

# Enhanced Computer Assisted Detection of Polyps in CT Colonography

**Robert Sadleir B.Eng.**

Robert.Sadleir@dcu.ie

This thesis is submitted in partial fulfilment of the requirement to transfer to a  
Doctor of Philosophy (Ph.D.) research program at Dublin City University

Supervised by **Prof. Paul F. Whelan**

School of Electronic Engineering  
Dublin City University

May 2006

I hereby certify that this material, which I now submit for assessment on the programme of study leading to the award of Doctor of Philosophy (Ph.D.) is entirely my own work and has not been taken from the work of others save to the extent that such work has been cited and acknowledged within the text of my work.

Signed: \_\_\_\_\_  
Robert Sadleir (B.Eng.)

ID No.: \_\_\_\_\_

Date: \_\_\_\_\_

# Acknowledgements

I wish to thank my supervisor, Prof. Paul F. Whelan for his help, encouragement and enthusiasm throughout the course of my research in the *Vision Systems Group* (VSG). I would also like to thank my family for their continued help and support throughout my studies, with special thanks to my sister Irene for her help with proofreading.

I would like to acknowledge funding support from the Irish Cancer Society (ICS), the Health Research Board of Ireland (HRB) and the School of Electronic Engineering at Dublin City University. I would also like to acknowledge the Science Foundation Ireland (SFI), who funded the acquisition of some of the data sets used in this research.

Finally, I would like to thank my colleagues from the VSG and the Mater Misericordiae Hospital (MMH), particularly Dr. Derek Molloy, Dr. Ovidiu Ghita, Tarik Chowdhury, Dr. Kevin Robinson, Dr. Helen Fenlon, Dr. John Bruzzi, Dr. Padraic Mac Mathuna and Dr. Alan Moss.

# Contents

<b>Acknowledgements</b>	<b>iii</b>
<b>Contents</b>	<b>iv</b>
<b>List of Figures</b>	<b>ix</b>
<b>List of Tables</b>	<b>xviii</b>
<b>Glossary of Acronyms</b>	<b>xix</b>
<b>Abstract</b>	<b>xxi</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Background . . . . .	1
1.2 Motivation . . . . .	3
1.3 Contributions . . . . .	4
1.4 Organisation . . . . .	5
<b>2 Standard CTC Review</b>	<b>7</b>
2.1 Introduction . . . . .	7
2.2 CTC Technique . . . . .	7
2.2.1 Bowel Preparation . . . . .	9
2.2.2 Patient Positioning . . . . .	10
2.2.3 Incidental Extracolonic Findings . . . . .	11
2.2.4 Radiation Dose . . . . .	12
2.2.5 Image Interpretation: 2-D versus 3-D . . . . .	15
2.2.6 Contrast Enhanced CTC . . . . .	16
2.3 Issues Associated with CTC . . . . .	17
2.3.1 Risk Factors . . . . .	17
2.3.2 Image Resolution . . . . .	18
2.3.3 Patient Preference . . . . .	18
2.3.4 Reader Training . . . . .	19
2.4 Alternative Screening Techniques . . . . .	20
2.4.1 Fecal Occult Blood Test . . . . .	21
2.4.2 Barium Enema . . . . .	21
2.4.3 Flexible Sigmoidoscopy . . . . .	22
2.4.4 DNA Testing . . . . .	22
2.5 Commercial CTC systems . . . . .	23
2.6 Discussion . . . . .	24

## Contents

<b>3</b>	<b>Computer Assisted CTC Review</b>	<b>26</b>
3.1	Introduction . . . . .	26
3.2	Colon Segmentation . . . . .	26
3.2.1	Lumen Segmentation . . . . .	28
3.2.2	(Inner) Wall Segmentation . . . . .	32
3.2.3	Digital Subtraction . . . . .	35
3.2.4	Segmentation Summary . . . . .	40
3.3	Centreline Calculation . . . . .	41
3.3.1	Centreline Calculation Summary . . . . .	53
3.4	Prone/Supine Registration . . . . .	58
3.4.1	Centreline Based Registration . . . . .	58
3.4.2	Region-Growing Based Registration . . . . .	60
3.5	CAD-CTC . . . . .	60
3.5.1	The Role of CAD at CTC . . . . .	61
3.5.2	Task Division: Primary versus Secondary Techniques . . . . .	62
3.5.3	Colorectal Polyp Detection Techniques . . . . .	63
3.5.3.1	Wall Thickness Analysis . . . . .	63
3.5.3.2	Isosurface Vertex Curvature Analysis . . . . .	64
3.5.3.3	Voxel Curvature Analysis . . . . .	70
3.5.3.4	Surface Normal Analysis . . . . .	74
3.5.3.5	Perception Based Methods . . . . .	76
3.5.4	A Contrast-Enhanced Approach . . . . .	78
3.5.5	Colorectal Mass Detection Techniques . . . . .	78
3.5.6	CAD-CTC Summary . . . . .	80
3.6	Discussion . . . . .	86
<b>4</b>	<b>Preprocessing Contributions</b>	<b>87</b>
4.1	Introduction . . . . .	87
4.2	Choice of Coordinate System . . . . .	87
4.3	Volumetric Data Interpretation . . . . .	88
4.3.1	Resampling . . . . .	91
4.4	Colon Segmentation . . . . .	91
4.4.1	Data set Cropping . . . . .	93
4.4.2	Rectum Point Detection . . . . .	95
4.5	Centreline Calculation . . . . .	95
4.5.1	Caecum Point Detection . . . . .	96
4.5.1.1	Revised Caecum Point Positioning . . . . .	97
4.5.2	Skeleton generation . . . . .	98
4.5.3	Skeleton reduction (centreline generation) . . . . .	99
4.5.4	Optimisation . . . . .	99
4.5.5	The Real Case . . . . .	102
4.6	Colon Subsegmentation . . . . .	103
4.6.1	Overview . . . . .	103
4.6.2	Flexure detection . . . . .	104
4.7	Discussion . . . . .	108
<b>5</b>	<b>CAD-CTC Contributions</b>	<b>110</b>
5.1	Introduction . . . . .	110
5.2	Initial Considerations . . . . .	111
5.3	Isosurface Extraction & Vertex Classification . . . . .	113

## Contents

5.3.1	The Standard Marching Cubes Algorithm . . . . .	113
5.3.2	Topology Errors (Holes) . . . . .	119
5.3.3	Improved Normal Calculation . . . . .	119
5.3.4	Vertex Curvature Classification . . . . .	120
5.3.4.1	Discrete Implementation . . . . .	123
5.3.5	Additional Vertex Properties . . . . .	127
5.3.5.1	Colon Segment Information . . . . .	128
5.3.5.2	Vertex Density Information . . . . .	129
5.3.6	Mesh Representation . . . . .	131
5.3.7	Neighbour Identification . . . . .	133
5.3.8	Modified MCA Summary . . . . .	133
5.4	Initial Polyp Candidate Identification . . . . .	134
5.5	False Positive Reduction . . . . .	137
5.5.1	Alternative Polyp Representation . . . . .	138
5.5.1.1	Polyp Candidate Projection . . . . .	138
5.5.1.2	Polyp Candidate Fingerprinting . . . . .	142
5.5.1.3	Reformatted Polyp Candidate Visualisation . . . . .	143
5.5.2	Polyp Candidate Expansion . . . . .	146
5.5.3	Global Polyp Candidate Features . . . . .	147
5.5.3.1	Expansion . . . . .	149
5.5.3.2	Eccentricity . . . . .	149
5.5.3.3	Protrusiveness . . . . .	150
5.5.4	Polyp Candidate Feature Summary . . . . .	150
5.5.5	Polyp Candidate Classification . . . . .	151
5.5.5.1	Minimum Distance Classifier . . . . .	152
5.6	Discussion . . . . .	152
<b>6</b>	<b>System Implementation</b> . . . . .	<b>154</b>
6.1	Implementation Software . . . . .	154
6.2	Implementation Strategy . . . . .	155
6.2.1	Staggered Execution . . . . .	155
6.2.2	Well Defined Structure . . . . .	155
6.3	Data Management . . . . .	157
6.3.1	Data Set Files . . . . .	158
6.3.2	Properties File . . . . .	158
6.3.3	Miscellaneous Data Files . . . . .	159
6.4	User Interfaces . . . . .	159
6.4.1	The Main System Interface . . . . .	160
6.4.2	The Slice Viewer . . . . .	160
6.4.2.1	Polyp Marking . . . . .	161
6.4.2.2	Seed Point Selection . . . . .	163
6.4.2.3	Viewpoint Selection . . . . .	163
6.4.3	The Surface Viewer . . . . .	164
6.4.3.1	Polyp View . . . . .	164
6.4.3.2	Candidate View . . . . .	164
6.4.3.3	Fly Through . . . . .	166
6.4.4	The Fingerprint Viewer . . . . .	166
6.5	Discussion . . . . .	166

## Contents

<b>7</b>	<b>Testing and Results</b>	<b>169</b>
7.1	Test Data . . . . .	169
7.1.1	BUSM Data . . . . .	170
7.1.2	MMH Data . . . . .	171
7.1.2.1	Patient Preparation . . . . .	171
7.1.2.2	Data Acquisition . . . . .	171
7.1.3	WRAMC Data . . . . .	172
7.2	Test Platform . . . . .	172
7.3	Centreline Evaluation . . . . .	173
7.3.1	Test Data . . . . .	173
7.3.1.1	Clinical Cases . . . . .	173
7.3.1.2	Phantom Data Set . . . . .	173
7.3.2	Accuracy Evaluation . . . . .	174
7.3.2.1	Endpoint Detection Accuracy . . . . .	176
7.3.2.2	Overall Centreline Accuracy . . . . .	176
7.3.3	Performance Analysis . . . . .	177
7.3.4	Centreline Correctness . . . . .	179
7.4	Subsegmentation Evaluation . . . . .	180
7.4.1	Flexure Detection Accuracy . . . . .	180
7.4.2	Subsegmentation Correctness . . . . .	183
7.5	CAD-CTC Evaluation . . . . .	183
7.5.1	Test Data . . . . .	183
7.5.2	Primary Search Criteria . . . . .	186
7.5.2.1	Initial Evaluation . . . . .	186
7.5.2.2	Overall Findings . . . . .	189
7.5.3	Automated Scoring . . . . .	189
7.5.4	Original versus Expanded Polyp Candidates . . . . .	190
7.5.4.1	Initial (Limited) Individual Feature Evaluation . . . . .	190
7.5.4.2	Extended Individual Feature Evaluation . . . . .	190
7.5.4.3	Candidate Density as a Discriminating Feature . . . . .	194
7.5.5	Combined Feature Evaluation . . . . .	198
7.5.6	Classification of Standard Candidates . . . . .	198
7.5.7	Classification of Expanded Candidates . . . . .	198
7.5.8	Classifier Performance Characterisation . . . . .	201
7.5.9	Additional Candidate Features . . . . .	201
7.6	Discussion . . . . .	201
<b>8</b>	<b>Conclusions and Future Work</b>	<b>205</b>
8.1	Research Contributions . . . . .	207
8.1.1	Major Contributions . . . . .	207
8.1.2	Minor Contributions . . . . .	208
8.2	Future Work . . . . .	209
8.2.1	Improved Colon Lumen Segmentation . . . . .	209
8.2.2	Enhanced False Positive Reduction . . . . .	209
8.2.3	Further Evaluation . . . . .	210
8.2.4	CTC Education . . . . .	210
	<b>Bibliography</b>	<b>212</b>

## Contents

<b>A</b>	<b>DICOM Interpretation</b>	<b>230</b>
A.1	API Implementation . . . . .	233
A.2	API Overview . . . . .	234
A.2.1	Data Interpretation . . . . .	235
A.2.2	Data Management . . . . .	235
A.2.3	Data Access/Manipulation . . . . .	235
A.2.4	Data Storage . . . . .	238
A.3	Sample Applications . . . . .	238
A.3.1	Simple DICOM viewer . . . . .	238
A.3.2	Sequence Viewer . . . . .	239
A.3.3	Volume Rendering . . . . .	240
A.3.4	Anonymising Data . . . . .	240
A.3.5	Image Processing . . . . .	240
A.4	Support Material . . . . .	241
A.5	Current Status . . . . .	243
A.6	NeatMed Users . . . . .	245
A.7	Summary . . . . .	246
<b>B</b>	<b>Summary of Available Data</b>	<b>247</b>
<b>C</b>	<b>Sample Candidates</b>	<b>259</b>
C.1	Colorectal Polyps . . . . .	259
C.2	Rectal Tube . . . . .	263
C.3	Haustral Folds . . . . .	264
C.4	Miscellaneous . . . . .	265

# List of Figures

1.1	Examples of structures and objects that mimic colorectal polyps in CTC examinations (a) the Haustral folds (b) residual fecal material and (c) the rectal catheter used for bowel insufflation. The shaded patches indicate false positive detections made by a computer aided detection algorithm. . . . .	3
1.2	An overview of the developed CAD-CTC system where minor and major contributions are highlighted. . . . .	6
2.1	A case (MMH-010/supine) where a large amount of residual faecal material has caused a complete blockage of the proximal descending colon. This blockage is imaged using a modified scout x-ray where white pixels represent endoluminal air (a). An examination of the relevant 2-D axial slice from the data set of the same patient (b) confirms that this is a case of blockage and not collapse. In both cases the location of the blockage is indicated by a red box. . . . .	10
2.2	Two examples of localised residue: (a) Residual fluid near the hepatic flexure of a patient (MMH-138/prone). Note the long horizontal transition at the fluid/air interface. Also note that the fluid/soft tissue boundary cannot be identified, if a submerged polyp is present then it will be undetectable. (b) Residual fecal material in the cecum resembling multiple polyps. . . . .	11
2.3	An example of collapse (MMH-381/supine) where the lumen is divided into two distinct sections represented by the red (rectum) and green (remainder) regions in (a) and a volume rendering of the same colon again illustrating the collapse at the rectosigmoid junction. . . . .	12
2.4	A good example of a case where a polyp is obscured in one position (supine) (a) and visible in the other (prone) (b). It should be noted that this patient (WRAMC_VC-241) received orally ingested contrast material. If contrast was not used then both the polyp and the residual material would be of similar density. This would effectively hide the polyp in the supine scan. . . . .	13
2.5	Two examples demonstrating the effect of radiation dose on image quality (a & c) with a tube current of 100mAs and (b & d) with a tube current of 13mAs. Both images were generated using the VSG phantom and are courtesy of T. A. Chowdhury (VSG). The window centre and window width for both images are -550 HU and 400 HU respectively and the 13 mAs image has 2.9 times more noise than the 100 mAs image. . . . .	14
2.6	A polyp viewed using a 2-D axial image (a) and a 3-D surface rendering (b), and a fold viewed using a 2-D axial image (c) and a 3-D surface rendering (d). All images were obtained from MMH-357/prone. Note the similarities between the polyp and the fold in the 2-D images (a) & (c) . In general 2-D is used for the primary evaluation of a data set and 3-D is used for problem solving. This approach would be relevant in the case of the images presented here. . . . .	16

## List of Figures

3.1	A CTC data set contains numerous gas filled regions that include the colon lumen, sections of the small intestine, and the stomach. As a result the task of colon lumen segmentation is not a trivial task. This is especially true where the colon lumen consists of more than one segment. A segmentation consisting of all gas filled abdominal regions for a sample data set (MMH-175/supine) contains 364,732 triangles (a) whereas a segmentation of the colon lumen alone contains 189,263 triangles. . . . .	27
3.2	A histogram of a randomly selected CTC data set (MMH-120/prone). The three labeled peaks represent air (A), fatty tissue (B), and lean tissue (C). Note that a threshold value of -800 HU adequately separates air and tissue and is suitably positioned to minimise shine through artifacts. . . . .	30
3.3	An overview of the digital subtraction process. (a) A cropped axial image illustrating a section of the colon lumen containing contrast enhanced residual fluid. (b) The original image following a simple threshold based subtraction. Note the residue artifacts at the air/contrast enhanced material boundary and edge artifacts at the fluid/soft tissue boundary. (c) Removal of the residue artifacts using a subtraction mask derived from the location of strong edges in the image. (d) Reconstruction of the colonic mucosa using targeted Gaussian filtering at the contrast enhanced material/soft tissue boundary. . . . .	37
3.4	Distance Field generation: (a) A 2-D representation of the human colon. (b) A distance field initiated from a user defined seed point (green cross) in the rectum. (c) A distance from surface field initiated automatically from the set of surface points. White pixels are furthest from the source, whereas black pixels are closest to the source. . . . .	42
3.5	One endpoint (red cross) of the colon can be identified as the point with the largest distance value relative to the seed point (green cross); this process is independent of seed point location. It is evident that the same endpoint is identified with a seed point in (a) the caecum, (b) the hepatic flexure and (c) the mid transverse colon. . . . .	44
3.6	The distance from source field alone cannot be used for centreline calculation as the shortest path (i.e. the path with the shortest Euclidean distance) between the start and end points is identified (blue path). This problem, commonly referred to as corner cutting is illustrated for initial seed point in (a) the rectum and (b) the caecum. . . . .	47
3.7	Surface renderings of the colon of the same patient (MMH-166) imaged using (a) prone and (b) supine positioning. Note significant deviation in the morphology of the colon between the two images. . . . .	59
3.8	The six curvature classes: (a) peak or cap, $H < 0$ , $K > 0$ , $R = 1.0$ , (b) ridge, $H > 0$ , $K > 0$ , $R = 0.75$ , (c) saddle, $K < 0$ , $R = 0.5$ , (d) valley or rut, $H > 0$ , $K < 0$ , $R = 0.25$ , (e) pit or cup, $H > 0$ , $K > 0$ , $R = 0.0$ , (f) plane, $H = 0$ , $K = 0$ , $R$ is indeterminable (Koenderink 1990). . . . .	65
3.9	Two common sources of false positives as seen at CTC. (a) A 2-D axial slice illustrating a protruding ileocecal valve (WRAMC_VC-286/prone) and (b) a 3-D surface rendering of the rectal tube (MMH-166/supine). Note that the patient in (a) is in the prone position hence the cecum is on the right side of the image. . . . .	68
3.10	The curvedness property of a surface. Three surfaces with increasing curvedness are illustrated (a) - (c). . . . .	71
3.11	An illustration of how surface normals can be used to identify convex regions associated with potential polypoid structures (a) a flat or concave structure, surface normals converge intraluminally (b) a convex structure, surface normals converge extraluminally. . . . .	74

## List of Figures

4.1	A surface rendered model of the colon displayed in conjunction with the Java3D coordinate system. . . . .	88
4.2	The merging process where individual DICOM files are converted into a custom volumetric data set. . . . .	90
4.3	The trilinear interpolation process. The interpolated value at $v_{xyz}$ is based on the values of its eight neighbours ( $v_{000} \rightarrow v_{111}$ ) and its location with respect to these eight neighbours. . . . .	92
4.4	An illustration of the region growing technique utilised for colon lumen segmentation using a simple 2-D example. (a) The original seed point is selected, (b) the identified voxels after two iterations of the region growing algorithm and, (c) All region voxels have been identified after nine iterations of the algorithm. . . . .	94
4.5	The segmentation process: The user selects a seed points to initiate the region growing process (a). The extent of the segmentation can be monitored using a simulated scout x-ray where the segmented region has been highlighted (b). . . . .	95
4.6	The rectum point can be identified by (a) selecting the slice at the base of the data set, (b) applying a border of one voxel to the selected slice and initiating region growing from the revised origin (-1, -1, -1) and, (c) thresholding the image and obtaining the centroid of resulting blob. . . . .	96
4.7	An overview of the centreline calculation process. (a) The original CTC data set, (b) the segmented colon lumen, (c) the skeleton and, (d) The final colon centreline. . . . .	97
4.8	Magnified sections of the colon lumen from Figure 4.7 (c) and (d) illustrating how extraneous loops can be caused by the Haustral folds (a) and how the removal of the these loops yields the final centreline (b). . . . .	100
4.9	A 2-D example illustrating the skeleton reduction process. (a) Original binary object (region of interest indicated using white); (b) Skeleton and distance from surface field generation; (c) Inverted distance from surface field (skeleton pixels only); (d) weighted graph (originating from the right); (e) The minimum cost path found using backward propagation through the distance field. . . . .	101
4.10	A 2-D illustration of the LUT index generation process. . . . .	102
4.11	In cases where the colon consists of multiple segments, a centreline segment must be calculated for each individual colon segment and the resulting centreline segments must be connected to yield the final centreline that extends between the endpoints located in the rectum and the caecum. . . . .	103
4.12	The anatomy of the colon. (a) An illustration of the ideal colon where all colonic segments and flexures have been identified. (b) An example of an actual human colon which has been segmented from a CTC data set. . . . .	104
4.13	Flexure detection theory: (a) The centreline of the ideal colon (see Fig. 4.12 (a)) and (b) a graph of adjusted values of $\theta_p$ for each point $p$ in the centreline. . . . .	105
4.14	Six iterations of the centreline filtering process illustrating the transition from the original distorted centreline (a) to an approximation of the original centreline where all of the kinks and loops have been removed (f). . . . .	106
4.15	An illustration of the flexure detection process. (a) The original colon centreline. (b) A filtered version of the original centreline ( $W = 319$ ). (c) A plot of $\theta_p$ against $p$ for the filtered centreline. Note that the hepatic flexure is the last point to be uniquely identified. [Flexure key: $\nabla = P_{rec}$ , $\triangle = P_{sig}$ , $\circ = P_{spl}$ & $\diamond = P_{hep}$ ] . . . . .	107

## List of Figures

4.16	The output of the subsegmentation process consists of a colon model where voxels are labeled as (a) the rectum, (b) the sigmoid colon, (c) the descending colon, (d) the transverse colon and, (e) the ascending colon. . . . .	108
5.1	The 15 possible configurations of eight binary voxels arranged in a cubic formation. The relevant surface patches have been inserted as documented in the original marching cubes algorithm specification (Lorensen & Cline 1987). . . . .	115
5.2	A sample voxel configuration where $v_1, v_5, v_6$ & $v_8$ are inside the isosurface and all other voxels are outside. An LUT index of 0xB1 is generated which results in the specified edge list. Note that the order of the edges is important as it defines the front face for the associated triangle in Java3D. . . . .	116
5.3	The calculation of $\delta_{iso}$ at three sample boundaries. In each case, the value of $\delta_{iso}$ is calculated using linear interpolation and the density at the point represented by $d_{iso}$ is -800, i.e. the isosurface density $d_{iso}$ . . . . .	116
5.4	The three masks (a, b & c) that are used by Lorensen & Cline (1987) to calculate the edge magnitudes in the $x, y$ & $z$ directions respectively. Note that the scaling factors are omitted as the implementation discussed here is intended for use with isometric data. . . . .	117
5.5	An illustration of an isosurface model for the colonic mucosa before (a) and after (b) the use of vertex and normal interpolation. . . . .	118
5.6	An example of the ambiguous case that results when the complement of configuration three (Figure 5.1 (d)) occurs next to configuration six (Figure 5.1 (g)). A hole is evident at the interface between these two cubes. . . . .	119
5.7	The eight additional configurations of eight binary voxels arranged in a cubic formation as defined by Nielson (2003). These extra configurations remove the need to generate complementary cases and as a result solve the topology problem associated with the original MCA. . . . .	120
5.8	The same isosurface extracted using the 15 core neighbourhood configurations (a) and the extended 23 neighbourhood configurations (b). . . . .	121
5.9	The three 26-neighbour masks representing the Zucker-Hummel edge operator where $m_a = 1.0$ , $m_b = \frac{\sqrt{2}}{2}$ & $m_c = \frac{\sqrt{3}}{3}$ . . . . .	122
5.10	The Gaussian filter $g_0(x)$ (a) and its first and second derivatives $g_1(x)$ (b) and $g_2(x)$ (c). The standard deviation $\sigma$ is set to 1.0 as recommended by Lohmann (1998). . . . .	125
5.11	The Deriche filter $d_0(x)$ (a) and its first and second derivatives $d_1(x)$ (b) and $d_2(x)$ (c). The values for $\alpha_1$ and $\alpha_2$ are set to 0.7 and 0.1 respectively as suggested by Campbell & Summers (2004) and the normalisation coefficients $c_0, c_1, c_2$ & $c_3$ are calculated using Equation 5.3.19 as suggested by Monga & Benayoun (1995). . . . .	127
5.12	An example of the segment skipping problem (a) the expected case where colon wall vertices are associated with centreline segments in the same segment (b) the actual case where vertices associate themselves with the closest centreline segment resulting in incorrect segment assignments. . . . .	129
5.13	The anatomical vertex labeling process. (a) Shows the wall at the hepatic flexure and the region growing process after three iterations. (b) Shows the completed region growing process where a smooth transition exists between the transverse colon (purple) and the ascending colon (orange). (c) Shows an exterior view of the labeled colon and, (d) shows an endoluminal view of the splenic flexure. . . . .	130

## List of Figures

5.14	An illustration of the density sampling process. (a) The interconnected vertices that result from the surface extraction algorithm. (b) The corresponding density samples (or virtual biopsies) are located at a fixed distance (the sampling depth) from each vertex in the opposite direction of the normal at that point. . . . .	131
5.15	An illustration of how indexing can be used to reduce the amount of data required to specify a mesh structure. A pyramid consists of five vertices and six triangles. The unoptimised representation yields 54 values (a) and the optimised alternative yields only 33 (b). Note that in Java the data type for vertex (float) and index (int) both require four bytes of storage space. . . . .	132
5.16	The neighbour indexing process: The neighbouring vertices for each mesh vertex are stored in a list to streamline the process of searching for vertex neighbours. . . . .	133
5.17	External (a) and internal (b) surface renderings of a sample CTC data set (MMH-765/supine) demonstrating the type of output generated by the initial candidate detection stage. Detected surface patches that have highlighted in red. This is an extreme example where there are a total of 562 initial polyp candidates. . . . .	135
5.18	Hysteresis thresholding based region growing in a polygonal mesh. An initial seed point is selected that belongs to the polyp cap class $V_s$ (a). Neighbouring vertices belonging to either the polyp cap or polyp collar classes are added to the region (b) & (c). The region growing process completes when no more suitable neighbours are available (d). Note that the region growing process must start from a vertex of the polyp cap class. . . . .	137
5.19	Examples of a true polyp and false positive detection generated by the initial candidate detection stage. (a) A 2-D axial slice containing a polyp. (b) A 3-D surface rendering of the region containing the polyp. (c) The section of the polyp identified by the initial candidate detection stage. (d) & (e) Irregular shaped false positive detections caused by regions of the colon surface with polyplike properties. . . . .	139
5.20	Alternative polyp representations. (a) Polyp candidate projection, where each candidate vertex is projected onto the plane in the direction of the normal to the surface of the plane. (b) Polyp candidate fingerprinting, an extension to polyp candidate projection, where a flattened representation of the surface, that preserves the original vertex layout, is generated. . . . .	140
5.21	The smallest angle between a reference vector $\vec{r}$ and another vector $\vec{n}$ can be calculated using the dot product. The resulting angle can be either clockwise (a) or anticlockwise (b) with respect to the reference vector. . . . .	142
5.22	Calculating the anticlockwise angle from a reference vector $u$ to another vector $v$ . (a) If the cross product $(u \times v)$ is in the same direction as the normal to the surface of the plane then the required angle $\theta$ is obtained using the dot product. (b) However, if the cross product is in the opposite direction to the normal then the anticlockwise angle is given by $2\pi - \theta$ . . . . .	143
5.23	An illustration of the different methods used for obtaining the radial component of the polar coordinates for a vertex being mapped to a 2-D plane (a) Polyp candidate projection (b) Polyp candidate fingerprinting. . . . .	144
5.24	The six different fingerprint representations for a polyp candidate (a) depth map, (b) density, (c) curvedness, (d) Gaussian curvature, (e) mean curvature and (f) shape index. In all cases black and white are used to represent the minimum and maximum values for the relevant property. . . . .	145

## List of Figures

- 5.25 The candidate expansion process (a) The original candidate regions consisting of a connected set of vertices (b) the minimum bounding ellipse of the original vertices and the set of additional vertices that it encapsulates. The combined set of vertices represents the expanded candidate region. . . . 147
- 5.26 Expanded versions of the fingerprints from Figure 5.24 (a) depth map, (b) density, (c) curvedness, (d) Gaussian curvature, (e) mean curvature and (f) shape index. The red points are the original polyp candidate points and the green points are the expanded polyp candidate points. . . . . 148
- 6.1 An illustration of the data flow through the completed system. Each system operation is represented by an individual block and the interconnections represent the data shared between the different operations. . . . . 156
- 6.2 The underlying file structure of the final system. Each patient has an individual directory under which the prone and supine data sets are stored. Each data set directory contains the relevant intermediate files generated by the system as well as a directory containing the original data set files in DICOM format. . . . . 157
- 6.3 An illustration of the custom data set format. The header is extensible and supports the addition of new fields. The data can be assessed using a range of methods. These methods provide direct access to the voxel data and the associated flag values. . . . . 158
- 6.4 The format of a file representing an array of `float` primitives that was created using a custom `OutputStream` class. First the constant header is written, then four zero bytes are written for the length field. All of the `float` values are written to the file and upon completion the length field is overwritten with the actual number of entries. . . . . 160
- 6.5 The main user interface for the completed system allowed selected data sets to be processed using selected operations. . . . . 161
- 6.6 Four possible visualisation modes of the slice viewer. (a) Polyp marking (interactive), (b) Segmentation (interactive), (c) Centreline calculation (display only), and (d) density calculation (display only). . . . . 162
- 6.7 The different modes of operation of the surface viewer. (a) The default external view of the segmented colon lumen. (b) The polyp view where each actual polyp can be viewed. (c) The candidate view where each candidate can be visualised and classified. (d) The fly through mode where the user is automatically guided through the surface model of the colon using the set of points generated during the centreline calculation stage. . . . . 165
- 6.8 The different modes of operation of the fingerprint viewer. (a) a sample polyp candidate (in this case an actual polyp) displayed using the 3-D surface viewer. (b) The uninterpolated fingerprint points associated with the candidate. (c) The fingerprint points interpolated with nearest neighbour interpolation. (d) The fingerprint points interpolated using inverse distance weighting based interpolation. (e) The superimposition of the original candidate points. (f) The superimposition of the minimum bounding ellipse and the expanded candidate points. . . . . 167
- 7.1 A series of sample images selected from data sets obtained from the BUSM (a) (BUSM-001/supine), the MMH (b) (MMH-640/supine) and from the WRAMC (c) (WRAMC-239/supine). There is usually a small amount of residue in the data sets obtained from the BUSM and the MMH. The samples in (a) & (b) were selected specifically to illustrate this point. This residue is approximately the same density as soft tissue. In the case of the WRAMC data, there is a significant amount of contrast enhanced residue. This is evident from the sample illustrated in (c). . . . . 170

## List of Figures

7.2	An illustration of the four computer generated synthetic phantom data sets that were used for centreline evaluation. . . . .	175
7.3	A comparison of the actual and calculated centrelines from an endoluminal perspective. (a) The two centrelines closely track each other but never converge. (b) The calculated centreline converges with the actual centreline. Note that in both cases the standard deviation between the actual and calculated centrelines is less than 1 mm. . . . .	177
7.4	A comparison of the actual and calculated centreline, for each of the phantoms, from an external perspective. . . . .	178
7.5	Examples of the centreline displayed in conjunction with transparently rendered colons for colons with 3, 4, 2 and 1 segments, respectively. . . . .	181
7.6	The two cases (out of 20) where centreline calculation failed. (a) A shine through artefact that resulted in the caecum point being detected in the mid ascending colon. (b) Disjoined centreline segments were incorrectly connected due to a large collapse in the sigmoid colon. . . . .	182
7.7	Examples of the output from the subsegmentation process for colons with 3, 4, 2 and 1 segments respectively. Note: These are from the same data sets as the centrelines that are illustrated in Figure 7.5. . . . .	184
7.8	An illustration of the same polyp as it appears in the prone (a) and (c) and supine (b) and (d) data sets of a patient (MMH-393). It is evident that there are significant variations in the perceived polyp size and polyp morphology between the two positions. . . . .	185
7.9	Internal and external renderings of the same data set (MMH-708/supine) where initial candidate detection has performed using the <i>H/K</i> features (a) & (c) and <i>R/S</i> features (b) & (d). It is clear from a visual inspection of these renderings that the <i>R/S</i> features generate far fewer polyp candidate regions. . . . .	188
7.10	A comparison of polyp coverage of the same set of polyps selected from MMH-017/supine using <i>H/K</i> features (a) - (d) and <i>R/S</i> features (e) - (h) initial candidate identification techniques. Note: The polyp sizes from left to right are 5.25, 3.0, 3.0 and 5.5 mm respectively. . . . .	189
7.11	A comparative evaluation of the original versus expanded shape index in a single patient study: (a) A scattergram representing the distribution of the original candidates. (b) A scattergram representing the distribution of the expanded candidates. (c) A ROC analysis comparing the performance of the original and expanded candidates. Note: At 100% sensitivity the expanded candidates provide a 32% reduction in false positive detections. . . . .	191
7.12	A comparative evaluation of the original versus expanded sphericity in a single patient study: (a) A scattergram representing the distribution of the original candidates. (b) A scattergram representing the distribution of the expanded candidates. (c) A ROC analysis comparing the performance of the original and expanded candidates. Note: At 100% sensitivity the expanded candidates provide a 30% reduction in false positive detections. . . . .	192
7.13	A comparative evaluation of the original versus expanded density in a single patient study: (a) A scattergram representing the distribution of the original candidates. (b) A scattergram representing the distribution of the expanded candidates. (c) A ROC analysis comparing the performance of the original and expanded candidates. Note: The density feature does not provide a good discrimination between true and false positives. . . . .	193

## List of Figures

7.14	A large polyp located on a fold (a) and its associated fingerprint (b) (from MMH-483/supine). The irregular boundary of the polyp causes the MBE to include points that belong to the fold. This dilutes the polyp's features e.g. the mean shape index drops from 0.92 in the original candidate to 0.81 in the expanded candidate. . . . .	194
7.15	A comparative evaluation of the original versus expanded shape index in a five patient study: (a) A scattergram representing the distribution of the original candidates. (b) A scattergram representing the distribution of the expanded candidates. (c) A ROC analysis comparing the performance of the original and expanded candidates. Note: At 100% sensitivity the expanded candidates caused a 137% increase in false positive detections. . . . .	195
7.16	A comparative evaluation of the original versus expanded sphericity in a five patient study: (a) A scattergram representing the distribution of the original candidates. (b) A scattergram representing the distribution of the expanded candidates. (c) A ROC analysis comparing the performance of the original and expanded candidates. Note: A sensitivity of 100% was not achievable in this experiment with expanded candidates. . . . .	196
7.17	A comparative evaluation of the original versus expanded density in a five patient study: (a) A scattergram representing the distribution of the original candidates. (b) A scattergram representing the distribution of the expanded candidates. (c) A ROC analysis comparing the performance of the original and expanded candidates. Note: The density feature does not provide a good discrimination between true and false positives. . . . .	197
7.18	A plot of polyp detection sensitivity versus increasing candidate expansion size threshold. Note that the system performance remains largely unaffected up to a size value of 5.6 mm. . . . .	200
7.19	A plot of false positive detection rate versus increasing candidate expansion size threshold. Note that the false positive detection rate at the optimum candidate expansion size threshold (5.6 mm) is 12.0 per data set. . . . .	200
7.20	A ROC curve comparing polyp detection sensitivity and false positive detection rate for clinically significant polyps $\geq 5$ mm in size. Note that at 97% sensitivity, original candidates generated 20.4 false positives per data set whereas the use of expanded candidates generated 12.0 (a 41.2% reduction). The use of the additional features calculated during the candidate expansion process enabled a further reduction in the false positive rate to a minimum of 8.6 per data set (a 57.8% reduction). . . . .	202
A.1	An overview of how the NeatMed API should be used in the development of medical imaging applications. . . . .	231
A.2	A representation of how data is stored and queried at the most basic level of the <code>DICOMImage</code> class. . . . .	236
A.3	A volume rendering of the human colon generated from an abdominal CTC data set using a combination of the NeatMed API and the VTK. . . . .	241
A.4	The threshold image processing operation. A sample axial CTC slice (a) before and (b) after application of the threshold. . . . .	242
A.5	Advanced image processing facilitated by NeatMed (a) Centreline calculation at CT colonography. (b) Segmentation of the biliary tree from MRCP data (Image courtesy of K. Robinson (VSG)). . . . .	243
A.6	The HTML Javadoc documentation for the NeatMed API. This documentation can be accessed online at <a href="http://www.eeng.dcu.ie/vsl/DICOM/">http://www.eeng.dcu.ie/vsl/DICOM/</a> . . . . .	244
B.1	MMH data sets. . . . .	256
B.2	MMH data sets (continued). . . . .	257
B.3	MMH data sets (continued). . . . .	258

## List of Figures

C.1	A 6 mm polyp from MMH-054/prone. The mean shape index of the original candidate is 0.924 and the mean shape index of the expanded candidate is 0.920 (a 0.4% reduction). . . . .	259
C.2	A 3 mm polyp from MMH-208/prone. The mean shape index of the original candidate is 0.896 and the mean shape index of the expanded candidate is 0.879 (a 1.9% reduction). . . . .	260
C.3	A 4 mm polyp from MMH-161/prone. The mean shape index of the original candidate is 0.859 and the mean shape index of the expanded candidate is 0.856 (a 0.3% reduction). . . . .	260
C.4	A 6 mm polyp from MMH-245/prone. The mean shape index of the original candidate is 0.899 and the mean shape index of the expanded candidate is 0.886 (a 1.4% reduction). . . . .	261
C.5	A 7 mm polyp from MMH-357/supine. The mean shape index of the original candidate is 0.892 and the mean shape index of the expanded candidate is 0.862 (a 3.3% reduction). . . . .	261
C.6	A 3 mm polyp from MMH-708/supine. The mean shape index of the original candidate is 0.923 and the mean shape index of the expanded candidate is 0.902 (a 2.2% reduction). . . . .	262
C.7	A 4 mm polyp from MMH-735/prone. The mean shape index of the original candidate is 0.919 and the mean shape index of the expanded candidate is 0.900 (a 2.0% reduction). . . . .	262
C.8	A polyp candidate from MMH-166/prone that is attributable to the rectal tube. The mean shape index of the original candidate is 0.866 and the mean shape index of the expanded candidate is 0.784 (a 9.4% reduction). . . . .	263
C.9	A polyp candidate from MMH-263/prone that is attributable to the rectal tube. The mean shape index of the original candidate is 0.882 and the mean shape index of the expanded candidate is 0.806 (a 8.6% reduction). . . . .	263
C.10	A polyp candidate from MMH-081/supine that is attributable a haustral fold. The mean shape index of the original candidate is 0.854 and the mean shape index of the expanded candidate is 0.816 (a 4.4% reduction). . . . .	264
C.11	A polyp candidate from MMH-139/supine that is attributable a haustral fold. The mean shape index of the original candidate is 0.882 and the mean shape index of the expanded candidate is 0.819 (a 7.5% reduction). . . . .	264
C.12	A polyp candidate from MMH-708/supine that is attributable to a patch of the colon surface. The mean shape index of the original candidate is 0.863 and the mean shape index of the expanded candidate is 0.778 (a 9.8% reduction). . . . .	265
C.13	A polyp candidate from MMH-526/Prone that is attributable to a patch of the colon surface. The mean shape index of the original candidate is 0.844 and the mean shape index of the expanded candidate is 0.762 (a 9.7% reduction). . . . .	265

# List of Tables

2.1	The results of large scale studies (involving $\geq 100$ patients) that compare the performance of manually read CTC to conventional colonoscopy for the detection of colorectal polyps. Note that the extremely low overall sensitivity reported by Cotton et al. (2004) is due to the high frequency of small polyps in their study, combined with the low sensitivity for the detection of these small polyps. . . . .	8
3.1	An overview of previously published centreline calculation algorithms. A rank has been calculated for each algorithm by scaling the reported execution time based on the processor speed. This is intended to provide an indication of algorithm performance that is independent of the hardware platform used. Instances where the required information is unavailable are indicated by '-'. The accuracy of each technique, compared to a known phantom centreline (P) and/or radiologist markings (R), is presented where available. The order in which the techniques are listed is based on the associated centreline calculation speed. . . . .	55
3.2	A summary of previously published CAD-CTC techniques organised chronologically (by submission date for journals and by meeting date for conferences). Instances where the required information is unavailable are indicated by '-'. . . . .	83
6.1	The properties supported by the properties file. These are essentially the global variables of the system and can be accessed by any of the system components. . . . .	159
7.1	The data sets used for centreline evaluation. The average number of segments in this test set was 2.25 and the average number of lumen voxels was 5,521,687. . . . .	174
7.2	The endpoint detection error distances for each of the four phantom studies. The mean endpoint detection error was approximately 2 mm. . . . .	176
7.3	The results of the overall accuracy evaluations for the four phantom studies. The average standard deviation between the actual and calculated centrelines was less than 1 mm. . . . .	177
7.4	Detailed results comparing the standard and optimised approaches to centreline calculation. Note: The optimisation technique only affected the skeleton generation stage of the process. The time required for all other stages remained the same as the unoptimised approach. . . . .	179
7.5	The results of a comparative evaluation of the <i>R/S</i> and <i>H/K</i> features for initial polyp candidate detection using 10 randomly selected CTC data sets with a total of 27 polyps. . . . .	187
B.1	A brief summary of the data that was used in this study. . . . .	248
B.2	A detailed summary of the patient data, provided by the MMH, that used for system training and evaluation. . . . .	249

# Glossary of Acronyms

## Technical Acronyms

<b>Acronym</b>	<b>Explanation</b>
1-D	<i>One Dimensional</i>
2-D	<i>Two Dimensional</i>
3-D	<i>Three Dimensional</i>
3-DTT	<i>Three Dimensional Topological Thinning</i>
AOF	<i>Average Outward Flux</i>
API	<i>Application Programmers Interface</i>
CAD	<i>Computer Assisted Diagnosis</i>
CAR	<i>Computer-Assisted Reader</i>
CT	<i>Computed Tomography</i>
CV	<i>Curvedness</i>
CTC	<i>Computer Tomography Colonography</i>
DBCE	<i>Double Contrast Barium Enema</i>
DICOM	<i>Digital Imaging and Communications in Medicine</i>
DSC	<i>Digital Subtraction Colonography</i>
FOBT	<i>Fecal Occult Blood Test</i>
FROC	<i>Free-Response Receiver Operator Characteristic</i>
GUI	<i>Graphical User Interface</i>
HFP	<i>Head First Prone</i>
HFS	<i>Head First Supine</i>
HU	<i>Hounsfield Unit</i>
HT	<i>Hough Transform</i>
HTML	<i>Hypertext Mark-up Language</i>
ICV	<i>Ileocecal Valve</i>
IDE	<i>Integrated Development Environment</i>
IDW	<i>Inverse Distance Weighted</i>
JDK	<i>Java Development Kit</i>
JVM	<i>Java Virtual Machine</i>
JPEG	<i>Joint Photographic Experts Group</i>
JRE	<i>Java Runtime Environment</i>
MAP	<i>Maximum a Posteriori (classifier)</i>
MBE	<i>Minimum Bounding Ellipse</i>
MCA	<i>Marching Cubes Algorithm</i>
MDC	<i>Minimum Distance Classifier</i>
MPLD	<i>Multi-Plane Linear Discriminant (classifier)</i>
MPR	<i>Multiplanar Reformatted</i>
MRC	<i>Magnetic Resonance Colonography</i>
MRCPC	<i>Magnetic Resonance Cholangiopancreatography</i>

**Acronym Explanation**

MRI	<i>Magnetic Resonance Imaging</i>
PET	<i>Positron Emission Tomography</i>
PVE	<i>Partial Volume Effect</i>
RLE	<i>Run Length Encoding</i>
ROC	<i>Receiver Operator Characteristic</i>
SNO	<i>Surface Normal Overlap</i>
SI	<i>Shape Index</i>
VC	<i>Virtual Colonoscopy</i>
VR	<i>Virtual Reality</i>
VTK	<i>Visualisation Toolkit</i>

**Non-technical Acronyms****Acronym Explanation**

ACRIN	<i>American College of Radiology Imaging Network</i>
ACR-NEMA	<i>American College of Radiology - National Electrical Manufacturers Association</i>
BUSM	<i>Boston University School of Medicine</i>
CSIRO	<i>Commonwealth Scientific and Industrial Research Organisation</i>
MMH	<i>Mater Misericordiae Hospital</i>
NIH	<i>National Institutes of Health</i>
SUNY	<i>State University of New York</i>
UCLA	<i>University of California, Los Angeles</i>
UC	<i>University of Chicago</i>
UHG	<i>University Hospital Gasthuisberg (Leuven, Belgium)</i>
VSG	<i>Vision Systems Group</i>
WRAMC	<i>Walter Reed Army Medical Centre</i>
WUSM	<i>Washington University School of Medicine</i>

# Abstract

## **Enhanced Computer Assisted Detection of Polyps in CT Colonography**

Robert Sadleir

Under the supervision of Prof. Paul F. Whelan  
at Dublin City University

Submitted in partial fulfilment of the requirements for the degree  
of Doctor of Philosophy Dublin City University, 2006

This thesis presents a novel technique for automatically detecting colorectal polyps in *CT colonography* (CTC). The objective of the documented *computer assisted diagnosis* (CAD) technique is to deal with the issue of false positive detections without adversely affecting polyp detection sensitivity. The thesis begins with an overview of CTC and a review of the associated research areas, with particular attention given to CAD-CTC. This review identifies excessive false positive detections as a common problem associated with current CAD-CTC techniques. Addressing this problem constitutes the major contribution of this thesis. The documented CAD-CTC technique is trained with, and evaluated using a series of clinical CTC data sets. These data sets contain polyps with a range of different sizes and morphologies. The results presented in this thesis indicate the validity of the developed CAD-CTC technique and demonstrate its effectiveness in accurately detecting colorectal polyps while significantly reducing the number of false positive detections.

# Chapter 1

## Introduction

### 1.1 Background

Colorectal cancer is the second most common cause of cancer-related death worldwide. For example, the latest statistics published by the *National Cancer Registry of Ireland* (NCRI 2005) show that, in Ireland, colorectal cancer accounted for 14,571 new cancer cases, and resulted in 7,439 deaths, during the period 1994 - 2001. Although the cause of colorectal cancer is as yet unknown, the landmark paper by Burkitt (1971) suggests that this type of cancer is linked to the lifestyle adopted by people in the developed world.

Colorectal cancer is preceded by the development of benign colorectal polyps and can be prevented in the majority of cases if precursor polyps are identified early in their course and are successfully resected. There are a number of screening techniques currently available, the most sensitive of these is conventional colonoscopy. This is an endoscopic procedure that is used to view the interior of the large intestine of a patient using an instrument known as a colonoscope. During the examination the patient lies on their side with their knees drawn up toward the abdomen. The colonoscope is advanced from the rectum to the cecum and then withdrawn. A more detailed examination of the colonic mucosa can be performed during withdrawal. The colonoscope captures endoluminal images of the colon and rectum. These images are displayed on a monitor and evaluated by a gastroenterologist. Any occlusive residual fluid can be removed using suction and any significant polyps that are encountered can be removed using an electrocautery snare. Thorough cleansing of the bowel is required prior to a colonoscopy examination. This generally involves a special low residue diet for two days before the examination combined with the use of laxatives to completely empty the colon before colonoscopy. The amount of

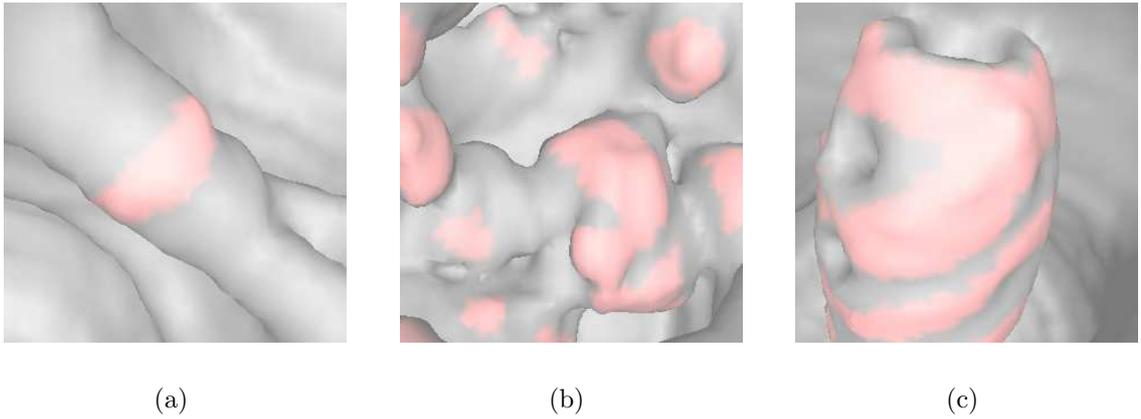
pain experienced by the patient is reduced using an analgesic (pain killer), and a sedative is also used to reduce patient discomfort.

There are a number of problems associated with conventional colonoscopy. The patient must undergo a rigorous bowel preparation immediately prior to the colonoscopy examination. This involves the use of laxatives and is both uncomfortable and embarrassing for the patient. The patient needs to be sedated before the colonoscope is inserted. As a result, the patient is unconscious for a period of time as well as being unable to function normally while recovering from the effects of the anesthetic. Conventional colonoscopy has a further embarrassment factor that is due to the nature of the examination. Although patients are under the effects of an anesthetic and have a high level of post-procedural amnesia, they still cite the invasive nature of the examination as being a major deterrent. Finally, there are certain risks associated with performing a colonoscopy examination. Due to the highly invasive nature of the procedure, it is possible to inadvertently perforate the colon which can in extreme cases, result in patient mortality. In their retrospective evaluation of 26,162 consecutive colonoscopy examinations, Tran et al. (2001) found that 21 of the cases required operations for perforation and that one patient died as a result of perforation.

*CT colonography* (CTC) (originally known as *virtual colonoscopy* (VC)<sup>1</sup>) is a minimally invasive alternative to conventional colonoscopy for colon imaging. The technique was originally demonstrated by Vining et al. (1994), however the concept of imaging colonic polyps using radiological techniques was proposed over ten years earlier by Coin et al. (1983). Using this technique a *computed tomography* (CT) scan of the patient's abdomen is performed. The patient must undergo the same bowel preparation that is required for conventional colonoscopy and, in addition, the patient's colon must be inflated with room air or CO<sub>2</sub>. The resulting CT data set is then examined either in *two-dimensions* (2-D) using the raw axial CT images or alternatively in *three-dimensions* (3-D) by generating a *virtual reality* (VR) model of the colon using the axial CT slices. This model can be evaluated in much the same way that a gastroenterologist examines images of the colon obtained using a colonoscope. If any abnormalities are discovered at CTC, the patient is referred for follow-up using conventional colonoscopy. This screening approach has the potential to substantially reduce the number of patients requiring conventional colonoscopy.

---

<sup>1</sup>Johnson et al. (1998) presented the justification behind the choice of nomenclature.



**Figure 1.1:** Examples of structures and objects that mimic colorectal polyps in CTC examinations (a) the Haustral folds (b) residual fecal material and (c) the rectal catheter used for bowel insufflation. The shaded patches indicate false positive detections made by a computer aided detection algorithm.

An additional benefit associated with CTC is that it can be used to image the entire colon in cases where conventional colonoscopy fails (Macari et al. 1999).

## 1.2 Motivation

The most active area of CTC research deals with the computer aided detection of colorectal polyps. A number of techniques dealing with this problem have been described in the literature, these techniques are summarised and compared where possible in Section 3.5. The task of automatically identifying polyps at CTC is extremely difficult due to the variety of polyp morphologies that may occur, and because other structures that are present in the colon at CTC may mimic the shape of polyps and generate false positive detections, see Figure 1.1. Dachman et al. (2002) cite nine potential causes of false positives: Prominent fold, small bowel/stomach, ileocecal valve, stool, rectal tube, flexular pseudotumor, motion artifact, anorectal junction and extrinsic compression. However, in their study they found that prominent folds (45%) and residual stool (25%) were the two main sources of false positives. Although the problem of automatically detecting colorectal polyps has been largely solved, the problem of excessive false positive detections still warrants attention.

## 1.3 Contributions

Yoshida & Dachman (2005) describe computer aided detection of polyps at CTC in terms of three main tasks:

1. **Colon lumen segmentation** - Separating the voxels representing the colon lumen from the entire CTC data set.
2. **Polyp candidate detection** - Identifying regions of the colon surface that have polyplike properties.
3. **False positive reduction** - Reducing the list of polyp candidates based on a more detailed analysis of the suspicious regions.

The major contribution of this thesis is to address the task of false positive reduction by introducing a novel CAD-CTC system that incorporates a series of innovative enhancements. The reduction in false positive detections is achieved by including an additional stage between stages two and three listed above. The purpose of the additional stage is to:

1. Generate an alternative polyp candidate representation. This involves creating a 2-D representation of the 3-D patch of the colon surface that is associated with a polyp candidate.
2. Expand the polyp candidate in a structured manner, based on an analysis of the alternative candidate representation, so that the resulting candidate region is more representative of an actual colorectal polyp.

It will be shown that the inclusion of this additional stage can significantly reduce the number of false positive detections generated by the CAD-CTC system. Other contributions deal with the associated tasks that are required to realise an innovative CAD-CTC system that incorporates a novel automated reporting scheme. These contributions include:

1. **A medical imaging interpretation toolkit** - A custom toolkit that can be used to interpret medical image data encoded according to both the DICOM and Analyze standards (Sadleir et al. 2004a). This toolkit can be used to simplify the task of medical imaging application development.

2. **A real time colon centreline calculation algorithm** - A novel centreline calculation technique based on 3-D topological thinning that can be used to accurately extract the central path through the colon in a computationally efficient manner. (Sadleir & Whelan 2005).
3. **A novel technique for colon lumen subsegmentation** - A technique that divides the centreline so that it represents the five main anatomical regions of the colon: Rectum, sigmoid colon, descending colon, transverse colon and ascending colon (Sadleir et al. 2002). Using this information, the location of an automatically detected polyp can be localised to a particular colon segment, enabling improved reporting of results to the examining radiologists.

