Applied Mathematics*, vol. 56, no. 4, pp. 1199–1246, 1996.
- [31] J. J. Koenderink, *Solid Shape*. Cambridge, USA: MIT Press, 1990.
- [32] O. Monga, S. Benayoun, and O. Faugeras, "From partial derivatives of 3-D density images to ridge lines," in *Proceedings IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 1992, pp. 354–359.
- [33] W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3D surface construction algorithm," in *SIGGRAPH '87: Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*. New York, NY, USA: ACM, 1987, pp. 163–169.
- [34] E. Meinhardt, E. Zacur, A. F. Frangi, and V. Caselles, "3D edge detection by selection of level surface patches," *Journal of Mathematical Imaging and Vision*, vol. 34, no. 1, pp. 1–16, Oct. 2008.
- [35] R. Kimmel and J. Sethian, "Computing geodesic paths on manifolds," in *Proc. Natl. Acad. Sci. USA*, 1998, pp. 8431–8435.
- [36] R. M. Haralick, K. Shanmugam, and I. Dinstein, "Textural features for image classification," *IEEE Trans Systems, Man and Cybernetics*, vol. 3, no. 6, pp. 610–621, 1973.
- [37] P. Domingos, "MetaCost: A general method for making classifiers cost-sensitive," in *Proceedings of the Fifth International Conference on Knowledge Discovery and Data Mining*, 1999, pp. 155–164.
- [38] C. Elkan, "The foundations of cost-sensitive learning," in *In Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*, 2001, pp. 973–978.