An overview of the developed CAD-CTC system is presented in Figure 1.2.

## 1.4 Organisation

The remainder of this thesis is organised as follows: Chapter 2 provides an overview of the manually read CTC technique. Chapter 3 reviews relevant research dealing with the automation of different aspects of CTC<sup>2</sup>. Chapter 4 describes the preprocessing operations that were developed to facilitate high performance CAD-CTC with automated reporting. Chapter 5 provides an overview of a CAD technique that was developed to deal specifically with the reduction of false positive detections. Chapter 6 provides a system overview and discusses some of the implementation issues that were encountered. Chapter 7 describes the testing strategy that was employed and presents the test results that were obtained. Finally, Chapter 8 summarises the main points of this thesis and suggests future research directions.

---

<sup>2</sup>In situations where multiple papers are discussed in a common thread the order in which journal papers are discussed will be determined by their respective submission dates. Conference papers will be discussed based on the date of the conference. This approach is adopted to ensure continuity when discussing the work of a particular author or group. In addition, references will be augmented to include the group that the author belonged to at the time of publication. In cases where the group name is represented by an acronym, the Glossary of Acronyms, located in the preamble to this thesis, should be consulted.

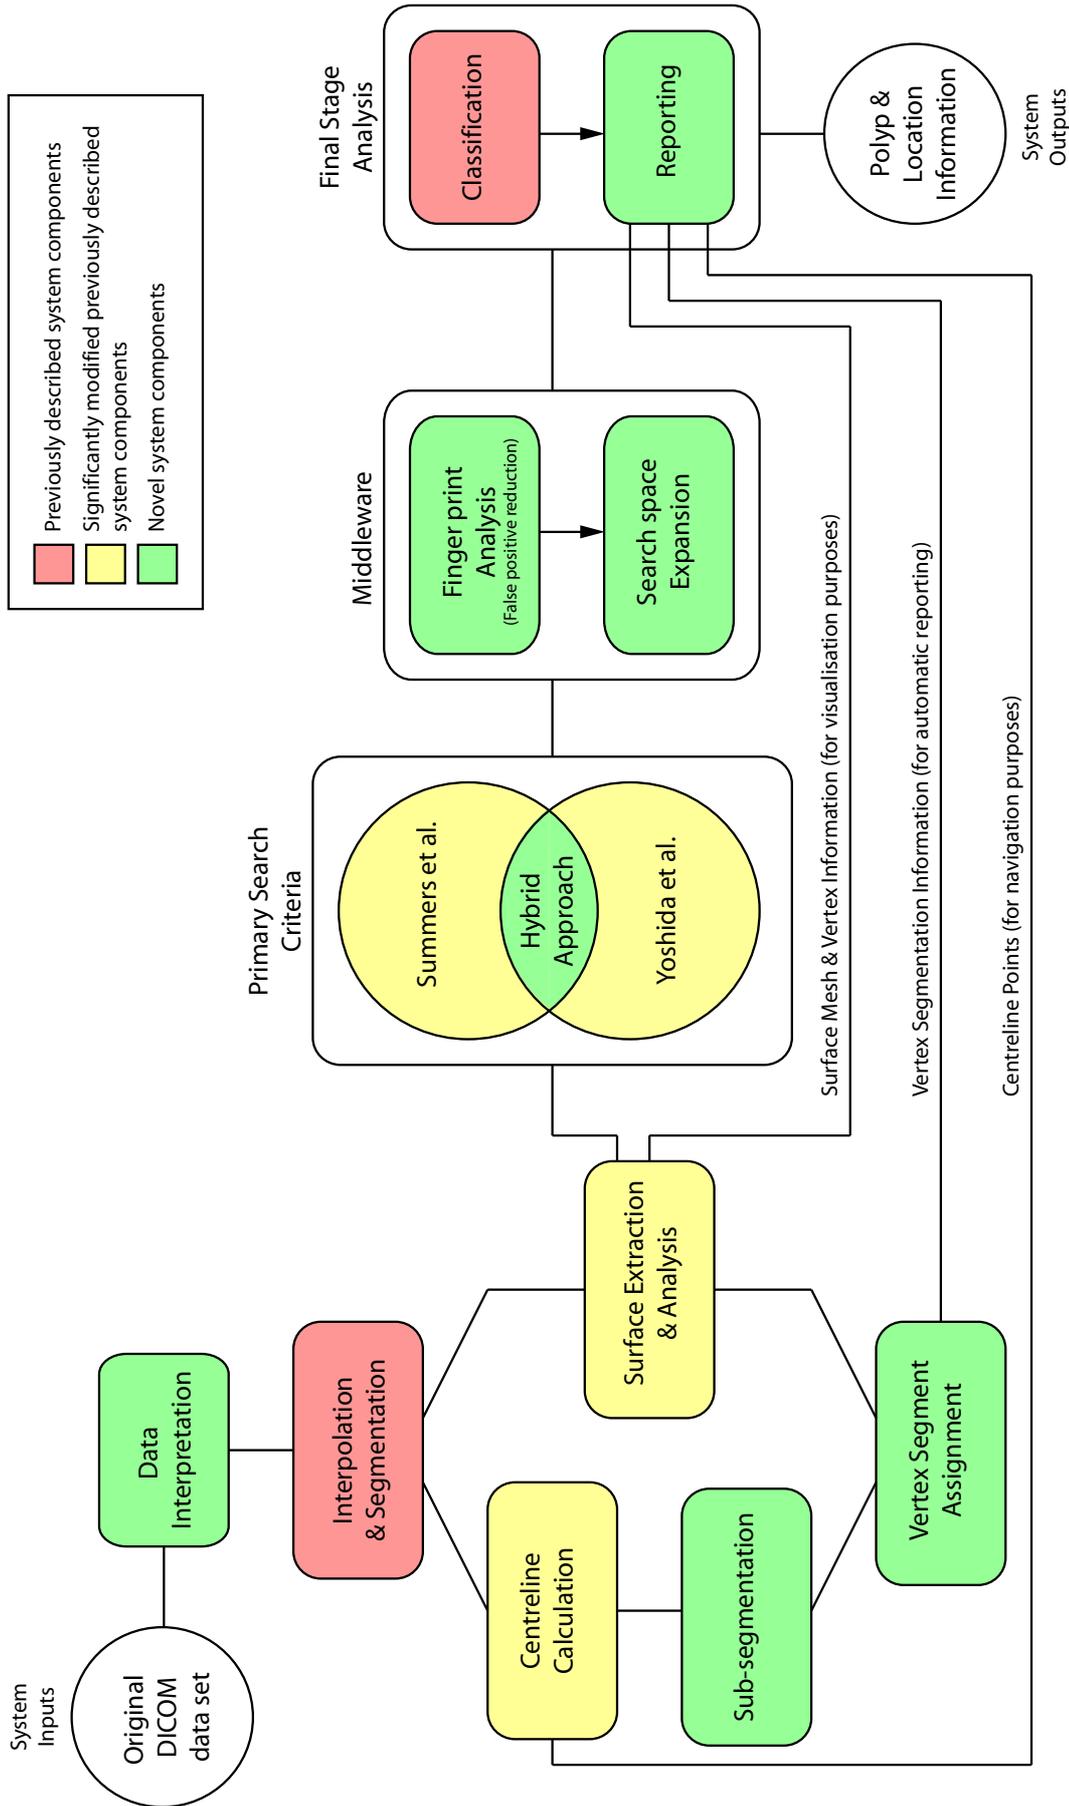


Figure 1.2: An overview of the developed CAD-CTC system where minor and major contributions are highlighted.

# Chapter 2

## Standard CTC Review

### 2.1 Introduction

The CTC protocol has been in a continuous state of flux since it was first demonstrated in 1994. Researchers have investigated many different aspects of the technique in order to establish an optimum approach to radiological imaging of the colon. This chapter reviews a selection of the important early studies, the outcomes of which have defined the way that the CTC protocol has evolved. The review of CTC development is continued in Chapter 3, where research dealing with the automation of different aspects of standard, or manually read, CTC is discussed.

### 2.2 CTC Technique

The first large-scale evaluation of CTC was carried out by Fenlon et al. of the Departments of Radiology and Gastroenterology at the Boston University School of Medicine between March 1997 and January 1999. Their subsequent report (Fenlon et al. 1999) provides a good overview of the standard CTC technique and will be summarised in the following paragraphs to form an introduction to this chapter.

The Boston study involved 100 high risk patients consisting of 60 men and 40 women with a mean age of 62 years. The patients underwent the same bowel preparation as that used for conventional colonoscopy as CTC was performed immediately before conventional colonoscopy for the purposes of the study. Prior to CT scanning the colon was insufflated with room air to maximum patient tolerance and a smooth muscle relaxant was administered to minimise peristalsis and reduce patient discomfort. The level of insufflation was assessed by means of a scout x-ray. Patients were scanned with a helical CT scanner with a table speed of 6.25 mm per second at a dose of 110 mAs and 110 kV. The resolution of the resulting axial images was  $512 \times$

**Table 2.1:** The results of large scale studies (involving  $\geq 100$  patients) that compare the performance of manually read CTC to conventional colonoscopy for the detection of colorectal polyps. Note that the extremely low overall sensitivity reported by Cotton et al. (2004) is due to the high frequency of small polyps in their study, combined with the low sensitivity for the detection of these small polyps.

Group	Cases	Polyps	Sensitivity			
			Small $\leq 5$ mm	Medium 6 - 9 mm	Large $\geq 10$ mm	Total
Fenlon et al. (1999)	100	115	55%	82%	91%	71%
Yee et al. (2001)	300	223	59.1%	80.1%	90%	69.7%
Cotton et al. (2004)	615	827	7.64%	22.7%	52%	13%
Pickhardt et al. (2003)	1233	550	-	89%	92%	89%

512 pixels. All patients were scanned in the prone and supine positions, a strategy that is employed to reduce the occluding effects of residual material in the colon. All scans were obtained during a single breath hold to reduce motion artifacts and all patients were reviewed by two experienced abdominal radiologists who were blinded to the results of conventional colonoscopy and all other previous evaluations.

Of the study cohort, 49 patients presented with colorectal neoplasia at conventional colonoscopy. These consisted of 115 polyps and three colorectal carcinomas. Of the polyps, 53 were 1 to 5 mm in size, 40 were 6 to 9 mm and the remaining 22 were  $\geq 10$  mm. It was noted that 89% of the conventional colonoscopy examinations were incomplete from the rectum to the caecum, various reasons were given for incomplete examinations. One complication (a nonfatal perforation of the sigmoid colon) resulting from colonoscopy was reported. In total, 71% of the 115 polyps reported at conventional colonoscopy were identified using CTC. The sensitivity of CTC compared to the gold standard (i.e. conventional colonoscopy) increased as the size threshold for significant polyps was increased. A per-polyp sensitivity of 91% was reported for polyps  $\geq 10$  mm in size. There were a total of 19 false positive detections over all size ranges. The authors noted that between 10 and 20 percent of colonic polyps, and up to 5% of colorectal cancers, can be missed by conventional colonoscopy. As a result, the actual number of false positives found in their study may have been slightly lower than the value reported.

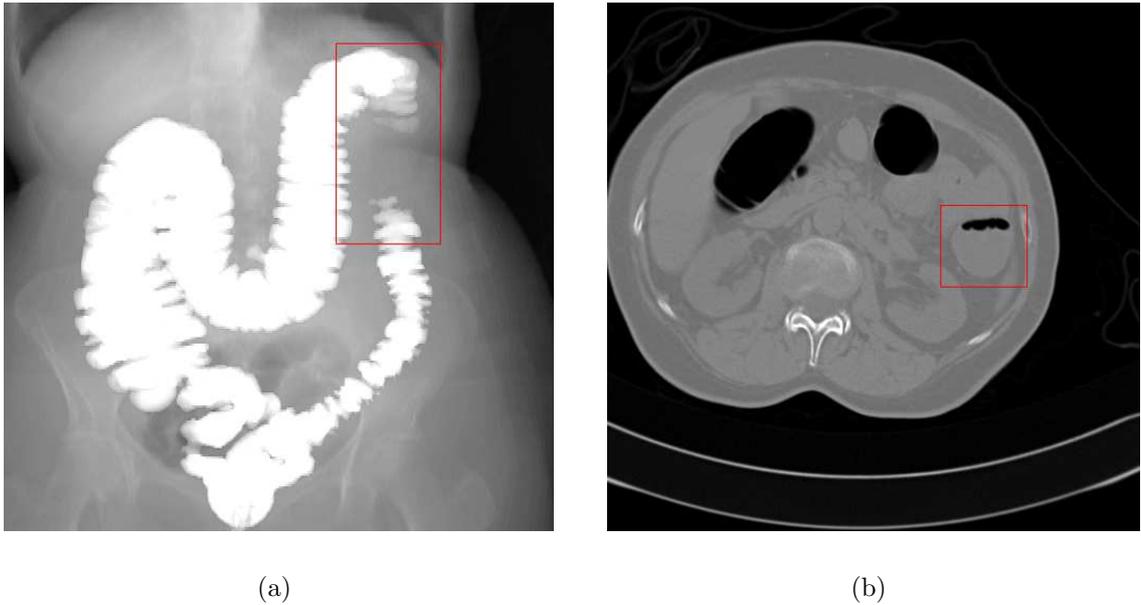
The Boston study was the first large scale study to demonstrate the feasibility of CTC for detecting significant colorectal polyps. The results of this study, and those of other significant comparative studies that have been published more recently, are summarised in Table 2.1. These results, with the exception of those presented by

Cotton et al. (2004), are promising and demonstrate that CTC is a serious candidate for large scale cancer screening. A number of potential issues associated with CTC were identified in an earlier publication by Fenlon & Ferrucci (1997). The remainder of this Section discusses the CTC protocol in more detail and reviews some of the issues raised by Fenlon & Ferrucci.

### 2.2.1 Bowel Preparation

At present CTC results depend largely on a clean well distended colon. Poor distention and any residual stool or intraluminal fluid may have a severe impact on the ability to detect polyps using CTC. Bowel preparation usually involves ingestion of a bowel cleansing agent prior to the CTC examination. Macari et al. (2001) evaluated the effect of different bowel preparations on the residual fluid during CTC. The study group consisted of 42 patients who received either a polyethylene glycol electrolyte solution (gastroenterologists choice) or a phospho-soda preparation (radiologists choice) the day prior to CTC. Macari et al. found that on average, a phospho-soda preparation resulted in significantly less residual intraluminal fluid. These results would suggest that the ideal bowel preparation depends on the type of colon examination being performed and that it is different for CTC and conventional colonoscopy examinations. A wet preparation is acceptable in the case of conventional colonoscopy as any residue can be removed by means of suction to reveal submerged regions. Conversely a dry, low residue, preparation is preferable in the case of a CTC examination. In both cases the bowel preparation is unpleasant for the patient. Recent developments have been made in the area of bowel preparation in which the use of a much more patient friendly non-purgative approach is being considered (see Section 3.2.3). Examples of suboptimal bowel preparation are presented in Figure 2.1, where a large blockage is evident in the proximal descending colon just below the splenic flexure, and in Figure 2.2, where residual material obscures the colon surface (a) and mimics colorectal polyps (b).

Immediately prior to CT scanning, a small bore rectal catheter is used to insufflate the colon with room air or carbon dioxide ( $\text{CO}_2$ ) to the maximum patient tolerance (generally about 1000 cc). McDermott et al. (2001) compared air and  $\text{CO}_2$  insufflation techniques in groups of 54 and 45 patients, respectively. Their study found no significant differences in image quality or in patient preference between the two groups. Insufficient bowel insufflation may result in collapsed regions

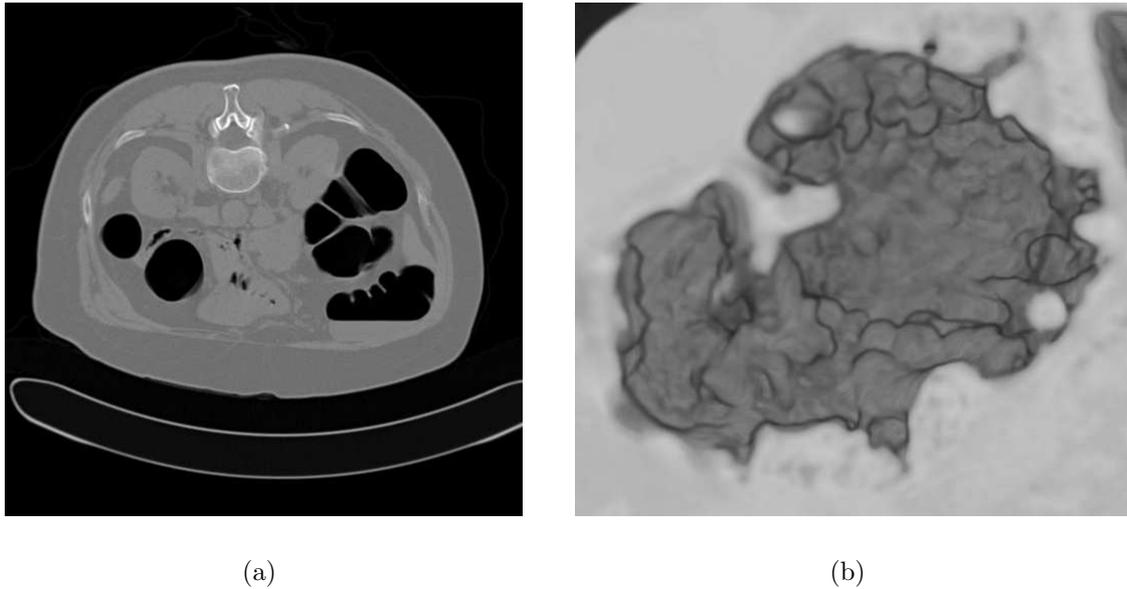


**Figure 2.1:** A case (MMH-010/supine) where a large amount of residual faecal material has caused a complete blockage of the proximal descending colon. This blockage is imaged using a modified scout x-ray where white pixels represent endoluminal air (a). An examination of the relevant 2-D axial slice from the data set of the same patient (b) confirms that this is a case of blockage and not collapse. In both cases the location of the blockage is indicated by a red box.

of the colon and thus degrade the quality of the exam. Collapse is common in the sigmoid colon, an example of this is illustrated in Figure 2.3.

### 2.2.2 Patient Positioning

Chen et al. (1999) investigated how patient positioning influenced bowel distention and polyp detection in a cohort of 23 patients (46 data sets). Distention was graded for each colonic segment in each scan on a four point scale from 0 (completely collapsed) to 3 (distended optimally). Partially or completely collapsed segments were found to be present in 11 of the prone data sets and 16 of the supine data sets. When both scans were considered together only three patients were found to have partially or completely collapsed segments. A total of 27 colonoscopy confirmed polyps were present in the 46 data sets (seven  $< 5$  mm and 20  $\geq 5$  mm). Eighteen of these polyps were identified from the prone images, and 13 were identified from the supine images. When both prone and supine data sets were considered together, 21 of the 27 polyps were identified. This study demonstrates that prone positioning is superior to supine positioning, but that dual patient positioning is necessary for adequate overall colonic distention, and results in greater sensitivity for polyp

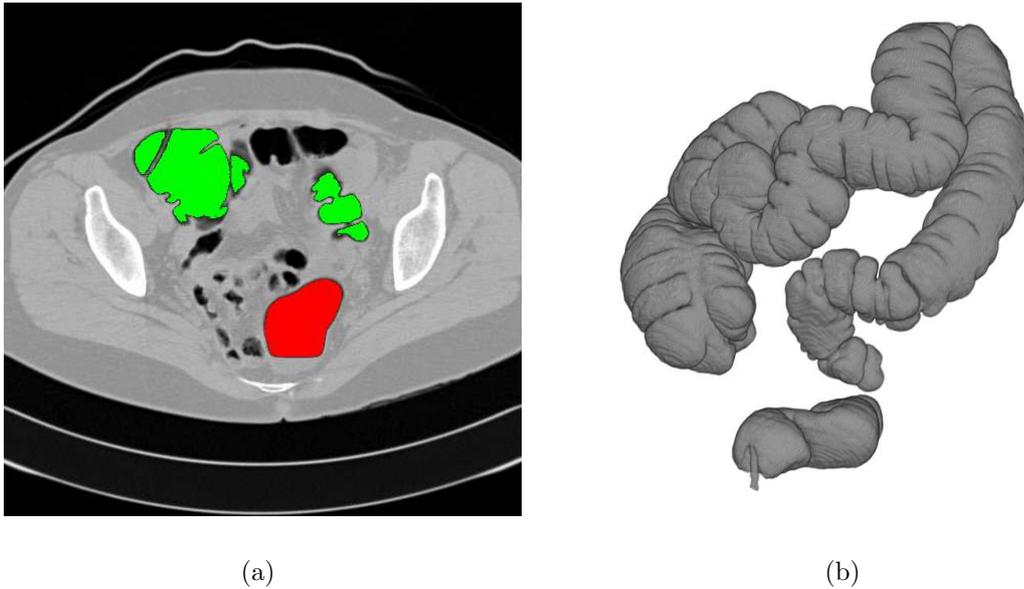


**Figure 2.2:** Two examples of localised residue: (a) Residual fluid near the hepatic flexure of a patient (MMH-138/prone). Note the long horizontal transition at the fluid/air interface. Also note that the fluid/soft tissue boundary cannot be identified, if a submerged polyp is present then it will be undetectable. (b) Residual fecal material in the cecum resembling multiple polyps.

detection. The superiority of prone positioning over supine was also noted by Morrin et al. (2000) in their study evaluating the utility of contrast enhancement at CTC, see Section 2.2.6. An example of how dual positioning facilitates the visualisation of a large previously submerged polyp is illustrated in Figure 2.4.

### 2.2.3 Incidental Extracolonic Findings

A conventional colonoscopy examination is restricted to an evaluation of the interior of the colon. A CTC examination is not restricted in this way, and can be used to visualise any organs or bone structures located outside the colon that have been captured by the abdominal CT scan. Hara et al. (2000) report on a novel study where they evaluate the frequency and importance of incidental extracolonic findings at CTC. Their study involved a cohort of 264 patients and any extracolonic findings were classified as being of high, moderate or low clinical significance. In total there were 151 findings in 109 (41%) of the patients, and 34 of these findings were considered to be of high importance. The financial impact associated with following up all significant findings was calculated, and the average additional cost per patient was found to be relatively small. The benefits associated with the identification of highly

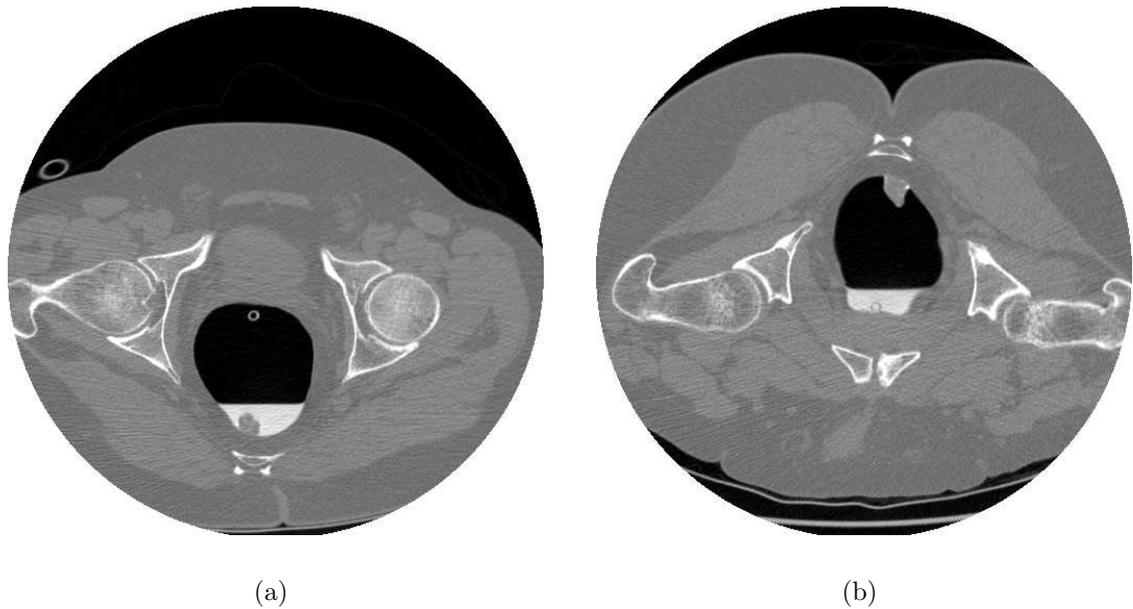


**Figure 2.3:** An example of collapse (MMH-381/supine) where the lumen is divided into two distinct sections represented by the red (rectum) and green (remainder) regions in (a) and a volume rendering of the same colon again illustrating the collapse at the rectosigmoid junction.

significant findings include a reduction in mortality due to undiagnosed conditions. However, the identification of multiple unimportant findings could result in expensive additional diagnostic examinations. This study by Hara et al. highlights the relatively high frequency of clinically significant extracolonic abnormalities, and suggests that, in the future, the search for abnormalities may become an integral part of the CTC protocol. This will represent a challenge for both radiologists and for those involved in the development of computer aided diagnostic software intended for use in CTC examinations.

#### 2.2.4 Radiation Dose

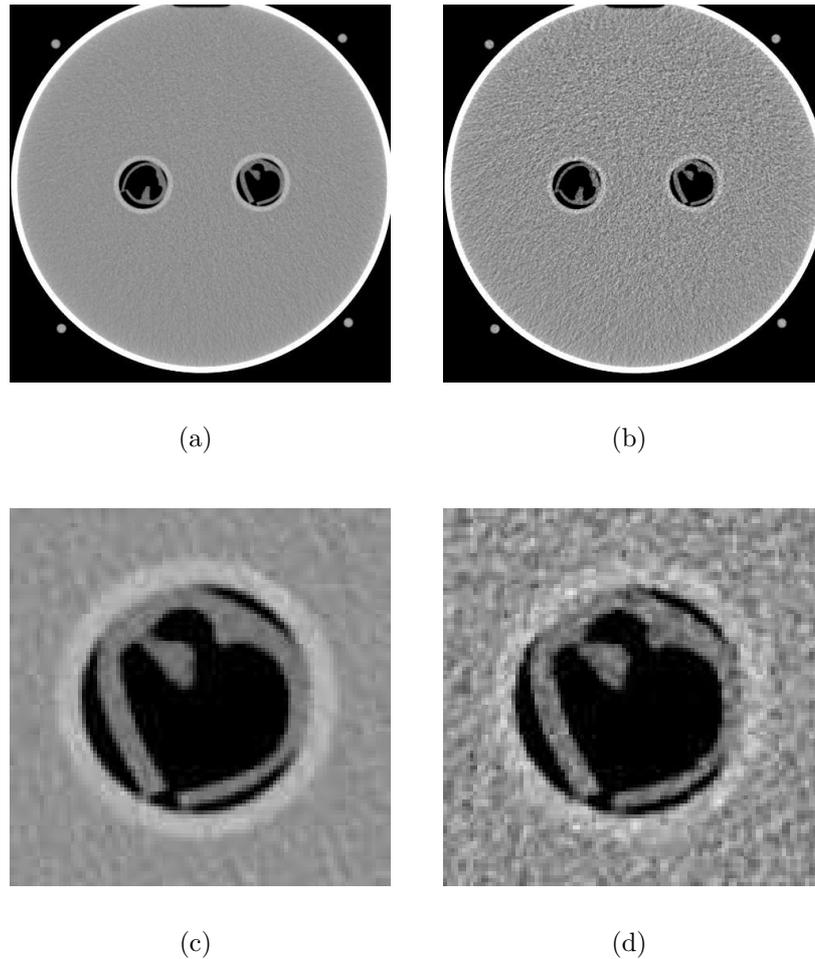
Ionising radiation is a known carcinogen, and consequently, the use of CT for colorectal cancer screening may actually result in the patient developing cancer at a later stage. This risk was quantified by Jensch et al. (2004) in a study that evaluated protocols obtained directly from institutions as well as those found in the literature as a basis for their calculations. They concluded that if applied once to a population aged 50, CTC performed in both the prone and supine positions would result in one fatal cancer in every 5,000 individuals (0.02%) over a long period of time (possibly decades). The problem of dose is compounded by the need for dual pa-



**Figure 2.4:** A good example of a case where a polyp is obscured in one position (supine) (a) and visible in the other (prone) (b). It should be noted that this patient (WRAMC\_VC-241) received orally ingested contrast material. If contrast was not used then both the polyp and the residual material would be of similar density. This would effectively hide the polyp in the supine scan.

tient positioning to adequately image the entire colon surface (discussed in Section 2.2.2).

Macari et al. (2002) performed a study to establish the effectiveness of low dose CTC in the detection of significant polyps. Their scanning protocol used a tube current of 50 mAs. This represents at least a 50% reduction in dose compared to Fenlon et al. (1999), Yee et al. (2001) and Pickhardt et al. (2003) who used 110 mAs, 120 - 150 mAs and 100mAs respectively. In the study by Macari et al., 105 patients received CTC followed immediately by conventional colonoscopy (the gold standard), and the polyp detection sensitivity compared to the gold standard was recorded for all polyp sizes (classified into three distinct size ranges). For the small ( $\leq 5$  mm), medium (6 - 9 mm) and large ( $\geq 10$  mm) size ranges Macari et al. reported polyp detection sensitivities of 12%, 70% and 93% respectively. These results compare very favourably with those presented in Table 2.1 for polyps in the medium and large size categories, indicating that low dose CTC is indeed feasible. The effect of dose reduction on image quality is illustrated in Figure 2.5. It is evident from these images that the air/phantom boundary is still clearly delineated despite the introduction of a large amount of noise into the image.



**Figure 2.5:** Two examples demonstrating the effect of radiation dose on image quality (a & c) with a tube current of 100mAs and (b & d) with a tube current of 13mAs. Both images were generated using the VSG phantom and are courtesy of T. A. Chowdhury (VSG). The window centre and window width for both images are -550 HU and 400 HU respectively and the 13 mAs image has 2.9 times more noise than the 100 mAs image.

Chowdhury et al. (2004) carried out a phantom study to evaluate the effect different radiation doses had on imaged polyp characteristics. Their phantom contained 11 spherical polyps ranging from 3 - 15 mm in size and three flat polyps. The phantom was scanned using a range of effective tube currents: 150, 90, 80, 70, 60, 50, 40, 30, 20 and 13 mAs. The authors noted that a reduction in radiation dose increased the amount of noise in the resulting images and reduced the polyp surface definition. However, the reduction in polyp surface definition was found to be minimal. The images in Figure 2.5 illustrate the observations made by Chowdhury et al..

### 2.2.5 Image Interpretation: 2-D versus 3-D

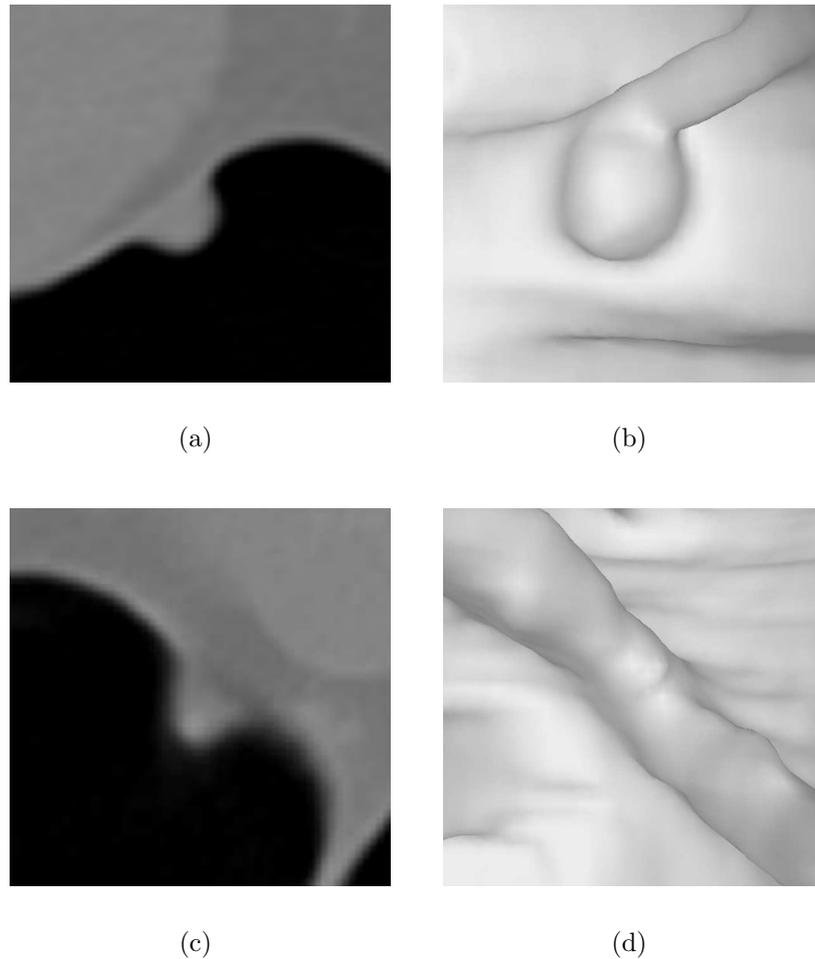
There is much debate as to whether a CTC data set should be examined as a series of 2-D slices, or as a reconstructed 3-D volume in which endoluminal views that are comparable to those obtained at conventional colonoscopy are generated. Hara et al. (1996) carried out a qualitative study to establish whether 2-D only, 3-D only, or a combination of 2-D and 3-D provided the best sensitivity for the detection of colorectal polyps. They used a single patient scanned in the supine position only who had no polyps detected at conventional colonoscopy. The resulting data set was copied five times and 11 synthetic polyps ranging in size from 1 to 10 mm were randomly placed throughout the five copies. The resulting data sets were blindly evaluated by two experienced radiologists using the three previously mentioned visualisation modes. A location and diagnostic confidence rating was specified for each polyp that was identified by the observers. None of the visualisation modes resulted in statistically significant differences in polyp detection sensitivity for either of the two observers. However, both observers did achieve the best results using the combined 2-D and 3-D approach.

In a more recent study Pickhardt et al. (2005) compared the accuracy of linear polyp measurement using 2-D *multiplanar reformatted* (MPR) and 3-D endoluminal views of the colon. The test data consisted of a colon phantom and ten patient data sets. The phantom contained a total of 10 synthetic polyps (four pedunculated and six sessile), ranging from 6 to 13 mm in size, and the patient data sets contained ten colonoscopy confirmed polyps ranging from 7 to 25 mm in size. Four experienced radiologists, who were blinded to the actual polyp sizes, measured the 10 real and 10 synthetic polyps using the commercially available Viatronix V3D-Colon CTC system<sup>1</sup>. The 3-D endoluminal displays facilitated the most accurate measurements in the phantom as well as in the real patient data.

The general consensus seems to be that 2-D should be used for initial detection as it facilitates a rapid evaluation of the colonic mucosa, and allows a search for extra colonic abnormalities to be carried out, whereas 3-D should be reserved for problem solving and general assessment of polyp morphology (e.g. size, as demonstrated by Pickhardt et al.), see Figure 2.6.

---

<sup>1</sup>Pickhardt (2003) found this system to be the best choice for primary 3-D examination of the colon in his comparative evaluation of three commercially available CTC systems (see Section 2.5).



**Figure 2.6:** A polyp viewed using a 2-D axial image (a) and a 3-D surface rendering (b), and a fold viewed using a 2-D axial image (c) and a 3-D surface rendering (d). All images were obtained from MMH-357/prone. Note the similarities between the polyp and the fold in the 2-D images (a) & (c). In general 2-D is used for the primary evaluation of a data set and 3-D is used for problem solving. This approach would be relevant in the case of the images presented here.

### 2.2.6 Contrast Enhanced CTC

CTC examinations are generally performed without the use of intravenously administered contrast material. Morrin et al. (2000) carried out a study to establish the potential benefits associated with the use of an intravenous contrast agent in CTC examinations. Their study group consisted of a cohort of 200 patients (115 of whom received contrast and 85 of whom did not). Contrast enhanced patients were scanned in the supine position without contrast and then in the prone position after contrast had been administered. Eighty one of the patients (48 of whom received contrast material) also underwent colonoscopy. These patients were used to gauge the effect contrast had on the detection of colonoscopy confirmed polyps

(noting of course that the shift in position between scans would have an impact on the results). The use of contrast enhancement was found to significantly increase the detectability of medium sized polyps (6 - 9 mm) from 58% to 75%. Contrast enhancement was also found to aid in the detection of large submerged polyps due to the difference in intensity between the polyps and the luminal fluid in which they were submerged. Interestingly the use of contrast only affected the detection of medium sized and large sized polyps ( $\geq 5$  mm), the detectability of small sized polyps ( $< 5$  mm) was not significantly affected. The use of contrast as suggested by Morrin et al. is not common in CTC protocols that are documented in the literature. Evidently the increase in sensitivity for the detection of medium and large sized polyps does not warrant the risks associated with the use of an intravenously administered contrast agent. Also it is important to note that the opacification of the bowel wall to detect submerged polyps is not the only solution to the problem of occluding residual material. Alternatively opacification of the faecal material using an orally administered contrast agent (see Section 3.2.3) would have a similar effect.

## 2.3 Issues Associated with CTC

There are a number of issues associated with CTC due to the nature of the examination as well as the infancy of the technique. These affect both the patient and the radiologist performing the examination. A number of key issues, specifically risk factors, image resolution, patient preference and reader training will be discussed in this Section.

### 2.3.1 Risk Factors

CTC is a minimally invasive alternative to conventional colonoscopy, but it still has a number of associated risks.

As mentioned earlier, (see Section 2.2.4) the exposure to ionising radiation is a known cause of cancer. Although the possibility of reducing the dose is being investigated, the fact remains that there is an associated risk of mortality in 1 out of every 5,000 screening candidates (Jensch et al. 2004). Luboldt et al. (1997) investigated 3-D *magnetic resonance imaging* (MRI) as an alternative to CTC for colon imaging and found that all relevant abnormalities including polyps were visible. In an MRI examination of the colon (MR colonography), water is used instead of

room air or CO<sub>2</sub>. The resolution of the data obtained from an MRI scan is lower than that obtained at CTC. Luboldt et al. cited a voxel size of  $1.25 \times 2.0 \times 2.0$  mm<sup>3</sup> in their study, whereas the average voxel size in the CTC data used in this research (see Appendix B) was  $0.848 \times 0.848 \times 1.5$  mm<sup>3</sup>.

While CTC is minimally invasive, it still requires the colon to be fully insufflated via a rectal catheter. Coady-Fariborzian et al. (2004) reported on the first known case of bowel perforation at CTC. This occurred in an elderly patient immediately after air insufflation of the colon and did not result in mortality. Although extremely rare, bowel perforation is a potential risk associated with CTC.

### 2.3.2 Image Resolution

The level of detail attainable at CTC is much lower than that achieved at conventional colonoscopy. This is particularly evident in relation to the detection of small polyps (i.e. those  $\leq 5$  mm in size) where the detection rates are significantly lower than those achieved at conventional colonoscopy. The reduced sensitivity for this size category is evident in all of the studies listed in Table 2.1. The fact that small polyps may be missed at CTC is not a major problem due to the slow progression of colon cancer, and the relatively frequent screening interval (every 5 years for individuals of 50 years and older). Another significant problem is the detection of flat or depressed colorectal neoplasia. These can be readily imaged at conventional colonoscopy as irregular patches of the colon surface, but they are very difficult to identify in CTC examinations due to the comparatively low image resolution. Flat polyps were originally thought to be unique to the Japanese population, however Rembacken et al. (2000) demonstrated from their evaluation of 1,000 U.K. colonoscopies, that this type of polyp is also frequent in a western population. In their study 36% of polyps were found to be flat.

### 2.3.3 Patient Preference

Reduced patient discomfort was one of the early promises of CTC, and it is clear that the CTC examination is much less invasive than conventional colonoscopy. A number of studies have been carried out to gauge actual patient opinion.

Hecht et al. (1997) carried out a study to compare CTC and conventional colonoscopy from the patients' perspective. All patients in their study cohort were queried about CTC and conventional colonoscopy both before and after screening

using either technique. The results of the questionnaire found that patients were more anxious about colonoscopy than CTC and anticipated markedly higher levels of discomfort and pain. The perceived and actual embarrassment ratings for both procedures was found to be the same, this was due mainly to the common bowel preparation and the invasive nature of both examinations. It should be noted that the post-procedural pain rating was only slightly higher for colonoscopy than for CTC, which is probably due to the fact that the patients were under general anesthetic for colonoscopy and did not remember much of the procedure.

Svensson et al. (2000) carried out a similar patient acceptance study of CTC compared to conventional colonoscopy on a cohort of 110 patients. Following each procedure the patients were queried about overall impression, level of discomfort during air insufflation or colonoscope insertion/extraction and pain. It was found that 82% of the patients preferred the CTC procedure and only 6% of the study group rated CTC as very painful compared with 29% for conventional colonoscopy. Overall, CTC achieved higher levels of acceptance than conventional colonoscopy with less pain and better patient tolerance.

Morrin et al. (1999) present a convincing argument in favor of CTC stating that from the patient's perspective the examination typically takes less than 10 minutes, does not require sedation or abdominal compression and is well tolerated apart from mild discomfort due to air or carbon dioxide insufflation. One of the major reasons for reduced patient acceptance of CTC is the need to undergo the same bowel preparation as for conventional colonoscopy or barium enema. It may be possible to forego conventional bowel preparation prior to CTC if a new technique, known as digital bowel subtraction (see Section 3.2.3), is employed. Using this technique patients are required to take an oral contrast agent prior to CTC in order to highlight any residual stool in the colon. This residue can then be digitally removed from the resulting abdominal CT images leaving only soft tissue. If this new bowel preparation technique proves effective then patients will be offered a preliminary non invasive colorectal screening procedure that gives them a significant chance of avoiding conventional colonoscopy.

### 2.3.4 Reader Training

In their editorial Soto et al. (2005) pose the question of how much reader training is required for CTC. They recognise that a high standard of performance and in-

terpretation will be required for widespread acceptance of CTC, and that this is only achievable through the specialised training of radiologists in the technique. In an earlier study intended to establish the general consensus regarding CTC training guidelines, Soto et al. (2004) sent out a questionnaire to 20 international experts in the field. The outcome of their survey suggested that the most appropriate method for reader training would consist of lectures and supervised hands-on workstation training with 40 to 50 cases (20% of which should be normal). The suitability of this level of training was confirmed by Pickhardt et al. (2004) who noted that reader performance was relatively stable after the interpretation of the first 50 cases. Soto et al. (2005) believe that the required training should be integrated into the relevant radiology residency and fellowship programmes. However, until that happens radiologists interested in acquiring competency in CTC must undertake a specific training course or participate in dedicated fellowships. This highlights the fact that accessibility to suitable training resources is an issue that will need to be addressed in the very near future.

It is also important to consider the use of CAD-CTC to assist in reader training. CAD-CTC is intended to act as a second reader by highlighting suspicious regions of the colonic mucosa (see Section 3.5.1). However, it is feasible that it could also be used to help trainee radiologists to differentiate between what is normal and what warrants further attention. In this way CAD-CTC could be used as a virtual tutor. This would only be possible if the ability of CAD-CTC to differentiate between true polyps and false positives was comparable to that of an experienced radiologist.

## 2.4 Alternative Screening Techniques

As mentioned in Chapter 1, conventional colonoscopy is the most sensitive colorectal cancer screening technique that is currently available. CTC is just one alternative to conventional colonoscopy. Others worthy of note, that will be reviewed in this Section are fecal occult blood test, barium enema, flexible sigmoidoscopy and DNA based screening. Different issues associated with these techniques including sensitivity, specificity and patient preference will be addressed.

### 2.4.1 Fecal Occult Blood Test

A *fecal occult blood test* (FOBT) can be used to detect the presence of small amounts of blood in stool samples, and thus, can be used to detect bleeding lesions that are located anywhere in the gastrointestinal tract. The test involves placing a small amount of stool on a chemically treated card. If there is a sufficient amount of blood in the stool then the card will change colour. Rockey et al. (1998) established the frequency of colorectal neoplasia (and upper gastrointestinal lesions) in 248 patients above 50 years of age all with positive FOBTs. Conventional colonoscopy was performed in each case and any detected polyps or cancers were recorded. Of the 248 patients, 13 (5.2%) were found to have cancers; 29 (11.7%) were found to have large polyps ( $\geq 10$  mm) and a further 29 (11.7%) were found to have medium sized polyps (5 - 9 mm). These results show that a positive test does not necessarily indicate the presence of colorectal cancers or polyps, and demonstrate the considerable lack of specificity provided by the fecal occult blood test for colorectal cancer screening.

### 2.4.2 Barium Enema

A barium enema involves taking an x-ray of the abdominal region in order to image the large intestine. There are two variations of this examination: (1) the single-contrast technique where barium sulphate is injected into the rectum prior to x-ray in order to acquire a well defined profile view of the large intestine, and (2) the double-contrast technique where air is inserted into the rectum after the barium sulphate has been expelled, and prior to the acquisition of the x-ray. The *double contrast barium enema* (DCBE) provides more detailed results. Gluecker et al. (2003) carried out a study to compare patient preference for CTC, conventional colonoscopy and DCBE. Two groups of patients were recruited for this study. The first group (696 patients) underwent CTC and conventional colonoscopy while the second group (617 patients) underwent CTC and DCBE. The patients were subsequently queried regarding the inconvenience and discomfort associated with the preparation, the examination discomfort, their willingness to repeat either of the examinations and their examination preference. The results for group 2 indicated that the discomfort associated with the preparation for both CTC and DCBE was considered to be comparable (moderate to mild). The examination discomfort was

rated as mild or none for CTC whereas patients undergoing DCBE regarded the discomfort as severe to moderate. Patients in group 2 were also more willing to undergo reexamination with CTC rather than DCBE and when patients were asked for their overall examination preference, the majority (97%) opted for CTC.

### 2.4.3 Flexible Sigmoidoscopy

Flexible sigmoidoscopy is used to endoscopically examine the left sided colon, i.e. the portion of the colon from the rectum to the splenic flexure. If polyps are detected, the patient must be referred for colonoscopy. Prior to the examination the patient must undergo bowel preparation but does not require sedation. If the left sided colon is free of polyps this usually indicates that there are no polyps in the entire colon. Patel & Hoffman (2001) verify this in their study of the anatomical distribution of colorectal polyps at colonoscopy. They state that an examination limited to the left sided colon would miss approximately 23% of lesions. Flexible sigmoidoscopy is not a complex procedure and can be performed by non-physician endoscopists. Shapero et al. (2001) present the results of a study where nurse endoscopists performed the flexible sigmoidoscopy examination. Video recordings of the procedures were subsequently reviewed by trained physicians. The results of this study indicate that flexible sigmoidoscopy can be carried out safely and effectively by non-physician personnel. This is very important from the point of view of screening costs, as the manpower costs for non-physician operators are considerably less than those for specialist physician endoscopists.

### 2.4.4 DNA Testing

An emerging colon screening technique involves the examination of a patients stool for excreted cells that exhibit DNA mutations associated with colorectal cancer and polyps. DNA based analysis is non-invasive and requires the acquisition of a stool sample much like the FOBT. Unlike the FOBT, DNA is continually excreted whereas rectal bleeding can be intermittent.

Ahlquist et al. (2000) evaluated the performance of DNA based stool analysis in discriminating between patients with colonoscopy confirmed colorectal neoplasia and those with normal colons. Their study consisted of three groups, 22 patients with cancer, 11 with large polyps ( $\geq 10$  mm) and 28 patients with normal colons. A sensitivity of 91% was reported for cancers, a sensitivity of 82% was reported for

polyps and an overall specificity of 93% was reported for the technique. It should be noted that the sensitivity of DNA analysis for the detection of large polyps ( $\geq 10$  mm) reported by Ahlquist et al. is lower than that reported for CTC in the majority of the large scale studies summarised in Table 2.1.

In their review article Levin et al. (2003) list a series of advantages and disadvantages associated with testing for DNA mutations in stool. Advantages included the patient friendliness and high accuracy of the test as well as the reduced level of false positive detections. The disadvantages or limitations cited by Levin et al. included lack of automation and the high cost of the test.

## 2.5 Commercial CTC systems

A number of commercial CTC systems have appeared on the market in recent years. These systems originally provided a relatively basic platform that enabled manual CTC reading. However, the latest releases are now starting to incorporate automation tools that have originated from research (these types of automation tools will be reviewed in Chapter 3). The main commercial CTC systems that are available are:

1. **Barco, Voxar Colonscreen** runs on a PC platform and enables 2-D and 3-D visualisation of prone and supine CTC data, supports lumen tracking, facilitates lesion marking and measurement with targeted volume rendering and provides a virtual double contrast barium enema for colon overview.  
(see <http://www.voxar.com/>; Accessed May 26th 2006.)
2. **GE Medical Systems, Navigator** can be used to display prone and supine views synchronously. Bookmarking can be used to record the location of detected polyps. Three oblique images and a 3-D endoluminal view are provided for inspecting potential polyps.  
(see <http://www.gehealthcare.com/rad/ct/applications/colon/>; Accessed May 26th 2006.)
3. **MedicSight, ColonCAD<sup>2</sup>** incorporates a CAD software component and enables CAD findings to be viewed in conjunction with the original CTC images.

---

<sup>2</sup>A predecessor to ColonCAD, called ColonCAR, was evaluated by Taylor et al. (2006). The results of this evaluation are summarised in Section 3.5.1.

The tolerance of the CAD system could be altered by adjusting a sphericity threshold.

(see <http://www.medicsight.com/>; Accessed May 26th 2006.)

4. **Rendoscopy** provides automatic colon lumen segmentation and path planning even in the presence of collapsed segments. An unfolded view of the colon is provided to enable a complete evaluation of the colonic mucosa. Automatic reports can be generated that summarise the findings of the radiologist.

(see <http://www.rendoscopy.com/>; Accessed May 26th 2006.)

5. **Viatronix, V3D-Colon** provides automatic segmentation and centreline calculation and electronic bowel cleansing. The user interface also provides measurement tools and facilitates interactive navigation within the colon.

(see <http://www.viatronix.com/>; Accessed May 26th 2006.)

6. **Vital Images, CT Colon option for Vitrea 2** facilitates interactive simultaneous visualisation of 2-D and 3-D prone and supine images. The Vital Images software can also be used to locate polyps throughout the entire length of the colon.

(see <http://www.vitalimages.com/>; Accessed May 26th 2006.)

Pickhardt (2003) compares the rendering capabilities of three of these systems (2, 5 and 6) with respect to polyp conspicuity, 3-D effect and likeness to optical colonoscopy. In this study 12 radiologists and 13 gastroenterologists evaluated the appearance of eight polyps ranging in size from 6 to 15 mm. In addition, the endoluminal navigation capabilities and other features of the three systems were evaluated. Pickhardt concluded that of the three systems the only one that seemed feasible for an effective and time efficient primary 3-D evaluation was the Viatronix V3D-Colon system.

## 2.6 Discussion

A detailed overview of the manually read CTC technique has been presented. Important aspects of the technique including bowel preparation, patient positioning, radiation dose and image interpretation have all been discussed in detail. Issues associated with CTC have also been outlined and alternative colorectal cancer screening techniques have been considered. Finally, a range of currently available commercial

CTC offerings have been summarised. These products demonstrate how CTC is transitioning from the research domain into widespread clinical usage. The review of CTC will continue in Chapter 3 where the development of computer assisted CTC techniques will be discussed.

# Chapter 3

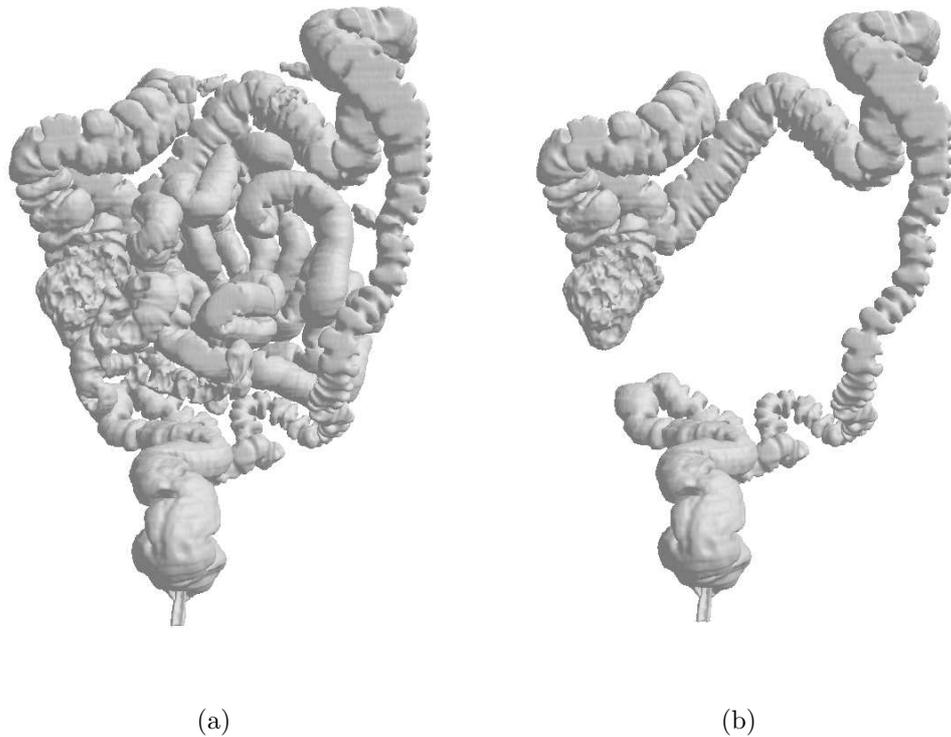
## Computer Assisted CTC Review

### 3.1 Introduction

Early CTC studies required a significant level of user interaction to navigate through and ultimately analyse CTC data sets. The use of automation at CTC has the potential to greatly increase the performance and productivity of radiologists by allowing them to concentrate on the most important aspect of a CTC examination i.e. the diagnosis of suspicious regions of the colonic mucosa. The task of automation at CTC is complicated by two main factors (1) the volume of data contained in a data set and (2) the complexity of the anatomy represented by the data set. Studies dealing with automation at CTC began to appear very soon after Vining et al. introduced the technique in 1994. Initial studies dealt with extracting the colon lumen and automating navigation through the extracted lumen to enable 3-D endoluminal flythroughs. More recent studies have dealt with the more difficult tasks of automatic polyp and mass detection. The speed and accuracy of published techniques has continued to increase, while the level of user interaction required has tended to decrease. This Chapter will discuss a range of automation techniques that are related to developments outlined in Chapters 4 & 5. This discussion deals specifically with the topics of colon segmentation, centreline calculation, data set registration and *computer aided diagnosis at CTC* (CAD-CTC).

### 3.2 Colon Segmentation

The region of the CTC data set occupied by the colon must be identified before any analysis of the colon can take place. This process is known as colon segmentation and can be achieved using a variety of techniques. An inadequate segmentation can hinder any subsequent examination of the colon. The task of colon segmentation



**Figure 3.1:** A CTC data set contains numerous gas filled regions that include the colon lumen, sections of the small intestine, and the stomach. As a result the task of colon lumen segmentation is not a trivial task. This is especially true where the colon lumen consists of more than one segment. A segmentation consisting of all gas filled abdominal regions for a sample data set (MMH-175/supine) contains 364,732 triangles (a) whereas a segmentation of the colon lumen alone contains 189,263 triangles.

is complicated by collapse and blockage. In both of these cases segmentation of the colon can yield multiple disjointed colon segments. Li & Santago (2005, Wake Forest) report that a fully segmented colon can have between 1 - 10 individual segments with an average of 2.4 segments. Under-segmentation occurs when the segmentation process fails to identify all of the individual colon segments. Another problem associated with segmentation is that the colon is not the only gas filled region that is contained in a CTC data set (Figure 3.1). Over-segmentation occurs when extracolonic regions are inadvertently included in the colon segmentation. Yoshida et al. (2002a, UC) found that 19.3% of their false positive detections made by their CAD algorithm were attributable to the small intestine or stomach. The following discussion on colon segmentation techniques is divided into two sections. The first section deals with techniques that identify the air in the colon lumen and the second section deals with techniques that extract an approximation of the colon wall from a CTC data set.

### 3.2.1 Lumen Segmentation

Chen et al. (2000a, SUNY) developed a contrast enhanced segmentation technique that segmented the entire colon lumen including regions that have been obscured by residual material. Following the removal of voxels external to the patient, a low level classification of the remaining data set voxels was performed using a vector quantisation algorithm. Voxels were thus divided into eight classes: air, soft tissue, muscle and bone or contrast enhanced material and four partial volume effect classes. Four of these classes were associated with the colon lumen: air, partial volume effect from air to soft tissue, enhanced material and partial volume effect from contrast enhanced material to soft tissue. This rough segmentation included the bone and the lungs. The lungs were identified based on their location at the top of the data set left and right of the centre and were removed using region growing. The remaining air voxels were attributed to the colon lumen, and the connected voxels representing the air to soft tissue interface were also included in the segmentation. Finally, the contrast enhanced material was included in the segmentation by region growing from an adjoining seed point located in the air filled colon. The segmentation results for 13 data sets were evaluated by a trained radiologist and the entire lumen was found to be successfully delineated in all cases (even in four cases with collapsed segments). However, sections of the small intestine and stomach were erroneously segmented in a small number of the test cases. Following a once off initialisation stage, the segmentation process required nine minutes to complete when deployed on a SGI Octane workstation with dual 250 MHz R10000 processors and 890 MB of RAM.

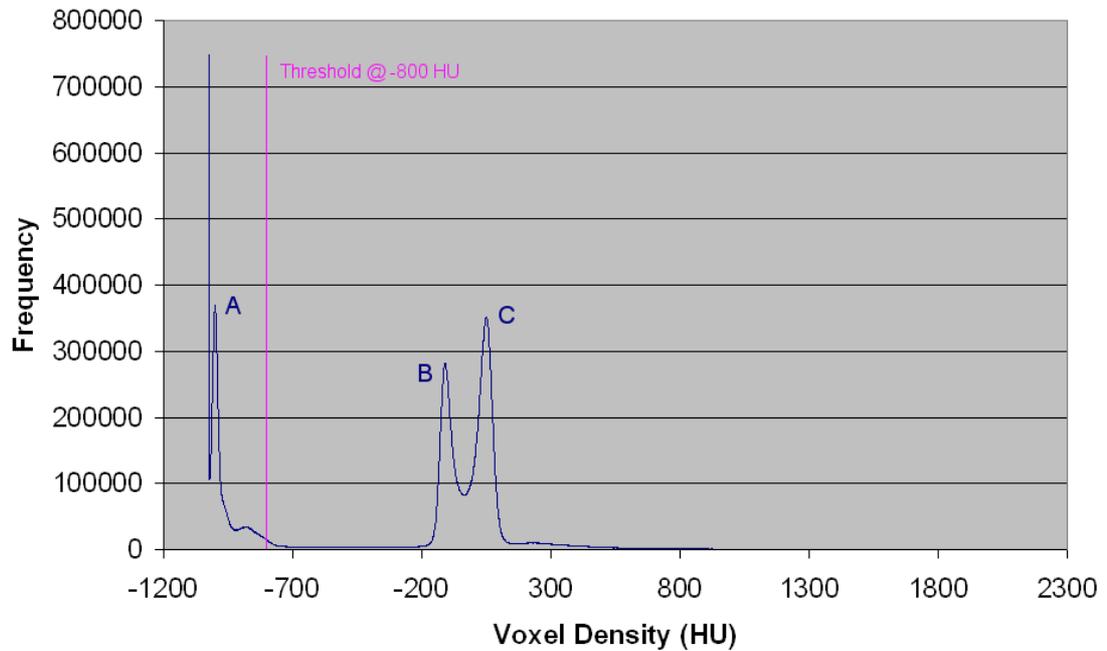
Wyatt et al. (2000, Wake Forest) developed a colon lumen segmentation technique with automatic seed point identification. Two bowel preparations were also evaluated as a part of this study: A standard bowel prep similar to that used in conventional colonoscopy and a contrast enhanced bowel preparation that incorporated a liquid diet. A region growing technique was used to segment the background air voxels. A threshold of  $-800$  *Hounsfield Units* (HU)<sup>1</sup> was then applied to identify the remaining gas filled regions located within the patient (justification for the use of this threshold value is illustrated in Figure 3.2). A distance transform was used to identify the voxel within the gas region located furthest from the surrounding tissue. Region growing was initiated at this voxel and the resulting volume was examined.

---

<sup>1</sup>The units of density in CT are named after the inventor of the CT technique, G. N. Hounsfield (Hounsfield 1973).

If it met certain shape and location criteria, it was added to the segmentation otherwise it was discarded. The next seed point was identified from the remainder of the gas region and the process continued until the segmentation was complete. An additional stage was included to deal with the contrast enhanced fluid resulting from the second bowel preparation. Large areas of low curvature were identified as the boundary between the gas and the contrast enhanced material. The gas region was then selectively dilated across this boundary and the region representing the contrasted enhanced material was identified using an adaptive region growing process. An adaptive approach was required due to variations in the density of the contrast enhanced material located throughout the colon. Twenty CTC data sets (ten with standard preparation and ten with contrast enhanced preparation) were used for test purposes. Five data sets from each group were used for determining algorithm parameters and the remainder were used for algorithm evaluation. The entire colon was identified in all data sets, however, the stomach and sections of the small intestine were erroneously identified as belonging to the colon lumen in several of the data sets. An evaluation of the extended algorithm showed that 40.36% to 80.15% (mean 65.74%) of the contrast enhanced material was correctly segmented when compared with manually segmented regions of contrast enhanced material. The algorithm was deployed on an SGI workstation with a MIPS R10000 processor and execution times of 65 minutes for the standard technique and 60 minutes for the contrast enhanced technique were reported. Note that sub-sampling was utilised in conjunction with the contrast enhanced technique due to the extra processing overhead required to deal with the contrast enhanced material.

Sato et al. (2001, SUNY) developed a colon lumen segmentation scheme that incorporated a patient-friendly bowel preparation protocol. During the day prior to scanning patients were instructed to eat only soft food. This diet was augmented with an orally ingested contrast agent to enhance the density of any residual material at CTC. An initial threshold based segmentation was performed using two thresholds  $T_{gas}$  and  $T_{fluid}$ . These threshold values were automatically identified based on an analysis of the data set histogram and were used to identify colon lumen voxels where  $\text{CO}_2$  voxels had densities  $< T_{gas}$ , and contrast enhanced fluid voxels had densities  $> T_{fluid}$ . A filtering stage was then employed to identify voxels representing the *partial volume effect* (PVE) between  $\text{CO}_2$  regions and contrast enhanced fluid regions. These PVE voxels were then classified as belonging to the colon lu-



**Figure 3.2:** A histogram of a randomly selected CTC data set (MMH-120/prone). The three labeled peaks represent air (A), fatty tissue (B), and lean tissue (C). Note that a threshold value of -800 HU adequately separates air and tissue and is suitably positioned to minimise shine through artifacts.

men. A more accurate segmentation was then obtained by initiating 26-connected region growing from the seed points selected from the PVE voxels, thus limiting the segmentation to the colon lumen and overcoming the over segmentation problems associated with the initial threshold based segmentation. In the final stage of their approach, a gradient magnitude based technique was used to expand the segmented colon lumen so that it more accurately represented the region enclosed by the colonic mucosa (i.e. the surface within the data set in the region of the colon lumen where the accumulated gradient magnitude has the maximum value). The PVE filter was then applied to the expanded region in order to yield the final segmentation. The documented technique required an execution time of less than five minutes to segment the colon lumen from typical CTC data sets when deployed on an SGI Onyx2 workstation. No results regarding segmentation accuracy were provided. However, cases of over segmentation were mentioned where the lungs or ribs were mistakenly included in the final result.

Iordanescu et al. (2005, NIH) reported on a study that dealt with automatic seed point selection for region growing based colon lumen segmentation. They noted that one seed point would be sufficient if the colon was well distended with no

collapse or blockage. Conversely, where this was not the case, multiple seed points would be required (one seed point for each segment). The documented algorithm automatically detected two seeds located at the most distal points within the colon i.e. in the rectum and in the caecum. Iordanescu et al. demonstrated that two seeds were sufficient to segment the entire colon lumen in the majority of cases (83.2%, or 243 out of 292, according to their report). Additionally, the selected locations for the seed points (i.e. at opposite ends of the colon) guaranteed that the seeds were located in different segments in cases where there was collapse or blockage. Their anatomy based seed point search identified the rectum as the air-filled region that was closest to the base of the data set and the caecum as the air filled region of greatest diameter that was closest to the side of the data set. The technique was evaluated using a cohort of 146 patients scanned in both the prone and supine positions. All automatically detected seed points were compared against manually placed seed points. The automatic and manually placed seed points matched in 82.2% or 240 of the 292 cases.

Li & Santago (2005, Wake Forest) presented a region growing based colon lumen segmentation technique that also incorporated surface extraction. The authors started by automatically identifying seed points in the uninterpolated CTC data sets by analysing each of the axial slices. Seed point identification was achieved using size and shape filtering of the air regions that were present in a slice. Size filtering involved removing connected air regions with areas above and below specified threshold values, e.g. background air regions and regions associated with the small intestine (note that this stage was applied to the 2-D axial images). Shape filtering involved evaluating the main axes of the remaining air regions, those with axes whose lengths were within a specified range were retained. Shape filtering was also used to remove areas with large aspect ratios e.g. the lung bases. The seeds were ultimately placed in the centres of air regions that passed the size and shape filters. This approach generated more seed points than were needed to completely segment the colon lumen. This is not a significant problem as unnecessary seeds were absorbed during the segmentation process. Segmentation was achieved using 3-D region growing initiated from the automatically detected seed points. The detected segments were subjected to an elongation test to identify and remove sections associated with the small intestine thus yielding the final segmentation of the colon lumen. A surface model of the colon lumen was also generated using a modified

marching cubes algorithm. A subjective evaluation of colon coverage was carried out by two radiologists using 100 data sets (50 prone and 50 supine). The average coverage was reported as 81.8% for the supine scans and 55.4% for the prone scans. A combined average coverage of 87.5% was ultimately reported with 6% of the segmentation attributable to extracolonic structures.

Chowdhury et al. (2005*a*, VSG) described a fully automated colon lumen segmentation technique that did not necessitate the administration of an orally ingested contrast agent. The first stage of their approach involved generating a binary representation of the data set that accurately separated gas and non-gas regions. This was achieved by applying a threshold value that was calculated based on an analysis of the global data set histogram. The lung bases and the air regions outside the patient were then eliminated using seeded 3-D region growing. The remaining lumen candidates were labeled so they could be uniquely identified. During the labeling process the endpoints of each candidate segment were identified. The endpoints were used to identify the orientation of, and to calculate the average cross-sectional area (volume/length) of, the candidate segments. The length between the endpoints of a segment was approximated using the Dijkstra shortest path algorithm (Dijkstra 1959). A well distended colon was subsequently identified as having an average cross-sectional area and a length that were both above specified thresholds. In addition, the orientation and structure of the colon had to conform to an ideal model. In the case of a collapsed colon, large segments that were above an average cross-sectional area threshold, were linked according to the ideal model to yield a rough segmentation of the colon. The rough segmentation was then refined by adding smaller sections based on their relative locations, lengths and curvature (highly curved regions were considered to be part of the small intestine). An evaluation of the technique using 115 data sets found that 100% of lumen voxels were identified in the case of well-distended colons and 95% of lumen voxels were identified in the case of colons with collapse or blockages. The number of voxels attributable to extracolonic structures was minimal (approximately 1%).

### 3.2.2 (Inner) Wall Segmentation

In many cases it is the colon wall and not the colon lumen that is of interest. This is particularly true where the segmentation of the colon is intended for use with CAD-CTC for polyp detection. A number of techniques that deal with the task of

colon wall segmentation have been published.

Masutani et al. (2001, UC) described a technique that was intended to segment only the colon wall from a CTC data set. Their anatomy orientated approach, based on earlier work in the area of segmentation for CT angiography, utilised a number of different image processing techniques including thresholding, region growing, labeling, gaussian smoothing, edge detection and 3-D mathematical morphology. Dual thresholding and labeling were used to identify the air outside the patients body. The high density bone structures were identified using dual thresholding (with higher threshold values than those used in the case of air). The bone structures were labeled and inspected individually and those greater than a predetermined size threshold were flagged for removal. The lung bases were identified, also using dual thresholding and labeling, as the two volumes closest to points at the top left and top right of the data set, that were located halfway between the front and the rear of the patient. These three structures (external air, bones and lungs) were dilated and the resulting regions were masked from the data set. A Gaussian smoothing operator was then applied to the data set to ensure that the colon walls would be of a suitable thickness (3 - 4 mm). The final stage of the technique involved the application of a gradient magnitude operator followed by a final thresholding stage. The colon wall was ultimately selected as the largest connected component following labeling (smaller regions were also included if their size was at least one quarter the size of the largest connected component). Of the 14 data sets used in testing, 12 (85.7%) achieved fair to excellent segmentation (where fair indicated that the segmentation included a section of the small intestine and was missing less than 5% of the colon).

In a subsequent paper from the same group, Näppi et al. (2002, UC) reported on an improved version of the automatic colon wall segmentation scheme described by Masutani et al. (2001). The technique used to segment the air outside the patient was improved so that no external air associated with either the table or air trapped beneath the patient was retained. The size threshold for including smaller regions in the final segmentation was reduced from 25% of the size of the largest connected component to 1%. This increased the number of extracolonic components as well as the amount of actual colon wall retained. The improved algorithm was augmented with a region growing stage to remove the extracolonic components as these are not usually connected to the colon lumen. Region growing within the colon lumen was initiated at an automatically detected seed point in the rectum. This seed point

was readily identifiable due to the relatively invariant location of the rectum in CTC data sets. In cases of collapse between the rectum and the sigmoid colon, a new seed point search was initiated 25 slices above the base of the data set. In the case of further collapse, seed points were selected from the remaining unsegmented internal air filled regions. The region growing process terminated in each colonic segment when the number of voxels added to the segmented region decreased by more than 75% from the previous iteration. The overall segmentation process continued until a series of global properties regarding the segmentation were satisfied. The segmented lumen was then expanded so that it covered the previously detected colon walls, and the intersection of the two structures was defined as the ultimate segmentation result i.e. the colon walls. The improved version of the algorithm, referred to as ‘knowledge guided segmentation’ by the authors, was evaluated by four radiologists using a total of 88 data sets (44 patients scanned in both the prone and supine positions). The average coverage of the improved technique was found to be 95%. However, between 10% and 15% of the segmentation was attributable to extracolonic components. When used in conjunction with polyp detection, the improved segmentation technique resulted in a 15% reduction in the false positive rate in a per-patient analysis (i.e. the number of candidates attributable to extracolonic components was reduced).

In a later collaborative effort, Frimmel et al. (2005, Uppsala & Harvard Medical) developed an algorithm that segmented the colon wall based on an analysis of the colon centreline. The technique utilised a simplified version of a centreline calculation algorithm developed by the same group (Frimmel et al. 2004). Air inside the patient was identified using a threshold based segmentation. External air was subsequently eliminated using region growing initiated from points on the scanner tunnel boundary. Centreline segments were calculated for each of the remaining air volumes. Segments close to the patient boundary were considered to be extracolonic and were removed. An anatomical evaluation of the remaining centreline segments was carried out to yield the final centreline. A segmentation of the colon lumen was then obtained by dilating the final centreline to completely fill the surrounding air filled regions. Finally, an approximation of the walls was generated by dilating the initial segmentation beyond the confines of the air/soft tissue boundary of the lumen to a desired wall thickness and then subtracting the original segmented lumen from the resulting volume. The centreline-based segmentation algorithm was com-

pared against the knowledge-guided segmentation approach discussed earlier (Näppi et al. 2002). Both techniques identified comparable regions of the colon. However, the centreline-based segmentation identified 21% fewer voxels associated with extra-colonic structures. The sensitivity of the centreline-based segmentation technique was reported to be 96%.

It is evident from the previous descriptions that these approaches are only intended to segment an approximation of the inner walls of the colon. Techniques intended to segment both the inner and outer wall of the colon as a precursor to polyp detection (Hunt et al. 1997, Vining et al. 1998, Wake Forest) are discussed in Section 3.5.3.1.

### 3.2.3 Digital Subtraction

An extension of colon segmentation touched on by Chen et al. (2000a), Wyatt et al. (2000) and Sato et al. (2001) (see Section 3.2.1) involves the patient ingesting orally administered contrast material prior to CT scanning in order to facilitate the removal of residual fecal material in the colon using image processing techniques. This approach illustrated in Figure 3.3, which will be referred to as *digital subtraction colonography* (DSC), has a number of associated potential benefits:

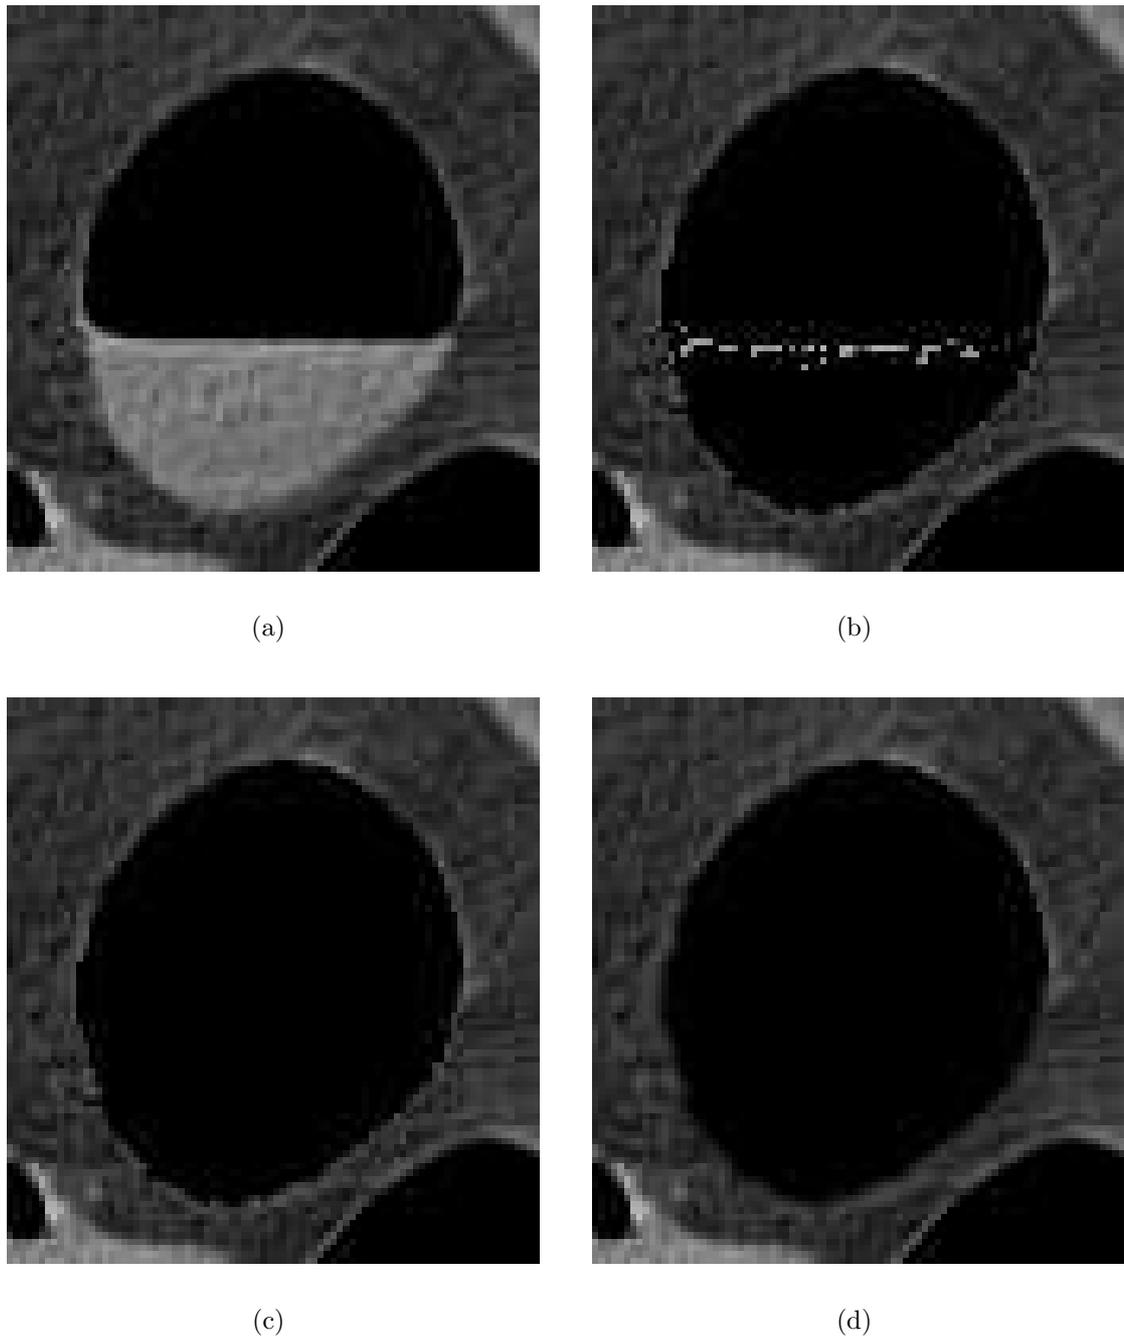
- DSC may remove the need to scan the patient using both prone and supine positioning. Dual patient positioning is used to deal with mobile residual fecal material. If this material is tagged then dual positioning may no longer be required. There are a number of benefits associated with scanning in a single position, most notably a 50% reduction in radiation dose.
- DSC may alleviate the problem of residual faecal material resembling polyps and in extreme cases causing blockages as illustrated in Figures 2.2 and 2.1 respectively.
- DSC doesn't require the use of laxatives and therefore DSC is more patient friendly and may deal with some of the issues raised in Section 2.3.3.

It should be noted that the concept of digital subtraction in radiological examinations of the colon is not entirely new. A similar technique had been used by radiologists to differentiate between polyps and residual fecal material in barium enema studies even before CTC was first demonstrated (Pochaczewsky 2002). The

following paragraphs review developments relating to DSC techniques for CTC, that have been reported in the literature.

A number of early techniques utilised a contrast enhanced bowel preparation in conjunction with a threshold based colon segmentation to achieve digital subtraction. Sheppard et al. (1999) reported on the feasibility of subtraction CTC (i.e. DSC) using an animal model with simulated polyps where a contrast enhanced bowel preparation was used. They employed a basic thresholding approach where voxels with values in excess of 200 HU were subtracted from the data set. The main problems associated with this approach were that the air/barium interface was still visible in the processed image and that no attempt was made to reconstruct a smooth representation of the colonic mucosa where the contrast enhanced material had been subtracted. In an evaluation of the digitally subtracted data sets, all simulated polyps greater than 5 mm were identified with a specificity of 88%. Callstrom et al. (2001) described a very simple threshold based technique that removed pixels with density values greater than 150 HU. Patients were instructed to take contrast material orally for 48 hours prior to the examination to ensure that the entire contents of the colon were labeled. Callstrom et al. evaluated their technique using a cohort of 56 patients. They found that polyp detection levels for stool labeling were comparable to levels achieved with conventional cathartic preparations.

Zalis & Hahn (2001, Massachusetts General) developed an early implementation of their contrast enhanced digital subtraction bowel cleansing CTC protocol that required no alterations to diet (with the exception of an orally administered contrast agent) and no laxatives. Bowel insufflation with room air was still required. Patients were administered an orally ingested contrast agent regularly for 48 hours prior to their CT scan. Following insufflation with room air, patients were scanned in both the prone and the supine positions. Images were analysed using a custom software system that operated independently of the display workstation. The resulting images were then processed using custom software to remove voxels representing the high contrast residual fecal material (and consequently voxels representing other high density regions, most notably bone). In addition, voxels representing the air/residual fecal material interface were also removed, thus eliminating the associated volume averaging artifacts associated with this boundary. The technique was evaluated using a cohort of five patients and colonic neoplasia were reported to be visible post processing using both 2-D axial and 3-D endoluminal views. No detailed results



**Figure 3.3:** An overview of the digital subtraction process. (a) A cropped axial image illustrating a section of the colon lumen containing contrast enhanced residual fluid. (b) The original image following a simple threshold based subtraction. Note the residue artifacts at the air/contrast enhanced material boundary and edge artifacts at the fluid/soft tissue boundary. (c) Removal of the residue artifacts using a subtraction mask derived from the location of strong edges in the image. (d) Reconstruction of the colonic mucosa using targeted Gaussian filtering at the contrast enhanced material/soft tissue boundary.

were provided due to the preliminary nature of the study. However, it is evident from the article that the recovered air/soft tissue boundary suffers from a severe edge artefact (Zalis & Hahn 2001, Figure 1 (d)). An example of this is evident at the recovered air/soft tissue boundary illustrated in Figures 3.3 (b) & (c).

In a later report on the subject, Zalis et al. (2003) described a revised version of their approach where pixels with attenuation values greater than 200 HU were subtracted from the images (these pixels represented contrast enhanced faecal material as well as bone and other high density material). The abrupt edges that resulted from the subtraction process were then identified using a Sobel edge detector. The response of the Sobel edge detector was found to be greater at the abrupt edges than at the naturally occurring edges which incorporated some degree of volume averaging. As an intermediate step the abrupt edges resulting from the subtraction process (which were identified using the Sobel edge detector) were replaced with a three pixel wide stair step contour. A natural air soft tissue transition was then simulated by applying a 2-D Gaussian filter to the pixels located in this contour region. The technique was evaluated by two experienced radiologists using a cohort of 20 patients scanned in the right decubitus position. Sixteen of the patients underwent digital subtraction bowel cleansing and the remainder (the control group of four patients) underwent traditional bowel cleansing. The readers gave a mean rating for image quality of the processed data sets of “interpretable with minimum subtraction artefact” (point 4 on a 5 point scale). Of the nine colonoscopy confirmed polyps, the first reader identified six and the second reader identified seven. The missed lesions were all  $\leq 10$  mm. Although not mentioned explicitly in the article, it is obvious that remnants of the boundary between the air and the contrast enhanced material are present in the processed images (Zalis et al. 2003, Figure 2 (b)). An example of a boundary or residue artifact is evident at the air/contrast-enhanced material boundary in Figure 3.3 (b).

Zalis et al. (2004) described a subsequent implementation of their digital subtraction bowel cleansing technique that dealt with the two main artifacts associated with the threshold based approach i.e. boundary artifacts (between air and contrast enhanced material) and edge artifacts, both of which are described in the previous paragraph and illustrated in Figure 3.3. Boundary artifacts were dealt with by specifying a mask for each axial image that consisted of voxels above a specified threshold and dilated edges above a certain threshold (i.e. the strong edges associated with

the air/contrast enhanced material interface). The resulting mask adequately covered all contrast enhanced material and could be used to effectively subtract it from the axial slice. The edge artifacts were dealt with using the surface reconstruction technique outlined earlier (Zalis et al. 2003). An evaluation of the technique was carried out using a cohort of 30 patients and one custom built phantom with 21 sessile shaped polyps ranging in size from 8 - 22 mm. A qualitative evaluation was carried out by two experienced radiologists comparing the 30 digitally subtracted data sets with 10 control cases where a standard bowel preparation was used. A small difference in quality was observed, however, this was deemed to be statistically insignificant. The phantom study was used to quantify the effect that digital subtraction had on polyp appearance. This was achieved by projecting rays from 10 of the synthetic polyps and sampling the density at regular intervals during the transition from colon wall to air. Samples were compared for an opacified, subtracted data set and an unopacified data set. Again, no statistically significant differences were observed.

Zalis et al. (2005a) evaluated the effect of DSC on polyp size measurement at CTC using the colon phantom from their earlier study (Zalis et al. 2004). A phantom was used in the study due to the ethical and logistical issues associated with performing repeated CT examinations on a human subject. Four different concentrations of contrast material were used, and the phantom was filled so that 11 of the 21 polyps were submerged or partially submerged for each of the scans. An additional scan was performed without any contrast material. This fifth and final scan was used as the control for the study. Two readers measured the sizes of the 11 relevant polyps in each scan. A subsequent comparison of the DSC measurements compared with the control measurements found no significant difference. The lowest mean measurement error was reported to be 0.6 mm. This was achieved in the case where the attenuation of the contrast-enhanced material was in the range 200 - 560 HU. The study demonstrated that DSC has no significant negative effect on reader performance at CTC.

Zalis et al. (2006) subsequently evaluated the effect of DSC on patient comfort and image readability. A total of 78 patients were used in this study. Ten patients, representing the control group, ingested polyethylene glycol electrolyte prior to imaging (i.e. a standard bowel preparation used in conventional colonoscopy), 25 patients ingested a barium sulfate solution (group 1), 21 ingested nonionic iodinated

contrast material (group 2) and 22 ingested a combination of nonionic iodinated contrast material and magnesium citrate (group 3). Each of the 78 patients underwent air insufflated CTC in the prone and supine positions where the same scanning parameters were used across the four study groups. A previously described DSC technique developed by the same group (Zalis et al. 2004) was used to remove the contrast enhanced material from the resulting CTC data sets. The level of patient discomfort, the bowel preparation quality, and the bowel preparation homogeneity were recorded for each of the patients enlisted in the study. The mean level of patient discomfort reported by the control group was significantly higher than the mean levels of discomfort reported by the patients from the three groups that underwent contrast enhanced bowel preparations. A subjective evaluation of bowel preparation quality found that the preparation used for group 3 patients was the most effective of the contrast enhanced approaches with reported readability scores comparable to those of the control group. A qualitative evaluation of bowel preparation homogeneity among the contrast enhanced patients found that the preparation used for group 3 patients yielded the most consistent residual material. This observation was verified quantitatively by determining the mean standard deviation of density values in the tagged fecal material across each of the three contrast enhanced groups. Again, the residue in group 3 patients was found to be the most consistent with a mean standard deviation of 78 HU. The authors point out that, although the use of DSC markedly reduces patient discomfort, the bowel preparation quality is still slightly inferior to the quality that is attainable with cathartic alternatives utilised in conventional colonoscopy. Consequently, further developments will be required to improve the bowel preparation quality provided by DSC and further evaluation will be required to establish the effect of DSC on polyp detection.

### 3.2.4 Segmentation Summary

Colon segmentation is the first processing task that is performed in a CTC system. An accurate segmentation is extremely important as it forms the basis for all subsequent processing and evaluation. It is clear from the previous review that the accuracy of colon segmentation techniques is improving and that the problems of over-segmentation and under-segmentation are proving to be less of an issue. The introduction of DSC represents a further improvement as it effectively eliminates the problems associated with residual faecal material in the colon which can resem-

ble polyps, or in extreme cases cause blockage. In addition, DSC also promises to make CTC a much more patient friendly examination by removing the need for a purgative bowel preparation and by reducing the radiation exposure significantly.

This Section has shown that there are various different approaches to colon segmentation and that these approaches ultimately yield either a representation of the colon lumen or a representation of the colon wall. The lumen is generally used when calculating the central path through the colon (see Section 3.3) or when extracting an isosurface representation of the colonic mucosa (e.g. when using the marching cubes algorithm (Lorensen & Cline 1987)). The wall representation is used in CAD algorithms that perform an analysis of the voxels surrounding the colonic mucosa in order to identify regions with polyplike characteristics (see Section 3.5.3.3). In any of these examples an accurate segmentation is required in order to ensure the that the correct outputs are obtained.

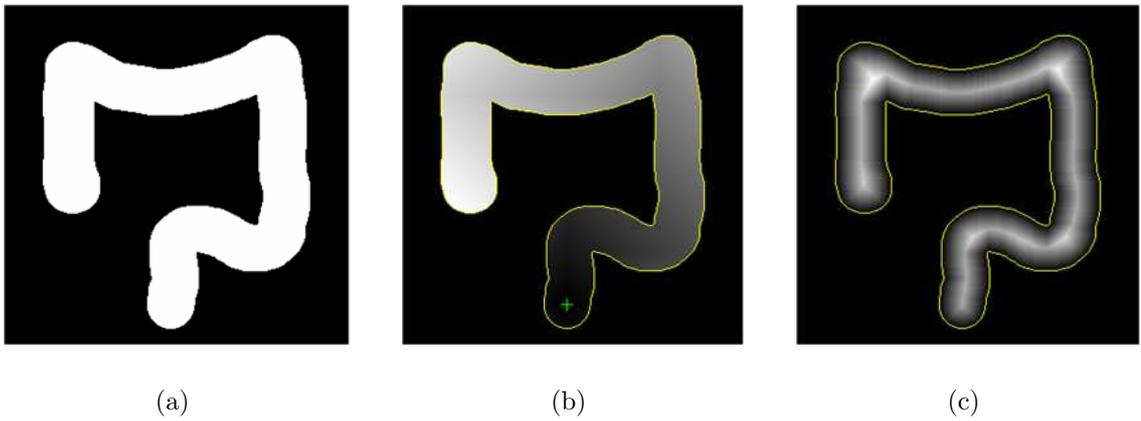
### 3.3 Centreline Calculation

A large number of centreline calculation algorithms are described in the literature. The common objective of each algorithm is to approximate the central path through the human colon for automating endoluminal navigation at CTC<sup>2</sup>. It is apparent from published results that this objective has been achieved with varying levels of accuracy, efficiency and user interaction. A review of previously published centreline calculation algorithms is presented in the following paragraphs, these algorithms are also summarised and compared (where possible) in Table 3.1.

Early techniques for centreline calculation required a significant level of user interaction. An example of this is described by McFarland et al. (1997, WUSM) where key points along the entire length of the colon were identified by a radiologist. The density of these points was greater in regions of increased tortuosity to insure that the resulting centreline resided within the colon lumen at all points. A cubic spline fit of the manually selected key points was ultimately calculated to approximate the colon centreline. A total of nine centrelines were generated by three radiologists for a single human colon in an average time of 18 minutes. A mean spatial error of 11 mm was reported and the worst fit centreline found to deviate outside the colon

---

<sup>2</sup>The centreline provides a compact representation of colonic structure and can be used in operations such as prone/supine registration (see Section 3.4) and colon lumen subsegmentation (see Section 4.6).



**Figure 3.4:** Distance Field generation: (a) A 2-D representation of the human colon. (b) A distance field initiated from a user defined seed point (green cross) in the rectum. (c) A distance from surface field initiated automatically from the set of surface points. White pixels are furthest from the source, whereas black pixels are closest to the source.

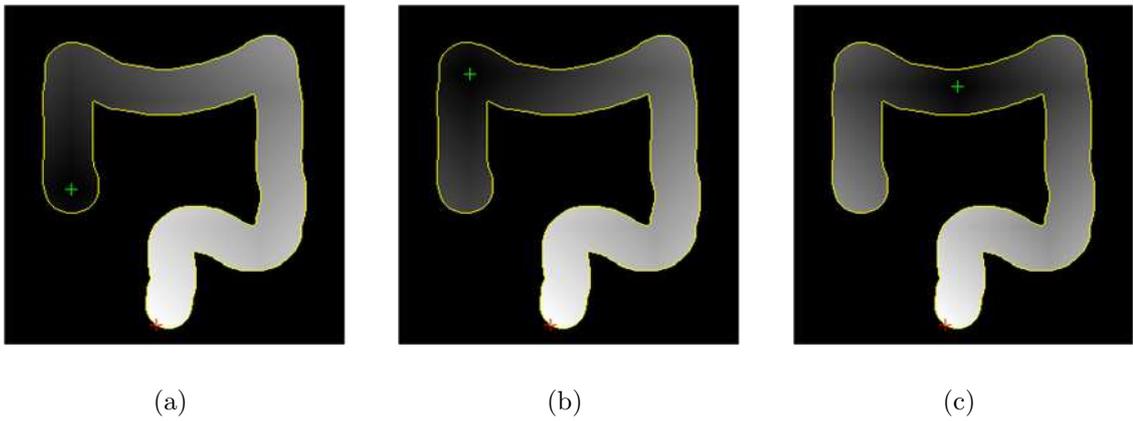
lumen on four occasions. Note that the spatial error was calculated by comparing each individual centreline with the average of the nine centrelines that were calculated during the study. More recently, automated centreline calculation algorithms have been published. These are much more efficient, generate more reproducible results and require considerably less user interaction. Automated techniques generally employ either distance fields (see Figure 3.4) or *3-D topological thinning* (3DTT).

Samara et al. (1998, UC) developed a centreline calculation algorithm that involved the generation of a distance from source field. This field was initiated from a user defined seed voxel in the rectum and effectively represented voxel wave fronts that radiated from the seed point. The centreline was subsequently approximated as the set of voxels that represented the centre of mass for each of the voxel wave fronts. Two approximations of the centreline were calculated from seed points that were located in both the rectum and the caecum. These antegrade and retrograde centrelines were then averaged to generate the final result. Centreline calculation using this technique was reported to take approximately one minute. Centreline points were indicated by two radiologists and the average RMS difference between indicated and calculated points was found to be 4.924 mm, this was less than the inter-observer RMS difference. In a revised version of this technique, Samara et al. (1999, UC) determined the orthogonal plane to each point in the averaged centreline. The centre of mass of voxels was calculated for each plane, and the resulting centroid points were then connected together to generate the final centreline. The revised centreline calculation technique required an average of five minutes for cen-

treline calculation and improved centreline accuracy for clinical cases by between 0.4 - 1.0 mm. In a phantom study the mean RMS errors between known and calculated centrelines was found to be between 1.0 - 1.5 mm.

Members of the Visualisation Lab from the State University of New York (SUNY) at Stony Brook have made a significant contribution to research in the area of centreline calculation for CTC. Their earliest publications described two different centreline calculation techniques. The first of these techniques (Hong et al. 1995, SUNY) was based on 3DTT (sometimes referred to as onion peeling) and involved the removal of successive layers of voxels from the colon lumen until only centreline points remained. A voxel was not removed if it was either one of the user defined endpoints, or if its removal results in a breach in connectivity between the two endpoints. The second technique (Hong et al. 1997) was based on potential fields. Two distance fields were generated, one represented the distance from target (located in the caecum or the appendix) and the second represented the distance from surface of the colon. These distance fields were combined to generate a potential field which was ultimately used to guide the user through the colon lumen. This technique facilitated automated navigation through the colon lumen. However, in cases where the central path was unable to escape from local minima, user intervention was required.

Zhou et al. (1998, UCLA and SUNY) developed a centreline calculation technique based on distance fields. This technique involved the generation of a distance from surface field using a segmented version of the colon lumen. The skeleton was identified from the distance field by identifying the set of clusters and the set of local maximum paths (LMpaths). A cluster was a set of at least edge connected local maximum voxels that had the same distance value, and a LMpath was defined as a special group of voxels that connected two clusters (see Zhou et al. (1998) for definition). User defined skeleton endpoints were then selected (ps and pt). The nearest points on the skeleton were then identified (ps1 and pt1). A distance from source field was then generated originating from ps1. The minimum cost path from pt1 back to ps1 (i.e. in the opposite direction of the distance field) was subsequently identified as the colon centreline. The technique was evaluated using the SUNY colon data set, previously used by Hong et al. (1997), and the entire centreline calculation process was reported to take 199 seconds. A more recent approach described by Zhou & Toga (1999, UCLA) involved generation of two distance fields



**Figure 3.5:** One endpoint (red cross) of the colon can be identified as the point with the largest distance value relative to the seed point (green cross); this process is independent of seed point location. It is evident that the same endpoint is identified with a seed point in (a) the caecum, (b) the hepatic flexure and (c) the mid transverse colon.

that represented a distance from surface field (i.e. a boundary-seeded distance field) and a distance from source field (i.e. a source-seeded distance field). A source point for the main source-seeded field was automatically identified by initiating an intermediate source-seeded distance field from an arbitrary point in the segmented lumen, the point with the highest distance value was subsequently identified as the source point (see Figure 3.5). In their paper, Zhou & Toga redefined the concept of a cluster as “a connected set of object voxels with the same source seeded distance value”. The centreline was then calculated using a two stage process that consisted of path identification and path centreing. The centreline was initiated from the cluster with the highest source seeded distance value (this fact was implied in their results section). The face neighbour of the starting point with the lowest source seeded distance value was used as the next point in the centreline and this process continued until the endpoint was reached. This first approximation of the centreline suffered from the corner cutting problem illustrated in Figure 3.6. The second stage was required to centre the initial centreline. Here each point in the approximated centreline was replaced by the medial point in the associated cluster. The resulting set of medial points represented the final centreline. This technique was also evaluated using the SUNY data set and the time required for centreline calculation was reported to be 519 seconds.

Chiou et al. (1998, 1999, SUNY) developed a centreline calculation technique based on the work by Zhou et al. (1998). They started by calculating distance fields for the segmented colon lumen representing the distance from the surface

(DFSrf) and distance from a user defined target point (DFTrg). Local maximum points/clusters were identified from the DFSrf field. These local maxima were subsequently referred to as attractors. Most of these attractors contributed to the centreline. However, some attractors were associated with irregularities in the colon and resulted in superfluous centreline loops. Therefore, the task of centreline calculation involved identifying and ultimately connecting the set of principal attractors. This was achieved by identifying the most powerful attractor and locating any attractors in its influence zone, i.e. its children. These children were then stored in a tree structure and the process was repeated for the next most powerful attractor in the tree until both the start and end principal attractors were identified, these were the attractors with the shortest Euclidean distances to the user defined start and end points. The principal attractors were ultimately connected to generate the final centreline. The technique was evaluated using the CTC data set originally used by Hong et al. (1997). The time required for centreline calculation was reported to be in the order of a few minutes, an exact figure was not provided.

Bitter et al. (2001, SUNY) described a general skeletonisation technique for volumetric objects that could also be applied to the colon for centreline calculation. This work improved on their earlier CEASAR<sup>3</sup> (Bitter et al. 2000) and TEASAR<sup>4</sup> (Sato et al. 2000) algorithms and consisted of two main stages: Data set reduction and centreline calculation.

1. The colon lumen was segmented from the abdominal CTC data set and the resulting data set was cropped. The cropping operation resulted in new volume that contained only the colon lumen plus a small boundary. A Euclidian distance from boundary field was calculated for the segmented colon lumen. This distance from boundary field was then used to calculate gradient vectors for each voxel in the segmented lumen. All local minima and local maxima in the gradient vector field were then flagged. The flagged voxels were connected along the gradient field vectors The resulting set of voxels represented a small subset of the original colon voxels (15 - 30 percent) which included all voxels that would ultimately represent the centreline.
2. A distance from source field originating from an arbitrary point in the set of

---

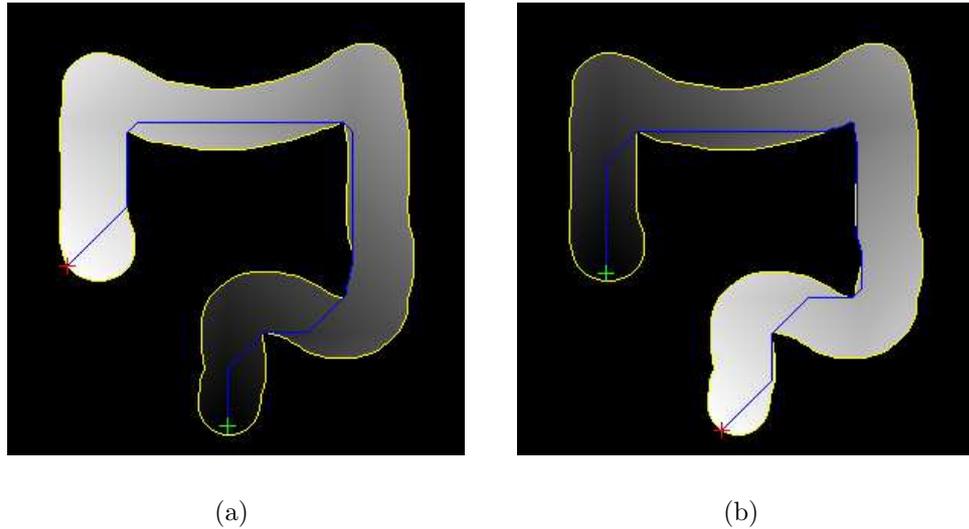
<sup>3</sup>A smooth, accurate and robust centreline extraction algorithm.

<sup>4</sup>Tree-structure extraction algorithm for accurate and robust skeletons.

flagged voxels was then used to identify one of the centreline endpoints, this was the voxel with the greatest distance value (see Figure 3.5). This voxel was referred to as the root voxel. A penalised distance from root field was then calculated. Here each voxel was assigned an accumulated cost, which was a combination of the Euclidean distance from the source and the penalty, and had an inverse relationship with the distance from the surface. The accumulated distance was also calculated for each voxel. This value was used to identify the other end of the centreline (i.e. the furthest point from the root voxel). The centreline was then identified as the minimum cost path from the endpoint back to the root. This graph reduction technique was an adaptation of the Dijkstra shortest path graph algorithm (Dijkstra 1959).

The technique was evaluated using five data sets, including a single CTC data set. The time required for centreline calculation of the CTC data set was reported to be 119 seconds.

Chen et al. (2000*b*, SUNY) described a multi-resolution approach to skeletonisation for virtual endoscopy. A simplified version of this approach was used to identify the centreline of the colon at CTC. This technique was based on the generation of a Euclidean distance field from a user defined root. This root voxel was defined as one end point of the centreline and the voxel with the maximum distance value was identified as the other end of the centreline (see Figure 3.5). Note that for the more general case of skeleton extraction, voxels with local maximum distance values were taken as skeleton endpoints. The minimum cost path was then identified by tracing from the endpoint back to the root. This path represented the shortest path between the two endpoints of the colon and, as such, suffered from the corner cutting problem (see Figure 3.6) and had to be centred. Voxels were selected from this initial centreline at regular intervals. The tangent was obtained for each of the selected voxels and the plane perpendicular to the tangent was determined. The centre of mass of lumen voxels that were contained in this cross-section was taken as a revised centreline voxel. The final centreline was generated by calculating the bi-cubic B-spline fit of the revised centreline points. Chen et al. included a data reduction technique based on the *one dimensional* (1-D) discrete wavelet transform. This technique was used to significantly reduce the size of the original CTC data set and thus greatly increase the performance of the technique. However, the au-



**Figure 3.6:** The distance from source field alone cannot be used for centreline calculation as the shortest path (i.e. the path with the shortest Euclidean distance) between the start and end points is identified (blue path). This problem, commonly referred to as corner cutting is illustrated for initial seed point in (a) the rectum and (b) the caecum.

thors point out that an over-reduction in size may compromise connectivity in the segmented colon lumen and compound the centreline calculation process. The technique was evaluated using three data sets that included a single CTC data set. In the case of the CTC data set, the centreline calculation process required a total of 36 seconds.

Wan et al. (1999, 2002, Boeing & SUNY) developed a technique similar to that described by Bitter et al. (2001), which operated on the entire segmented colon lumen. Using this technique an inverse Euclidean distance from the surface field was generated. The Dijkstra algorithm (Dijkstra 1959) was then employed to find the minimal cost path through a directed graph representation of the segmented colon lumen. The individual vertex weights in this directed graph corresponded to the distance value of the node at the end of the vertex. A source node (the start point) was selected by the user (Wan et al. 1999) or automatically (Wan et al. 2002) and the minimum cost from each node to the start point was calculated. This represented a minimum cost spanning tree. A distance from source field was also calculated and this was ultimately used to identify the corresponding endpoint. The centreline was finally identified as the path from the endpoint to the start point with the minimal cost. The technique was evaluated using four CTC data sets and a mean centreline calculation time of 14.7 seconds was reported. However, it should

be noted that the time required for the calculation of the distance from boundary and distance from source fields was not included in the documented results.

Deschamps & Cohen (2000, 2001, Philips & Université Paris IX Dauphine) described a navigation path identification technique for virtual endoscopy based on extracting minimum energy paths from 3-D images. This involved the calculation of a minimal action map that defined the minimal energy path between a user defined start point and any other point in the lumen. The energy for each path was calculated using a simplified snake energy model where the potential at a point in the image corresponded to the intensity value at that point. The path endpoint could either be user defined, or automatically detected as the last point added to the minimal action map, i.e. the furthest point from the start point (Truyen et al. 2001). The minimum energy path was subsequently identified from the minimal action map by backward propagation from the endpoint to the start point. At CTC the potential inside the colon lumen was effectively constant, and as a result the minimal energy path between the start and end points was the path between these two points with the shortest Euclidean distance. Consequently, this approach suffered from the corner cutting problem common to all distance field techniques (see Figure 3.4). This problem was alleviated by introducing a centreing potential that decreased with distance from the colon wall. The resulting minimal path, calculated in the presence of this centreing potential, provided a more accurate representation of the colon centreline. The technique was evaluated using a single CTC data set. The time required for centreline calculation was reported to be 30 seconds.

Ge et al. (1996, Wake Forest) described a centreline calculation technique that was based on 3-D skeletons. An initial skeleton was generated using 3DTT constrained by the centres of maximally inscribed spheres. The user was subsequently required to select the centreline endpoints from the skeleton. Thinning was applied a second time with these two endpoints as the only anchor points. The user was then required to identify points on extraneous skeletal loops. These points were subsequently removed prior to the third and final application of the 3DTT algorithm. The resulting skeleton represented the centreline of the colon. More recently, Ge et al. (1999) described an improved centreline calculation algorithm also based on 3DTT where the user was only required to select the two centreline endpoints located in the rectum and in the caecum. The task of extraneous loop reduction was performed automatically. In addition, optimisation techniques including sub

sampling and surface voxel tracking were utilised to improve performance. Ge et al. used a total of 20 CTC data sets to evaluate their technique. They reported that centreline calculation required an average of approximately 518 seconds. However, when subsampling was used the average time for centreline calculation reduced to approximately 60 seconds.

Sadleir & Whelan (2002, VSG) developed an approach based on 3DTT (Tsao & Fu 1981) which extended the work by Ge et al. (1999). An endpoint in the rectum was automatically identified by examining the slice at the bottom of the data set. This endpoint was subsequently used as the seed point for region growing and the generation of a distance from source field. Both operations occurred simultaneously and facilitated colon lumen segmentation and the detection of the second endpoint in the caecum. Standard 3DTT is an extremely slow process. This is mainly due to the extensive local neighbourhood analysis that must be performed prior to the removal of each surface voxel. In order to enhance the performance of 3DTT, Sadleir & Whelan generated every possible  $3 \times 3 \times 3$  binary neighbourhood ( $2^{26}$  combinations) and determined if the central voxel could be deleted without compromising the topology of the object being thinned. The results were then stored in a lookup table where the unique index to a particular result was generated as a function of the associated neighbourhood configuration. This table was calculated only once and required 67 MB of memory. Subsequently, thinning could be performed much more efficiently by consulting the lookup table as opposed to performing an extensive local neighbourhood analysis prior to the conditional removal of each surface voxel. The thinning operation resulted in a skeleton representation of the segmented colon lumen. The skeleton was then reduced to the centreline, i.e. the set of points between the two automatically detected endpoints that is furthest from the original colon wall, this was achieved by pruning any extraneous loops that were present in the skeleton. A more comprehensive evaluation of this technique was subsequently published by the same authors (Sadleir & Whelan 2005). This later work also provided a more detailed description of the automatic endpoint detection technique that was used for the detection of endpoints in the rectum and in the caecum. The technique was evaluated using a total of 12 CTC data sets that represented a mixture of prone and supine studies. The optimised approach to centreline calculation required an average of 3.438 seconds. This was reported to be approximately 45 times faster than the unoptimised approach which required an average of 155 seconds for centreline

calculation.

Paik et al. (1998, Stanford) proposed a novel variation of 3DTT for centreline calculation that did not require extensive local neighbourhood analysis. Their technique involved identifying the shortest path between two user defined endpoints, that consisted of only surface voxels. The object was then thinned using a three stage process.

1. The surface voxels were removed in parallel. No local neighbourhood testing was performed and all voxels were removed regardless of whether they were endpoints or their removal resulted in a connectivity breach.
2. A revised shortest path between the start and end voxels was then obtained. This was calculated using a union of the previous shortest path voxels and the new surface voxels. Voxels on the new surface were preferred over voxels on the old shortest path.
3. The iterative process continued until only the centralised path remained.

This approach effectively initialised the centreline on the surface of the segmented colon lumen and iteratively centred it using an extremely efficient 3-D thinning algorithm, until the final result is obtained. The technique described by Paik et al. was evaluated using five data sets that included one CTC data set. The time required for centreline calculation in the CTC data set was reported to be 759 seconds. However, it should be note that this value included the time required for colon lumen segmentation (129 seconds).

Iordanescu & Summers (2003, NIH) presented a method where the centreline was calculated from a polygonal mesh representation of the colon surface. This approach differed from conventional methods, as it operated on the colon surface rather than on the lumen. The authors cited applications of their approach including the registration of polyps between prone and supine data sets and the estimation of local colonic distention. The algorithm started by extracting an isosurface representation of the colon surface using the marching cubes algorithm (Lorenson & Cline 1987). Decimation was then used to reduce the amount of data (i.e. the number of vertices) thus increasing the performance of the technique. A reduction in vertices of approximately 80% was reported. An approximation of the centreline was obtained by iteratively averaging the vertex locations in the decimated colon model.

The thinned structure was sampled at regular intervals ( $\sim 10$  mm) to generate a line strip array representation of the thinned surface. This was achieved using a vertex based region growing technique. Pairs of points on the line strip array and their associated distance from rectum values were mapped to cylindrical segments of the decimated surface model. The edge vertices of each of the cylindrical segments were averaged and the resulting set of points represented the final centreline. A post-processing stage was used to deal with any points that fell outside the colonic lumen. Finally, points in the decimated colon were mapped back to the original colon to facilitate polyp matching between prone and supine data sets. The technique was validated quantitatively using a computer generated phantom and qualitatively using data sets from 10 patients scanned in both the prone and the supine positions (i.e. a total of 20 data sets). The average time required for centreline calculation was reported to be 108 seconds. The phantom study demonstrated the accuracy of the approach (the average deviation from the known centreline was reported to be  $1.1 \pm 1.0$  mm). The use of real patient data demonstrated the ability to accurately register 11 polyps in seven patients between prone and supine data sets using the distance from the rectum along the normalised centrelines.

Frimmel et al. (2004, Uppsala & UC) presented a centreline calculation technique that involved the analysis of a distance map of the segmented colon lumen. An initial set of local maxima were identified from the distance map based on an inspection of local four and six-connected neighbourhoods. This set was refined by associating an exclusion zone with each local maximum, where the value of the local maximum defined the radius of the exclusion zone. A threshold was then used to identify seed points from the remaining local maxima. These seed points were connected, according to a set of rules<sup>5</sup>, to form graphs representing minimum distance spanning trees within the colon lumen. The other local maxima (i.e. those below the threshold) were appended to the resulting graphs to form branches. All branches were then eliminated except those representing graph endpoints. Finally, any graphs that did not contain a seed point were removed as they were considered to be associated with extracolonic structures. The resulting set of graphs represented the final colon centreline. The proposed algorithm was evaluated against a computer generated phantom with a known centreline and centrelines identified

---

<sup>5</sup>These rules ensured that the resulting graph, representing the shortest path between the seed points, did not contain any loops, and did not go through the colon wall at any point.

from real patient data by three radiologists. In all cases the deviation of the calculated centreline from the gold standard was small. In the case of the phantom studies, where the true centreline was known, the mean displacement between the true and calculated centreline was typically of the order of 3 mm. In the clinical cases the average value for the mean displacement between the automatically calculated centreline and those indicated by the three radiologists was 3.8 mm. The performance of the technique was evaluated using a total of 40 CTC data sets and the average time required for centreline calculation was reported to be 10.5 seconds. Note: Frimmel et al. (2005, see Section 3.2.1) subsequently published a colon lumen segmentation technique using their centreline calculation algorithm.

Bouix et al. (2005, Harvard Medical School) described a fully automatic centreline extraction algorithm based on their previous work in the area of skeletonisation. The algorithm started by calculating the *average outward flux* (AOF) at each point in the segmented colon lumen. The AOF was lower for voxels further away from the surface of the colon and decreased to a strictly negative value at the centre. The centreline calculation process involved the sequential removal of simple voxels<sup>6</sup>, that were not endpoints and had an AOF value above a specified threshold. The order in which voxels were examined was determined by their Euclidean distance from the colon surface. Finally, the centreline was extracted from the resulting medial curve using the Dijkstra shortest path algorithm (Dijkstra 1959). This was achieved by identifying the longest of the shortest paths between any two endpoints of the medial curve. The accuracy of the technique was established using six synthetic shapes with known centrelines. The percentage overlap, the average distance and the maximum distance between points on the computed and known centrelines were calculated in each case. The resulting average values of 91%, 0.11 mm and 2.67 mm confirmed the validity of their approach. The performance of the technique was evaluated using a range of data sets that included one CTC data set. A centreline calculation time of 31 seconds was reported for the CTC data set.

Kang & Ra (2005) proposed a centreline calculation technique that was intended to maximise the visibility in virtual flythrough examinations of CTC data sets. Their algorithm started by obtaining a segmentation of the colon using a seeded region growing algorithm. An initial centreline was then obtained from the segmented

---

<sup>6</sup>i.e. points whose removal did not (1) create a hole (2) create a cavity or (3) disconnect a connected component.

lumen using the distance field technique developed by Wan et al. (2002). An approximation of the thickness of the colon at each point along the centreline was obtained using a 3-D distance transform calculated using the 3-4-5 chamfer based method. The endoluminal visibility was subsequently maximised by moving the view point back from the current centreline point so that the full diameter of the colon at that point could be visualised within the limited field of view. Potential problem cases where the viewpoint was moved outside the lumen or behind an obstruction were avoided by ensuring that there was a clear line of sight between the viewpoint and the current centreline point. The path traced by the viewpoint, representing the improved visibility centreline, was smoothed using a low-pass-filter. Finally the up-vectors were selected to minimise the rotation between consecutive images. This new centreline facilitated the visualisation of the colon surface in highly curved sections of the colon lumen. The visualisation of regions obscured by the haustral folds was made possible by the use of bidirectional (antegrade & retrograde) endoluminal navigation. The documented approach was tested using three patient data sets and enabled the visualisation of an average of 97% of the colon surface for a camera field of view of 60 degrees. This represented an average improvement of 4.74% over the conventional technique with the same field of view. The average time required for centreline calculation in the three data sets was reported to be 18.667 seconds.

### 3.3.1 Centreline Calculation Summary

The centreline is an important feature of the colon lumen as it can be used to automatically guide the viewer on a flythrough examination of the interior of the colon. The centreline is not limited to automating endoluminal navigation and has been used for other purposes e.g. to register polyps between prone and supine data sets. The centreline is generally calculated from a binary representation of the segmented colon lumen, with the exception of the method described by Iordanescu & Summers (2003) that operated on a polygonal mesh representation of the colonic mucosa. Lumen based techniques use either distance fields or 3DTT to identify either the centreline of the colon lumen or the skeleton of the colon lumen (which is a superset of the centreline). In cases where the skeleton is obtained it must be reduced to the centreline by pruning unnecessary extraneous loops. The end points of the centreline can be manually selected or automatically detected. In the case of automatic detection the rectum point is more readily detectable due to its invariant position. The

caecum point can then be detected by identifying the furthest point from the rectum by means of a distance field initiated from the rectum point. However, this method of caecum point detection is only possible in the case of well distended colons, without blockage, that do not include other air filled regions located in the abdomen due to over segmentation. Recently published centreline calculation techniques require less than 60 seconds to reduce the colon lumen to its centreline representation (see Table 3.1) and as a result these techniques can be used to facilitate a real time flythrough examination of the colon in a clinical environment. The question of accuracy has been addressed by comparing the calculated centreline to that of a known centreline in computer generated phantom studies (e.g. Bouix et al. (2005)) and by comparing the calculated centreline to radiologist markings (e.g. Frimmel et al. (2004)). The approach by Kang & Ra (2005) represents a further development in centreline calculation research where the calculated centreline represents the path that provides the best visibility as opposed to the more central path through the colon<sup>7</sup>. The area of CTC research dealing with centreline calculation has evolved considerably from the semiautomated approach proposed by McFarland et al. (1997) to the fully automated, real time approaches that have been documented more recently.

---

<sup>7</sup>It should be noted at this point that the centreline calculation technique documented in this thesis (see Section 4.5) is intended to provide an accurate representation of the structure of the colon, i.e. it is not intended to maximise the visibility of the colon surface. The centreline is ultimately used to identify the colonic flexure points and thus enable the subsegmentation of the colon lumen (see Section 4.6).

**Table 3.1:** An overview of previously published centreline calculation algorithms. A rank has been calculated for each algorithm by scaling the reported execution time based on the processor speed. This is intended to provide an indication of algorithm performance that is independent of the hardware platform used. Instances where the required information is unavailable are indicated by '-'. The accuracy of each technique, compared to a known phantom centreline (P) and/or radiologist markings (R), is presented where available. The order in which the techniques are listed is based on the associated centreline calculation speed.

Group	Technique	Accuracy (mm)	Platform	CPU (Mhz)	RAM (Mb)	EPs	Time (s)	Rank
Sadleir & Whelan (2005)	Optimised 3D topological thinning using lookup tables.	-	Intel	1600	512	0	3.438	1
Frimmel et al. (2004)	Identification and connection of distance field local maxima followed by spline interpolation.	~3 (P), 3.8 (R)	Intel	800	512	0	10.5	2
Wan et al. (2001)	Minimum cost path through a distance from surface field with manual endpoint selection.	-	Intel	700	655	1	14.75 <sup>a</sup>	4
Wan et al. (2002)	Minimum cost path through a distance from surface field with automatic endpoint detection.	-	Intel	700	655	0	18.205 <sup>a&amp;b</sup>	5
Kang & Ra (2005)	A technique for maximising endoluminal visibility based on the distance field method proposed by Wan et al. (2002).	-	Intel	2000	0	-	18.667	8
Sadleir & Whelan (2002)	Optimised 3D topological thinning using lookup tables.	-	Intel	700	512	0	24.42	7
Deschamps & Cohen (2000, 2001)	Minimum energy path through a distance from source field with a centreline stage.	-	Sun/ Solaris	300	1024	1	30	3

Continued on next page...

Table 3.1: continued from previous page

Group	Technique	Accuracy (mm)	Platform	CPU (Mhz)	RAM (Mb)	EPs	Time (s)	Rank
Bouix et al. (2005)	Skeletonisation algorithm that utilises voxels properties including the average outward flux and the Euclidian distance from boundary.	0.11 (P)	Intel	2000	1024	0	31.0	11
Chen et al. (2000b)	Minimum cost path through a distance from source field followed by a centreing stage.	-	SGI	195 × 2 (R10000)	896	1	36	6
Samara et al. (1998)	Distance form source field with cluster centreing.	-	SGI	-	-	2	~60	-
Iordanescu & Summers (2003)	Averaging of colon surface vertices followed by a set of refinement stages to yield final result.	1.1 ± 1.0 (P)	Intel/Windows	1500	2048	0	108	13
Bitter et al. (2001)	Centreline calculated as minimum cost path from end to start point in a penalised distance field.	-	Intel/Windows	1000	-	0	119	-
Zhou et al. (1998)	Skeleton extraction using a distance form surface field followed by centreline identification.	-	SGI	-	-	2	199	-
Sato et al. (2000)	Centreline calculated as minimum cost path from end to start point in a penalised distance field.	-	SGI	194 (R10000)	-	0	208	9
Bitter et al. (2000)	Centreline calculated as minimum cost path from end to start point in a penalised distance field.	-	SGI	194 (R10000)	4096	0	276	10
Samara et al. (1999)	Distance from source field with cluster centreing and centreline refinement stage.	4 - 5 (R), 1.0 - 1.5 (P)	SGI	-	-	2	~300	-

Continued on next page...

Table 3.1: continued from previous page

Group	Technique	Accuracy (mm)	Platform	CPU (Mhz)	RAM (Mb)	EPs	Time (s)	Rank
Ge et al. (1999)	3-D topological thinning used in conjunction with surface voxel tracking.	-	SGI	NA	-	2	518(60 <sup>c</sup> )	-
Zhou & Toga (1999)	Distance from source field with cluster centres used as centreline points.	-	SGI	- (R10000)	-	0	519	-
Paik et al. (1998)	A combination of a distance from source field and streamlined 3D topological thinning.	-	SGI	180 (R5000)	256	2	759 <sup>d</sup>	12
McFarland et al. (1997)	Radiologist marking & spline interpolation.	11 (R)	-	-	-	2	1080	-
Chiou et al. (1998, 1999)	Identify and connect principal attractors, i.e. local maxima in distance from surface field.	-	SGI	-	3072	2	Exact figure unavailable	-
Hong et al. (1997)	Uses a potential field, derived for distance from source and distance from surface fields, to identify the central path through the colon lumen.	-	-	-	-	-	-	-
Hong et al. (1995)	Reduce the segmented colon lumen to a single path of connected voxels using onion peeling (i.e. 3DTT).	-	-	-	-	-	-	-

Abbreviations: *EPs* The number of colonic endpoints that must be manually specified. An EP count of zero indicates that the associated centreline calculation technique is fully automated.

<sup>a</sup>Does not include the time required to generate the distance fields (approximately 20 seconds).

<sup>b</sup>The same results as Wan et al. (2001), however, the time required for extraneous branch removal (i.e. 3.46 seconds) is included.

<sup>c</sup>Obtained by subsampling the data set.

<sup>d</sup>Includes time required for segmentation stage.

## 3.4 Prone/Supine Registration

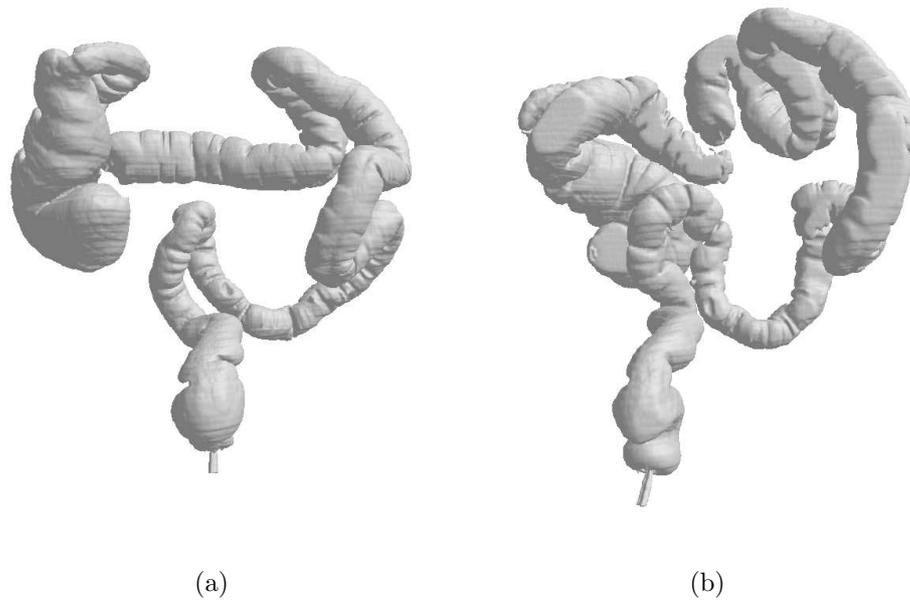
The registration of prone and supine CTC data sets can be used for a number of purposes e.g. to match polyp candidates between prone and supine studies or to identify the presence of mobile residual stool (Näppi et al. 2005a, UC). The task of accurately registering prone and supine data sets is challenging in cases where the colon is severely deformed or contains several collapsed regions. An example of the significant differences between scans of the same patient using prone and supine positioning is illustrated in Figure 3.7. This section reviews a number of methods that address the registration problem. The process of prone/supine registration is discussed at this point due to its relevance to subsegmentation, a process that will be introduced in Section 4.6.

### 3.4.1 Centreline Based Registration

Acar et al. (2001, Stanford) described a method for automatically registering points between the prone and supine centrelines that were calculated using a previously published algorithm (Paik et al. 1998). Their technique involved matching local extrema in the  $x$ ,  $y$  &  $z$  directions. The first stage in the process involved scaling the supine centreline points in a linear manner so that the supine centreline was the same length as the prone centreline. The extrema were then calculated and neighbouring extrema were clustered to simplify the matching problem. The clusters were then ordered and matched based on their distance along the centreline and their perpendicular distance from the horizontal plane. The scaling was then modified in a nonlinear fashion so that all of the extrema matched and not just the two endpoints as described earlier. Using experimentally obtained optimum parameters for their system, Acar et al. compared the locations of 25 reference points (five reference points in five patients). They found that the mean registration error decreased significantly from 51.0 mm down to 14.2 mm after the application of their registration algorithm.

Li et al. (2004, Stanford) extended the approach described by Acar et al. (2001) to include polyp registration. Three polyp properties were used for this purpose:

1. The distance of the polyp along the central path (after path registration). An approximation of the relevant path point was found by identifying the path point with the shortest Euclidean distance to the polyp.



**Figure 3.7:** Surface renderings of the colon of the same patient (MMH-166) imaged using (a) prone and (b) supine positioning. Note significant deviation in the morphology of the colon between the two images.

2. The location of the polyp, represented as a 3-D orientation vector, on the intersecting plane that was perpendicular to the central path. A transformation was used to reduce variations caused by colonic distortions between prone and supine scans.
3. The shape characteristics of the polyp. The CAD intensity score from the surface normal overlap polyp detection algorithm (Paik et al. 2004) was used for this purpose. A higher score indicated a higher likelihood of the candidate being an actual polyp.

A match score for a pair of prone/supine polyp candidates was then calculated from a joint probability density function that was approximated using the density functions for each of the three polyp characteristics listed above. In each case the individual density functions were estimated using a reference standard (in this case all data sets except for the one being evaluated). A total of three experiments were carried out to evaluate the polyp registration algorithm. The first experiment used only known polyps and resulted in a 78% success rate. In the second experiment when nonpolyps were included the success rate dropped to 50%. The final experiment determined the number of manual comparisons required to identify polyp pairs. This was found to be 1.78 with polyp registration compared to 4.28 comparisons

without registration.

### 3.4.2 Region-Growing Based Registration

Näppi et al. (2004a, UC) described an automated method for establishing a correspondence between the segmented colons of prone and supine data sets of a particular patient. Their approach involved automatically identifying a series of landmarks (anus, rectosigmoid junction, sigmoid junction, splenic flexure, hepatic flexure and caecum) in both data sets. Then, each voxel in the segmented colon lumen was assigned a value that represented its distance from the caecum. The distance values of the voxels in the prone and supine lumens were subsequently normalised based on the distance values at the landmarks. The prone and supine lumens were then divided into segments where a segment was located between two landmarks. Small segments (below a predefined minimum length) were merged, and large segments (above a predefined maximum length) were split into smaller overlapping segments. This registration technique was developed to reduce false positives at CAD-CTC. If a polyp candidate is flagged in only one of two corresponding segments then it is disregarded as mobile stool. Only polyps that are present in both corresponding segments were retained. An evaluation of their technique on 121 prone and supine CTC data sets (where 28 patients had 42 confirmed polyps 5 - 25 mm in size) demonstrated a reduction in the number of false positives by 20%.

In a subsequent publication Näppi et al. (2005a) provided a more detailed description of their approach dealing with landmark detection, segment generation, true positive retention and parameter estimation. A more thorough evaluation of their approach using the same patient cohort as their earlier study and a slightly different CAD scheme demonstrated that prone/supine correspondence reduces false positives by 19% (or from 3.0 per data set to 2.4 per data set). The authors note that only 41% of potential landmarks were detected using their technique thus the ability to eliminate potential false positives was greatly reduced.

## 3.5 CAD-CTC

At present clinical CTC examinations are read manually. This involves a radiologist examining all of the images that constitute the data set using 2-D, 3-D or a combination of both viewing techniques (see Section 2.2.5). Manually reading CTC

data sets is a slow process that is prone to error. CAD algorithms can be used in CTC examination to act as a second reader. The utilisation of CAD in this manner has the potential to increase the diagnostic performance of readers in the detection of colorectal polyps and decrease the variability in the diagnostic accuracy among readers (Yoshida & Dachman 2005, UC).

### 3.5.1 The Role of CAD at CTC

CAD-CTC is intended to be an assistive tool that will aid radiologists in the detection of colorectal polyps and masses. A number of studies have been carried out to establish how CAD-CTC will affect the diagnostic performance of radiologists involved in colorectal cancer screening.

Summers et al. (2002, NIH) compared the performance of two radiologists who examined a cohort of 40 patients (20 normal and 20 with at least one polyp  $\geq 10$  mm). Each study was examined with and without the help of their CAD system (described in Section 3.5.3.2) and the results were ultimately compared with same day conventional colonoscopy findings. Of the 65 polyps present in the study, 25 were  $\geq 10$  mm (large polyps). The study found that CAD increased the polyp detection sensitivity for all polyps by 16% and for large polyps ( $\geq 10$ mm) by 33%.

Mani et al. (2004, Stanford) evaluated the performance of three radiologists with and without the aid of CAD results. A total of 41 data sets were used containing a total of 30 polyps. The CAD system employed was that described by Paik et al. (2004) (discussed in Sec. 3.5.3.4). Unassisted readers detected on average 63% of polyps  $\geq 10$  mm. This increased to 74% when CAD was used.

Näppi et al. (2005*b*, UC) reported on the complementary effect of their CAD-CTC system (described in Sec 3.5.3.3) when used by experienced radiologists. In their study the colon was examined by three observers (an experienced radiologist, an inexperienced radiologist and a gastroenterologist). The observers evaluated a series of 20 CTC data sets consisting of 10 normal data sets and 10 abnormal data sets containing 11 polyps ranging in size from 5 - 12 mm. The data sets were evaluated with and without the benefit of CAD. The CAD scheme coloured the colon lumen so that folds were assigned a pink colour, the background was assigned a dark red colour and regions detected by the CAD algorithm were given a green colour. This colouring scheme was referred to by the authors as shape-scale colouring. The performance of the three observers was evaluated both with monochromatic colour-

ing and shape-scale colouring using *receiver operator characteristic* (ROC) curves, and in all cases the performance of the observer increased with the aid of the CAD scheme. An earlier report from the same group (Okamura et al. 2004) also demonstrated improved performance by human readers when assisting CAD was used.

Taylor et al. (2006, UK) compared the performance of a cohort of radiologists, with the performance of *computer-assisted reader* (CAR) software, for the detection of colorectal polyps. The CAR software, ColonCAR version 1.2, identified potential polyps using sphericity and protrusiveness measurements calculated from the axial slices of a CTC data set. The software was trained using prone and supine data sets from 45 patients with 100 colonoscopy confirmed polyps  $\geq 3$  mm in size, and evaluated using prone and supine data sets from an additional 25 patients with 32 polyps  $\geq 5$  mm in size. These additional data sets were also independently evaluated by three radiologists who had extensive CTC experience. In all cases a polyp was deemed to be detected if it was identified in either the prone data set or the supine data set. The ColonCAR software identified 15 (75%) of the 20 polyps 6 - 9 mm in size and 11 (92%) of the polyps  $\geq 10$  mm in size. The overall sensitivity was 81% and there were an average of 13 false positive detections per patient (i.e. per prone and supine data set pair). The first and second radiologists achieved individual overall sensitivities of 88% (28/32) and 91% (29/32) respectively. However, the third radiologist achieved a much lower overall sensitivity of 56% (18/32). As a result, the average overall sensitivity achieved by the radiologists was 78%. This was slightly lower than the sensitivity reported for the ColonCAR software. The authors observed that the ColonCAR software identified all seven polyps missed by the first and second radiologists, and 12 of the 14 polyps missed by the third radiologist. Consequently, if the ColonCAR software was used as a second reader, the three radiologists would have achieved sensitivities of 100%, 100% and 94% respectively.

### 3.5.2 Task Division: Primary versus Secondary Techniques

Early CAD-CTC systems provided high sensitivity, however, the level of false positive detections was unacceptably high. Göktürk et al. (2001, Stanford) recognised this problem and proposed the use of a two stage CAD-CTC system to reduce the number of false positive detections to a reasonable level. A primary stage is used to identify a list of initial polyp candidates and a secondary stage refines this list,

increasing specificity while maintaining the sensitivity of the primary stage. Most of the more recent polyp detection techniques utilise this two stage approach and this will be evident in the review of published CAD-CTC algorithms that follows.

### 3.5.3 Colorectal Polyp Detection Techniques

There has been a great deal of activity in the area of CAD-CTC research. Summers (2003) identified the main contributing institutions as:

- The National Institutes of Health (NIH)
- Department of Radiology, University of Chicago (UC)
- Department of Radiology, Stanford University
- Wake Forest University School of Medicine
- University Hospital Gasthuisberg (UHG), Leuven, Belgium
- State University of New York (SUNY)
- Vision Systems Group (VSG), Dublin City University

The following discussion will review the achievements of these and other groups involved in CAD-CTC research. Each of the reviewed studies will be summarised and compared where possible in Table 3.2. By comparing the performances listed in this Table with those for manually read CTC (see Table 2.1) it is evident that CAD has much to offer radiologists, particularly in the role of a second reader at conventional CTC.

#### 3.5.3.1 Wall Thickness Analysis

A number of early CAD-CTC techniques involved an assessment of the colon wall to identify regions of abnormal thickness. These techniques were developed by the Wake Forest group i.e. the group that originally demonstrated the CTC technique (Vining et al. 1994).

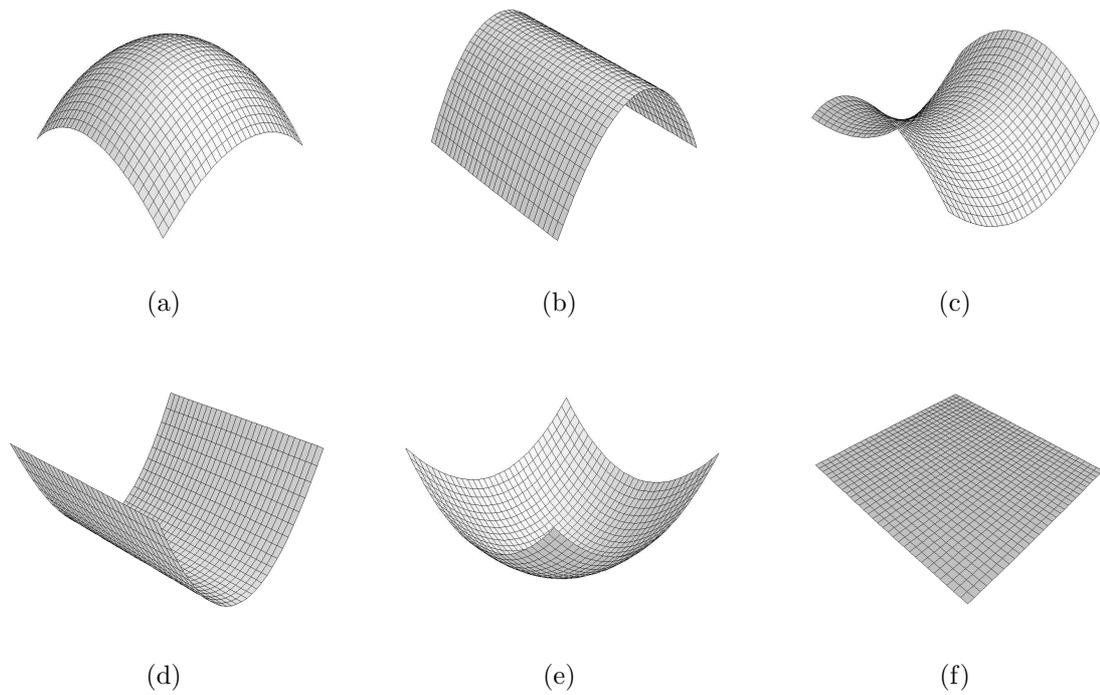
Hunt et al. (1997, Wake Forest) provided the earliest description of a software system that automatically highlights potential colorectal lesions in CTC data. Their approach involved measuring the colon wall thickness. This was achieved by initially fitting a triangular mesh to the air/soft tissue interface representing the colonic

mucosa (inner surface). Then at each vertex in the mesh, by traveling along the associated normal vector away from the colon lumen, the corresponding outer surface vertex was identified where either a soft tissue/air or a soft tissue/fat interface was encountered. A maximum search depth was specified in the event that neither type of interface was found in the vicinity of the inner vertex. In regions of abnormal wall thickness the convexity of the inner surface was measured and convex patches greater than a certain size threshold were ultimately highlighted as potential polyps. Hunt et al. reported that an evaluation of their technique on a cohort of five patients with known polyps yielded no false negative detections. However, the authors also reported that their technique generated numerous false positive detections. The number of actual polyps or false positive detections was not documented in their paper.

An extended version of this technique was presented by the same group in a later publication (Vining et al. 1998, Wake Forest). As before, an approximation of the internal and external surface of the colon was estimated. Vertices with abnormal wall thickness, increased curvature and increased convexity were considered to be abnormal (i.e. representative of polyps and masses). Patches of abnormal vertices with a surface area  $\geq 50 \text{ mm}^2$  were flagged as initial polyp candidates. A shape filter was then used to eliminate elongated candidates caused by the haustral folds to yield the final list of polyp candidates. Colour coding was used in conjunction with their *FreeFlight* virtual endoscopy system (Vining et al. 1997) to highlight the presence of potential colorectal lesions. In a cohort of 10 patients with a total of 11 polyps ranging in size from 7 - 40 mm the system was found to have a sensitivity of 73% with an average of 49.1 false positive detections per data set (range 9 - 90).

### 3.5.3.2 Isosurface Vertex Curvature Analysis

A number of CAD-CTC techniques that are described in the literature involve the extraction of an isosurface, in the form of a polygonal mesh, that represents the colonic mucosa. Each vertex in the mesh is then assigned the shape characteristics that correspond to its location in the original data set. The resulting mesh is then evaluated and suspicious regions with polyplike characteristics are flagged for further evaluation. The main group involved in this research is based at the National



**Figure 3.8:** The six curvature classes: (a) peak or cap,  $H < 0$ ,  $K > 0$ ,  $R = 1.0$ , (b) ridge,  $H > 0$ ,  $K > 0$ ,  $R = 0.75$ , (c) saddle,  $K < 0$ ,  $R = 0.5$ , (d) valley or rut,  $H > 0$ ,  $K < 0$ ,  $R = 0.25$ , (e) pit or cup,  $H > 0$ ,  $K > 0$ ,  $R = 0.0$ , (f) plane,  $H = 0$ ,  $K = 0$ ,  $R$  is indeterminable (Koenderink 1990).

Institutes of Health<sup>8</sup>. The six different curvature classes and their curvature characteristics (*Gaussian curvature* ( $K$ ), *mean curvature* ( $H$ ) and *shape index* ( $R$ )) are illustrated in Figure 3.8.

Summers et al. (2000, NIH) evaluated the feasibility of an early shape-based automated polyp detector using a single CTC data set that was modified to include 10 simulated polyps all approximately 10 mm in size. Two main experiments were carried out to evaluate different polyp detection criteria. The second experiment, designed to enhance the profile of polyps, diminish the profile of haustral folds and improve conspicuity between polyps and adjacent folds, yielded the best results with a sensitivity of 80% and no false positive detections. The first stage of this algorithm involved identifying candidate polyps. This was achieved by locating regions of the colon that were spherical in shape and that projected into the colon lumen. This was referred to as the primary shape criterion. The process of calculating the relevant vertex curvature information is described in detail elsewhere Summers (2000, Chapt.

<sup>8</sup>The work of the NIH group in automated polyp detection at CTC extends from their earlier work dealing with the computer assisted detection of polypoid lesions of the airways at virtual bronchoscopy (Summers et al. 1998).

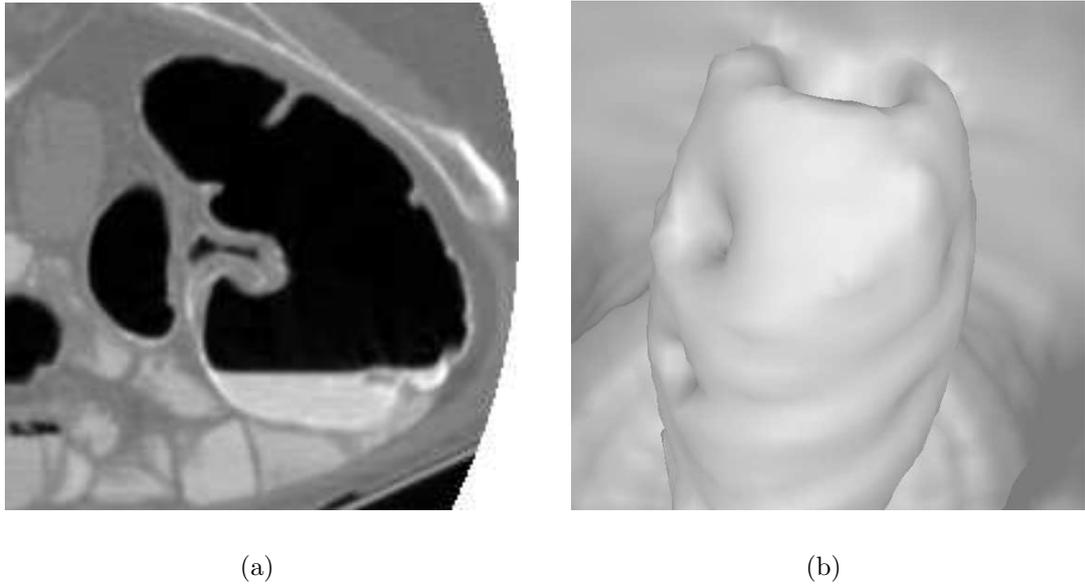
45) and will be summarised later in Section 5.3.4.

In a subsequent publication, Summers et al. (2001) evaluated the feasibility of an improved version of their automated polyp detector using a cohort of 20 patients. In their improved algorithm a range of different filters were introduced to account for the greater variety of polyp morphologies found in real patient data. In addition, a false positive reduction technique based on sampling the CT attenuation of suspected polyps was also included. A sensitivity of 71%, with 3.5 false positives per data set, was reported when this algorithm was applied to well-distended segments of the colon. Ling et al. (2002) evaluated the performance of this CAD technique in a phantom study where polyp size, scan orientation, scanner type and radiation dose were all varied. A total of eight synthetic polyps were used in the study ranging in size from  $10 \times 10$  to  $5 \times 5$  (diameter  $\times$  height in mm), the orientation of the phantom to the scan plane was varied between 0, 45 and 90 degrees. Single and multislice helical scanners were used in the study and the dose was varied between 0.7 and 10 mGy. They found that clinically significant polyps (i.e. those 10 mm in size) could be detected with 100% sensitivity in all orientations, doses, collimations and modes that were examined.

Jerebko et al. (2003*a*, NIH) used the output generated by the CAD-CTC system described earlier (Summers et al. 2001) in their study evaluating neural networks and binary trees for use in false positive reduction. In each case the classifier input consisted of a number of shape based features including minimum and maximum principal curvatures, the number of candidate vertices and the sphericity of the candidate. Prone and supine studies of 40 patients (i.e. 80 data sets) containing 39 polyps ranging in size from 3 - 25 mm were used in the evaluation. Neural networks were found to provide the best results with a sensitivity of 90% and a specificity of 95% (or 16 false positives per study). In a followup study Jerebko et al. (2003*b*) described a multiple neural network classification scheme that used 12 features (eight classifiers with four features per classifier where feature reuse occurred). A majority vote from the committee of classifiers was ultimately used to classify a polyp based on the feature set provided. The system was evaluated using a cohort of 40 patients scanned in both positions resulting in a total of 23 polyps ranging in size from 5 - 25 mm. The use of multiple classifiers in this way was found to reduce the false positive rate by 36% from 4.6 per study to 2.9. An overall sensitivity of 82.9% was reported for the system with a specificity of 95.3%.

Yao et al. (2004, NIH) introduced a novel method for volumetric polyp segmentation intended to address the task of false positive reduction at CAD-CTC. Their multistage algorithm started by identifying initial candidates using the standard NIH isosurface filtering technique. Four features: mean curvature range, vertex count, maximum dimension and sphericity were used in the filtering stage. Single  $64 \times 64$  2-D slices centred on the initial candidate clusters were extracted for further analysis. The intensities of voxels within these images that had polyplike properties (i.e. voxels close to convex boundaries) were increased using a process referred to as knowledge guided intensity adjustment. Fuzzy C-means clustering was then used to group voxels into one of three classes: *lumen air*, *polyp tissue* or *nonpolyp tissue*. Ultimately the polyp segmentation process was performed using active contours where the contour was attracted to high gradients in the edge map of the membership function for *polyp tissue*. This process was applied to adjoining slices until the entire 3-D volume of the polyp had been identified. In an evaluation of their technique using 40 data sets containing 65 polyps (3 - 30 mm) a 76.3% correspondence was found between automatic and manually segmented regions. When used in conjunction with the aforementioned CAD-CTC technique (Summers et al. 2001), automatic segmentation reduced the number of false positives by 30% at a sensitivity of 100%. A further reduction in false positives was achieved by eliminating polyp candidates based on an analysis of their volumetric features. This resulted in a total reduction in false positives of 50%, however, the sensitivity was slightly compromised, reducing to 80%.

The basic isosurface curvature CAD-CTC technique was further extended by Summers et al. (2004) to deal with a common source of false positive detections i.e. the *ileocecal valve* (ICV) (see Figure 3.9 (a)). The documented approach recognised the ICV at CTC based on two main characteristics: its low density ( $< -124$  HU) and its large size ( $>1.5$  cm<sup>3</sup>). Using this technique in a cohort of 60 patients, 14 of the 27 false positives attributable to the ICV were eliminated whereas none of the true positives were affected. In a subsequent study, Iordanescu & Summers (2004, NIH) dealt with another common source of false positives, i.e. the rectal tube required for bowel insufflation (see Figure 3.9 (b)). Their approach involved an analysis of the slices toward the base of the data set. Potential rectal tube shaft regions, visible as small pockets of air, were identified using morphological operators with circular structuring elements. A set of shaft candidate voxels was then generated



**Figure 3.9:** Two common sources of false positives as seen at CTC. (a) A 2-D axial slice illustrating a protruding ileocecal valve (WRAMC\_VC-286/prone) and (b) a 3-D surface rendering of the rectal tube (MMH-166/supine). Note that the patient in (a) is in the prone position hence the cecum is on the right side of the image.

by matching a template to each potential rectal tube shaft region and retaining those with the highest similarity (the template simulated the actual rectal tube cross-section). Candidate voxels were connected to form axes and the axis that traversed mostly air was identified as rectal tube cylinder axis. This axis was extended to the tip of the tube and then expanded to yield the final region occupied by the rectal tube by means of conditional dilation. Iordanescu & Summers tested their technique using 40 patients scanned in both the prone and the supine positions. The rectal tube was completely identified in 72% of cases (58 data sets). Removal of the rectal tube reduced the number of initial (primary stage) false positive detections at CAD by 9.2%. The number of (secondary stage) false positives after the application of the classifier was reduced by 7.9%. In both cases the true positive detections were unaffected.

In a further extension of the NIH technique, Summers et al. (2005) incorporated support for dealing with contrast-enhanced CTC data sets. This was achieved using a bowel segmentation technique that included a contrast-enhanced fluid subtraction algorithm (see Section 3.2.3). The segmentation process utilised a region-growing algorithm that could traverse from air to opacified fluid and back again, and ultimately yielded a surface representation of the colonic mucosa located at transitions from

air to colon wall, fluid to colon wall and fluid boundary to colon wall. A threshold refinement stage was included to reduce the level of over-segmentation (i.e. inclusion of the small intestine). Curvature values were calculated for all vertices in the resulting surface and calculated values were inverted for vertices associated with contrast enhanced fluid. Surface patches that met the previously discussed shape criteria (elliptical curvature of the peak sub-type) were ultimately selected as polyp candidates. An evaluation of the technique was carried out using 17 patients with 22 submerged polyps ranging in size from 5 - 51 mm. Leakage into the small bowel was quantified as small or none in the majority of cases (15/17) based on a visual inspection of the segmented data sets, and 19 of the 22 polyps (86%) were automatically detected. One of the three polyps that were missed was located on the air/fluid boundary.

Jerebko et al. (2005) extended the original NIH technique (Summers et al. 2001) by utilising a committee of support vector machines for the detection of colorectal polyps. The input to their system was an isosurface representation of the inner surface of the colon wall. Regions of the isosurface were identified as initial polyp candidates using features including mean curvature, maximum dimension, vertex count, sphericity, wall thickness, average density and surface area. Four additional candidate features including the curvature of the candidate neck and an approximation of candidate texture were also calculated for use in the classification stage. The effectiveness of subsets of these features for polyp classification was quantified and classifiers with an effectiveness above a certain threshold qualified for inclusion in the committee. This resulted in a committee consisting of four support vector machine classifiers each using four features drawn from a core set of nine (feature reuse was permitted). The ultimate decision regarding the class of a candidate was based on the majority vote of these four classifiers. The system was trained using 40 data sets containing a total of 47 polyps  $> 4$  mm and evaluated, using a modified leave-one-out method, with a further 80 data sets containing 20 polyps. The committee of support vector machines was found to have a sensitivity of 81%, with an average of 2.6 false positives per data set. This represented a notable improvement over the performance of a single support vector machine using all nine input features and the average performance of each committee member when applied individually.

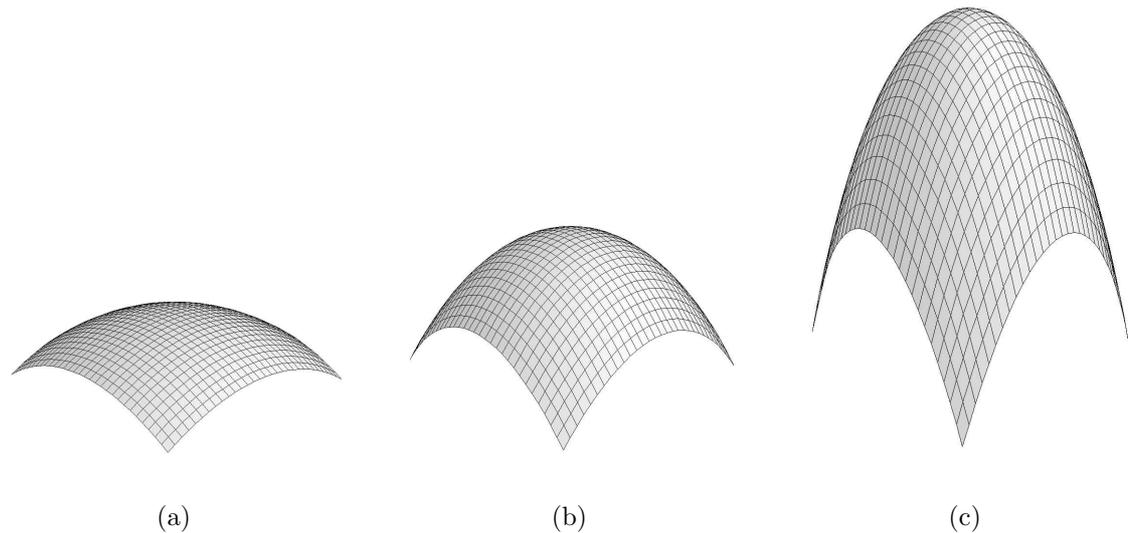
The Wake Forest Group also developed a CAD-CTC technique based on isosurface curvature analysis (Li et al. 2005). The first stage in the documented approach

involved obtaining a surface representation of the colon using their aforementioned segmentation technique (Li & Santago 2005, see Section 3.2.1). Each vertex in the resulting mesh retained references to its directly connected neighbours in order to facilitate region growing within the mesh. The two principal curvatures, and subsequently the mean and Gaussian curvatures, were calculated directly from the mesh. Polyp candidates were then identified, using region growing, as surface patches above a specified size threshold that meet certain requirements regarding mean and Gaussian curvature. Features including minimum and maximum candidate radius, candidate surface area, roundness, elongation factor and distance from catheter were used in conjunction with *multi-plane linear discriminant function* (MPLD) and *maximum a posteriori* (MAP) classifiers to reduce false positives and yield the final list of polyp candidates. Both classifiers were evaluated using a cohort of 34 subjects. ROC and *free-response receiver operator characteristic* (FROC) analyses indicated that there was no statistically significant difference between the two classifiers. At a patient level sensitivity of 90% the MLDF classifier yielded 1.3 false positives per case and the MAP yielded 1.2 false positives per case.

### 3.5.3.3 Voxel Curvature Analysis

Voxels belonging to the colon wall can be used instead of isosurface vertices for approximating the curvature characteristics of the colonic mucosa. This removes the need to extract an isosurface representation of the colon wall from the CTC data set. A number of techniques that utilise this alternative have been developed at the University of Chicago. The Chicago group use the shape index (see Figure 3.8) and curvedness (S) (see Figure 3.10) to identify initial polyp candidates.

Yoshida et al. (2002a, UC) developed a CAD technique that analysed a segmented model of the colon lumen obtained from a linearly interpolated isometric data set. The segmentation of the colon was achieved using a previously described technique developed by the Chicago group (Masutani et al. 2001, see Section 3.2.2) and yielded a representation of the entire colon wall and not just the inner surface of the colonic walls. Two features: the shape index and the curvedness were calculated for all voxels belonging to the segmented colon wall. The shape index could be used to assign a voxel to a particular shape class (rut, saddle, ridge & cap). The curvedness gave the magnitude of the curvature for the voxel. Voxels with a shape index between 0.9 and 1.0 (representing cap structures) and a curvedness between



**Figure 3.10:** The curvedness property of a surface. Three surfaces with increasing curvedness are illustrated (a) - (c).

0.08 and  $0.2 \text{ mm}^{-1}$  (corresponding to effective polyp sizes from 12.5 to 5 mm) were identified as initial candidate voxels. Hysteresis thresholding with relaxed shape index and curvedness thresholds was used to identify polyp candidate regions. Relaxed thresholds were used to reflect the fact that polyps do not always appear as perfect cap-like shapes. Candidate regions that were close together were merged and fuzzy C-means clustering was employed to remove voxels attributable to noise. A subsequent refinement involving a rule based linear discriminant analysis was used to further reduce the number of false positives. Finally minimum and maximum values were identified for the CT values and the gradient of the CT values of actual polyps. These did not provide a high level of discrimination between true and false positives, but were used to eliminate outliers from the ultimate list of polyps using a simple threshold. An evaluation of this technique using 82 data sets containing a total of 18 polyps, resulted in a sensitivity of 89% with 2 false positives per data set and a sensitivity of 100% with 2.5 false positives per data set. Upon an analysis of the false positives, the majority were found to be caused by prominent folds, incorrectly segmented small intestine and stomach, the ileocecal valve, residual stool and the rectal tube. Other less frequently occurring false positives were attributed to flexural pseudo tumors, motion artifacts, the anorectal junction and extrinsic compression.

In a subsequent paper, Yoshida & Näppi (2001) provided a detailed report on an improved version of their technique that incorporated a more accurate segmentation

technique Näppi et al. (2002, see Section 3.2.2). Further enhancements included the use of two new features: the gradient concentration which provides a measure of sphericity and the directional gradient concentration which provides a similar measure that accounted for the hemispherical shape of polyps. A total of six features and nine statistics (i.e. 54 feature statistics) were used to represent each candidate. Linear discriminant analysis and quadratic discriminant analysis were employed to reduce the number of false positive detections. The improved technique was tested using 86 data sets with a total of 21 polyps ranging in size from 5 to 30 mm. An overall evaluation of performance demonstrated a sensitivity of 95% with 1.2 false positives per data set and a sensitivity of 100% with 2.0 false positives per data set.

In a related study Näppi & Yoshida (2002, UC) carried out an unbiased estimation of the performance of their CAD technique using the *leave-one-out* method for classifier evaluation. The effect of different combinations of the 54 feature statistics on system performance was also evaluated. Their results indicated that the optimum combination of feature statistics consisted of the mean shape index, the mean directional gradient concentration and the variance of the CT value. An evaluation of their CAD technique using the optimum feature statistics and the same patient data from their previous study (Yoshida & Näppi 2001) resulted in a per-patient sensitivity of 95% with 1.7 false positives per data set and a per-patient sensitivity of 100% with 2.4 false positives per data set. A round-robin evaluation of the per-polyp performance of the same technique using a similar patient cohort (Yoshida et al. 2002b) resulted in a maximum sensitivity of 90% with a total of 2.0 false positives per data set.

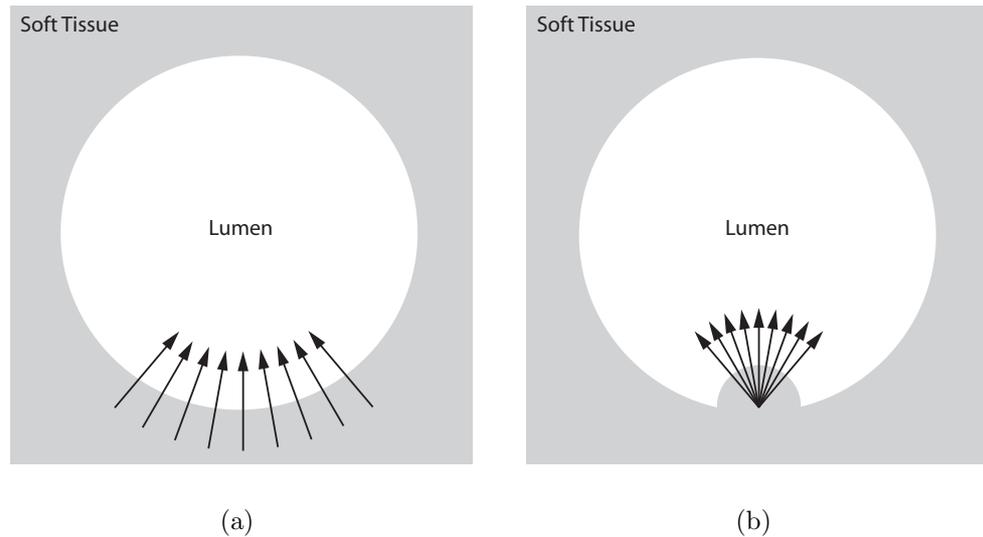
Näppi & Yoshida (2003) noticed two problems associated with the fuzzy clustering stage of their algorithm: (1) insufficient polyp coverage and (2) the inability to completely extract large polyps (i.e. those  $\geq 10$  mm). In an effort to solve this problem, an alternative to fuzzy clustering, referred to as feature-guided analysis, was proposed. This involved the iterative conditional morphological dilation of the initial polyp candidates to yield expanded candidate regions. The expanded regions provided a more comprehensive representation of the polyp candidate, comparable to the region that would be associated with the candidate upon visual inspection. In addition, an alternative to the directional gradient concentration, referred to as the modified gradient concentration, was used to identify candidates with spherical (pedunculated) and hemispherical (sessile) shape characteristics. Experimentally

determined thresholds were used to limit expanded candidate voxels to those with polylike properties. The resulting voxels were used to calculate candidate features including mean shape index, mean modified gradient concentration and the variance of the CT values. Ultimately candidate classification was performed, as in the previous version of the technique, using quadratic discriminant analysis. A round-robin evaluation comparing feature-guided analysis with fuzzy clustering was carried out using a cohort of 72 patients (14 of whom had 21 polyps ranging in size from 5 to 25 mm). At a maximum per-polyp sensitivity of 95% feature-guided analysis generated 1.51 false positives per data set compared with 5.19 for fuzzy clustering. The inclusion of the modified gradient concentration did not have a significant impact on the performance of either technique.

Wang et al. (2005, SUNY) also developed a CAD-CTC technique that utilised the surface curvature characteristics of a voxel representation of the colonic mucosa. Local and global shape index values were calculated for each voxel and regions with polypoid characteristics were identified using a clustering algorithm. The authors reported that clustering typically generated over 100 suspicious regions. A filtering stage was subsequently employed to eliminate small regions below a given size threshold and regions whose global general shape was not indicative of a spherical cup or trough structure<sup>9</sup>. The filtering stage was reported to eliminate up to 80% of suspicious regions and its output represented the set of initial polyp candidates. Subsequently, the authors utilised a false positive reduction scheme that approximated, and analysed the volumetric region occupied by each polyp candidate. This region, referred to as the *ellipsoid region of interest* (eROI), was approximated by fitting an ellipsoidal model to the inner and outer border points of the candidate, where the inner border points were located at the candidate/normal tissue boundary and the outer border points were located at the candidate/air boundary. A total of six candidate features were identified during this process: The volume of the candidate region, the ratio of the smallest and largest axes of the eROI (the axis ratio), a measure of density distribution within the candidate region (the growth ratio), the mean density for the candidate, a measure of the level of candidate coverage provided by the border points (the coverage ratio) and a measure of the symmetry among the border points (the radiation ratio). Ultimately, a suitably trained two-

---

<sup>9</sup>The SUNY group uses an approach to surface curvature calculation that considers a polyp to be concave structure, whereas other groups invariably use approaches to surface curvature calculation that consider a polyp to be a convex structure.



**Figure 3.11:** An illustration of how surface normals can be used to identify convex regions associated with potential polypoid structures (a) a flat or concave structure, surface normals converge intraluminally (b) a convex structure, surface normals converge extraluminally.

level linear discriminant classifier was used to map the feature values to the range  $[0, 1]$  and generate the final decision regarding the class of a polyp candidate. The prone and supine data sets from 153 patients were used for system evaluation. A total of 45 of these patients had 61 colonoscopy confirmed polyps  $\geq 4$  mm in size. A leave one out strategy was employed for system evaluation and a polyp was deemed to be detected if it was found in either the prone data set or the supine data set. An overall sensitivity of 100% was reported with 2.68 false positives per data set. A false positive rate of 3.44 was reported for polyps 4 - 10 mm in size and a false positive rate of 2.0 was reported for polyps 10 - 30 mm in size.

#### 3.5.3.4 Surface Normal Analysis

At a more basic level the curvature characteristics of the colon wall can be evaluated based on an inspection of the surface normals. A number of groups have used this approach to identify potential polyps from CTC data sets. Figure 3.11 illustrates the distribution of surface normals for normal regions of the colon and region that contain convex polyplike structures.

Kiss et al. (2002, UHG) provided a description of a technique that utilised both surface normal analysis and sphere fitting in order to automatically detect potential colorectal polyps. Polyp candidates were ultimately identified as the local maxima

of sphere centres in the processed data set. Their approach was tested using 18 data sets (nine with and nine without polyps) containing a total of 20 polyps with various sizes. 100% sensitivity with 1.88 false positives per data set was reported for polyps  $\geq 10$  mm.

Paik et al. (2004, Stanford) developed a novel CAD algorithm, called the *surface normal overlap* (SNO) method, for automatically detecting both colonic polyps and lung nodules in CT images. The following discussion will concentrate on the polyp detection aspect of their approach, as it is most relevant to this review. The SNO method started by obtaining an isometric segmentation of the colon that was representative of the region occupied by the colonic mucosa and had a thickness of approximately 5 mm. This was achieved using a combination of clamping, interpolation, region growing, masking and morphological dilation. The surface normals of the segmented colon were subsequently calculated using a modified Canny edge detector. Then each voxel belonging to the segmented colon was assigned a score proportional to the number of intersecting surface normals in its vicinity. An exact intersection was not required thus facilitating the detection of roughly globular hemispherical objects representative of polyps. Regions with a high overlap score that were above a specified size threshold were ultimately flagged as potential colorectal polyps. A theoretical analysis of the SNO technique was carried out involving the creation of shape models representing polyps (hemispheres) and haustral folds (half-cylinders). The models were deformed using parameters obtained from an analysis of real patient data in order to more accurately reflect the morphology of the actual anatomical structures being approximated. In this analysis, the SNO method was compared with a previously published *Hough transform* (HT) based technique (Erberich et al. 1997). The SNO method was found to provide greater discrimination between polyps and folds particularly when the structures began to deviate from their ideal cases (i.e. hemispheres and half-cylinders). In addition, the response of the SNO method was found to increase with polyp size, whereas, the opposite was found to be the case with the HT method. The parameters for the Canny edge detector were optimised using 110 hemispherical phantoms generated using simulated CT scanning software. The optimisation involved the identification of parameters that placed the centre of the hemisphere detected using the SNO method closest to the true hemisphere centre. The system was tested using a cohort of eight patients selected from a database of 116. All patients were scanned in the supine position

only. The test set contained a total of seven polyps  $\geq 10$  mm and 11 polyps in the range 5 - 9 mm. An evaluation of the performance of the system using a leave-one-out approach found that a maximum sensitivity of 100% was achieved for polyps  $\geq 10$  mm with an average of 7.0 false positives per data set.

Chowdhury et al. (2005b, VSG) developed a CAD-CTC technique that utilised surface normal analysis to identify initial polyp candidates. Their system operated on a voxel based representation of the colon surface and calculated the required surface normals using the 3-D Zucker-Hummel edge operator (Zucker & Hummel 1981). Intersections of surface normals, where the number of intersections was above a specified threshold, were identified as polyp candidate centres, and each centre was associated with a cluster of surface voxels. A convexity test was used to eliminate non-convex surface voxels and clusters were divided into multiple candidates in cases where discontinuities occurred. A total of seven features were computed for each candidate cluster. These included features that were calculated using sphere and ellipsoid fitting. A suitably trained nearest neighbour classifier was ultimately used to determine the class of the initial polyp candidates. The system was trained using 64 polyps and 354 nonpolyps, and evaluated using real patient data, real patient data augmented with synthetic polyps, and a colon phantom. The real patient data consisted of 36 patients with 57 colonoscopy confirmed colorectal neoplasia (polyps and masses). An overall sensitivity of 70% was reported with a false positive rate of 4.05 per data set. The authors also reported that the documented CAD-CTC technique required an average of 3.9 minutes to process each data set. In a subsequent study, Chowdhury et al. (2005c) evaluated the effect of radiation dose on the operation of their CAD-CTC technique. A synthetic colon phantom containing a total of 48 polyps was scanned using a range of effective tube currents: 100, 70, 60, 40, 30, 20 and 13 mAs. A sensitivity of 87% was achieved for all tube currents with the exception of 13 mAs where a slightly lower sensitivity of 85% was achieved. This study demonstrated that a substantial reduction in radiation dose is possible without a significant impact on the automatic detection of polypoid structures.

### 3.5.3.5 Perception Based Methods

A number of techniques developed by the Stanford group attempted to identify polyps at CAD-CTC based on the methods and techniques that radiologists use to identify polyps at conventional, manually read, CTC.

Göktürk et al. (2001, Stanford) recognised the low specificity associated with early CAD-CTC systems as a significant problem and proposed a two stage system to address the issue of false positive reduction. Initial polyp candidates were identified using a Hough transform based technique previously developed by the Stanford group (Paik et al. 2000), this yielded a list of polyp candidates and constituted the primary stage of their system. The secondary stage was intended to replicate, and to attempt to improve on, the way that a radiologist examines a data sets using axial, sagittal and coronal views. This involved the generation of a novel representation of the candidates that consisted of 100 randomly orientated triples of mutually orthogonal planes rotated about the centre of each  $32 \times 32 \times 32$  voxel candidate subvolume. The planar images were cropped to eliminate regions associated with extraneous tissue and the geometric features of the resulting cropped images were calculated and combined to generate a feature vector or shape signature for the candidate. The ultimate classification of a candidate as either a polyp or nonpolyp was achieved using a support vector machine classifier trained using a gold standard consisting of manually segmented polyps. The system was tested using a cohort of 48 patients with a total of 40 usable polyps ranging in size from 2 - 15 mm (six polyps were excluded due to their proximity to the data set boundary). A ten-fold cross-validation evaluation of the two stage system showed that, on average, the specificity of the system, at 100% sensitivity, increased from 0.35 to 0.75 when the secondary stage was included.

Acar et al. (2002, Stanford) also developed a CAD-CTC system based on the way a radiologist identifies polyps while performing a 2-D evaluation of a CTC data set. The Hough transform based technique developed by Paik et al. (2000), was used to identify initial polyp candidates. Volumes of interest of size  $21 \times 21 \times 21$  voxels and centred at each initial candidate voxel were extracted. Subsequently, a smoothed mean edge displacement field for each volume was calculated for the axial, sagittal and coronal directions. A characterisation of the resulting edge displacement fields resulted in three parameters: the divergence, the circulatory behavior and the spread of the streamlines about the centre. An average value for each parameter was obtained for the axial, sagittal and coronal directions thus creating a 3-D feature vector for each candidate volume. A binary classification of candidates was then carried out using a Mahalanobis distance based linear classifier in order to yield the final list of polyp candidates. Evaluation data was generated by extracting volumes of

interest for all usable polyps and 50 nonpolyps randomly selected from each patient in their cohort of 48 (as in the study by Göktürk et al. (2001)). These volumes were divided into 10 sets and a system evaluation was carried out using the leave-one-out approach. The results of this evaluation showed that the use of the edge displacement fields (secondary stage) increased the mean specificity of the Hough transform approach (primary stage) from 0.19 to 0.47 at a sensitivity of 100%.

### 3.5.4 A Contrast-Enhanced Approach

Luboldt et al. (2002, Frankfurt) evaluated the feasibility of a threshold based CAD technique used in conjunction with contrast enhanced CTC data. Following a standard bowel preparation, their study cohort of 18 patients underwent CTC followed by conventional colonoscopy. Immediately before CT scanning, an iodine based contrast agent was administered to increase the conspicuity of polyps and masses. A segmentation of the colon lumen was obtained and dilated by two iterations to overcome the partial volume effect. An experimentally determined threshold of 90 HU was applied to the surface voxels of the segmentation (i.e. those representing the colon wall). Any voxels found to be above this threshold were colour coded to indicate the presence of potential colorectal neoplasia. The intensity of the colour coding ranged from 0% to 100% for voxels in the range 90 to 160 HU. The resulting colour coding information was ultimately superimposed onto a virtual barium enema and a virtual endoscopic display for evaluation purposes. The lower threshold value of 90 HU was found to highlight 80% of the significant masses ( $\geq 10$  mm) and only 10% of the entire colon wall. A qualitative comparison of the two display modes found the virtual endoscopic perspective to be superior. The results show that the density value of the colon wall is a useful parameter for CAD-CTC systems if intravenous contrast is used. A similar study dealing with the use of intravenously administered contrast at manually read CTC, carried out by Morrin et al. (2000) (see Section 2.2.6), also yielded positive results.

### 3.5.5 Colorectal Mass Detection Techniques

The CAD studies discussed in the previous Section dealt specifically with the detection of colorectal polyps. A more difficult problem, due to reduced contrast enhancement, smooth transition to the colon wall and variable shape characteristics, is the detection of colonic masses or carcinomas. The detection of colorectal masses

is generally treated as a separate issue as masses do not have the same geometrical properties as polyps, and as such cannot be reliably detected using polyp detection algorithms. The following paragraphs review colorectal mass detection algorithms that have been published in the literature.

Näppi et al. (2004b) described an approach to colorectal mass detection based on their earlier work in the area of polyp detection Näppi & Yoshida (2003, see Section 3.5.3.3). The authors characterised masses as either projecting (lobulated, polypoid or circumferential) or non-projecting (those associated with a thickening of the mucosal wall). Separate detection techniques were proposed for each of these categories of colorectal mass. Projecting masses were detected using a fuzzy merging method that combined polyp candidate regions to yield potential mass regions. The merging process was governed by three criteria:

1. The CT values along the path between the merge candidates were reasonably consistent.
2. There were no boundaries (i.e. large deviations in the gradient of the CT values) between the two merge candidates.
3. The distance between two merge candidates was less than a predetermined threshold of 50 mm.

If a merged region met a certain size requirement it was included as a projecting mass candidate otherwise its component regions were referred for analysis using the polyp detection technique referenced above. Non-projecting masses were detected using an analysis of wall thickness based on the method described by Vining et al. (1998, reviewed in Section 3.5.3.1). This involved identifying solid, clearly delineated regions of the colon wall, with a high density relative to the surrounding tissue and an abnormal thickness ( $\geq 5$  mm) that did not arise from adjoining bones or other internal organs. Only regions with diametrically opposed counterparts were considered as non-projecting mass candidates<sup>10</sup>. The final list of non-projecting mass candidates was obtained by omitting regions below a volume threshold and regions deemed to be associated with residual fluid. The final mass regions were extracted using a level set method where the projecting and non-projecting candidate regions

---

<sup>10</sup>This restriction was imposed as only non-projecting masses of the circumferential variety were present in the available patient data.

were used as seeds. A round robin by-patient evaluation of the technique using a cohort of 82 patients with 14 masses detected all but one of the masses with an average false positive rate of 0.21 per patient. The single mass that was not detected resided in a boundary region i.e. the entire mass was not imaged in the CT study.

Luboldt et al. (2005, Frankfurt) presented results from their preliminary study dealing with colonic mass detection. Their algorithm was based on the observation that voxels associated with a mass have a greater density than those associated with normally occurring colonic features. The first stage of the algorithm involved identifying voxels that represent the air filled colon lumen. Subsequent steps included identifying border voxels<sup>11</sup> and border voxel analysis involving an inspection of border voxel density values and resulting region volumes. Regions above a specified size threshold were then coloured to indicate the presence of a potential mass. The algorithm was tested using a cohort of five patients with known colorectal carcinomas. All five carcinomas were detected by the automated mass detection algorithm. Seventeen false positives were also identified by the algorithm.

### 3.5.6 CAD-CTC Summary

It is clear that CAD-CTC techniques are evolving. Early single stage algorithms (e.g. Hunt et al. (1997)) are yielding to more recent dual stage algorithms that consist of initial candidate detection followed by false positive reduction (e.g. Göktürk et al. (2001)). It is evident from Table 3.2 that the performance of recent CAD systems compares well to the performance of radiologists (see Table 2.1), particularly in the detection of clinically significant polyps. The problem of false positive detections by CAD-CTC algorithms has been highlighted. Kiss et al. (2002) reported false positives caused by the colonic wall (45.58%), the haustral folds (31.97%), residual stool or fluid (12.25%), the inflation tube (7.48%) and the ileocecal valve (2.71%). Yoshida et al. (2002a) also provides a breakdown of false positive detections generated by their algorithm with the most common causes being prominent folds (29.7%), the small intestine or stomach (19.3%), the ileocecal valve (13.8%), residual stool (13.1%) and the rectal tube (9.6%). The reduction of false positives caused by the ileocecal valve and the rectal tube has been investigated by Summers et al. (2004) and Iordanescu & Summers (2004) respectively. The utility of orally

---

<sup>11</sup>i.e. those voxels that were within 4 mm of the air/soft tissue boundary representing the colonic mucosa, this was intended to be an approximation of the colon wall.

administered contrast (see Section 3.2.3) to reduce the number of false positives caused by residual material has been investigated by Summers et al. (2005). Registration of polyps between prone and supine data sets has also been used to reduce the number of false positives caused by mobile residual stool (Näppi et al. 2005a).

Various different approaches to initial (or primary stage) candidate detection have been investigated. These approaches can be divided into two main groups: surface curvature based and surface normal based. A variety of different false positive reduction schemes (secondary stages) have also been proposed. These include the use of techniques based on human perception (Göktürk et al. 2001, Acar et al. 2002). The use of a variety of different classifiers has also been shown to provide better discrimination between true and false positive detections. CAD-CTC is evolving into two separate but related research streams that deal with the detection of colorectal polyps and colorectal masses independently. The tasks of polyp and mass detection are mutually exclusive operations and should not be combined due to the difference in morphology between the two types of anomaly. This is demonstrated by Acar et al. (2002) where a single 32 mm mass present in their study cohort was omitted from the training/evaluation data. This approach is valid, as the inclusion of masses would obviously have an adverse effect on the training and thus the operation of their polyp detection system.

There is one significant issue that is apparent from a review of the relevant literature. This relates to the inconsistent reporting of system performance. For example, Näppi & Yoshida (2003) use the terms by-patient and by-polyp to describe the sensitivity of their system. A 100% by-patient sensitivity occurs when at least one polyp is detected from either the supine or prone data set of a patient with one or more polyps. If a patient has four significant polyps visible in both the prone and supine data sets and a CAD-CTC algorithm detects only one of the eight polyps then a by-patient sensitivity of 100% can be reported even though only 12.5% of the polyps have been detected. The intended role of CAD-CTC is as an assistive tool i.e. as a second reader to increase polyp detection sensitivity (see Section 3.5.1). By-patient reporting would only be relevant if CAD-CTC were intended as the primary reader. Conversely a by-polyp sensitivity of 100% occurs when all polyps are detected from either the prone or supine scans of a patient. If a patient has one significant polyp that is only detected in one of the two orientations, then a by-polyp sensitivity of 100% can be reported even though only 50% of the polyp instances have been

identified. Other groups use different reporting strategies. Testing strategies also differ greatly between studies. Test sets can consist of complete patient data sets or alternatively sets of sub-sampled volumes containing (all) true and (selected) false positives that were identified by a primary stage algorithm. Acar et al. (2002) adopt the later test set generation strategy in their study due to computational concerns. There are various other differences between CAD-CTC techniques that are summarised in Table 3.2: The number of patients (8 - 72), the scanning orientations used, the number of polyps (7 - 65), the bowel preparation and the CT equipment used. All of these factors make a direct comparison between different techniques difficult.

**Table 3.2:** A summary of previously published CAD-CTC techniques organised chronologically (by submission date for journals and by meeting date for conferences). Instances where the required information is unavailable are indicated by ‘-’.

Group	Study Date	Cohort Size	Scans	P (mm)	P #	Sens.	FPR / Spec.	Comments
Chowdhury et al. (2005 <i>b</i> )	'05, Sept.	36	P & S	≤ 5 > 5 All	34 16 50	62% 88% 70%	- - 4.05	A per polyp sensitivity is reported. Note: Separate training and evaluation data sets were used in this study.
Wang et al. (2005)	'05, Mar	153	P & S	4 - 10 10 - 30 All	50 22 61	100% 100% 100%	3.44 2.0 2.68	A per polyp sensitivity is reported. The false positive reduction stage eliminated 93.1% of initial candidates without affecting sensitivity.
Taylor et al. (2006)	'04, Dec	25	P & S	6 - 9 ≥ 10 All	20 12 32	75% 92% 81%	- - 13	Per polyp sensitivity reported. Patient level false positive rate reported i.e. false positive rate is 6.5 per data set.
Li et al. (2005)	-	34	P & S (× 2)	-	-	90%	1.2 [1.3]	FP rates given for MAP & [MLDF] classifiers. Prone and supine scans obtained with and without oral contrast (i.e. 4 scans). Patient level sensitivity reported.
Jerebko et al. (2005)	'04, Mar.	40	P & S	< 10 > 10 All	5 15 20	81% [73%]	2.6 [4.0]	Comparative results provided for a committee of classifiers and [a single classifier]. A per-data set FPR reported.
Summers et al. (2005)	'04, Feb.	17	P & S	5 - 9 10 - 51 All	14 8 22	- - 86%	- - -	Orally administered contrast agent used. All polyps were submerged or partially submerged. No FP information reported.
Iordanescu & Summers (2004)	'03, Oct.	40	P & S	-	19	-	1.1 [1.0]	FP rate before and [after] elimination of rectal tube FPs. Note: FP reduction did not affect sensitivity.
Paik et al. (2004)	'03, Sept.	8	SO	≥ 10	7	100% 90% 80%	7.0 6.0 4.6	A dual purpose algorithm capable of detecting lung nodules as well as colorectal polyps. 90% sensitivity & a FPR of 5.6 reported for lung nodule detection.

Continued on next page...

Table 3.2: Continued from previous page.

Group	Date	Cohort	Scans	P (mm)	P #	Sens.	FPR	Comments
Summers et al. (2004)	'03, Aug.	20 ( $ds_1$ ) 40 ( $ds_2$ )	P & S	$\geq 10$	26 18	81% 50%	4.6 [3.0] 2.0 [0.8]	FP rate before and [after] elimination of FPs caused by ICV. $ds_2$ used for evaluation and $ds_1$ used for training. Sensitivity unaffected by FP reduction.
Yao et al. (2004)	'03, Feb.	20	P & S	3 - 30	65	100% 80%	11.5 5.5	An evaluation of density and volumetric features of polyp candidate used for false positive reduction.
Jerebko et al. (2003b)	'02, Sept.	40	P & S	5 - 25	21	83% [58%]	2.9 [4.5]	Comparative results for an 8 member committee of 4 input classifiers and [a single 13 input classifier].
Näppi & Yoshida (2003)	'02, June	72	P & S	5 - 10 11 - 12 25 All	17 3 1 21	- - - 95% [100%]	- - - 1.5 [1.3]	By-polyp and [by-patient] results reported. By-polyp: a polyp is considered detected if found in either orientation. By-patient: a patient is considered to be diagnosed correctly if any polyps are detected.
Jerebko et al. (2003a)	'02, Mar	40	P & S	3 - 25	39	90% [82%]	16 [17]	A comparative study evaluating two types of classifier: neural networks and [recursive binary trees].
Luboldt et al. (2002)	'01, Nov.	18	P   S	$< 10$ $\geq 10$ All	28 7 35	- 80% -	- - -	Although no specificity rate is reported, the authors do state the their algorithm flags an estimated 10% of the colon surface as abnormal.
Yoshida et al. (2002b)	'01, Nov.	71	P & S	5 - 9 10 - 25 All	15 6 21	- - 90% [100%]	- - 2.0 [2.0]	By-polyp and [by-patient] results reported. These terms are defined above (see comment for Näppi & Yoshida (2003) study.)
Näppi & Yoshida (2002)	'01, Sept.	43	P & S	$\geq 5$	12	95% [100%]	1.7 [2.4]	Data set-based and [case-based] results reported i.e. a true positive occurs when any polyps are detected in one orientation or [in either orientation].
Acar et al. (2002)	'01, July	48	P   S	2 - 15	40	100% 95% 90% 80%	47% 56% 66% 76%	10 test sets generated using 40 polyps from a total of 46 available (6 omitted due to proximity to boundary) and 50 random nonpolyps per patient. Average specificities for 10 experiments reported.

Continued on next page...

Table 3.2: Continued from previous page.

Group	Date	Cohort	Scans	P <sub>(mm)</sub>	P#	Sens.	FPR	Comments
Yoshida & Näppi (2001)	'01, June	43	P & S	≥ 5	12	95% 100%	1.2 2.0	Patient level results reported i.e. A true positive occurs if any polyp is detected in either the prone or supine data sets.
Göktürk et al. (2001)	'01, June	48	P   S	2 - 15	40	100% 95%	69% [19%] 74% [35%]	See comment for Acar et al. (2002) above. Square brackets indicate specificity for the Hough transform based primary stage.
Kiss et al. (2002)	'01, Mar.	18	-	5 - 9 ≥ 10 Total	5 10 15	40% 100% 60%	8.17 1.88 8.17	Four flat lesions and 1 polyp < 5 mm were present but not detected. FPs attributed to: colon wall, haustral folds, residual stool/fluid, inflation tube & ICV.
Yoshida et al. (2002a)	'01, Feb	41	P & S	5 - 10	9	100% 89%	2.5 2.0	As with their later study (Yoshida & Näppi 2001) patient level results are reported i.e. a true positive is recorded if a polyp is detected in either orientation.
Summers et al. (2001)	'00, May	20	SO	≥ 10 5 - 9 < 5 Total	28 12 10 50	64% - - 40%	6 [3.5] - - -	A density based false positive reduction scheme is employed. The false positive rate is given before [and after] density based false positive reduction. Per-polyp sensitivity reported.
Summers et al. (2000)	'99, Aug.	1	SO	10	10	80%	0.0	All polyps were simulated as spherical structures 10 mm in diameter with the appropriate attenuation, pixel dimensions and partial volume effect.
Vining et al. (1998)	'98, Sept.	10	-	7 - 40	11	73%	49.1	Although a high level of false positives was reported (range 9 - 90) a ranking scheme was used and true positives generally had a high rank.
Hunt et al. (1997)	'97, May	5	-	-	-	100%	-	A 100% sensitivity was reported for significant polyps however Hunt et al. also reported numerous (but unquantified) false positive detections.

Abbreviations: Polyp sizes (P<sub>(mm)</sub>), Polyp count (P#), Sensitivity (Sens.), Specificity (Spec.)

### 3.6 Discussion

This chapter has provided a review of a selection of automation techniques that have been developed for use with CTC. Segmentation techniques automate the identification of the colon lumen/wall in CTC data sets with a high degree of accuracy. Centreline calculation techniques can then be used to identify the central path through the extracted colon lumen. The combination of these techniques enables an automated flythrough of the colon with minimal user interaction. Prone and supine registration can be used by a radiologist, or a CAD-CTC algorithm, to compare candidates between both positions thus enabling a more accurate diagnosis. Finally, CAD-CTC systems can be used to aid the radiologists in the detection of colorectal polyps and masses. Although significant progress has been made in each of these areas, there is still scope for improvement. Chapters 4 & 5 of this thesis outline a series of contributions that deal with centreline calculation and CAD-CTC. The documented techniques build on, and draw from, much of the previously published research that has been reviewed in this Chapter.

# Chapter 4

## Preprocessing Contributions

### 4.1 Introduction

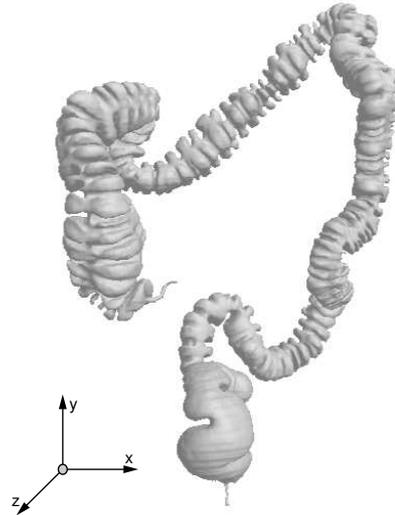
It is evident from Figure 1.2 that the CAD-CTC system described in this thesis consists of a range of well defined and interrelated tasks. These tasks can be divided into two main groups:

- **Preprocessing tasks:** Those that are required to enable automatic polyp detection, automated reporting and system testing. These tasks can be divided into two further groups:
  - Mandatory tasks that must be performed and whose outcomes are necessary for system operation (and testing).
  - Optional tasks that complement system operation by providing additional functionality and information.
- **CAD specific tasks:** Those that deal directly with automated polyp detection and automatic reporting. These are essentially the core system tasks.

This chapter provides details of the preprocessing tasks, describing their function in relation to overall system operation. The discussion regarding CAD specific tasks is presented in Chapter 5.

### 4.2 Choice of Coordinate System

At this point it is important to be aware of the implementation language as the choice influences certain aspects of system development. The Java programming language from Sun Microsystems and its related *application programmers interfaces* (API)s are used for all aspects of system implementation (see Section 6.1)). The



**Figure 4.1:** A surface rendered model of the colon displayed in conjunction with the Java3D coordinate system.

choice of 3-D coordinate system is dictated by the implementation software. In the case of Java3D the positive x-axis is from the left of the screen to the right, the positive y-axis is from the bottom of the screen to the top and the positive z-axis projects outward from the screen. A default patient orientation is selected so that the patient is upright and facing outward with respect to this coordinate system i.e the positive x-axis is from the patient’s right to their left, the positive y-axis is from the pelvis to the lung bases and the positive z-axis is from the posterior of the patient to the anterior. Adhering to this convention ensures that the orientation of the patient data is consistent, and independent of the position of the patient during scanning. This is beneficial when visualising patient data and in the case of global operations where the anatomy of the colon is examined. Figure 4.1 illustrates a surface rendered model of the colon displayed in accordance with the selected coordinate system.

### 4.3 Volumetric Data Interpretation

All CTC data sets are presented to the system as a sequence of images. Each image represents a single slice from the study volume and is encoded according to the *digital imaging and communications in medicine* (DICOM) standard. This data format is not ideal for processing as:

1. Voxel densities and other salient information is not readily accessible due to

the encoding scheme employed by the DICOM standard.

2. A large amount of unnecessary header information is present, and in many instances repeated, in each file.

As a result the interpretation of DICOM encoded data is slow with a significant processing overhead. In order to address this problem, a custom data set format is used that provides a compact representation of the DICOM encoded CTC data where all of the required information is directly accessible and all unnecessary information is discarded.

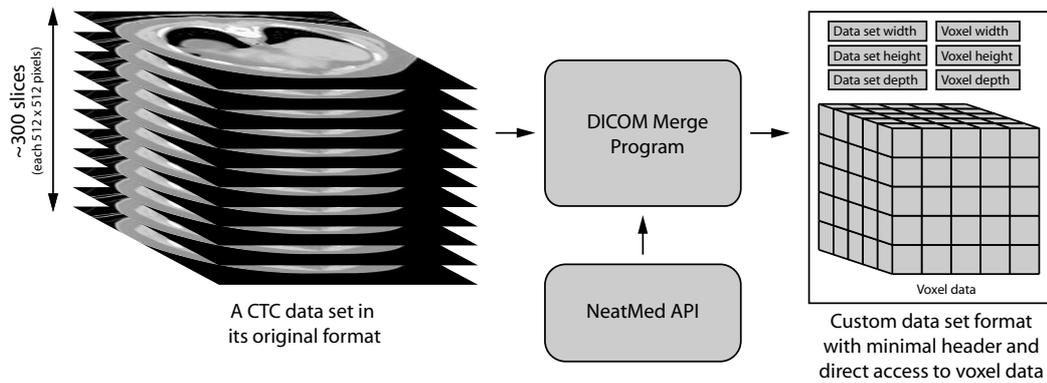
The DICOM images are interpreted using the NeatMed (API) (Sadleir et al. 2002a, 2004a). This API provides direct access to medical image data and associated header information that is encoded according to either the DICOM or Analyze standards<sup>1</sup>. A custom data set representation is initialised using a subset of the data elements from the original DICOM files and then populated with the relevant image data in accordance with the coordinate system outlined in Section 4.2. The DICOM data elements that are required for this process are:

- **The patient ID** (0x0010, 0x0020): A numerical value that uniquely identifies a particular patient. This was the only patient specific information used in this research<sup>2</sup>.
- **The patient position** (0x0018, 0x5100): A text value representing one of eight possible patient positions. Only two of these, *head first prone* (HFP) where the patient is lying down facing upwards and *head first supine* (HFS) where the patient is lying facedown, were utilised during data acquisition (see Section 7.1.2.2) and as such are relevant to this research.
- **The instance number** (0x0020, 0x0013): A numerical value representing the position of a particular slice within a CT volume. An instance number of 1 is assigned to the first image in a volume. Since all patient positions are head first this slice will always be located at the base of the lungs. The *y* coordinate

---

<sup>1</sup>A bespoke API was developed as no suitable alternatives were available at the commencement of the research programme outlined in this thesis. In addition, the use of an in-house API ensures maximum levels of flexibility and control when developing medical imaging applications. The operation of the NeatMed API is described in detail in Appendix A.

<sup>2</sup>All CTC data sets were anonymised upon receipt. This was achieved using a custom application developed using the NeatMed API (see Section A.3.4).



**Figure 4.2:** The merging process where individual DICOM files are converted into a custom volumetric data set.

for the slice with an instance number  $i$  is calculated as  $y = s - i$  where  $s$  is the number of slices in the data set<sup>3</sup>.

- **The slice location** (0x0020, 0x1041): A numerical value representing the location of the slice relative to an unspecified (scanner specific) reference point. The distance between two slices with consecutive instance numbers is used to determine slice spacing (i.e. the voxel spacing along the  $y$  axis).
- **The pixel spacing** (0x0028, 0x0030): A pair of numerical values representing the horizontal and vertical spacing between pixel values. These were ultimately used to represent the voxel spacing along the  $x$  and  $z$  axes respectively in the 3-D representation of the CTC data.
- **The image data** (0x7fe0, 0x0010): A sequence of numerical values representing the density values that constitute a particular slice.

Ultimately a custom data set contains only the density information stored as a 1-D array of 16-bit signed values, the data set dimensions and the voxel dimensions. The remainder of the information (i.e. the patient ID and the patient position) are made available as global system values as they are relevant to all subsequent instances of the data set within the system. The volumetric data interpretation process is illustrated in Figure 4.2.

<sup>3</sup>The DICOM standard does not provide a data element for the number of slices in a CT data set. This number can only be found by counting the number of files representing that data set

### 4.3.1 Resampling

The data sets are resampled in order to generate isometric volumes. This was a preemptive measure to avoid difficulties associated with anisotropic volume processing and visualisation. The resampled voxel size was selected to be the same as the minimum original voxel dimension<sup>4</sup>. This ensured that no information was lost and that the minimum amount of memory was required to store the resulting isometric data set. The use of resampling (or interpolation) to generate isometric CTC data sets has previously been documented in the literature (Yoshida & Näppi 2001). Resampling was performed using a trilinear interpolation scheme derived from the bilinear scheme described by Press et al. (1992, pg. 123). Using this technique an interpolated value  $v_{xyz}$  is calculated using the weighted values of its eight nearest neighbours. The weight for each neighbour is based on the distance from that neighbour to the interpolation position. This relationship is formalised in Equation 4.3.1 and the required values are illustrated in Figure 4.3.

$$\begin{aligned}
v_{xyz} = & v_{000}(1 - \Delta x)(1 - \Delta y)(1 - \Delta z) + \\
& v_{001}(1 - \Delta x)(1 - \Delta y)\Delta z + \\
& v_{010}(1 - \Delta x)\Delta y(1 - \Delta z) + \\
& v_{011}(1 - \Delta x)\Delta y\Delta z + \\
& v_{100}\Delta x(1 - \Delta y)(1 - \Delta z) + \\
& v_{101}\Delta x(1 - \Delta y)\Delta z + \\
& v_{110}\Delta x\Delta y(1 - \Delta z) + \\
& v_{111}\Delta x\Delta y\Delta z
\end{aligned} \tag{4.3.1}$$

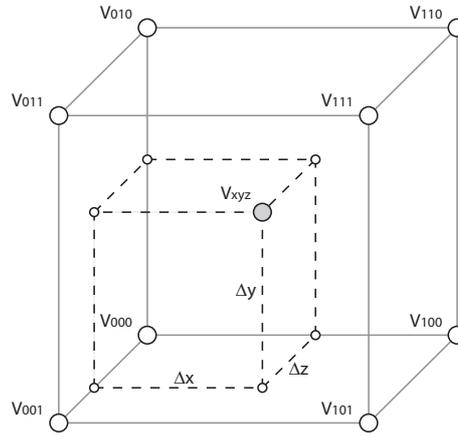
It should be noted that resampling is only required along the  $y$  axis. The full 3-D interpolation equation is specified here as it will be referred to at a later stage. In general the values for  $\Delta x$ ,  $\Delta y$  and  $\Delta z$  must be normalised, however this is only required in the case of non-integer or anisotropic sampling.

## 4.4 Colon Segmentation

The data set must be segmented following interpolation in order to identify the data set voxels that represent the colon lumen. The task of colon segmentation is

---

<sup>4</sup>Voxel sizes ranged from 0.486 - 0.848 mm (mean 0.699 mm) in the axial plane for the data sets used in this research.



**Figure 4.3:** The trilinear interpolation process. The interpolated value at  $v_{xyz}$  is based on the values of its eight neighbours ( $v_{000} \rightarrow v_{111}$ ) and its location with respect to these eight neighbours.

extremely important especially when used in conjunction with CAD. Yoshida et al. (2002a) report that approximately 20% of false positives generated by their CAD technique were attributable to over segmentation where sections of the small intestine and stomach were mistakenly identified as belonging to the colon. It is evident from the review of segmentation techniques described in the literature (see Section 3.2.1) that, although a high level of accuracy has been demonstrated, no technique can guarantee complete segmentation of the entire colon in all cases. For this reason a manual segmentation technique is utilised. This ensures that all subsequent processing and analysis is applied to the entire lumen and that no segments are omitted. It should be noted that the segmentation stage is the only part of the entire CAD-CTC system discussed in this thesis that is not completely automated. Manual segmentation has also been used in some of the more recent CAD-CTC studies e.g. Jerebko et al. (2005). Segmentation is performed using a manually seeded 3-D region growing algorithm. A 2-D example of the region growing process is presented in Figure 4.4. The operation of this algorithm can be summarised as follows:

1. A threshold is selected to define the air/soft tissue boundary representing the colon wall. A value of -800 HU is generally selected for this purpose. This threshold provides an adequate segmentation while minimising the possibility of shine through between disconnected regions, see Figure 3.2.
2. The user then selects a seed point located within the lumen that is to be segmented. If the seed point is less than the threshold value then it is flagged

and its coordinates are placed in a queue (queue1). The region growing process is initiated from this point, and involves the following steps:

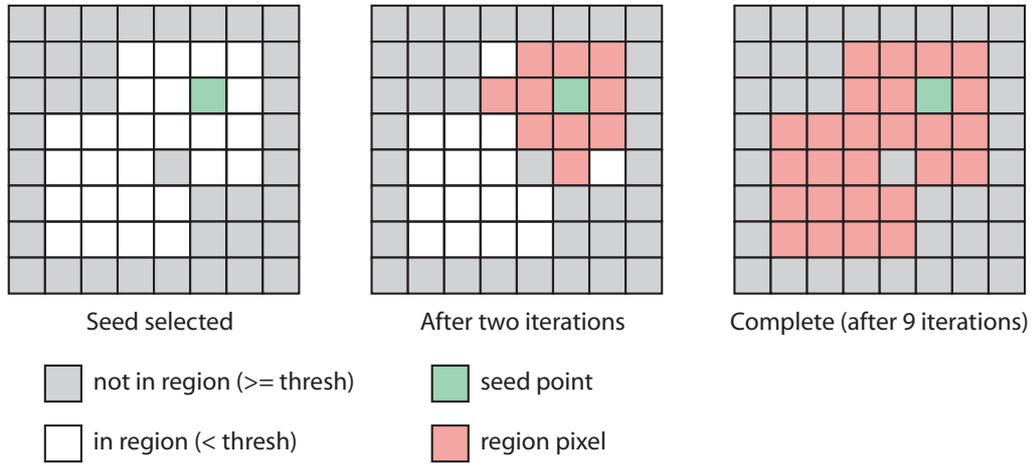
- (a) Each of the voxels stored in the queue (queue1) are removed in sequence and their neighbours are examined. Any neighbouring voxel that is not flagged and has a value less than the threshold, is flagged and its coordinates are placed in a second queue (queue2).
  - (b) When queue1 is empty it is replaced by queue2 and the process outlined in (a) continues.
  - (c) The region growing process completes when both queues have been emptied.
3. Upon completion, the set of flagged voxels represents a connected volume, with a density less than the specified threshold, which contains the original seed point.

Where the colon consists of a number of different segments due to collapse or blockage, the process listed above must be repeated for each segment. The 6-connected (face connected) approach is adopted in the region growing process. Other options including 18-connected (face & edge) and 26-connected (face, edge & corner) were decided against as they could inadvertently introduce shine through artifacts in the resulting segmentation. An illustration of how the user initiates and monitors region growing is presented in Figure 4.5.

Two other related tasks are performed at this stage in addition to colon lumen segmentation. These are data set cropping previously described by Paik et al. (1998) and Bitter et al. (2001), and rectum point detection. Data set cropping reduces the data set so that it only includes the colon lumen surrounded by a small border. In general, this significantly reduces the amount of space required to store a CTC data set while at the same time reducing the amount of time required to reload the data set. Rectum point detection is important for the centreline calculation process described in Section 4.5.1.

#### 4.4.1 Data set Cropping

A CTC data set contains a vast amount of information that does not relate to the colon or its immediate surroundings. This information is not important for sub-



**Figure 4.4:** An illustration of the region growing technique utilised for colon lumen segmentation using a simple 2-D example. (a) The original seed point is selected, (b) the identified voxels after two iterations of the region growing algorithm and, (c) All region voxels have been identified after nine iterations of the algorithm.

sequent processing and as such it is deemed unnecessary<sup>5</sup>. Data set cropping can reduce the amount of memory required to store a CTC data set by discarding any information outside the 3-D cuboid ROI that encloses the colon lumen (and a small border region). During the colon segmentation process the minimum and maximum values of  $x$ ,  $y$  &  $z$  for all lumen voxels are recorded. For subsequent processing only the lumen voxels and their immediate neighbours are of interest, and for this purpose a border is defined indicating the thickness of this region. The data set is then cropped around the lumen voxels with the necessary border region included. The revised origin  $(x_0, y_0, z_0)$  is given by:

$$\begin{aligned} x_0 &= x_{min} - border \\ y_0 &= y_{min} - border \\ z_0 &= z_{min} - border \end{aligned} \tag{4.4.1}$$

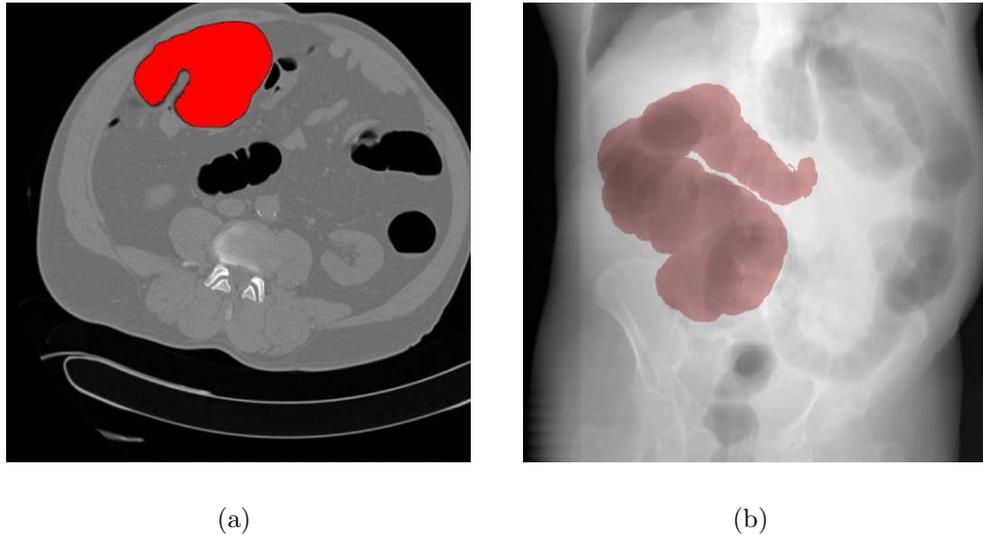
And the revised data set dimensions are:

$$width = x_{max} + border - x_0 \tag{4.4.2}$$

$$height = y_{max} + border - y_0 \tag{4.4.3}$$

$$depth = z_{max} + border - z_0 \tag{4.4.4}$$

<sup>5</sup>It should be noted that this information is important from a conventional diagnostic perspective due to the presence of extracolonic abnormalities that are imaged at CTC (Hara et al. 2000).



**Figure 4.5:** The segmentation process: The user selects a seed points to initiate the region growing process (a). The extent of the segmentation can be monitored using a simulated scout x-ray where the segmented region has been highlighted (b).

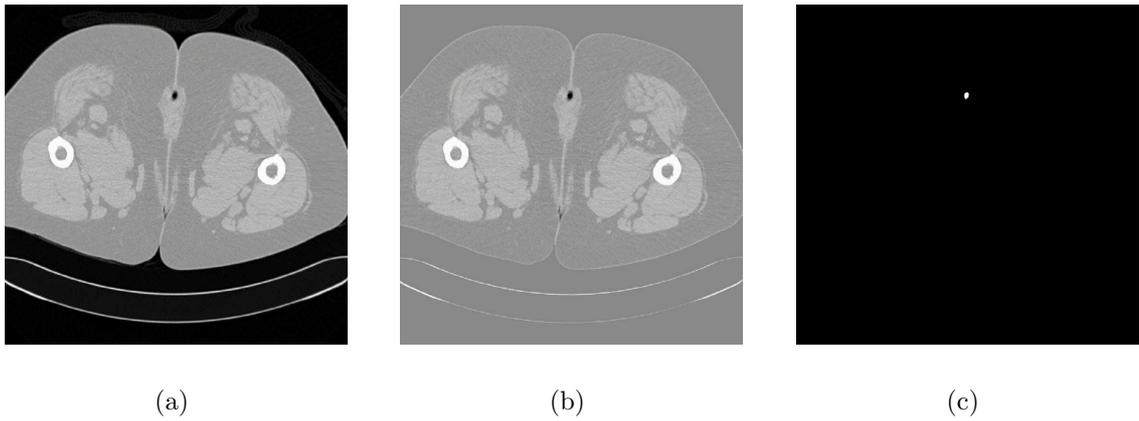
The data set defined by the revised origin and dimensions is then saved for processing by further stages of the system.

#### 4.4.2 Rectum Point Detection

The rectum point is an important landmark in a CTC data set, and is a point that can be readily detected during the segmentation process. Because of the way a CTC data set is populated (see Section 4.2), it can be assumed that the data set is captured from the pelvis at the point where the inflation device is inserted to the base of the lungs, thus ensuring that the entire colon is imaged. During the segmentation process the lumen point with the lowest  $y$  coordinate is recorded as the rectum point. If multiple points are recorded with the lowest  $y$  coordinate, then the centremost point is used as the rectum point. Alternatively the rectum point can be determined independently of the segmentation process based on an examination of the pelvic region of the data set (i.e. the base of the data set). This alternative approach to rectum point detection is outlined in Figure 4.6.

### 4.5 Centreline Calculation

The centreline of the colon is generally used to automate endoluminal navigation at CTC. It can also be used for other purposes including the registration of prone

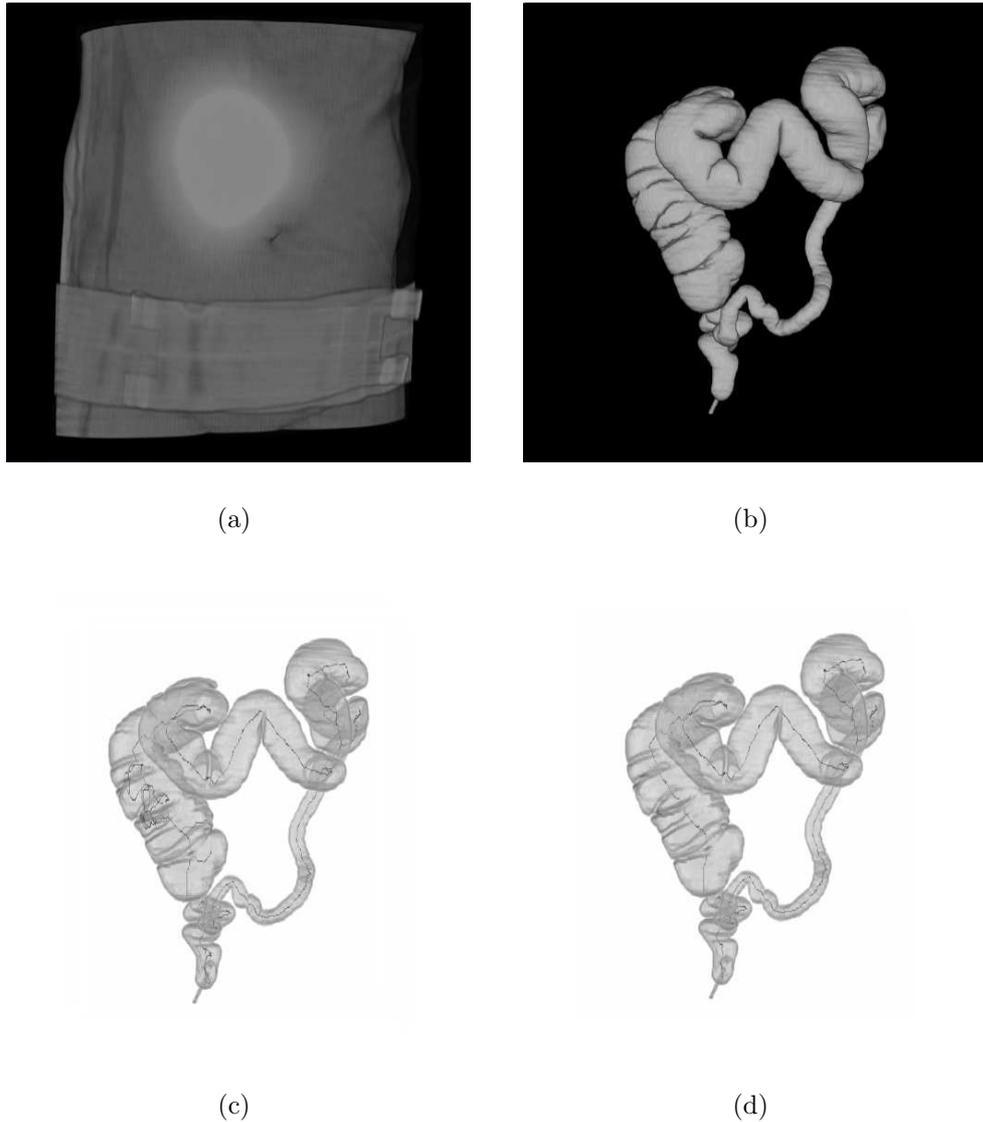


**Figure 4.6:** The rectum point can be identified by (a) selecting the slice at the base of the data set, (b) applying a border of one voxel to the selected slice and initiating region growing from the revised origin  $(-1, -1, -1)$  and, (c) thresholding the image and obtaining the centroid of resulting blob.

and supine data sets (see Section 3.4) and identifying the locations of anatomical landmarks associated with the colon (Sadleir et al. 2002). An efficient, accurate and fully automated colon centreline calculation algorithm, that is based on work previously published by Ge et al. (1999), will be presented. This algorithm will initially be discussed for the ideal case i.e. where the colon lumen is presented as a single connected hollow tube extending from rectum to caecum. The more realistic case where the colon is presented as multiple disconnected segments will then be addressed.

### 4.5.1 Caecum Point Detection

The identification of the rectum point was discussed in Section 4.4.2. The rectum point represents one of the centreline endpoints. The other endpoint, located in the caecum, must be identified before the centreline calculation process can be initiated. The caecum is the most distant region from the rectum within the colon lumen. Based on this observation the caecum point can be identified by initiating a distance from source field from the rectum point. The caecum point is subsequently identified as the point with the greatest distance value. A 2-D example of the caecum point detection process is illustrated in Figure 3.4 (a) where a distance from source field is initiated from the automatically detected rectum point (green cross) and terminates at the caecum point (red cross). This approach to detecting the second endpoint of the colon lumen was used in many of the distance field based centreline calculation



**Figure 4.7:** An overview of the centreline calculation process. (a) The original CTC data set, (b) the segmented colon lumen, (c) the skeleton and, (d) The final colon centreline.

techniques outlined in Chapter 3.3.

#### 4.5.1.1 Revised Caecum Point Positioning

The automatically detected caecum point is not guaranteed to be located at the lowest point of the caecum. This is a direct result of using a distance field for detecting the caecum point. A revised caecum point can be located by initiating a search from the initial caecum point. The revised caecum point is defined as the point in the resulting search space that has the lowest y coordinate. In cases where multiple points are identified, the centremost point is used. This revised approach provides a much more suitable indication of the caecum point.

### 4.5.2 Skeleton generation

The skeleton of the segmented colon lumen is obtained using 3DTT (Tsao & Fu 1981). The thinning process involves the removal of layers of surface voxels from a binary object. A surface voxel, i.e. one that is not completely surrounded by other object voxels, is only removed if it does not affect certain structural properties of the object being thinned. These properties describe the object in terms of overall connectivity, holes and endpoints. Layers of surface voxels are removed iteratively and the process terminates when no more surface voxels can be removed without compromising the structural properties outlined above. The resulting set of voxels represents the skeleton of the original object.

The initial set of surface voxels are identified by performing a raster scan of the entire volume. Subsequent surface voxels can be identified more easily by examining the six directly connected (face) neighbours of deleted voxels. A total of three tests must be performed prior to the deletion of a particular surface voxel, these tests deal with the following requirements:

1. **Endpoint retention** - The number and location of endpoints in the skeleton must be the same as in the original object. In the case of centreline calculation there are only two endpoints of interest, these are located in the rectum and caecum.
2. **Connectivity preservation** - The number of distinct binary objects in the scene must be the same before and after the application of the topological thinning algorithm.
3. **Hole prevention** - The removal of a surface voxel must not introduce a new hole into the object being thinned.

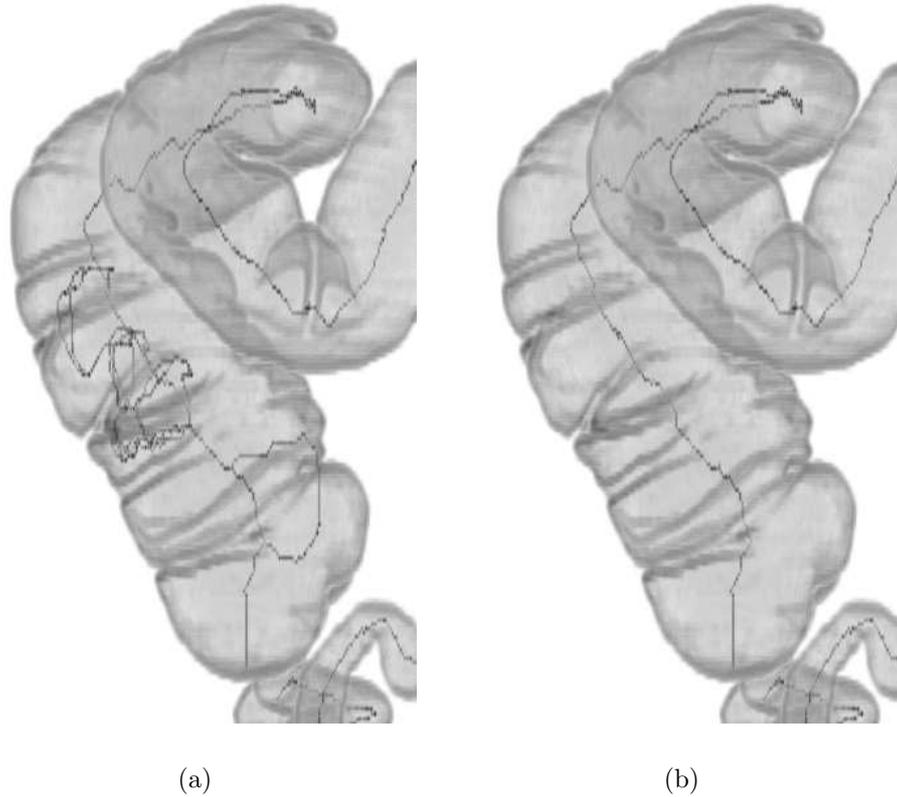
If the deletion of a particular surface voxel respects all of these requirements then it is removed, otherwise it is retained possibly to be deleted at a subsequent iteration. The endpoint retention test is trivial and involves comparing the coordinates of the voxel under examination with those of the two predefined endpoints. The other two tests, connectivity preservation and hole prevention, are much more complicated and require a significant amount of time to perform. In each of these cases the test involves extensively examining the configuration of the 26 voxels in the  $3 \times 3 \times 3$  region surrounding the deletion candidate.

### 4.5.3 Skeleton reduction (centreline generation)

The skeleton that results from the topological thinning stage incorporates the colon centreline. Unfortunately, it also includes a number of extraneous loops that occur due to holes in the original segmented colon lumen (see Figure 4.8 (a)). These holes are a common occurrence usually associated with the Haustral folds. The skeleton branches caused by holes are closer to the surface of the colon than those associated with the actual centreline. The proximity of a particular skeleton voxel to the original surface is related to the thinning iteration at which that voxel was uncovered. By assigning the relevant thinning index to each voxel as it is exposed, it is possible to generate a distance field that represents the approximate distance of each lumen voxel from the colon surface. Following an inversion of the distance values it is possible to generate a weighted graph where the centreline can be identified as the minimum cost path between the two predefined endpoints. The minimum cost path can be found in an extremely efficient manner using a simplified version of the Dijkstra shortest path algorithm (Dijkstra 1959). The resulting set of points represents the final centreline, i.e. the path between the two predefined endpoints that is farthest from the surface of the original segmented colon lumen (see Figure 4.8 (b)). The skeleton reduction process is illustrated in Figure 4.9 using a 2-D example.

### 4.5.4 Optimisation

The method for centreline calculation outlined in the previous paragraphs forms the basis for an optimised implementation and will be used as a standard for comparison purposes. The enhancements deal specifically with the skeleton generation stage, particularly the tests for connectivity preservation and hole prevention. Both of these tests must be performed repeatedly, at least once for each object voxel, and represent the most time consuming aspect of the centreline calculation process. In each case the  $3 \times 3 \times 3$  neighbourhood surrounding the voxel under examination is extensively analysed in order establish whether its removal will affect the relevant global properties of the object, i.e. overall connectivity and number of holes. The result of this analysis is dependent on the values of the 26 neighbours surrounding the voxel under examination. Each of these voxels can belong to one of two classes, either object or background, consequently the number of possible neighbourhood



**Figure 4.8:** Magnified sections of the colon lumen from Figure 4.7 (c) and (d) illustrating how extraneous loops can be caused by the Haustral folds (a) and how the removal of these loops yields the final centreline (b).

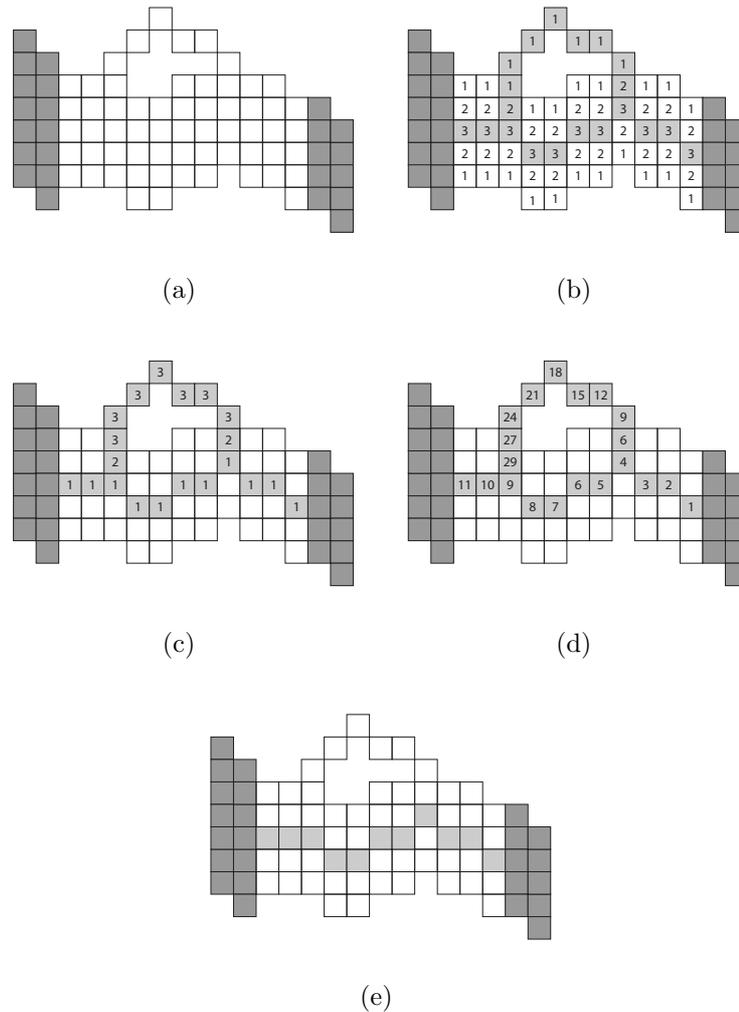
configurations is 67,108,864 ( $2^{26}$ ).

In order to streamline the centreline calculation process all possible neighbourhood configurations are generated. Each neighbourhood is tested for connectivity preservation and hole prevention. The results of these tests are combined into a single pass/fail binary value and stored in a LUT structure. Each result is addressed by a unique index,  $I$ , that is obtained by multiplying the value of each voxel in the neighbourhood ( $V_0 - V_{25}$ ) with a weight in the range  $2^0 - 2^{25}$  and then summing the results:

$$I = \sum_{n=0}^{25} 2^n V_n \quad (4.5.1)$$

This is essentially a 3-D convolution operation where the value of the central voxel is not used in the calculation. A simplified 2-D example of the index generation process is illustrated in Figure 4.10.

The LUT can subsequently be used to streamline the standard centreline calcu-



**Figure 4.9:** A 2-D example illustrating the skeleton reduction process. (a) Original binary object (region of interest indicated using white); (b) Skeleton and distance from surface field generation; (c) Inverted distance from surface field (skeleton pixels only); (d) weighted graph (originating from the right); (e) The minimum cost path found using backward propagation through the distance field.

lation algorithm. During the thinning phase of the algorithm, the computationally intensive tests for connectivity preservation and hole prevention are no longer necessary. Instead, the relevant index is quickly generated using the equation above and the LUT is queried to determine whether or not the central voxel can be removed. The thinning process continues until only skeleton voxels remain and the skeleton is reduced using the technique described earlier to obtain the final centreline.

It is extremely important to note that the LUT is calculated only once. It is subsequently stored in a file on the hard disk and reloaded (not recalculated) each time it is required for centreline calculation. The time required to load the LUT is significantly less than the time required to generate it.

**General Case:**

$P_0$	$P_1$	$P_2$
$P_7$	$x$	$P_3$
$P_6$	$P_5$	$P_4$

 $\ast$ 

$2^0$	$2^1$	$2^2$
$2^7$	$0$	$2^3$
$2^6$	$2^5$	$2^4$

 $=$  LUT Index

**Example:**

0	0	0
1	x	1
1	1	1

 $\ast$ 

$2^0$	$2^1$	$2^2$
$2^7$	$0$	$2^3$
$2^6$	$2^5$	$2^4$

 $=$  248  
 (8 + 16 + 32 + 64 + 128)

1 Foreground Pixel  
 0 Background Pixel  
 x Central Pixel (value irrelevant)  
 \* Convolution Operation

(a)

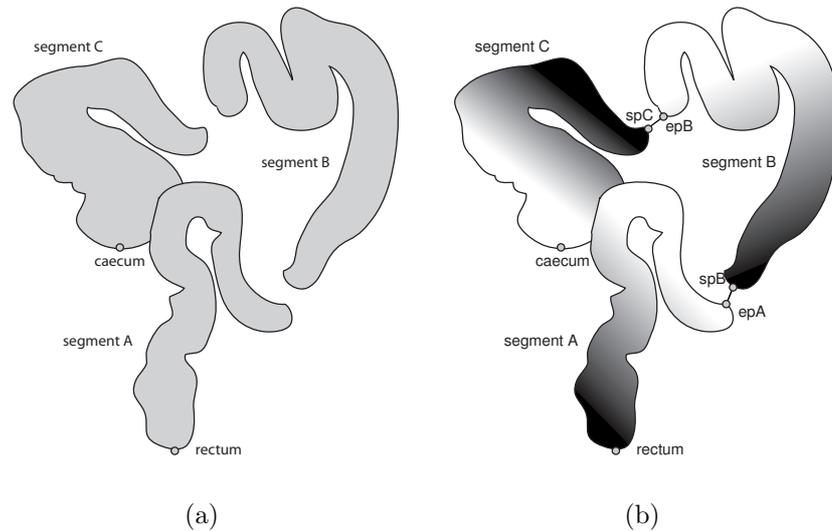
**Figure 4.10:** A 2-D illustration of the LUT index generation process.

### 4.5.5 The Real Case

Typically, the segmented colon lumen consists of two or in some cases more individual segments. The centreline calculation algorithm described above deals with the case where the colon is represented by a single hollow lumen with two endpoints located in the rectum and the caecum. The same algorithm is valid in the case where the colon consists of multiple disconnected segments. In such cases, a centreline segment must be calculated for each colon segment. The resulting centreline segments must ultimately be connected in order to yield the final centreline.

Starting with the rectum point, the endpoint for the rectum segment is identified using the distance from source field approach, outlined in Section 4.5.1. The centreline segment is then calculated for the rectum segment between these two endpoints. Each of the remaining segments are identified using a raster scan of the entire volume. A distance from source field is initiated from the point where the raster scan encounters each segment. The point in the resulting distance field with the largest distance value is identified as the first endpoints for that segment. The second endpoint is identified by initiating a new distance from source field from the first endpoint and identifying the point in the resulting distance field with the largest distance value. This endpoint detection technique is essentially an extension of the

approach described by Bitter et al. (2001). The remaining centreline segments are then calculated for each of the additional colon segments. The individual centreline segments are ultimately connected together, based on the proximity of their end-points, to yield the final centreline. The process for centreline calculation in cases where the colon consists of multiple disconnected segments is illustrated in Figure 4.11.



**Figure 4.11:** In cases where the colon consists of multiple segments, a centreline segment must be calculated for each individual colon segment and the resulting centreline segments must be connected to yield the final centreline that extends between the endpoints located in the rectum and the caecum.

## 4.6 Colon Subsegmentation

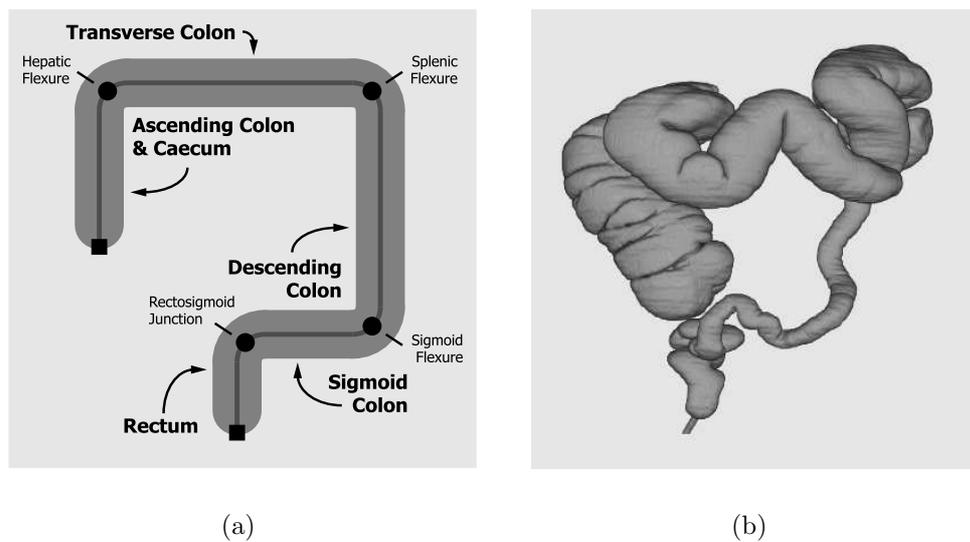
This section describes a novel use of the colon centreline to automatically identify the locations of anatomical landmarks associated with the colon. This novel technique, which will be referred to as subsegmentation, has been developed to complement CAD by providing meaningful information about the location of automatically detected colorectal neoplasia.

### 4.6.1 Overview

The human colon is divided into five anatomical regions, these are the: caecum and ascending colon, transverse colon, descending colon, sigmoid colon and rectum. From the ideal model of the colon (see Figure 4.12 (a)), it is apparent that each of these

regions represents either a horizontally or vertically aligned colonic segment. Furthermore, the junction between each of these regions is characterised by a significant change in direction, or flexure. There are four flexures in total, these are the: hepatic ( $P_{hep}$ ), splenic ( $P_{spl}$ ) and sigmoid ( $P_{sig}$ ) flexures and the rectosigmoid junction ( $P_{rec}$ ), which will also be classified as a flexure for the purpose of this discussion.

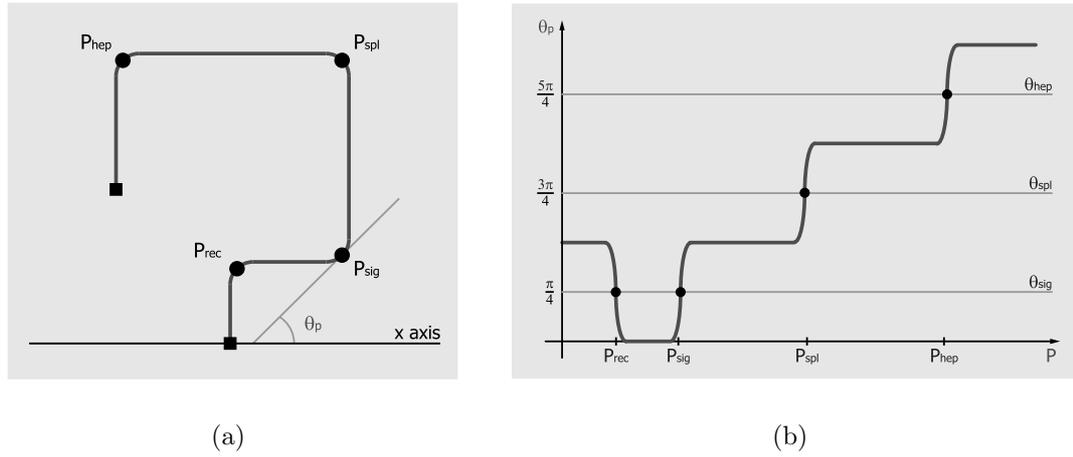
In reality, the shape of the human colon can deviate significantly from the ideal model (see Figure 4.12 (b)). Kinks and loops are introduced over time, and the colon invariably becomes distorted with age. The general shape of the colon is retained to some extent, and can still be identified, however, the task of automatic flexure detection is complicated by the frequency and magnitude of the colonic distortions. The remainder of this Section describes a series of techniques that can be used to reliably and accurately detect the four flexure points from real patient data, even in the presence of significant colonic distortions. Accurate identification of these flexures is required to facilitate colon lumen subsegmentation.



**Figure 4.12:** The anatomy of the colon. (a) An illustration of the ideal colon where all colonic segments and flexures have been identified. (b) An example of an actual human colon which has been segmented from a CTC data set.

## 4.6.2 Flexure detection

The centreline consists of a set of 3-D points. These points can be readily analysed to determine changes in direction, thus facilitating automated flexure detection. As mentioned earlier, real human colons deviate significantly from the ideal model and,



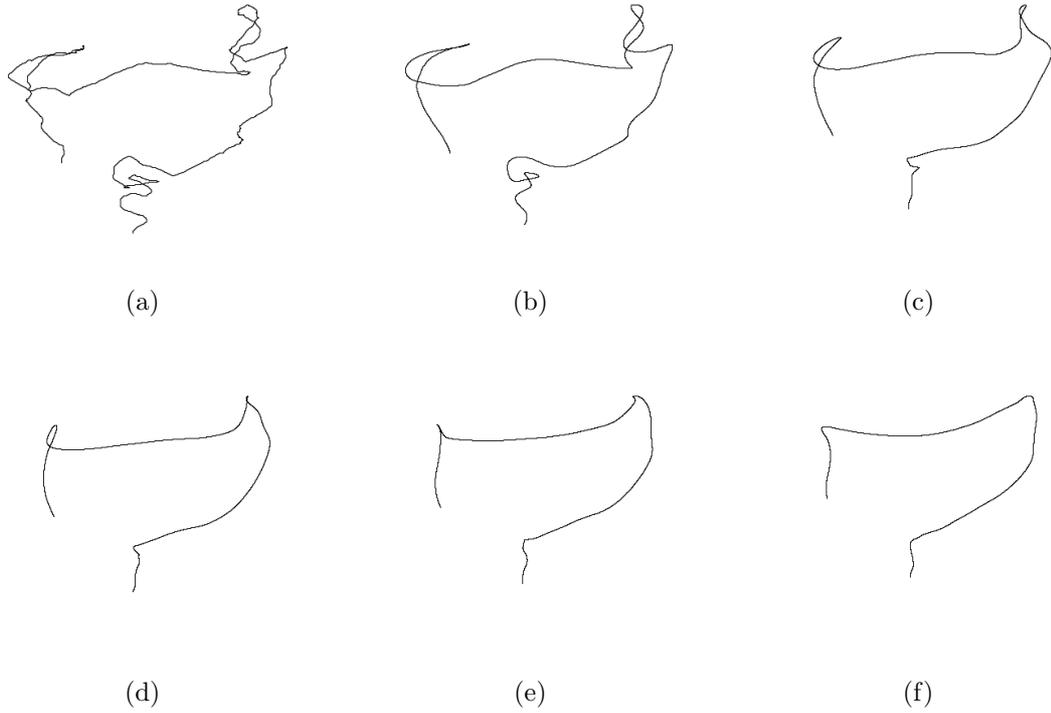
**Figure 4.13:** Flexure detection theory: (a) The centreline of the ideal colon (see Fig. 4.12 (a)) and (b) a graph of adjusted values of  $\theta_p$  for each point  $p$  in the centreline.

as a result, steps must be taken to simplify the centreline in order to assist automatic flexure detection.

The colon is generally observed from a coronal perspective and although deviations may be present in the sagittal plane, these do not have a significant influence on the location of the colonic flexures. Sagittal deviations in the centreline points can be ignored by simply projecting the centreline points from x-y-z space (3-D) onto the x-y plane (2-D). Reducing the dimensionality of the centreline points in this manner decreases the amount of information required to describe the centreline while retaining much of the important shape information.

The projected centreline still retains the distortions described earlier. These distortions must be eliminated, or at least their magnitude must be reduced, so that the flexure points can be reliably detected. If the colon centreline is treated as signal then the kinks and loops represent noise. It is apparent that the frequency of this noise is significantly higher than the frequency of the underlying signal, i.e. the undistorted colon centreline. High frequency noise of this type can be suppressed by passing the data points through a symmetrical lowpass filter. In the case of the colon centreline, lowpass filtering is achieved by calculating the average location of each centreline point over a certain range  $W$  (the filter width) i.e.  $(x_p, z_p) \rightarrow (\bar{x}_p, \bar{z}_p)$ . The filtering process is illustrated in Figure 4.14, it is obvious that the filtered centreline (f) conforms much better to the ideal model than the original centreline (a).

The four flexures can be reliably identified by examining the retrograde (rectum to caecum) direction of the filtered 2-D representation of centreline at each point.



**Figure 4.14:** Six iterations of the centreline filtering process illustrating the transition from the original distorted centreline (a) to an approximation of the original centreline where all of the kinks and loops have been removed (f).

An approximation of the centreline direction can be obtained by calculating the discrete first derivative ( $\frac{dz}{dx}$ ), i.e the slope of the tangent to the centreline. The tasks of filtering and differentiating can be combined by calculating the change in the  $x$  position ( $\Delta x_p$ ) and the change in  $z$  position ( $\Delta z_p$ ) over a number of points, the number of points must be the same as the lowpass filter width  $W$ .

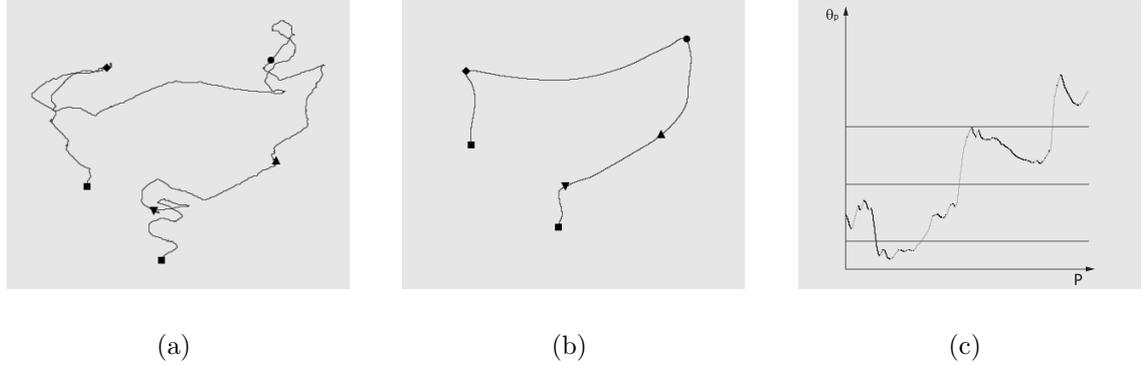
$$\Delta x_p = \sum_{i=p-\frac{W-1}{2}+1}^{p+\frac{W-1}{2}} x_i - x_{i-1} \quad (4.6.1)$$

$$\Delta z_p = \sum_{i=p-\frac{W-1}{2}+1}^{p+\frac{W-1}{2}} z_i - z_{i-1} \quad (4.6.2)$$

An approximation of the angle between the tangent to the centreline and the  $x$  axis ( $\theta_p$ ) can then be calculated for each of the centreline points.

$$\theta_p = \tan^{-1} \left( \frac{\Delta z_p}{\Delta x_p} \right) \quad (4.6.3)$$

The resulting value for  $\theta_p$  will be in the range  $-\frac{\pi}{2} \rightarrow \frac{\pi}{2}$ . It is possible to adjust this value to the range  $0 \rightarrow 2\pi$  by examining the discrete second derivative of the centreline ( $\frac{d^2z}{dx^2}$ ). The values of  $\theta_p$  for the ideal colon centreline are plotted in Figure 4.13 (b).

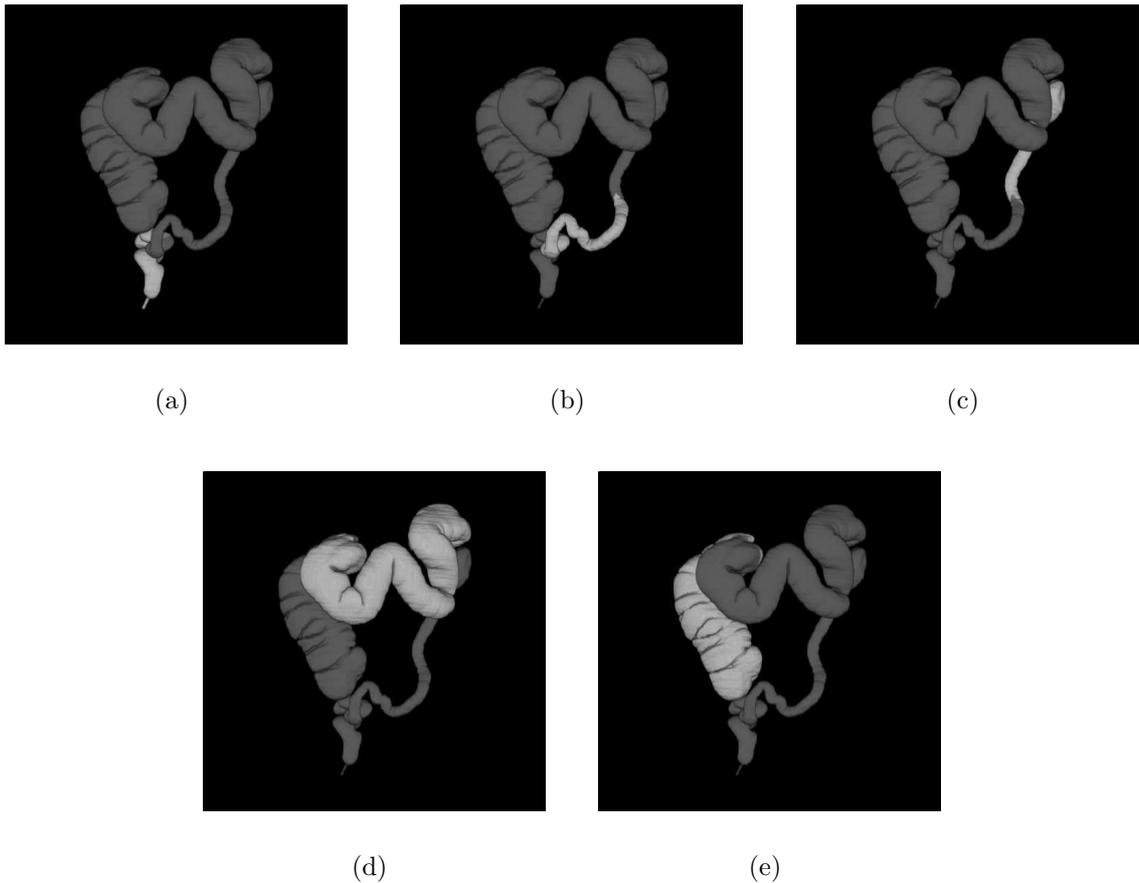


**Figure 4.15:** An illustration of the flexure detection process. (a) The original colon centreline. (b) A filtered version of the original centreline ( $W = 319$ ). (c) A plot of  $\theta_p$  against  $p$  for the filtered centreline. Note that the hepatic flexure is the last point to be uniquely identified. [Flexure key:  $\nabla = P_{rec}$ ,  $\triangle = P_{sig}$ ,  $\circ = P_{spl}$  &  $\diamond = P_{hep}$ ]

It is apparent from an examination of the ideal centreline (see Figure 4.13 (a)), that the flexure points are located where the centreline transitions from horizontally to vertically aligned colonic segments and vice versa. Consider the splenic flexure, this is located where the direction of the centreline changes from  $\frac{\pi}{2}$  (ascending colon) to  $\pi$  (transverse colon). The ideal splenic flexure is located midway between these two angles, at the point where the  $\theta_p$  curve intersects with the line at  $\frac{3\pi}{4}$ , i.e. the splenic angle ( $\theta_{spl}$ ). Similarly the hepatic angle ( $\theta_{hep}$ ) is at  $\frac{5\pi}{4}$  and the sigmoid angle ( $\theta_{sig}$ ), which identifies both the rectosigmoid junction and the sigmoid flexure, is at  $\frac{\pi}{4}$ . It is possible to differentiate between the sigmoid flexure and the rectosigmoid junction by examining the discrete second derivative of the centreline at each of these points. In the case of the rectosigmoid junction, the value of  $\theta_p$  transitions from  $\frac{\pi}{2} \rightarrow 0$ , therefore the value of  $\theta_p$  is decreasing and the discrete second derivative of the centreline will be negative at this point. Note that the opposite is true for the sigmoid flexure.

In the final system the width of the lowpass filter  $W$  is initialised at a low value and then increased until all flexure points are identified. The location of each flexure point is recorded immediately after it has been uniquely detected, this ensures that regions with a large number of strong deviations do not affect those which are

relatively smooth, thus ensuring maximum sensitivity for individual flexure point detection. The flexure detection process terminates when the last flexure point has been uniquely identified (see Figure ??). The automatically detected flexure points are ultimately used to facilitate the subsegmentation of the colon lumen (see Figure 4.16). The subsegmentation information is ultimately used to accurately report the location of automatically detected polyps (e.g. an automatically detected polyp could be localised to the transverse colon).



**Figure 4.16:** The output of the subsegmentation process consists of a colon model where voxels are labeled as (a) the rectum, (b) the sigmoid colon, (c) the descending colon, (d) the transverse colon and, (e) the ascending colon.

## 4.7 Discussion

This chapter has introduced the range of preprocessing tasks that are performed by the CAD-CTC system prior to automatically detecting polyps. A number of these tasks are core to system operation i.e. data set interpretation, data set resampling and colon lumen segmentation, while others are value added tasks that ultimately

enhance the operation of the system i.e. centreline calculation and subsegmentation. The outcomes generated by the preprocessing stages form the basis of system operation and are used to enable automatic polyp detection and automatic reporting (discussed in Chapter 5), polyp visualisation and endoluminal navigation (discussed in Chapter 6) and overall system testing (discussed in Chapter 7).

# Chapter 5

## CAD-CTC Contributions

### 5.1 Introduction

Colorectal polyps appear as projections into the air insufflated colon lumen at CTC. The task of automating the polyp detection process has been discussed at length in the literature and a number of different polyp detection techniques have been proposed (see Section 3.5). The majority of these techniques analyse either surface curvature characteristics or the distribution of surface normals in order to identify suspicious regions of the colonic mucosa that have polyplike properties. One of the main problems associated with these automated techniques is the high number of false positive detections. This problem is a result of numerous factors including the complicated morphology of the colon surface and the presence of residual material with polyplike shape and density characteristics. More recently documented CAD-CTC techniques have addressed this problem to some extent by utilising a two stage approach to polyp detection. The first, or primary, stage is intended to identify all polyp candidates with a very high level of sensitivity. The high sensitivity of the primary stage is invariably achieved by sacrificing specificity. A second stage is intended to refine the list of candidates using a secondary processing technique and/or classifier to yield the final list of polyp candidates with an acceptably low number of false positives.

This Chapter introduces a novel CAD-CTC technique that identifies initial polyp candidates based on the curvature characteristics of an isosurface representation of the colonic mucosa. The primary stage evaluation consists of a hybrid approach that combines earlier surface curvature analysis techniques, particularly those developed at the NIH (Summers et al. 2000) and the University of Chicago (Yoshida & Näppi 2001). The shape of candidates generated by primary stage algorithms is not

necessarily representative of the underlying polyp morphology. This is particularly true in the case of false positives where the candidate shape can be quite irregular. A range of true and false positives that illustrate this issue are presented in Appendix C. Upon inspection of these examples it is clear that actual polyps have a roughly circular or elliptical shape, whereas false positives have a much more irregular shape. Based on this observation, a novel intermediate stage has been developed to increase the differentiability between true polyps and false positives. This intermediate stage restructures initial candidates, expanding them to fill minimum sized elliptical regions. It will be shown that this process causes false positives to include more vertices with false positive characteristics, whereas, polyps retain vertices with polylike characteristics. As such, the expansion process essentially optimises the candidates prior to classification.

The following discussion will deal with issues such as initial design considerations, isosurface extraction, surface normal calculation, calculation of vertex curvature characteristics, anatomical vertex labeling (i.e. assigning a vertex to a particular colon segment), vertex density measurements, initial polyp candidate identification and false positive reduction using the aforementioned intermediate stage for structured candidate expansion.

## 5.2 Initial Considerations

At the most basic level, a CAD-CTC algorithm performs an analysis of the local shape characteristics of the colonic mucosa in order to identify suspicious regions that warrant further attention. As a result, when developing a CAD-CTC algorithm, two important decisions must be made: (1) what representation of the colon surface should be used, and (2) which local shape characteristics should be evaluated. These decisions were influenced by previously published techniques from two of the main contributors to surface curvature based CAD-CTC research:

- Summers et al. (2000) used a polygonal mesh representation of the colon surface in their CAD-CTC system. They evaluated vertex level shape characteristics including mean curvature and Gaussian curvature as well as global candidate sphericity in order to identify initial polyp candidates. The required curvature characteristics were calculated for voxels in the original volume and subsequently assigned to the vertices of the isosurface mesh.

- Yoshida & Näppi (2001) used a voxel based representation of the colon surface. Their initial candidate detection scheme involved calculating shape index and curvedness values for each voxel that belonged to the colon surface. Initial polyp candidates were subsequently identified using a hysteresis threshold based region growing technique. Note that other groups, e.g. Wang et al. (2005), have also used the shape index for initial polyp candidate detection.

The main difference between these two techniques is the method used for representing the colon surface. Summers et al. used a polygonal mesh representation of the colon surface, whereas Yoshida & Näppi used a voxel based representation of the area occupied by the colon surface (obtained using their segmentation technique (Masutani et al. 2001, Näppi et al. 2002)). The mesh representation has a number of benefits over the voxel based alternative:

1. The mesh representation provides a much more compact representation of the colon surface than the voxel based alternative.
2. It is possible to achieve sub-voxel precision using the mesh representation using the interpolation techniques outlined in Section 5.3.
3. The mesh representation can be readily visualised using standard computer graphics hardware i.e. surface rendering can be used, as opposed to volume rendering which is much more computationally intensive.

Based on these observations, the mesh approach was adopted for representing the colon surface<sup>1</sup>. This approach provides a compact and accurate representation of the surface of the colon that can be easily visualised.

The sets of curvature characteristics used by the NIH and Chicago groups, although different, can all be calculated in a similar manner (see Section 5.3.4). The choice of surface curvature characteristics to be used for initial candidate detection was based on an evaluation of these two sets of characteristics. The results of this evaluation indicated that the shape index and curvedness measurements used by Yoshida & Näppi provided comparable sensitivity for initial candidate detection with the lowest level of false positive detections. The details of this study are outlined in Section 7.5.2.

---

<sup>1</sup>Note that CAD-CTC techniques that utilise a polygonal mesh representation of the colon surface have also been documented by the Wake Forest group e.g. Hunt et al. (1997), Vining et al. (1998) and Li et al. (2005). However, other groups invariably used a voxel based representation of the colon surface.

The initial stage of the CAD-CTC system outlined in this Chapter, i.e. the component that deals with surface extraction and initial polyp candidate detection, is essentially a hybrid technique that draws on the best aspects of approaches published by two of the main contributors to surface curvature based CAD-CTC research. The isosurface approach to colon surface representation was selected over the voxel based alternative as it provides a compact representation of the colon with sub-voxel precision. In addition, an isosurface model can be readily visualised using standard computer graphics hardware. Conversely, a comparative evaluation of the two sets of surface curvature characteristics for the detection of initial polyp candidates found that shape index and curvedness provided the best results, hence their selection for use in the CAD-CTC system.

### 5.3 Isosurface Extraction & Vertex Classification

The task of identifying the triangular mesh representing the air/soft tissue interface associated with the colonic mucosa is performed using a modified version of the marching cubes surface extraction algorithm (Lorensen & Cline 1987). In addition to geometry extraction the basic algorithm is extended so that curvature, density and information regarding anatomical location is obtained for each vertex. Isosurface extraction using the marching cubes algorithm forms the basis of the polyp detection scheme outlined in this thesis, and due to its importance and relevance to subsequent Sections it will be briefly outlined at this point. The modifications and enhancements will be discussed following a review of the standard algorithm.

#### 5.3.1 The Standard Marching Cubes Algorithm

The *marching cubes algorithm* (MCA) is used to extract a surface represented by an isosurface value ( $d_{iso}$ ) from a volumetric data set. The value of  $d_{iso}$  is dependent on the surface being extracted and in the case of colon surface  $d_{iso}$  has a value of -800 HU (see Section 4.4). The MCA begins by thresholding the data set, assigning a 1 to voxels  $\geq d_{iso}$  (inside the isosurface) and a 0 to voxels  $< d_{iso}$  (outside the isosurface). The isosurface associated with the colonic mucosa cannot be uniquely identified using a simple threshold operation due to the the number of gas/soft tissue interfaces that are present in a CTC data set (i.e. those due to the lung bases, the small intestine, the stomach and the exterior of the patient). In this

case, the segmentation information (see Section 4.4) is used in conjunction with the isosurface value to identify the region associated with the colon surface. A cubic mask of size  $2 \times 2 \times 2$  is then passed through the volume and at each mask location the configuration of the eight underlying voxels is examined and the relevant surface patches are generated.

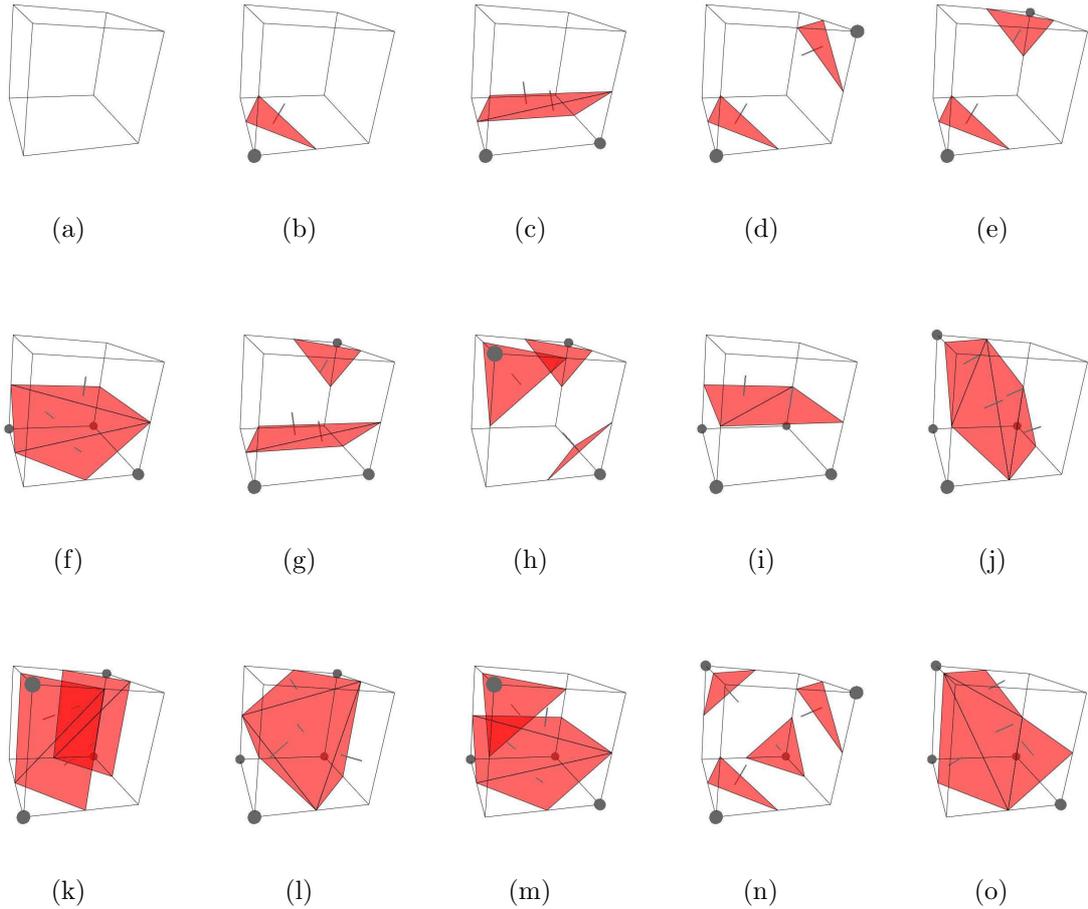
There are 256 ( $2^8$ ) possible configurations of eight binary voxels and although possible, the task of manually specifying the surface patches associated with each configuration is both tedious and prone to error. Lorensen & Cline observed that this task could be greatly simplified by considering all rotations and complementary cases for each configuration. Using this approach, the number of possible configurations is reduced from 256 down to just 15, as illustrated in Figure 5.1. This significant reduction in the number of configurations makes the task of manually specifying surface patches associated with particular voxel configurations much more manageable and a lot less prone to error.

Once the relevant surface patches have been specified for the 15 base configurations the original 256 configurations are regenerated by applying a series of rotations to the original 15 configurations and their complements. In each case the associated list of vertices (i.e. surface patches) and their complements are also rotated. The resulting information is used to populate a 256 element lookup table associating voxel configurations with edge lists (i.e. lists of vertices that are positioned relative to the local origin of the cube). The lookup table index is generated based on the voxel configuration. This process is illustrated in Figure 5.2.

By default a vertex will lie midway between the two complementary valued voxels ( $v_a$  &  $v_b$ ) that led to its creation. This is demonstrated in general by all of the vertices associated with each of the 15 cases illustrated in Figure 5.1 and in particular by the vertex resulting from  $v_1$  &  $v_2$  in Figure 5.2. In order to fit the extracted surface more accurately to the actual isosurface identified by  $d_{iso}$ , each vertex must be interpolated between  $v_a$  &  $v_b$  based on the relationship between their densities  $d_a$  &  $d_b$  and the isosurface density  $d_{iso}$ . This is achieved by calculating the normalised distance ( $\delta_{iso}$ ) from the voxel that is closest to the origin (either  $v_a$  or  $v_b$ ) to the isosurface:

$$\delta_{iso} = \frac{d_{iso} - d_a}{d_b - d_a} \quad (5.3.1)$$

The value for  $\delta_{iso}$  represents the intersection location relative to the reference voxel



**Figure 5.1:** The 15 possible configurations of eight binary voxels arranged in a cubic formation. The relevant surface patches have been inserted as documented in the original marching cubes algorithm specification (Lorenson & Cline 1987).

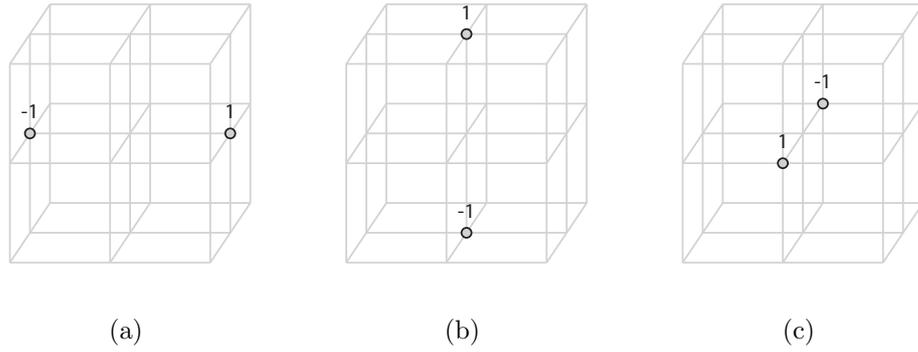
in terms of the intervoxel spacing (see Figure 5.3). The closest voxel to the origin is used as the reference point  $v_a$  to ensure  $\delta_{iso}$  has positive value and to standardise the interpolation process throughout the surface extraction algorithm. The  $\delta_{iso}$  value can then be used to calculate an interpolated vertex location that is more representative of the actual isosurface. Assuming that  $v_a$  is the closest voxel to the origin, the interpolated vertex location  $(x_i, y_i, z_i)$  located between  $v_a$  and  $v_b$  is calculated using:

$$\begin{aligned}
 x_i &= x_a + \delta_{iso}(x_b - x_a) \\
 y_i &= y_a + \delta_{iso}(y_b - y_a) \\
 z_i &= z_a + \delta_{iso}(z_b - z_a)
 \end{aligned}
 \tag{5.3.2}$$

where  $(x_a, y_a, z_a)$  and  $(x_b, y_b, z_b)$  are the locations of  $v_a$  and  $v_b$  respectively.

The final stage of the MCA involves generating unit normals for each vertex in





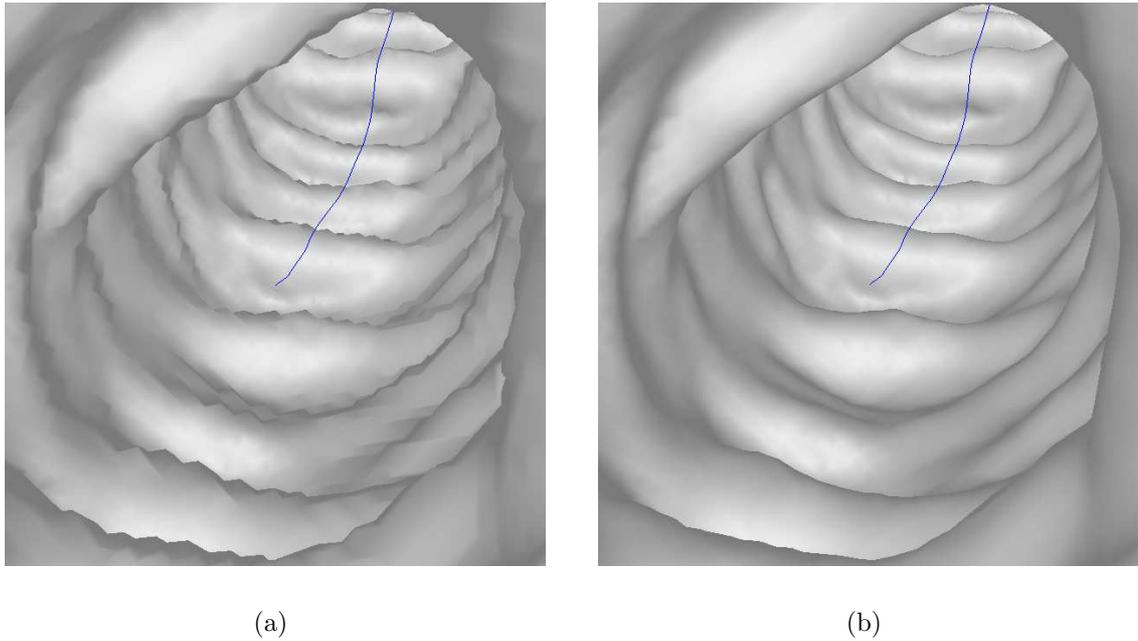
**Figure 5.4:** The three masks (a, b & c) that are used by Lorensen & Cline (1987) to calculate the edge magnitudes in the  $x$ ,  $y$  &  $z$  directions respectively. Note that the scaling factors are omitted as the implementation discussed here is intended for use with isometric data.

$$\begin{aligned}
 a_i &= (1 - \delta_{iso})a_a + \delta_{iso}a_b \\
 b_i &= (1 - \delta_{iso})b_a + \delta_{iso}b_b \\
 c_i &= (1 - \delta_{iso})c_a + \delta_{iso}c_b
 \end{aligned}
 \tag{5.3.3}$$

Where  $(a_a, b_a, c_a)$  and  $(a_b, b_b, c_b)$  are the normals associated with voxels  $v_a$  and  $v_b$  respectively. The effect of vertex and normal interpolation on the quality of the extracted surface is illustrated in Figure 5.5.

This completes the basic description of the MCA. The vertices and their associated normals can now be rendered. This basic form is only suitable for display purposes. In order to use the MCA for polyp detection a number of modifications and enhancements are required. These modifications, which are summarised below, are dealt with in the remainder of this section.

1. The standard MCA does not generate airtight surfaces. In certain cases holes may be inadvertently introduced into the generated mesh. The standard MCA must be updated so such surface discontinuities do not occur.
2. The standard MCA uses a very basic method to generate normals. A more accurate normal generation technique is proposed as normals will be used for surface analysis as well as surface visualisation.
3. The mesh generated by the standard MCA is wasteful of memory as it contains a vast amount of repeated information. A more streamlined alternative is used



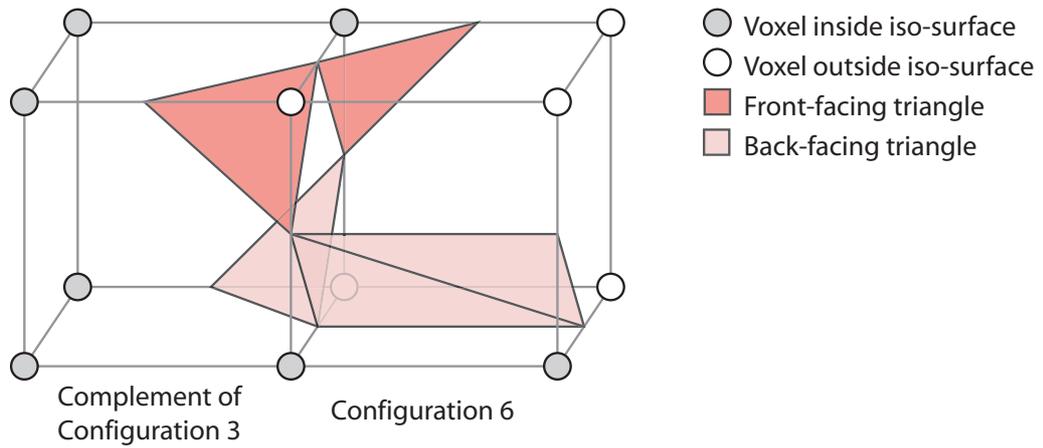
**Figure 5.5:** An illustration of an isosurface model for the colonic mucosa before (a) and after (b) the use of vertex and normal interpolation.

to reduce the amount of memory required to store vertex coordinates and associated information (e.g. normals, colours, etc.).

4. Curvature characteristics, density information and anatomical location markers are assigned to each mesh vertex to facilitate automatic polyp detection and automated reporting.
5. In the streamlined mesh, mentioned in 3. above, a vertex is no longer represented by an  $(x, y, z)$  coordinate. Instead it is represented by a unique index into a list of common coordinates. A neighbour list is generated for each vertex index that identifies all of the directly connected neighbouring vertices. This is extremely useful for region growing in a triangular mesh and crucial in identifying the surface of the mesh with polyplike properties.

### 5.3.2 Topology Errors (Holes)

Upon visual inspection of the output generated by the standard MCA, it is clear that topology errors (or holes) are present in the generated surface, see Figure 5.8 (a). These holes are due to ambiguous cases resulting from mismatches between the surface patches of adjoining cubes. An example of an ambiguous case is illustrated in Figure 5.6.



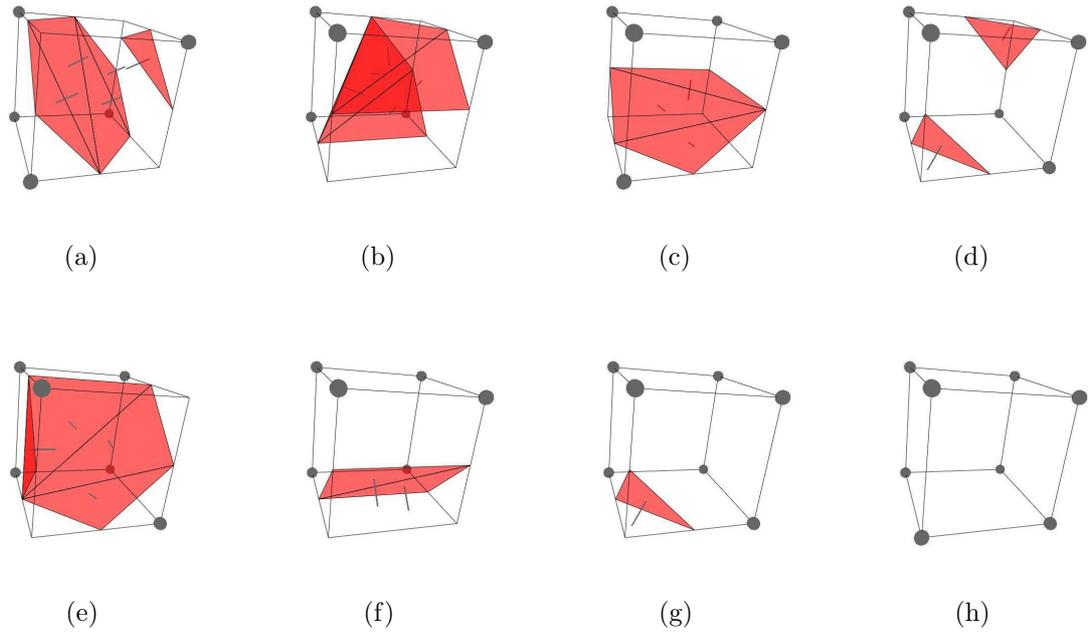
**Figure 5.6:** An example of the ambiguous case that results when the complement of configuration three (Figure 5.1 (d)) occurs next to configuration six (Figure 5.1 (g)). A hole is evident at the interface between these two cubes.

The ambiguous cases that result in unwanted holes are a direct result of the use of complementary cases in the standard MCA to reduce the number of core cube configurations that must be specified. By disregarding complementary cases and using only rotation to identify equivalent cube configurations the number of core configurations increases from 15 to 23. The eight additional cases and their associated surface patches (as defined by Nielson (2003)) are illustrated in Figure 5.7.

Altering the standard MCA to include these eight extra cases removes the necessity to generate complementary cases and thus results in the generation of airtight surfaces that do not contain unwanted holes. The result of using the revised algorithm is presented in Figure 5.8 (b) where the solitary hole that is evident at the top of Figure 5.8 (a) is no longer present.

### 5.3.3 Improved Normal Calculation

The edge detector that is used for normal generation by the standard MCA (see Figure 5.4) is very basic and provides only a rough estimation of 3-D edge direction. Vertex normals are usually only used for visualisation purposes i.e. to enable shading so that surfaces generate a more realistic response to lighting. In the enhanced MCA, normals have a dual role and are used for both visualisation and analysis purposes. As a result, a more accurate estimation of 3-D edge direction is required. The Zucker-Hummel edge operator (Zucker & Hummel 1981) was selected for this task due to its inherent smoothing effect. The use of this edge operator gives a more global indication of the normal at each vertex location. The three masks for the 3-D



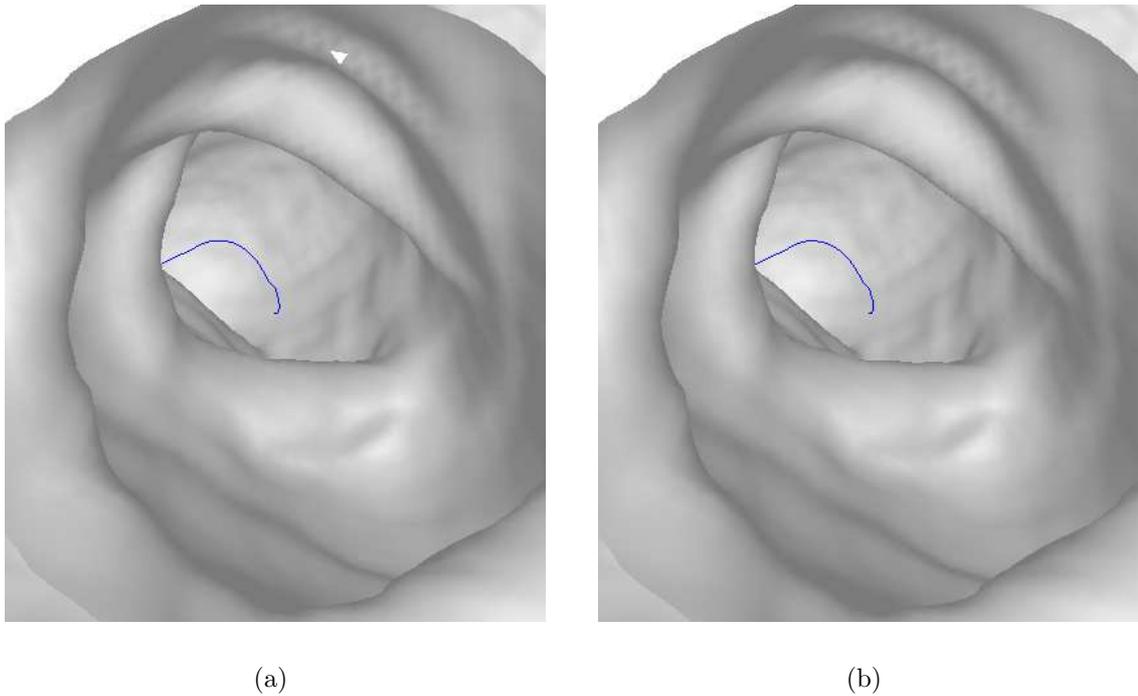
**Figure 5.7:** The eight additional configurations of eight binary voxels arranged in a cubic formation as defined by Nielson (2003). These extra configurations remove the need to generate complementary cases and as a result solve the topology problem associated with the original MCA.

Zucker-Hummel operator are illustrated in Figure 5.9.

### 5.3.4 Vertex Curvature Classification

It is clear from the reviews of the Chicago and NIH techniques in Sections 3.5 and 5.2 that these groups used a combined total of five local curvature characteristics to describe the surface of the colon. The mesh curvature is important for analysis purposes as polyps generally present as spherical projections into the colon lumen. Accordingly, each mesh vertex is classified in relation to the curvature of the mesh at that point. The curvature characteristics calculated for each vertex are:

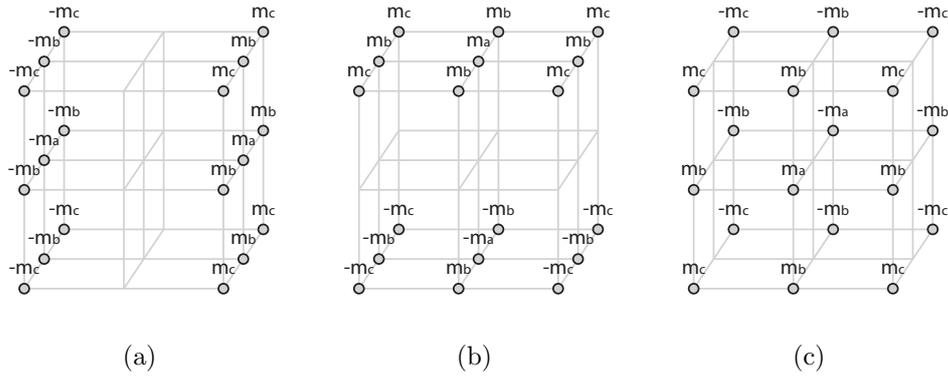
- **Mean curvature** - indicates whether the surface is convex (the mean curvature is negative) or concave (the mean curvature is positive) and the level of convexity or concavity. A flat surface has a mean curvature of 0.
- **Gaussian curvature** - indicates whether the surface is elliptic (the gaussian curvature is positive) or hyperbolic (the gaussian curvature is negative). A gaussian curvature of 0 is indicative of parabolic surface.



**Figure 5.8:** The same isosurface extracted using the 15 core neighbourhood configurations (a) and the extended 23 neighbourhood configurations (b).

- **Shape Index** - a value in the range 0 - 1 that defines the local shape to be a member of one of six shape classes: cup, rut, saddle, ridge, cap and plane where the shape index of a perfect cup is 0 and the shape index of a perfect cap is 1 (see Figure 3.8).
- **Curvedness** - indicates the curvature per unit length at the selected point, thus quantifying the level of surface curvature at the point represented by a vertex. A low curvedness value would indicate a flat region whereas a high value would indicate a highly curved region (see Figure 3.10).
- **Sphericity** - a property that is defined for a region (i.e. a set of vertices). The sphericity indicates how spherical a region is. A sphericity value of 0 would indicate a perfectly spherical region whereas a value of 2 would indicate a flat surface.

Sample surfaces and their respective mean curvatures, Gaussian curvatures and shape indices are illustrated in Figure 3.8, and the concept of curvedness is illustrated in Figure 3.10.



**Figure 5.9:** The three 26-neighbour masks representing the Zucker-Hummel edge operator where  $m_a = 1.0$ ,  $m_b = \frac{\sqrt{2}}{2}$  &  $m_c = \frac{\sqrt{3}}{3}$ .

Summers et al. (1998) originally demonstrated the use of isosurface curvature analysis to detect endobronchial lesions at virtual bronchoscopy. They subsequently extended this research to deal with the automatic detection of colorectal polyps at CTC, initially using simulated polyps (Summers et al. 2000) and then, in a more recent study (Summers et al. 2001), real polyps. Their technique involved calculating mean and Gaussian curvature values at all points that belong to the relevant isosurface, i.e. the colon wall for the detection of colorectal polyps, and then identifying regions that have polyplike characteristics or “elliptical curvature of the peak subtype”. These two measures, the Gaussian curvature ( $K$ ) and the mean curvature ( $H$ ) can be calculated as follows:

$$\begin{aligned}
 K = \frac{1}{h^2} [ & f_x^2(f_{yy}f_{zz} - f_{yz}^2) + 2f_yf_z(f_{xz}f_{xy} - f_{xx}f_{yz}) + \\
 & f_y^2(f_{xx}f_{zz} - f_{xz}^2) + 2f_xf_z(f_{yz}f_{xz} - f_{yy}f_{xz}) + \\
 & f_z^2(f_{xx}f_{yy} - f_{xy}^2) + 2f_xf_y(f_{xz}f_{yz} - f_{zz}f_{xy}) ]
 \end{aligned} \tag{5.3.4}$$

$$\begin{aligned}
 H = \frac{1}{2h^{3/2}} [ & f_x^2(f_{yy} + f_{zz}) - 2f_yf_zf_{yz} + \\
 & f_y^2(f_{xx} + f_{zz}) - 2f_xf_zf_{xz} + \\
 & f_z^2(f_{xx} + f_{yy}) - 2f_xf_yf_{xy} ]
 \end{aligned} \tag{5.3.5}$$

where  $f_x$  is the partial derivative of the image data with respect to  $x$ ,  $f_{xx}$  is the second partial derivative of the image with respect to  $x$ ,  $f_{xy}$  is the mixed partial derivative with respect to  $x$  and  $y$  etc. and  $h = f_x^2 + f_y^2 + f_z^2$ . The derivation of these formulae is dealt with by Thirion & Gourdon (1995)<sup>2</sup>. Once the values for  $H$  and  $K$

<sup>2</sup>A technical report on the same topic is also available (Thirion & Gourdon 1993).

have been determined, the principle, i.e. minimum and maximum, curvatures ( $\kappa_{min}$  &  $\kappa_{max}$ ) can be calculated for each point:

$$\kappa_{min}, \kappa_{max} = H \pm \sqrt{H^2 - K}, \kappa_{max} \geq \kappa_{min} \quad (5.3.6)$$

Summers et al. (2000) also suggest the use of another shape measurement to identify potential polyps. This measure, sphericity, indicates how spherical a region is and ranges from 0 for a perfect sphere to 2 for a flat region. The sphericity at a point is calculated using the principle and mean curvatures as follows:

$$s = \frac{\bar{\kappa}_{min} - \bar{\kappa}_{max}}{\bar{H}} \quad (5.3.7)$$

It is important to note that sphericity as defined by Summers et al. is a global measurement. It is calculated for an entire polyp candidate and not individual vertices using the average values for principle  $\bar{\kappa}_{min}$ ,  $\bar{\kappa}_{max}$  and mean  $\bar{H}$  curvatures.

The principle curvatures  $\kappa_{min}$  &  $\kappa_{max}$  can then be used to calculate values for the shape index  $R$  and the curvedness  $S$  (Koenderink 1990, pp 319 - 324).

$$R = \frac{1}{2} - \frac{1}{\pi} \arctan \left( \frac{\kappa_{min} + \kappa_{max}}{\kappa_{min} - \kappa_{max}} \right) \quad (5.3.8)$$

$$S = \frac{2}{\pi} \ln \sqrt{\frac{\kappa_{min}^2 + \kappa_{max}^2}{2}} \quad (5.3.9)$$

Yoshida & Näppi (2001) use these curvature values for identifying initial polyp candidates in their polyp detection algorithm.

#### 5.3.4.1 Discrete Implementation

It can be seen from Equations 5.3.4 and 5.3.5 that several first and second order derivatives must be calculated for a particular surface (the colonic mucosa in this case) before the associated values for the mean and Gaussian curvatures can be obtained. It is possible to calculate the first derivative of a 3-D digital image by convolving the image with the derivative of a Gaussian in one direction and smoothing it with the Gaussians with the same standard deviation in the other two directions. Similarly the second derivative can be found by convolving the image with the second derivative of a Gaussian in one direction and smoothing as outlined above. Mixed partial derivatives can be obtained using a similar method (Lohmann 1998).

The one dimensional Gaussian  $g_0(t)$  and its first and second derivatives,  $g_1(t)$  and  $g_2(t)$ , are given below:

$$g_0(t) = ce^{\frac{-t^2}{2\sigma^2}} \quad (5.3.10)$$

$$g_1(t) = \frac{-tc}{\sigma^2} e^{\frac{-t^2}{2\sigma^2}} \quad (5.3.11)$$

$$g_2(t) = \frac{c}{\sigma^2} e^{\frac{-t^2}{2\sigma^2}} \left( \frac{t^2}{\sigma^2} - 1 \right) \quad (5.3.12)$$

where  $c = \frac{1}{\sigma\sqrt{2\pi}}$ . Plots of these three functions are illustrated in Figure 5.10. Alternatively Deriche filters (Deriche 1987) can be used to calculate the required derivatives. The one dimensional Deriche filter  $d_0(t)$  and its first and second derivatives,  $d_1(t)$  and  $d_2(t)$ , are given below:

$$d_0(t) = c_0(1 + \alpha_1|t|)e^{-\alpha_1|t|} \quad (5.3.13)$$

$$d_1(t) = c_1\alpha_1^2 e^{-\alpha_1|t|} \quad (5.3.14)$$

$$d_2(t) = c_2(1 - c_3\alpha_2|t|)e^{-\alpha_2|t|} \quad (5.3.15)$$

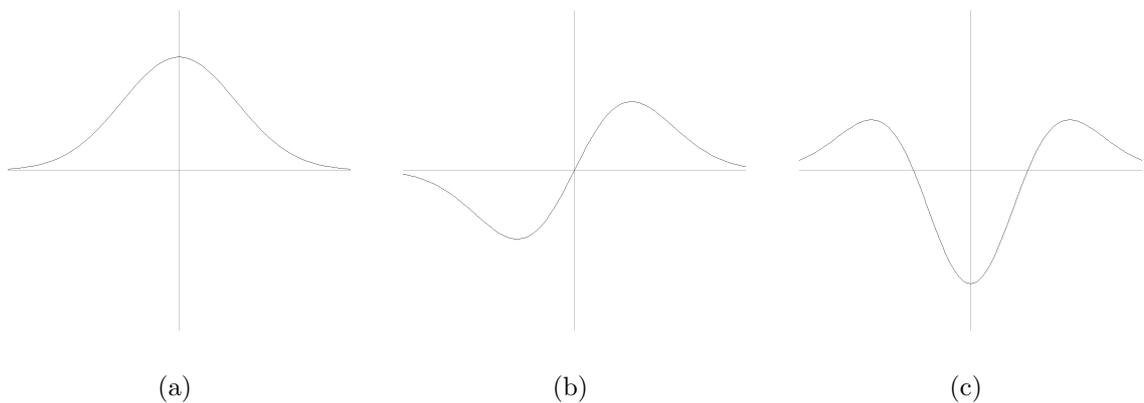
where  $\alpha_1$  and  $\alpha_2$  are smoothing parameters and control the specific shape of the filters. Plots of the Deriche function and its derivatives are illustrated in Figure 5.11. Note that the wider Deriche filters provide a greater degree of smoothing than the Gaussian filters.

The process of curvature calculation using the aforementioned technique (i.e Gaussian smoothing) is extremely computationally intensive and calculating the curvature characteristics at each voxel in a CTC data set is an extremely time consuming (and unnecessary) process. Only the curvature values for voxels in the proximity of the isosurface vertices associated with the colonic mucosa are required (i.e. those voxels associated with cases 1 - 21 of the extended MCA) to completely characterise the curvature of the surface model of the colon. This significantly reduces the amount of data processing required as the isosurface voxels represent a very small percentage of the voxels that are present in the data set. The partial derivatives can be calculated as follows:

$$\begin{aligned}
f_x &= (f_1(x)f_0(y)f_0(z)) * I, \\
f_{xx} &= (f_2(x)f_0(y)f_0(z)) * I, \\
f_{xy} &= (f_1(x)f_1(y)f_0(z)) * I, \\
f_y &= (f_0(x)f_1(y)f_0(z)) * I, \\
f_{yy} &= (f_0(x)f_2(y)f_0(z)) * I, \\
f_{yz} &= (f_0(x)f_1(y)f_1(z)) * I, \\
f_z &= (f_0(x)f_0(y)f_1(z)) * I, \\
f_{zz} &= (f_0(x)f_0(y)f_2(z)) * I, \\
f_{xz} &= (f_1(x)f_0(y)f_1(z)) * I.
\end{aligned} \tag{5.3.16}$$

where  $f_n(t)$  is a discretised version of the selected filter, i.e. either  $g_n(t)$  or  $d_n(t)$  and  $I(x, y, z)$  is the 3-D image in question (i.e. the CTC data set).

Lohmann (1998) recommends using Gaussian filters and suggests a filter width of 7 voxels and a standard deviation of 1.0 (see Figure 5.10). In contrast, Summers (2000) advises the use of Deriche filters. This choice is motivated by the fact that the mean and Gaussian curvatures are susceptible to noise, and that Deriche filtering can be used to smooth the data and calculate the derivatives simultaneously. Deriche filters were used for calculation of the required partial derivatives based on this observation.



**Figure 5.10:** The Gaussian filter  $g_0(x)$  (a) and its first and second derivatives  $g_1(x)$  (b) and  $g_2(x)$  (c). The standard deviation  $\sigma$  is set to 1.0 as recommended by Lohmann (1998).

Monga & Benayoun (1995) suggest using values of 1.0 and 0.7 for  $\alpha_1$  &  $\alpha_2$  in conjunction with the Deriche filters. This is based on the fact that higher order

derivatives require greater smoothing in order to obtain reasonable immunity to noise. However a recent analysis of the kernel method for surface curvature estimation by Campbell & Summers (2004) suggests that the smoothing parameters  $\alpha_1$  &  $\alpha_2$  should be 0.7 and 0.1 respectively as these demonstrated a reasonable accuracy and low variance in their study. To complete the calculation of the relevant derivatives Monga & Benayoun suggest selecting the normalisation coefficients for the Deriche filters  $c_0$ ,  $c_1$ ,  $c_2$  and  $c_3$  so that:

$$\begin{aligned} \int_{-\infty}^{+\infty} d_0(x)dx &= 1 \\ \int_{-\infty}^{+\infty} xd_1(x)dx &= 1 \\ \int_{-\infty}^{+\infty} d_2(x)dx &= 0 \quad \text{and} \quad \int_{-\infty}^{+\infty} \frac{x^2}{2}d_2(x)dx = 1 \end{aligned} \tag{5.3.17}$$

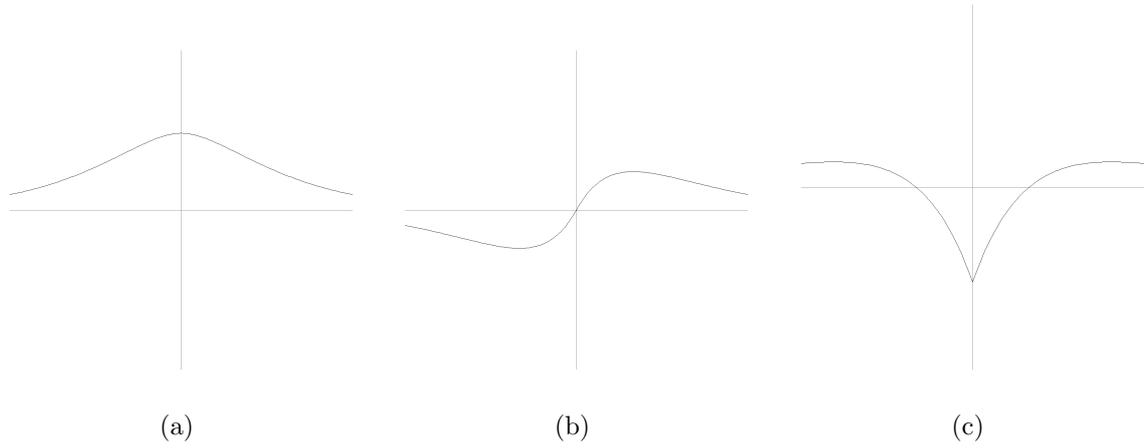
In the case of a discrete implementation with isotropic integer sampling these formulae become:

$$\begin{aligned} \sum_{-\infty}^{+\infty} d_0(x) &= 1 \\ \sum_{-\infty}^{+\infty} xd_1(x) &= 1 \\ \sum_{-\infty}^{+\infty} d_2(x) &= 0 \quad \text{and} \quad \sum_{-\infty}^{+\infty} \frac{x^2}{2}d_2(x) = 1 \end{aligned} \tag{5.3.18}$$

By substituting the Deriche filter and its first and second derivatives for  $f_0(x)$ ,  $f_1(x)$  &  $f_2(x)$  the following representations of the normalisation coefficients are obtained:

$$\begin{aligned} c_0 &= \sum_{-r}^{+r} \frac{1}{(1 + \alpha_1|x|)e^{-\alpha_1|x|}} \\ c_1 &= \sum_{-r}^{+r} \frac{1}{x^2\alpha_1^2e^{-\alpha_1|x|}} \\ c_2 &= \sum_{-r}^{+r} \frac{2}{x^2(1 - c_3\alpha_2|x|)e^{-\alpha_2|x|}} \\ c_3 &= \sum_{-r}^{+r} \frac{1}{\alpha_2|x|} \end{aligned} \tag{5.3.19}$$

where  $r$  is the half filter width (the filter width  $w$  being  $2r + 1$  voxels) and  $\alpha_1$  and  $\alpha_2$  have the values suggested by Campbell & Summers (2004).



**Figure 5.11:** The Deriche filter  $d_0(x)$  (a) and its first and second derivatives  $d_1(x)$  (b) and  $d_2(x)$  (c). The values for  $\alpha_1$  and  $\alpha_2$  are set to 0.7 and 0.1 respectively as suggested by Campbell & Summers (2004) and the normalisation coefficients  $c_0$ ,  $c_1$ ,  $c_2$  &  $c_3$  are calculated using Equation 5.3.19 as suggested by Monga & Benayoun (1995).

This process generates the relevant curvature values for the voxels that encapsulate the isosurface representing the colon wall. In order to calculate the curvature values for the vertices of the isosurface that represents the colon wall, an interpolation scheme must be used. The interpolated value of the relevant curvature property  $p_i$  (e.g. shape index ( $R_i$ ) or Gaussian curvature ( $K_i$ )) is calculated as follows:

$$p_i = (1 - \delta_{iso})p_a + \delta_{iso}p_b \quad (5.3.20)$$

where,  $p_a$  and  $p_b$  are the curvature values associated with the two voxels that lead to the generation of the vertex in question, and  $\delta_{iso}$  is the normalised distance from the voxel that is nearest the origin.

### 5.3.5 Additional Vertex Properties

Other properties are calculated for each vertex in addition the curvature characteristics. These are:

- The colon segment that the vertex belongs to (i.e. the rectum, sigmoid colon, descending colon, transverse colon or ascending colon, as identified using sub-segmentation (see Section 4.6)). The colon segment information can be used to facilitate automated reporting by localising an automatically detected polyp to a particular colon segment. This provides a much more meaningful indication of polyp location than an (x, y, z) coordinate.

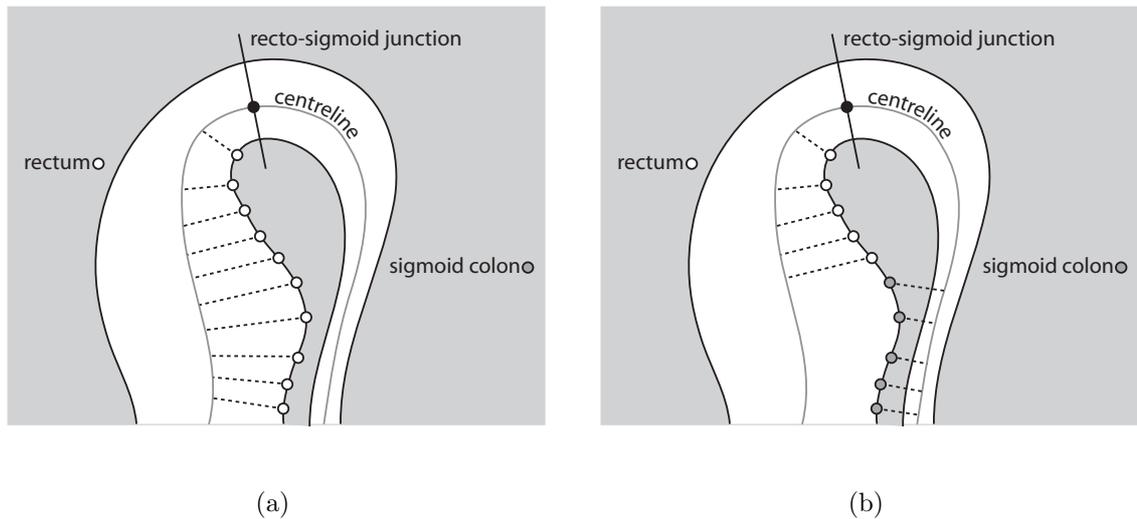
- The density, measured in HU, at the point in the colonic mucosa that corresponds to the vertex. The density information can be used to discriminate between true polyps and false positives that are caused by residual faecal material.

Both of these vertex properties are calculated during the surface extraction process using the modified marching cubes algorithm.

### 5.3.5.1 Colon Segment Information

The most straightforward method for obtaining the segment vertices is to find the closest centreline point to each individual vertex and assign the vertex to the same region as the closest centreline point. This approach may seem to be adequate, however, it was found to be susceptible to a problem where vertices in one segment were incorrectly flagged as belonging to another segment. This occurred in cases where, for example, a vertex actually belonged to the rectum but it was closer to a centreline point located in the sigmoid colon. In this case the rectum vertex would be incorrectly flagged as belonging to the sigmoid colon. This particular case is illustrated in Figure 5.12. A simple line of sight solution was considered where a vertex was assigned to the same segment as the closest centreline point in its field of view. However, this solution was not practical as the line of sight between a vertex and the centreline can be obscured by a Haustral fold or some other obstruction. In these cases the colon surface would contain a patch with undefined colon segment information.

An alternative solution ensures that each vertex is assigned the correct colon segment information and that no patches are left undefined. This involves reversing the centreline calculation process and performing automated seeded region growing within the limits of the colonic mucosa. This region growing process is initiated from the centreline points i.e. each centreline point represents an individual seed point. As the centreline expands, the voxels that are added to the resulting structure are assigned the same colon segment information as their nearest neighbour. This process effectively dilates the colon centreline so that it expands to fill the same region as the original segmented colon lumen. As a result, each individual centreline segment expands to fill its associated region in the colon lumen. In order to prevent leakage between the segments during the dilation process, planar structures three voxels thick are inserted at each flexure point that lies on the colon centreline (i.e those



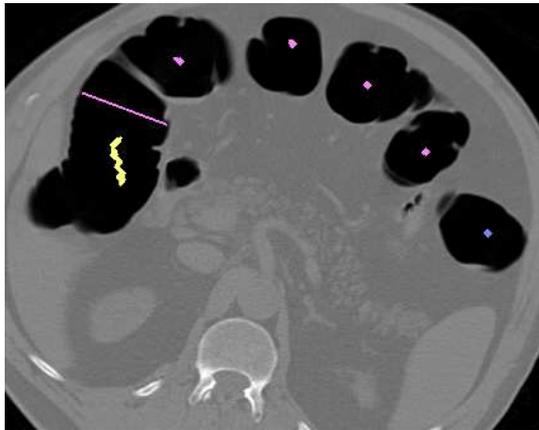
**Figure 5.12:** An example of the segment skipping problem (a) the expected case where colon wall vertices are associated with centreline segments in the same segment (b) the actual case where vertices associate themselves with the closest centreline segment resulting in incorrect segment assignments.

points detected during the subsegmentation process). The planes are orientated perpendicularly to the centreline at the relevant flexure points. This ensures that the border between each segment is smooth and well defined.

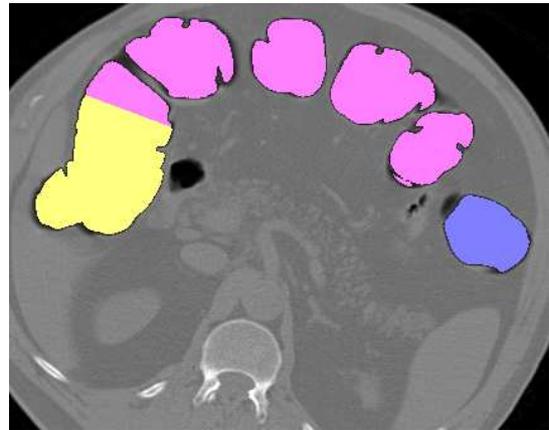
Once the dilation process is complete, each vertex is assigned to the same colon segment as its closest neighbouring lumen voxel. The closest lumen voxel to a vertex is detected by performing a search in the  $3 \times 3 \times 3$  voxel neighbourhood surrounding the vertex. The vertex is subsequently assigned to the same segment as the lumen voxel with the shortest Euclidean distance. An illustration of the anatomical vertex labeling process is presented in Figure 5.13. Colour coding is used to indicate the different colon segments.

### 5.3.5.2 Vertex Density Information

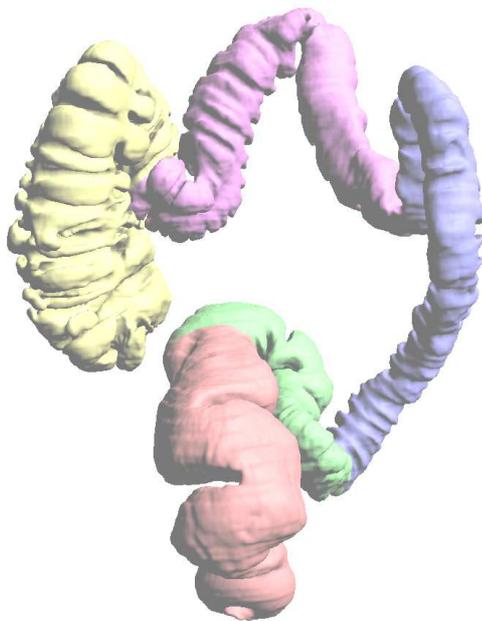
Density information has been used by others to discriminate between true polyps and false positives that are caused by residual faecal material (Summers et al. 2001) in CTC data sets. This process is especially effective when a contrast enhanced bowel preparation is used (see Section 3.2.3) as the conspicuity between polyps and residual fecal material is greatly increased. Density information can also be used to identify false positive detections that are attributable to the ICV (Summers et al. 2004).



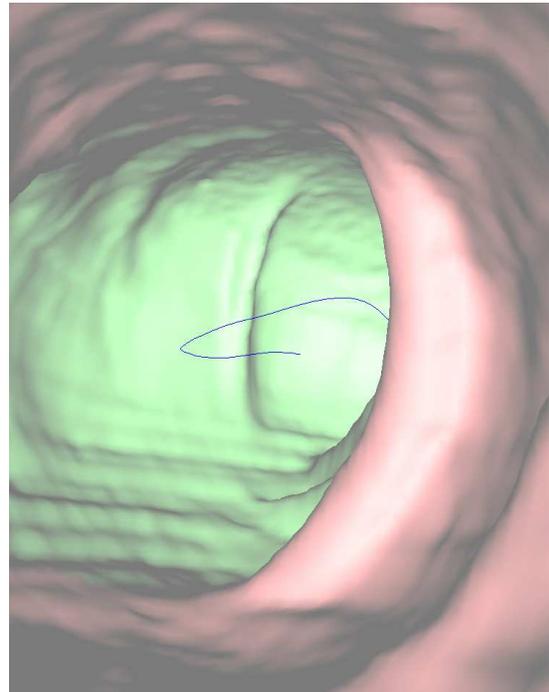
(a)



(b)



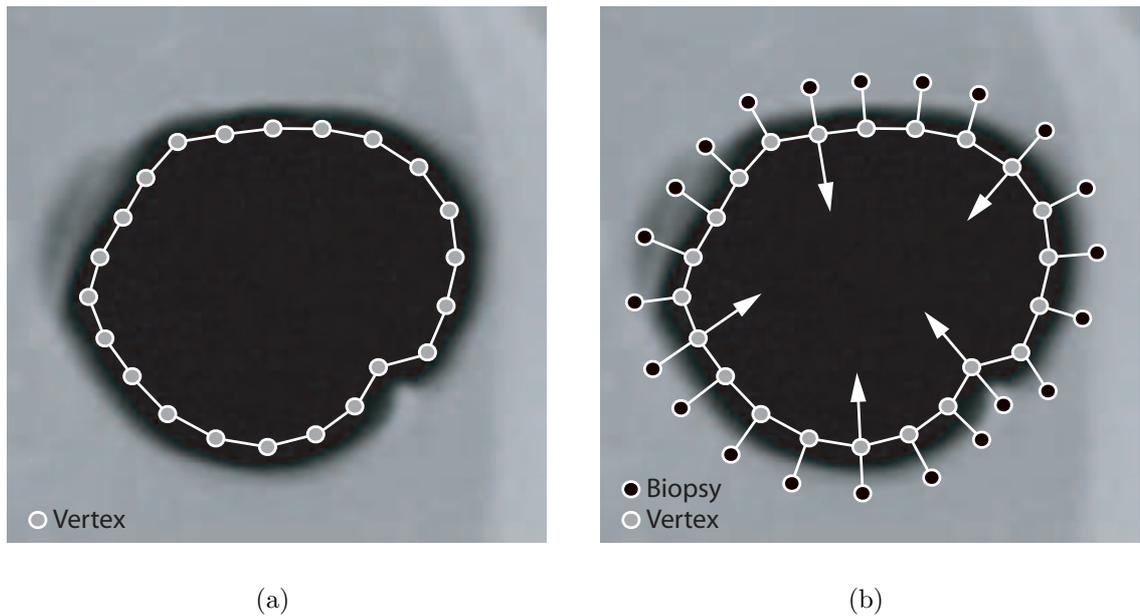
(c)



(d)

**Figure 5.13:** The anatomical vertex labeling process. (a) Shows the wall at the hepatic flexure and the region growing process after three iterations. (b) Shows the completed region growing process where a smooth transition exists between the transverse colon (purple) and the ascending colon (orange). (c) Shows an exterior view of the labeled colon and, (d) shows an endoluminal view of the splenic flexure.

The density at each vertex can be obtained by sampling the data set at a fixed distance from the vertex in the opposite direction of the normal at that point. The data at the required sample point is obtained using the trilinear interpolation technique described in Section 4.3.1. This is essentially the same as performing a virtual biopsy at every vertex in the mesh representing the colon surface. The distance of the sample from the vertex (i.e. the sampling distance) was selected to be 2 mm in order to overcome the partial volume effect at the colon walls. The density sampling process is illustrated in Figure 5.14.



**Figure 5.14:** An illustration of the density sampling process. (a) The interconnected vertices that result from the surface extraction algorithm. (b) The corresponding density samples (or virtual biopsies) are located at a fixed distance (the sampling depth) from each vertex in the opposite direction of the normal at that point.

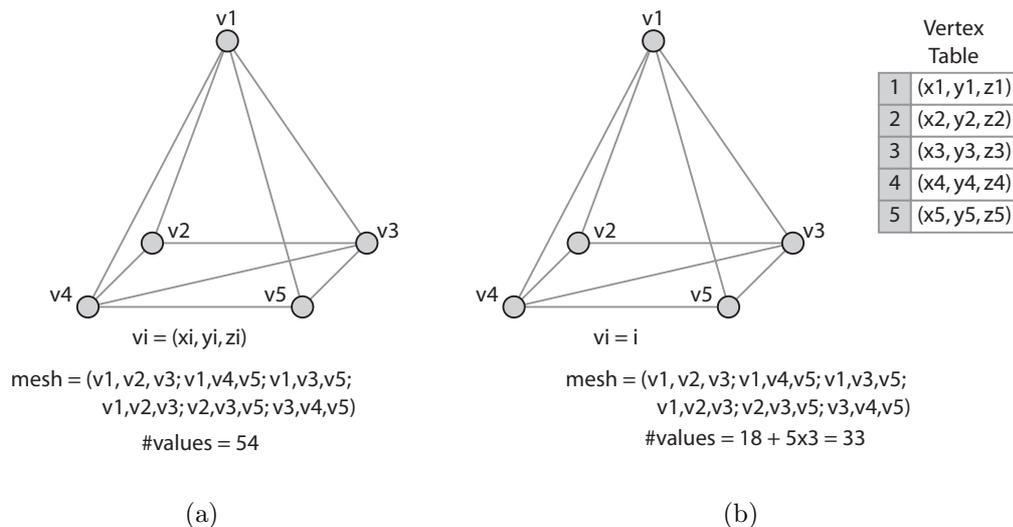
### 5.3.6 Mesh Representation

The mesh generated by the standard MCA consists of a disjointed set of triangular patches where there is a high degree of vertex repetition<sup>3</sup>. This representation, although suitable for visualisation purposes, is not ideal for analysis. It is also extremely wasteful of memory due to the high degree of vertex repetition.

An alternative mesh representation involves storing each vertex in an array structure where a particular vertex can be present only once. Using this approach, tri-

<sup>3</sup>An evaluation of several data sets used in this research found that each colon surface vertex was reused an average of 6 times.

angles that represent the mesh are specified as indices into the vertex array and not as actual vertices triplets. An array of unit normals is generated and populated in the same manner, and the same approach is used to store the vertex measurements discussed in Sections 5.3.4 and 5.3.5. This approach to mesh storage has the potential to significantly reduce the amount of memory that is required to store a fully characterised isosurface representation of the colonic mucosa. A simple example illustrating the difference between the standard mesh and the indexed mesh is presented in Figure 5.15.

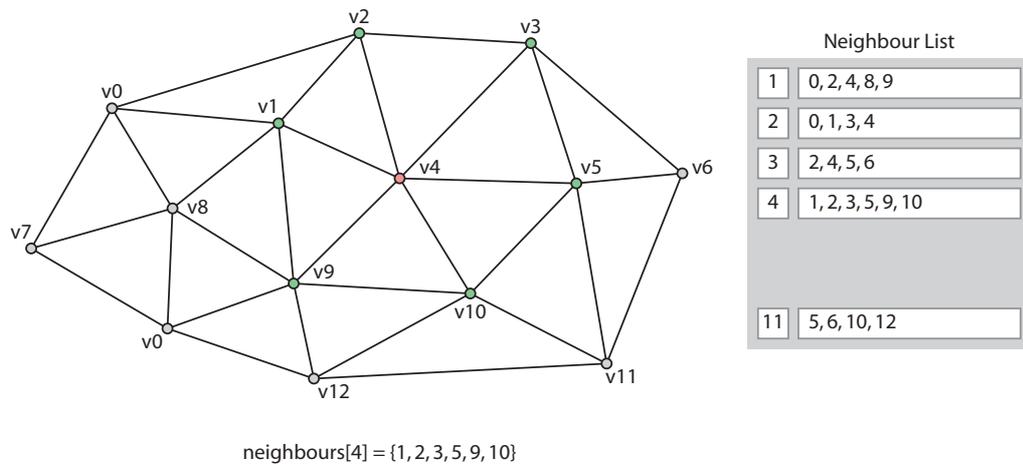


**Figure 5.15:** An illustration of how indexing can be used to reduce the amount of data required to specify a mesh structure. A pyramid consists of five vertices and six triangles. The unoptimised representation yields 54 values (a) and the optimised alternative yields only 33 (b). Note that in Java the data type for vertex (float) and index (int) both require four bytes of storage space.

The indexing process requires a modification to the standard MCA. As each vertex is encountered it is assigned an index and stored in the vertex list. This index is then used to represent the relevant vertex coordinates (i.e. an Integer primitive (four bytes) is used to reference three floating point primitives (12 bytes)). If a vertex is encountered that was already assigned an index then that index is used. Conversely, if a new vertex is encountered then a new index is generated. Only vertices associated with the same slice or two adjacent slices can be shared. As a result, only two slices need to be resident in memory at any one time. Noting that in the context of the marching cubes algorithm a slice is actually two voxels thick.

### 5.3.7 Neighbour Identification

The final modification to the standard MCA involves identifying all of the directly connected neighbouring vertices within the mesh. This step is required to facilitate region growing operations where region growing is used to identify initial polyp candidates that meet the required shape and size criteria. As each triangle, consisting of vertices  $e_0$ ,  $e_1$  and  $e_2$ , is identified, its vertex relationships are added to an array structure where:  $e_0$  is associated with  $e_1$  &  $e_2$  (i.e. two vertex pairs  $(e_0, e_1)$  and  $(e_0, e_2)$ ),  $e_1$  is associated with  $e_0$  &  $e_2$  and  $e_2$  is associated with  $e_0$  &  $e_1$ . A vertex pair is only added to the array structure if this vertex pair is not already present. An example illustrating the neighbour identification process is illustrated in Figure 5.16. Representing the vertex neighbours in this way reduces the task of neighbourhood identification from an extensive mesh search to a simple table lookup i.e. by specifying the index of one vertex the indices of all of the neighbouring vertices are returned. This method for recording the vertex relationships within the mesh has also been used by (Li et al. 2005).



**Figure 5.16:** The neighbour indexing process: The neighbouring vertices for each mesh vertex are stored in a list to streamline the process of searching for vertex neighbours.

### 5.3.8 Modified MCA Summary

The improved version of the MCA described in this Section represents a significant extension over the standard MCA. The standard version generates vertices and unit normals that represent a particular isosurface. The improved version fixes the topology problems associated with the standard algorithm, provides local shape measurements for each vertex in the extracted mesh, generates more accurate unit

normals and provides information regarding vertex neighbours, density at a particular vertex, as well as being able to assign a vertex to a particular anatomical segment of the colon. All of this is done in an extremely memory efficient manner<sup>4</sup>. Organising and characterising the colonic mucosa in this way provides a robust and flexible foundation for the development and implementation of multistage CAD-CTC algorithms.

## 5.4 Initial Polyp Candidate Identification

Initial polyp candidate identification involves segmenting the colonic mucosa into two categories: suspicious regions that may be representative of colorectal polyps and normal regions attributable to benign, naturally occurring colonic features. This operation represents the primary stage in a multistage CAD-CTC system typical of the more recently documented systems that were reviewed in Section 3.5. The primary stage is intended to quickly remove large sections of the colon surface that are obviously representative of benign structures, and its output typically consists of a distinct set of surface patches that represent any actual polyps that are present in the CTC data set as well as numerous false positive regions (see Figure 5.17). These surface patches can then be evaluated individually using a subsequent, more comprehensive, analysis stage. Thus, the purpose of the primary stage is twofold:

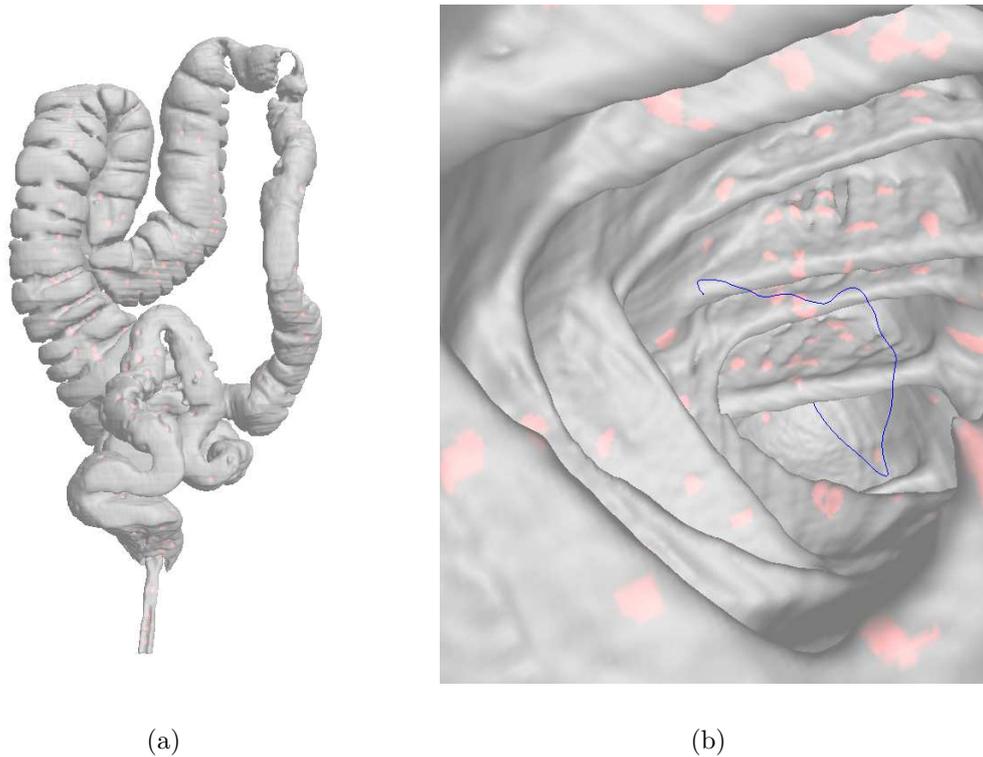
- **Data reduction** - to eliminate the need to perform a detailed, time consuming, analysis of regions of the colon surface that are obviously not attributable to colorectal polyps.
- **Data formatting** - so that the pertinent regions of the colon surface are represented as a series of surface patches that are more suitable for individual analysis.

It is extremely important that the primary stage retains all regions that are associated with clinically significant polyps. As a result, even though a significant amount of data reduction can be achieved, the number of false positive detections retained for subsequent evaluation can be extremely high.

Initial polyp candidate identification is performed based on an analysis of surface characteristics at vertices in the polygonal mesh that represents the colonic surface.

---

<sup>4</sup>The amount of memory required to hold all of this information is approximately 50% of the amount of memory required to store the original data set.



**Figure 5.17:** External (a) and internal (b) surface renderings of a sample CTC data set (MMH-765/supine) demonstrating the type of output generated by the initial candidate detection stage. Detected surface patches that have highlighted in red. This is an extreme example where there are a total of 562 initial polyp candidates.

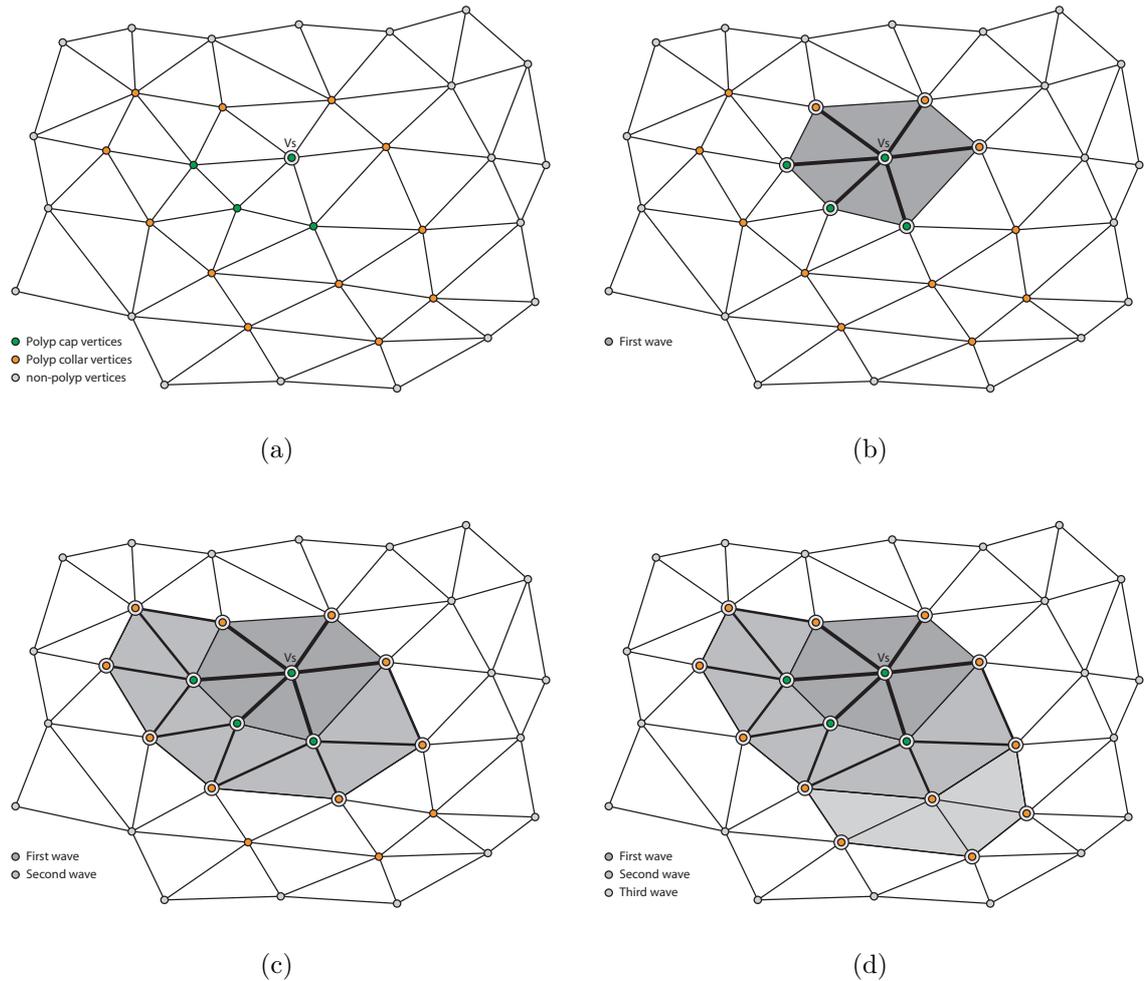
All vertices in the mesh are grouped into three distinct classes as defined by Yoshida & Näppi (2001):

1. **Polyp cap** - Vertices of the polyp cap class have a shape index greater than 0.9 and a curvedness in the range  $0.08 \text{ mm}^{-1}$  -  $0.33 \text{ mm}^{-1}$ .
2. **Polyp collar** - Vertices of the polyp collar class have a shape index greater than 0.8 and a curvedness in the range  $0.05 \text{ mm}^{-1}$  to  $0.5 \text{ mm}^{-1}$ . The requirement for the polyp collar class is motivated by the fact that peripheral regions of a polyp do not always conform to a perfect cap shape.
3. **Nonpolyp** - All other vertices that do not meet the criteria for polyp cap or polyp collar classes fall into the nonpolyp class.

Note that the maximum curvedness values for the polyp cap and polyp collar classes are higher than the values suggested by Yoshida & Näppi. This modification was made in order to facilitate the detection of smaller polyps in the 2 - 3 mm size range.

The identification of initial polyp candidates can be performed once the vertex level classification has been completed. This process involves the use of hysteresis thresholding (as described by Lohmann (1998)). Seed points are identified that belong to the polyp cap class. Vertex based region growing is initiated at these vertices and allowed to expand into vertices of the polyp collar class. The process continues until all of the surface patches with polyp like properties have been identified. Any patches that have a maximum patch diameter above 2 mm are recorded for further analysis while the remainder, which do not meet this requirement, are discarded. The minimum size constraint eliminates extremely small candidate regions and reduces the number of detections retained for further processing. It is clear that the polyp size specified above is much smaller than the minimum size for clinically significant polyps, i.e. 5 mm. This is due to the fact that only a portion of the surface of a polyp may be flagged as having polyplike characteristics. Setting the minimum size to 2 mm ensures that a 5 mm polyp will be retained for further processing even if only a small proportion the surface was flagged (i.e. approximately 15% in the case of a perfectly spherical polyp). An example of the output generated by the initial polyp candidate detection technique is illustrated in Figure 5.17. In this example, the total number of vertices in the colon surface is 714,716. The initial polyp candidate detection technique flags 13,526 (1.89%) of these before the application of the size filter, and 7,208 (1.01%) are retained after the application of the size filter. Thus, it is clear that the primary stage significantly reduces the amount of data retained for subsequent analysis.

The vertex based region growing process that is used to identify polyp candidate regions is much the same as the voxel based region growing process that is used for colon segmentation (see Section 4.4). The main difference between the two approaches is the concept of a neighbour. In the case of voxel based region growing there are several different possible neighbourhood configurations: face connected with six neighbours, edge connected with 18 neighbours and corner connected with 26 neighbours. Conversely, in the case of vertex based region growing the concept of a neighbour is very well defined. Here, the neighbours of a vertex are only those vertices that are directly connected as outlined in Section 5.3.7. The concept of vertex based region growing is illustrated in Figure 5.18.



**Figure 5.18:** Hysteresis thresholding based region growing in a polygonal mesh. An initial seed point is selected that belongs to the polyp cap class  $V_s$  (a). Neighbouring vertices belonging to either the polyp cap or polyp collar classes are added to the region (b) & (c). The region growing process completes when no more suitable neighbours are available (d). Note that the region growing process must start from a vertex of the polyp cap class.

## 5.5 False Positive Reduction

Following initial polyp candidate detection, a more detailed analysis of the flagged regions is generally carried out. This second stage analysis is intended to reduce the number of false positive detections to an acceptable level while, at the same time, retaining the high sensitivity achieved by the primary stage. This Section describes a novel false positive reduction technique that is based on a number of important observations:

1. A polyp is a spherical or ellipsoidal projection into the colon lumen. As a result the boundary of a polyp will be circular or elliptical. This is evident from the illustration of an actual polyp in Figure 5.19 (b) where the boundary of the

polyp is roughly elliptical.

2. In the case of an actual polyp, the candidate region may not completely encompass the region of the colon surface that is attributable to that polyp. This is evident from the illustration of an actual polyp and its associated candidate region in Figure 5.19 (c).
3. A candidate region attributable to a false positive detection can have an irregular boundary with numerous bays. This is evident from the illustration of sample false positive detections in Figures 5.19 (d) & (e).

This novel false positive reduction technique involves generating an alternative polyp candidate representation and expanding the resulting candidate region so that it is more representative of an actual polyp. Thus, an actual polyp should retain its original polyplike properties whereas a false positive should expand to include more nonpolyp features. This should make it easier to discriminate between true polyps and false positives.

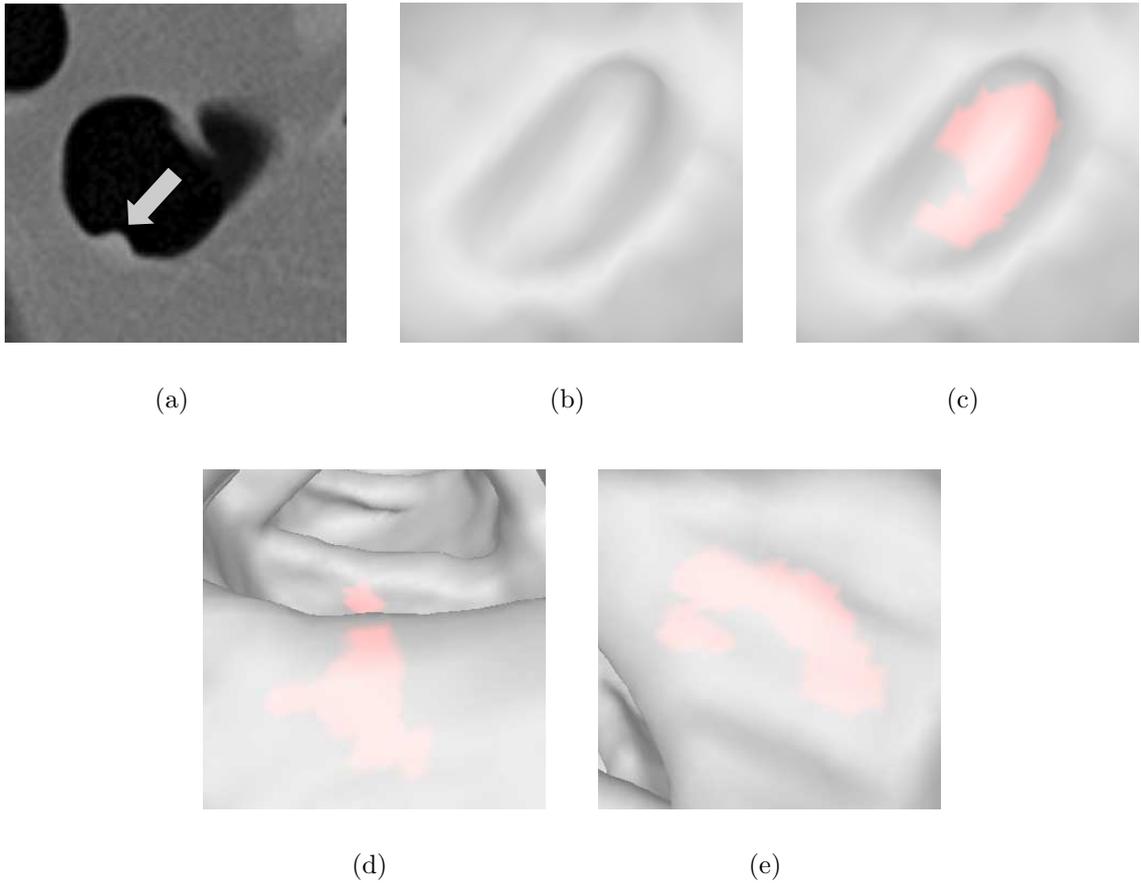
### 5.5.1 Alternative Polyp Representation

The alternative polyp representation involves generating a 2-D approximation of the 3-D surface patch that is associated with a polyp candidate. This alternative polyp representation can then be analysed using 2-D techniques. The generation of the alternative representation consists of two stages: The first stage, polyp candidate projection, involves projecting the vertices of the candidate onto a plane. The second stage essentially unwraps the candidate onto the plane so that the entire surface of the candidate can be visualised. The entire process will be referred to as fingerprinting as it creates a 2-D approximation of a 3-D surface.

#### 5.5.1.1 Polyp Candidate Projection

Polyp candidate projection involves projecting the individual vertices of the polyp candidate onto a plane that is roughly parallel to the surface of the candidate. The first stage in this process involves the definition of the projection plane. A plane can be defined by specifying:

- A point that belongs to that plane  $(x_p, y_p, z_p)$ . The central candidate vertex is used for this purpose, and is found by identifying the vertex with the minimum



**Figure 5.19:** Examples of a true polyp and false positive detection generated by the initial candidate detection stage. (a) A 2-D axial slice containing a polyp. (b) A 3-D surface rendering of the region containing the polyp. (c) The section of the polyp identified by the initial candidate detection stage. (d) & (e) Irregular shaped false positive detections caused by regions of the colon surface with polyplike properties.

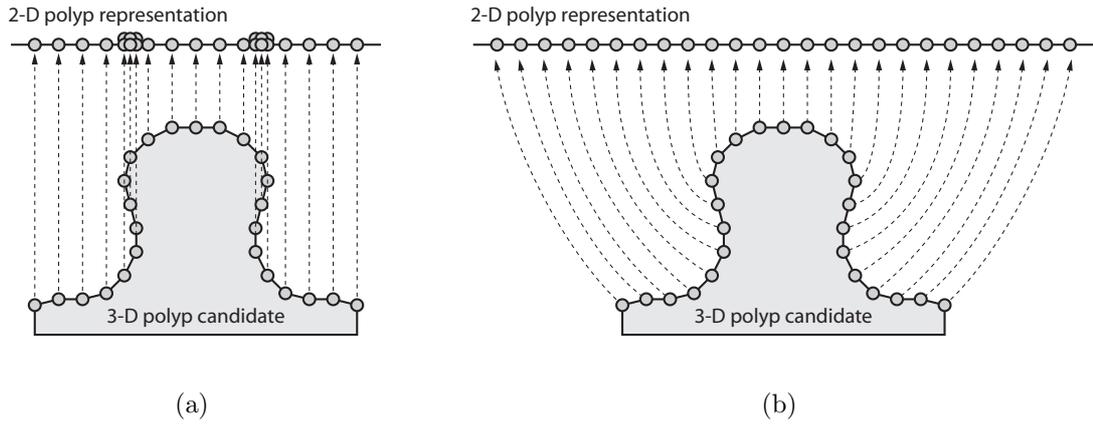
accumulated Euclidean distance to all other vertices in the candidate, by means of an exhaustive search.

- A vector specifying the direction of the normal to the surface of the plane  $(a, b, c)$ . The required vector is calculated by averaging the normals at each vertex in the polyp candidate. Thus yielding the average normal for the candidate region.

(see Equation 5.5.1). The resulting plane intersects the centremost candidate vertex and is roughly parallel to the surface of the candidate.

$$a(x - x_p) + b(y - y_p) + c(z - z_p) = 0 \quad (5.5.1)$$

The projection process involves generating a line perpendicular to the plane (i.e. in the same direction as the normal to the surface of the plane) that intersects each



**Figure 5.20:** Alternative polyp representations. (a) Polyp candidate projection, where each candidate vertex is projected onto the plane in the direction of the normal to the surface of the plane. (b) Polyp candidate fingerprinting, an extension to polyp candidate projection, where a flattened representation of the surface, that preserves the original vertex layout, is generated.

of the candidate vertices. The intersection that this line makes with the plane is recorded as the projected vertex location. The line is defined by a direction i.e.  $(a, b, c)$  from above and a point that belongs to the line  $(x_l, y_l, z_l)$  (i.e. the relevant vertex location). This is represented by the following parametric equation:

$$\begin{aligned} x &= x_l + at \\ y &= y_l + bt \\ z &= z_l + ct \end{aligned} \tag{5.5.2}$$

By substituting the values of  $x$ ,  $y$  and  $z$  from 5.5.2 into 5.5.1 the following value for  $t$  is obtained:

$$t = \frac{a(x_p - x_l) + b(y_p - y_l) + c(z_p - z_l)}{a^2 + b^2 + c^2} \tag{5.5.3}$$

Substituting this value for  $t$  back into 5.5.2 gives the coordinates for the intersection between the line originating from the associated polyp candidate vertex and the plane, i.e. the intersection point  $(x_i, y_i, z_i)$ . This process is repeated for all candidate vertices, as well as their immediate neighbours.

The intersection points consist of a set of coplanar 3-D coordinates and must be transformed into 2-D coordinates in order to facilitate analysis of the projected candidates using 2-D techniques. This is achieved by first transforming the points to 2-D

polar coordinates and subsequently transforming them to 2-D cartesian coordinates as follows:

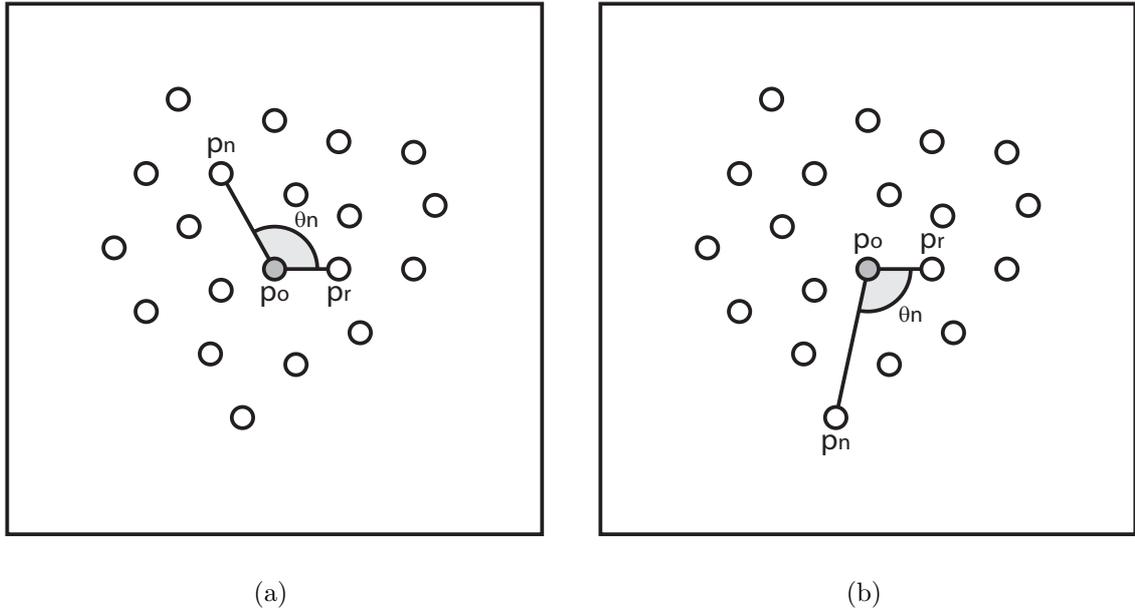
- Identify an origin point. The central point used to define the plane. The point  $(x_p, y_p, z_p)$  is probably the most suitable option.
- Identify another random intersection point with the plane that is not  $(x_p, y_p, z_p)$ . These two points are considered to define a vector in the direction of the positive x-axis (reference vector  $\vec{u}$ ).
- Calculate all of the vectors  $\vec{v}_n$  that connect the origin to each intersection point.
- Each intersection point can then be represented by:
  - An angle  $\theta_n$ , the angle between  $\vec{u}$  and  $\vec{v}_n$ , which is calculated from the dot product (5.5.4)
  - A length  $r_n$ , which is the magnitude of  $\vec{v}_n$

i.e. the polar coordinates  $(r_n, \theta_n)$

$$\theta_n = \cos^{-1} \left( \frac{\vec{u} \cdot \vec{v}_n}{|\vec{u}| |\vec{v}_n|} \right) \quad (5.5.4)$$

Further calculations are necessary because the dot product will only generate angles in the range  $0 \rightarrow \pi$ , i.e. the smallest of the angles between the two vectors (see Figure 5.21). The required angle, which represents the anti-clockwise angle  $\varphi_n$  between the two vectors (from the viewer's perspective), can be obtained using the cross product (see Equations 5.5.5 and 5.5.6) as follows:

- If the cross product is in the same direction as the surface normal then the anti-clockwise angle  $\varphi_n$  with respect to the reference vector  $\vec{u}$  is the angle obtained from the dot product i.e.  $\theta_n$ .
- However, if the cross product is in the opposite direction to the surface normal then the anti-clockwise angle  $\varphi_n$  with respect to the reference vector  $\vec{u}$  is equal to  $2\pi - \theta_n$ .



**Figure 5.21:** The smallest angle between a reference vector  $\vec{r}$  and another vector  $\vec{n}$  can be calculated using the dot product. The resulting angle can be either clockwise (a) or anticlockwise (b) with respect to the reference vector.

where the cross product of the reference vector  $\vec{u}$  and the vector under investigation  $\vec{v}_n$  is:

$$\vec{u} \times \vec{v}_n = \begin{vmatrix} u_y & u_z \\ v_{ny} & v_{nz} \end{vmatrix} i - \begin{vmatrix} u_x & u_z \\ v_{nx} & v_{nz} \end{vmatrix} j + \begin{vmatrix} u_x & u_y \\ v_{nx} & v_{ny} \end{vmatrix} k \quad (5.5.5)$$

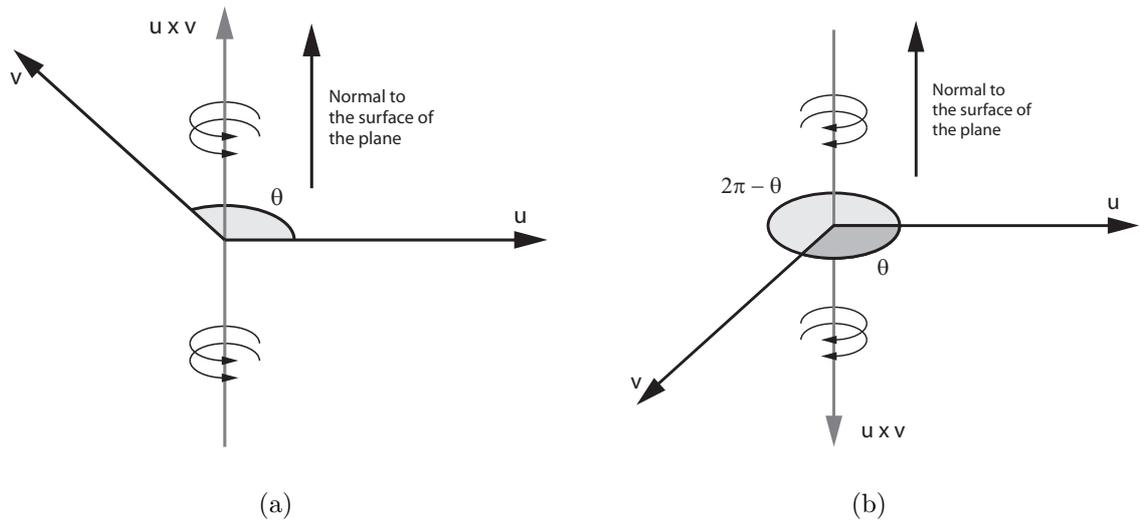
or more appropriately from an implementation perspective:

$$\vec{u} \times \vec{v}_n = (u_y v_{nz} - u_z v_{ny}) - (u_x v_{nz} - v_{nx} u_z) + (u_x v_{ny} - v_{nx} u_y) \quad (5.5.6)$$

Note that the orientation of the angles  $\theta_n$  and  $\varphi_n$  is from the viewer's perspective and that the viewing vector is in the opposite direction to the surface normal  $(a, b, c)$  discussed earlier. The angle adjustment process is illustrated in Figure 5.22. Once the angle has been adjusted, the intersection point can be converted from polar coordinates to cartesian coordinates for display purposes where,  $x_n = r \cos(\varphi_n)$  and  $y_n = r \sin(\varphi_n)$ .

### 5.5.1.2 Polyp Candidate Fingerprinting

It is evident from the example in Figure 5.20 (a) that polyp candidate projection results in vertex overlap in situations where the vertices being projected belong to a surface that is perpendicular to the projection plane. An alternative polyp representation, that provides a more even distribution of candidate vertices, can be achieved



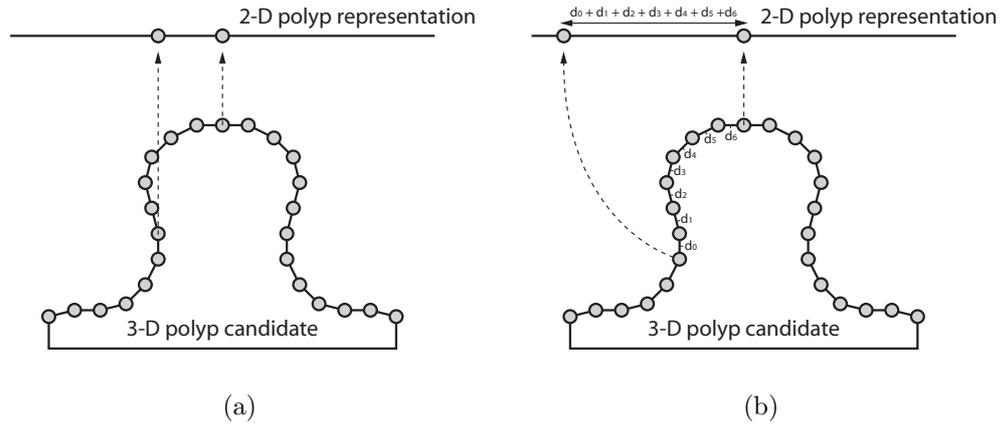
**Figure 5.22:** Calculating the anticlockwise angle from a reference vector  $u$  to another vector  $v$ . (a) If the cross product ( $u \times v$ ) is in the same direction as the normal to the surface of the plane then the required angle  $\theta$  is obtained using the dot product. (b) However, if the cross product is in the opposite direction to the normal then the anticlockwise angle is given by  $2\pi - \theta$ .

by making a minor modification to the polyp candidate projection process. The alternative representation, which will be referred to as polyp candidate fingerprinting, involves altering the radial component of polar coordinates generated during the polyp candidate projection process.

The revised radial component of the polar coordinates represents the geodesic distance from the relevant vertex to the reference point along the surface of the colonic mucosa. This process is illustrated in Figure 5.23 (b) and can be compared to the original approach, illustrated in Figure 5.23 (a), where the distance from the relevant intersection to the reference point was used. In the revised approach, the distance along the colon surface from a particular vertex to the reference vertex is approximated by identifying the shortest path between the two vertices in the polygonal mesh that represents the colon surface. Mapping the modified polar coordinates into cartesian coordinates, provides a flat representation of the polyp candidate vertices that can be visualised and analysed using conventional 2-D techniques.

### 5.5.1.3 Reformatted Polyp Candidate Visualisation

The final stage of the reformatting process involves visualising the 2-D representation of the polyp candidate. The reformatted polyp candidate consists of irregularly spaced coplanar data points. The visualisation of these data points as a continuous



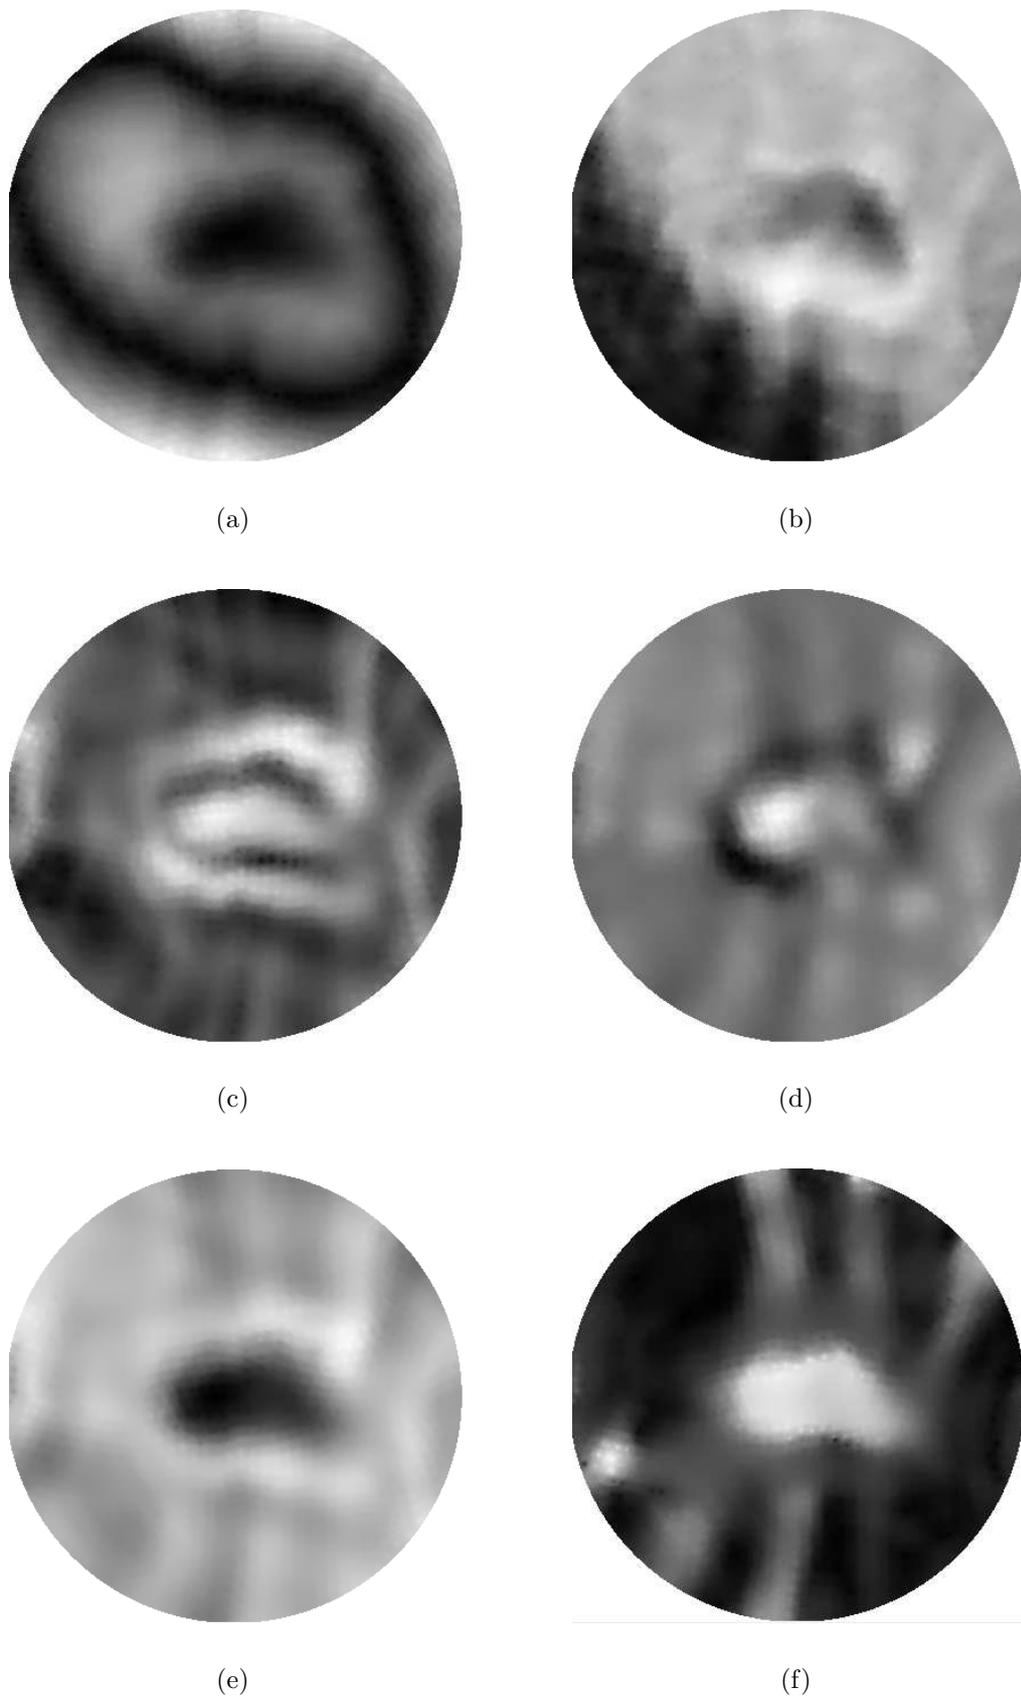
**Figure 5.23:** An illustration of the different methods used for obtaining the radial component of the polar coordinates for a vertex being mapped to a 2-D plane (a) Polyp candidate projection (b) Polyp candidate fingerprinting.

2-D digital image necessitates the use of some kind of interpolation scheme. Shepard (1968) proposed a 2-D interpolation technique for irregularly-spaced data that was developed for use with variables such as population density, housing conditions and other variables from the fields of planning and geography. The *inverse distance weighted* (IDW) interpolation scheme, proposed by Shepard, can be used to calculate the value  $f_1(P)$  at any point  $P$  in a randomly sampled set of data points as follows:

$$f_1(P) = \begin{cases} \frac{\sum_{i=1}^N (d_i)^{-u} z_i}{\sum_{i=1}^N (d_i)^{-u}} & \text{if } d_i \neq 0 \text{ for all } D_i \text{ (} u > 0 \text{)} \\ z_i & \text{if } d_i = 0 \text{ for some } D_i \end{cases} \quad (5.5.7)$$

where  $f_1(P)$  is the interpolated value to be calculated at point  $P$ ,  $d_i$  is the distance from  $P$  to the data point  $D_i$ ,  $z_i$  is the value at data point  $D_i$ ,  $u$  is an arbitrary positive real number called the power parameter (a value of 2 is typically used) and  $N$  is the number of points in the randomly sampled data.

This approach can be used to visualise different properties associated with a polyp candidate including it's depth map, density, curvedness, Gaussian curvature, mean curvature and shape index. In each case, the minimum and maximum values for each property are used to normalise the property in the range 0 - 255 for display purposes. The fingerprints associated with the polyp from Figure 5.19 are presented in 5.24.



**Figure 5.24:** The six different fingerprint representations for a polyp candidate (a) depth map, (b) density, (c) curvedness, (d) Gaussian curvature, (e) mean curvature and (f) shape index. In all cases black and white are used to represent the minimum and maximum values for the relevant property.

## 5.5.2 Polyp Candidate Expansion

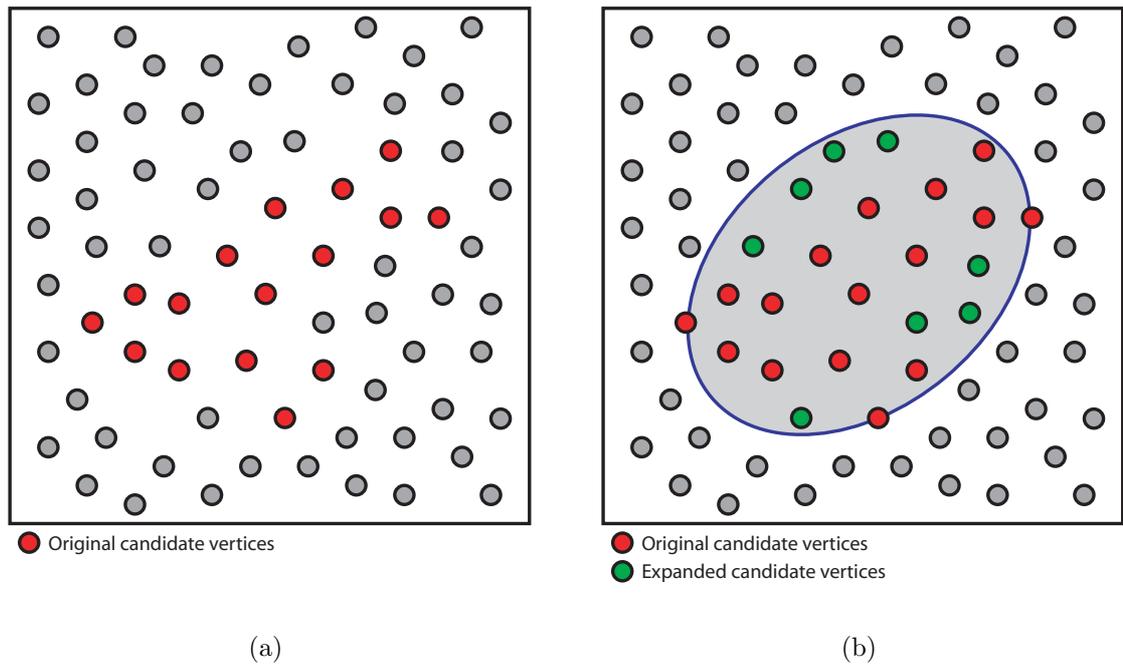
Polyp candidates generally occur as irregularly shaped surface patches that are not necessarily representative of the shape of the actual candidate region. An example of this can be seen in Figure 5.19 (c) where it is evident that the candidate region is a subset of the actual polyp surface. Polyps present as roughly spherical or ellipsoidal projections into the colon lumen. Consequently, the fingerprint representation of a polyp will have a circular or elliptical boundary. A novel expansion technique based on this observation has been developed to increase the surface coverage of a polyp candidate in a structured manner (i.e. so that the structure of the candidate conforms more closely to that of an actual polyp).

The expansion technique involves an analysis of the fingerprint points to find the *minimum bounding ellipse* (MBE) for the subset of points that constitute a polyp candidate. Any additional points enclosed by the MBE, that were not already flagged as polyp candidate points, are included in the candidate region. There are two main consequences associated with expanding the candidate in this way:

- Candidates that represent actual polyps expand to include additional points that belong to the polyp. As a result, the candidate region retains its polylike shape characteristics.
- False positive regions, on the other hand, expand to include additional points that do not belong to a polyp. The inclusion of these points effectively dilutes the overall shape characteristics of the candidate so it becomes less polylike.

Note that this approach is intended to expand the candidate region so that it is more representative of an actual polyp. It is different to approaches that use 3-D sphere or ellipsoid fitting to obtain a volumetric representation of polyp candidates (e.g. Wang et al. (2005)) or to quantify conformance of polyp candidates to spherical/ellipsoidal models (e.g. Kiss et al. (2002) and Chowdhury et al. (2005b)).

Expanding polyp candidates in this manner should ultimately improve the ability to discriminate between true polyps and false positive detections. This observation is verified in Chapter 7, where the results of a study comparing the use of regular and expanded candidates for polyp detection will be presented. An example of the polyp candidate expansion process is illustrated in Figure 5.25.



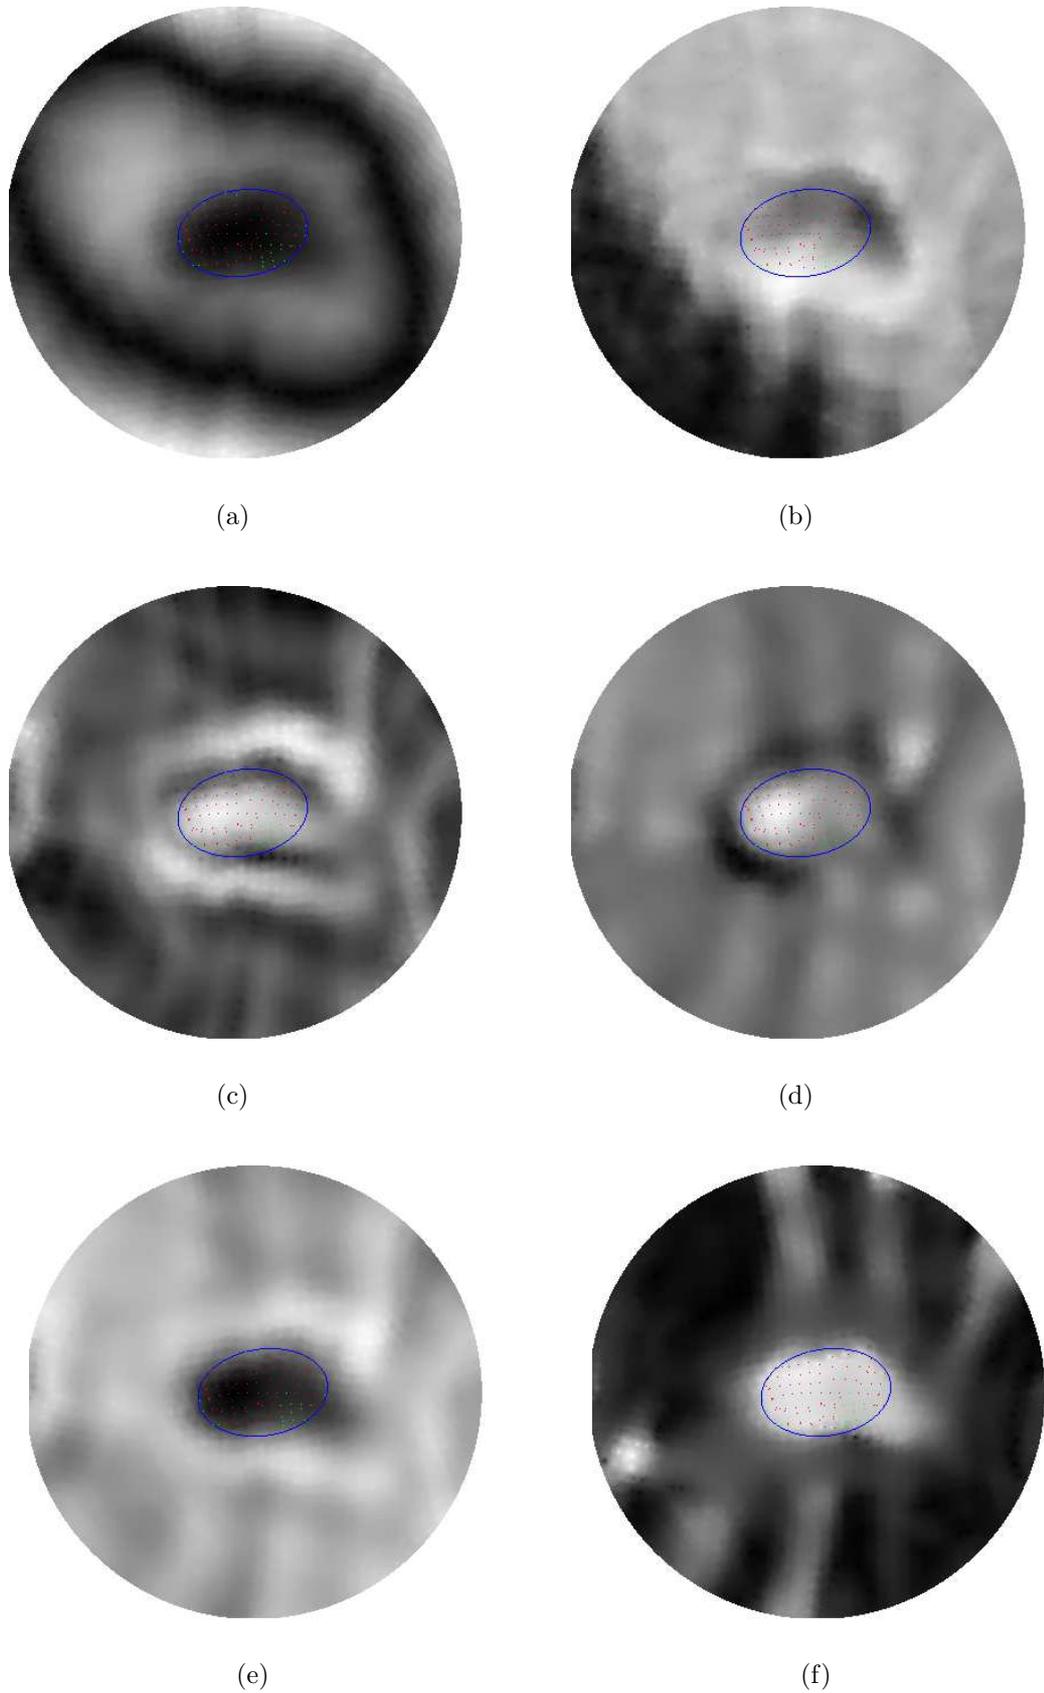
**Figure 5.25:** The candidate expansion process (a) The original candidate regions consisting of a connected set of vertices (b) the minimum bounding ellipse of the original vertices and the set of additional vertices that it encapsulates. The combined set of vertices represents the expanded candidate region.

There are various methods available for calculating the MBE for a set of coplanar points. Post (1994) proposed a brute force approach that involved calculating every ellipse that intersected with all subsets of 3, 4 or 5 points selected from the entire set of points. The ellipse with the smallest area that enclosed all of the original points was identified as the MBE. Welzl (1991) developed a simple randomised MBE algorithm, based on the Linear Programming algorithm developed by Seidel (1990). This approach was more efficient than the technique proposed by Post. An implementation of Welzl’s approach, developed by P. Sakov from the Marine Research division of the Commonwealth Scientific and Industrial Research Organisation (CSIRO)<sup>5</sup>, was ultimately used for MBE calculation. An illustration of the candidate expansion process is presented in Figure 5.26.

### 5.5.3 Global Polyp Candidate Features

Up to now a polyp candidate has been described in relation to its vertex level or local features. A statistical analysis of these local features can be used to yield a set of global features that describe the overall shape or density characteristics of the polyp

<sup>5</sup>Available for download at: <http://www.marine.csiro.au/~sakov/>; Accessed May 26th 2006.



**Figure 5.26:** Expanded versions of the fingerprints from Figure 5.24 (a) depth map, (b) density, (c) curvedness, (d) Gaussian curvature, (e) mean curvature and (f) shape index. The red points are the original polyp candidate points and the green points are the expanded polyp candidate points.

candidate. It is also possible to obtain additional global polyp candidate features during the fingerprinting and candidate expansion processes. Three of these features, expansion, eccentricity and protrusiveness, will be introduced below and their utility in discriminating between true polyps and false positives will be evaluated in Chapter 7.

### 5.5.3.1 Expansion

An actual polyp generally has a roughly spherical or elliptical boundary. As a result the number of extra points added to an actual polyp during the candidate expansion process should be minimal. Conversely, a false positive can have an extremely irregular boundary with numerous bays (see Appendix C). Consequently, in the case of a false positive, the number of extra points added during the candidate expansion process can be quite high.

The expansion is a global polyp candidate feature that is based on this observation. The expansion can be calculated as the ratio of the number of extra candidate points to the original number of candidate points. In the case of the example in Figure 5.25 (i.e. eight additional points and 16 original points) the expansion would be 0.5. The expansion gives an indication of the irregularity of the original polyp candidate.

### 5.5.3.2 Eccentricity

The MBE generated during the candidate expansion process provides an approximation of the shape of the polyp boundary. In the case of an actual polyp, the boundary should be circular or slightly elliptical. Conversely, in the case of a false positive the shape can be more elongated. This is particularly true in the case of false positives caused by the Haustral folds and those caused by the rectal tube (see Appendix C). Li et al. (2005) also proposed the use of candidate elongation as a discriminating feature for the detection of colorectal polyps.

The eccentricity of the MBE can be used as a global polyp candidate feature to distinguish between true polyps and false positive detections based on this observation. The eccentricity of an ellipse  $e$ , defined in Equation 5.5.8, provides a relationship between the semimajor axis  $a$  and the semiminor axis  $b$  of an ellipse where  $0 \leq e < 1$ .

$$e \equiv \sqrt{1 - \frac{b^2}{a^2}} \quad (5.5.8)$$

A perfect circle has an eccentricity of exactly 0, whereas an ellipse has eccentricity greater than 0 but less than 1. The more elongated an ellipse the higher the eccentricity value.

### 5.5.3.3 Protrusiveness

A polyp is defined as a protrusive growth that projects from the colonic mucosa into the colon lumen. A global measure of protrusiveness should help to distinguish between true polyps and flat false positives that have shape characteristics which conform to those of an actual polyp. The protrusiveness (or flatness) of a polyp candidate has previously been used as a discriminating feature for polyp detection (Taylor et al. 2006). The protrusiveness of a polyp candidate can be obtained during the fingerprinting process and is defined as the ratio of the maximum polyp candidate radius to the maximum vertex depth relative to the projection plane. The maximum polyp candidate radius is the Euclidean distance between the centremost vertex and furthest vertex belonging to the candidate. The maximum vertex depth is the distance between the projection plane ( $ax + by + cz + d = 0$ ) and the furthest vertex from the projection plane ( $x_f, y_f, z_f$ ) (see Equation 5.5.9).

$$D = \operatorname{argmax} \left( \frac{ax_f + by_f + cz_f + d}{\sqrt{a^2 + b^2 + c^2}} \right), f \in C \quad (5.5.9)$$

where  $C$  is the set of candidate surface points. Note that the protrusiveness can also be measured during the fingerprinting process by calculating the distance between each candidate vertex and its associated intersection point in the projection plane. A perfectly hemispherical candidate would have a protrusiveness value of 1.0, a candidate caused by flat region with polylike properties would have a protrusiveness less than 1.0 and a non-flat region caused by a sharp fold for example would have a protrusiveness greater than 1.0.

### 5.5.4 Polyp Candidate Feature Summary

A total of nine polyp candidate features have been introduced. Five of these features: shape index, curvedness, sphericity, mean curvature and Gaussian curvature (see Section 5.3.4 for definitions), are based on local measurements that are calculated at each vertex in the mesh representation of the colon surface. A density

measurement is also obtained at each vertex (see Section 5.3.5.2) and brings the number of measurements made using vertex level information to six. For each measurement, with the exception of sphericity, the associated global candidate feature is obtained by calculating the mean of the relevant local measurements for the set of vertices that represent the polyp candidate. Recall that sphericity is already a global candidate feature that is calculated using the average values for the principle (i.e. minimum and maximum), and mean curvatures for the set of vertices representing the candidate region (see Equation 5.3.7). Two versions of each global feature are generated for each candidate: One for the original candidate, and one for the expanded candidate. Three additional features: Expansion, eccentricity and protrusiveness, are also calculated during the candidate expansion process.

### 5.5.5 Polyp Candidate Classification

A polyp candidate is described by a feature vector  $x$  that consists of  $n$  individual measurements  $(x_1, x_2, \dots, x_n)$  and can be represented by a point in the feature space i.e.  $x \in R^n$ . The classification process involves determining the class of an object based on an evaluation of its feature vector. In the case of polyp candidate classification, the feature vector contains the pertinent candidate properties (e.g. a selection of those summarised in Section 5.5.4) and a candidate can be classified as either a true polyp or a false positive (i.e. the classification process involves a binary decision). Note that the feature selection process and the precise composition of the polyp candidate feature vector will be discussed in Section 7.5.4.1.

It is clear from the review in Section 3.5 that a range of different classifiers and classifier configurations can be used for false positive reduction in CAD-CTC systems. It is also clear from this review that the choice of classifier is motivated to some extent by the type of study that is being carried out. The main objective of the study outlined in this thesis, is to compare the use of regular polyp candidates and those that have been expanded using the technique outlined in Section 5.5.2. A linear classifier provides an adequate solution to the two-class classification problem of false positive reduction. It also provides a means of generating a reasonably unbiased comparison between multiple feature sets. Consequently, a *minimum distance classifier* (MDC), which is a type of linear classifier, was selected for use in the comparative evaluation of the aforementioned feature sets and, ultimately, a MDC using the optimum feature set was used for the purpose of false positive reduction in

the documented CAD-CTC system. A brief discussion concerning the training and operation of a MDC will follow.

### 5.5.5.1 Minimum Distance Classifier

A MDC (see Gonzalez & Woods (1992, pg. 580) for description) can be used to perform supervised classification and must be trained prior to use. The training process involves obtaining the mean feature vector  $m_j$  for each class  $\omega_j$  that is supported by the classifier i.e.

$$m_j = \frac{1}{N_j} \sum_{x \in \omega_j} x \quad j = 1, \dots, c \quad (5.5.10)$$

where  $c$  is the number of classes supported by the classifier,  $N_j$  is the number of training vectors that are available for class  $\omega_j$  and only feature vectors  $x$  that belong to objects of class  $\omega_j$  are used in the summation.

Once the MDC has been trained it can be used to assign an unknown object with feature vector  $x$  to one of the classes supported by the classifier. The process involves calculating the Euclidean distance  $D_j(x)$  between  $x$  and  $m_j$  for all  $c$  classes where:

$$D_j(x) = \|x - m_j\| \quad j = 1, \dots, c \quad (5.5.11)$$

or alternatively

$$D_j(x) = \sqrt{\sum_{i=1}^n (x_i - m_{ji})^2} \quad j = 1, \dots, c \quad (5.5.12)$$

where  $n$  represents the size of the feature vector. The unknown object is subsequently assigned to the same class  $\omega_j$  as that of the closest reference vector.

## 5.6 Discussion

This chapter has provided a complete description of the CAD element of the CTC system that is outlined in this thesis. It is clear that the documented CAD system incorporates novel elements and draws on previously published techniques and, as such, represents a hybrid approach. The use of a polygonal mesh provides a memory efficient way of representing the colon surface that also provides subvoxel

accuracy. Modifications to the standard marching cubes algorithm have been introduced. These modifications facilitate the generation of an enhanced representation of the colon surface that includes a range of vertex level measurements as well as guaranteeing airtight surfaces and enabling mesh based region growing.

A novel false positive reduction scheme has also been introduced. This scheme involves the generation of an alternative representation of a polyp candidate prior to polyp candidate classification. The alternative representation expands the polyp candidate to fit an elliptical boundary. The expansion process is intended to make the candidate conform more accurately to an actual polyp. This causes actual polyps to retain their polyplike characteristics, while false positives expand to include more false positive features. The testing and evaluation of all aspects of this CAD scheme and the preprocessing techniques that were introduced in Chapter 4 will be discussed in Chapter 7 following a discussion regarding system implementation in Chapter 6.

# Chapter 6

## System Implementation

The task of implementing the concepts discussed in Chapters 4 and 5 requires some consideration. This is due to a number of issues relating to overall system performance and memory usage. Of particular concern was the significant computational overhead associated with certain operations (e.g. the modified marching cubes algorithm) and the large volume of data being processed by the system as a whole. With these issues in mind, a modular approach to system implementation was adopted. This strategy enabled a phased system implementation, which was beneficial in relation to testing and development, and ultimately streamlined the system integration process.

### 6.1 Implementation Software

All system software was implemented using the Java programming language from Sun Microsystems. This choice was influenced by a number of contributing factors:

- Java is a multi-platform programming language and can be deployed on a number of different platforms without modification. This is especially important in a collaborative research environment.
- Java provides good support for both 2-D and 3-D graphics through well documented, dedicated APIs.
- The object orientated features of Java make it suitable for large scale structured software development.
- The performance of recent versions of Java, particularly in the case of data intensive operations typical of those performed in a CTC system, has been demonstrated to be comparable to that of C++ (Sadleir & Whelan 2005).

## 6.2 Implementation Strategy

Each system operation is implemented as a separate stand alone Java application. Data is shared between related operations by means of a common filesystem. In general, the only parameter that is passed to each application upon execution is the name of the directory in which the common data is stored. In the case of the classifier, multiple directories representing multiple data sets must be specified for training and evaluation purposes. An illustration of the dataflow between the modules of the completed system is presented in Figure 6.1. This approach to system implementation has a number of associated benefits.

### 6.2.1 Staggered Execution

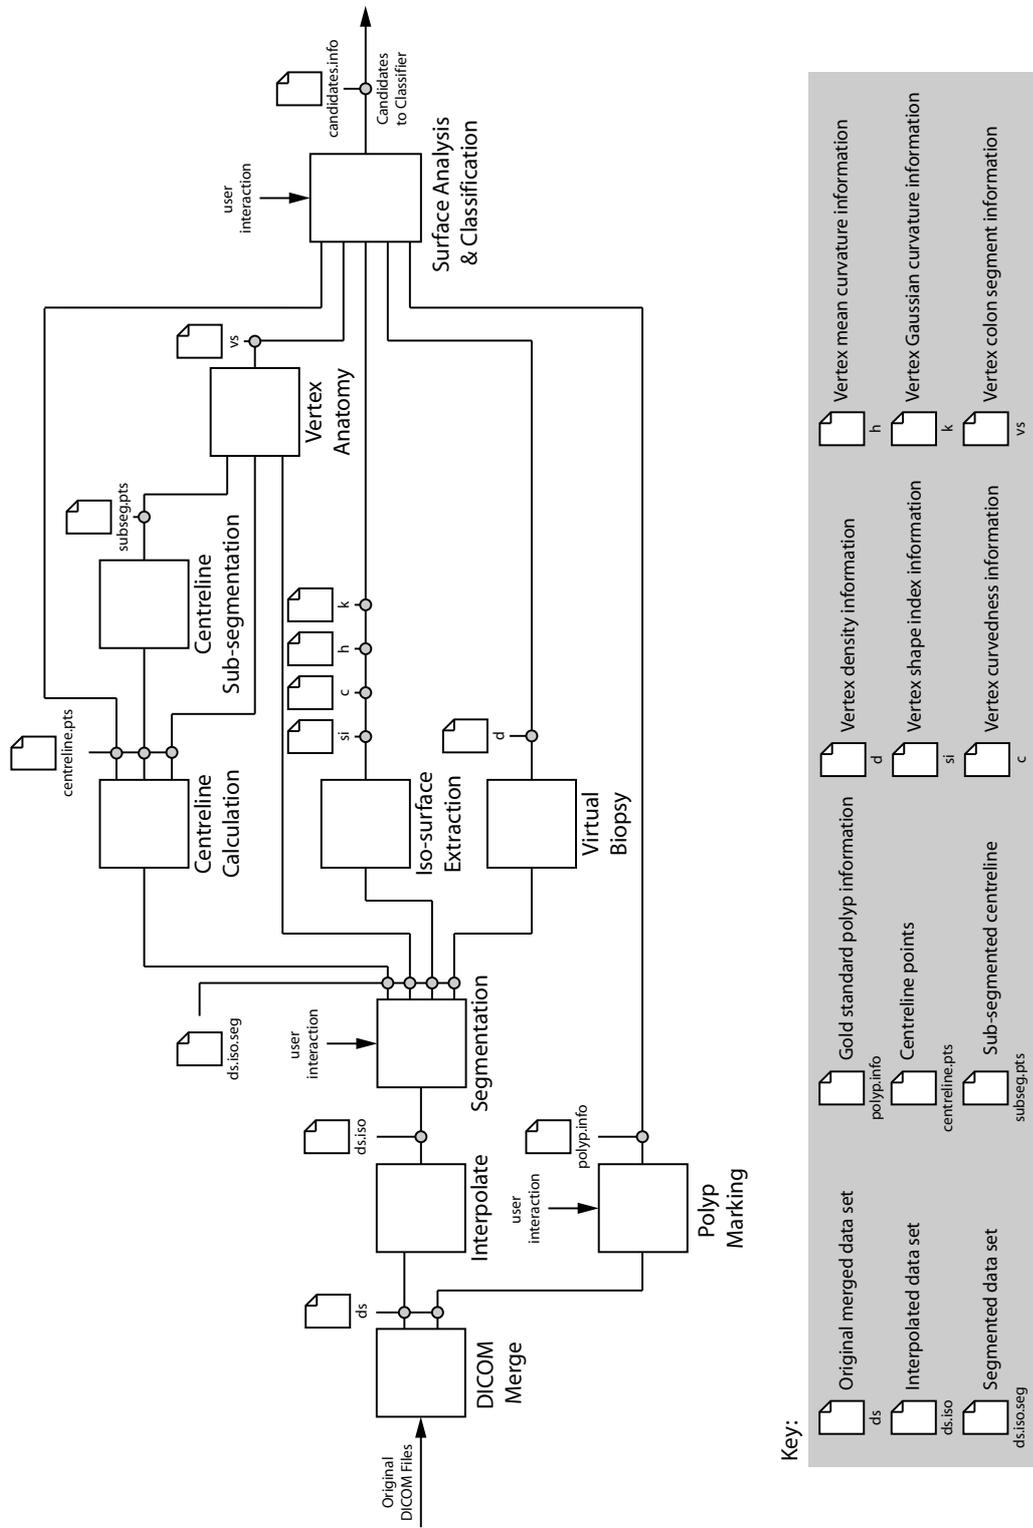
Executing all of the system operations sequentially within a single application would require significant memory resources due to the large amount of intermediate data that is generated. Bearing in mind that a single data set is typically between 100 and 150 MB in size, the available system memory would quickly be exhausted<sup>1</sup>. In addition to the considerable memory requirements, there would be a significant processing overhead if system execution took place in this manner. A modular system implementation with staggered execution is a much more attractive alternative, particularly in the case where a phased approach to system implementation is employed.

### 6.2.2 Well Defined Structure

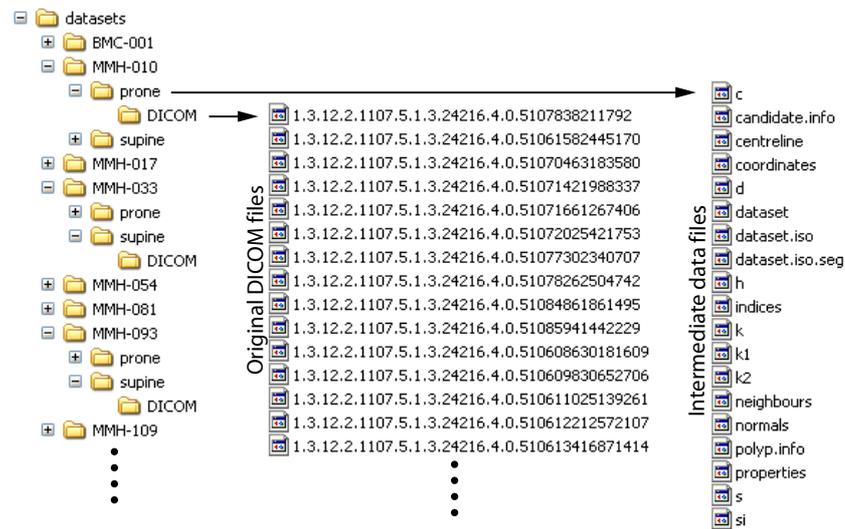
Generally speaking, each of the operations has a well defined set of inputs and outputs (e.g. centreline calculation requires the segmented colon lumen as an input and generates a set of points representing the central path through the colon as an output). Implementing the operations as separate applications helps to formalise the boundaries, requirements and deliverables of the different system operations. This is evident from the system overview illustrated in Figure 6.1.

---

<sup>1</sup>The amount of memory that can be allocated to a particular instance of the Java Virtual Machine (JVM) is limited. This value varies from platform to platform. In the case of the test platform the maximum was found to be 1612 MB. However the use of Java3D reduced this to under 1024MB.



**Figure 6.1:** An illustration of the data flow through the completed system. Each system operation is represented by an individual block and the interconnections represent the data shared between the different operations.

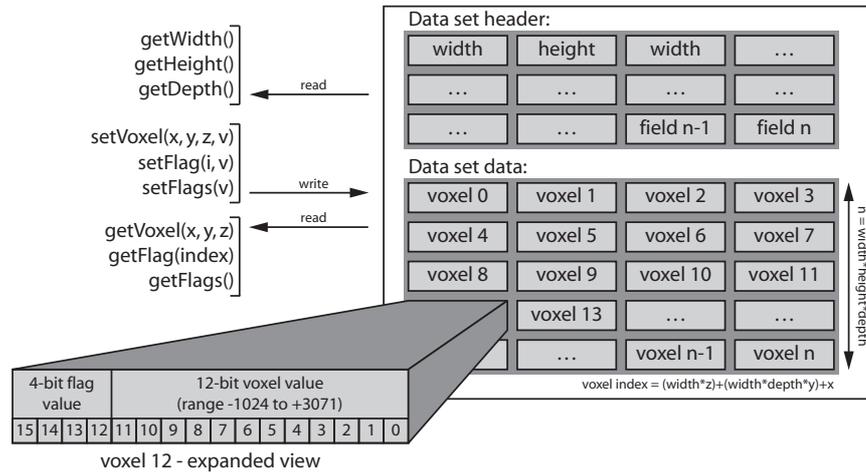


**Figure 6.2:** The underlying file structure of the final system. Each patient has an individual directory under which the prone and supine data sets are stored. Each data set directory contains the relevant intermediate files generated by the system as well as a directory containing the original data set files in DICOM format.

### 6.3 Data Management

A significant amount of intermediate data is generated by the system due to its modular structure. The common information shared between different system modules is stored in data files. These data files are generally related to a particular data set and, as such, a directory structure reflecting this relationship is utilised. The root system directory contains a subdirectory for each patient, which in turn contains two subdirectories for the prone and supine data sets respectively. A data set directory contains all of the intermediate files that are shared between the system modules as well as a further subdirectory containing the original DICOM files. Using this file structure, the data associated with a particular data set can be identified once the patient number and patient position are specified. The expanded file structure for the MMH-010/prone data set is illustrated in Figure 6.3.

The issue of data access is extremely important and has a significant impact on overall system performance. This is due to the large amount of information stored in each of the data files and the number of read and write operations that must be performed by the system as a whole. In order to deal with this issue, a number of custom data files were introduced. In the majority of cases these data files exploited the high speed `ObjectInputStream` and `ObjectOutputStream` classes provided by Java. The different data files that are provided, deal with the storage of data sets,



**Figure 6.3:** An illustration of the custom data set format. The header is extensible and supports the addition of new fields. The data can be assessed using a range of methods. These methods provide direct access to the voxel data and the associated flag values.

data set related properties and miscellaneous data that can be stored using arrays of primitives.

### 6.3.1 Data Set Files

A data set file contains both header and voxel information. The variable length header holds the dimensions of the data set, and can be used to store additional information if required. The voxel information is stored in a 1-D array of `short` primitives in which each voxel value requires only 12-bits of storage space. The remaining bits are used to store four binary flags or a single 4-bit flag value, thus utilising all of the available storage space. The flag values can be used to specify additional information regarding a voxel e.g. in relation to the segmentation process, a flag value of 1 indicates that a particular voxel belongs to the colon lumen. A data set file is used to represent the original data set, the interpolated data set and the segmented data set. Saving and loading data sets takes just seconds.

### 6.3.2 Properties File

Each data set has a properties file where certain information that is required by multiple modules is stored. The properties file is essentially a `Hashtable` structure where predefined property names are paired with the relevant property values. Each time a property value is set the entire property file is saved. The supported properties are listed in Table 6.1.

**Table 6.1:** The properties supported by the properties file. These are essentially the global variables of the system and can be accessed by any of the system components.

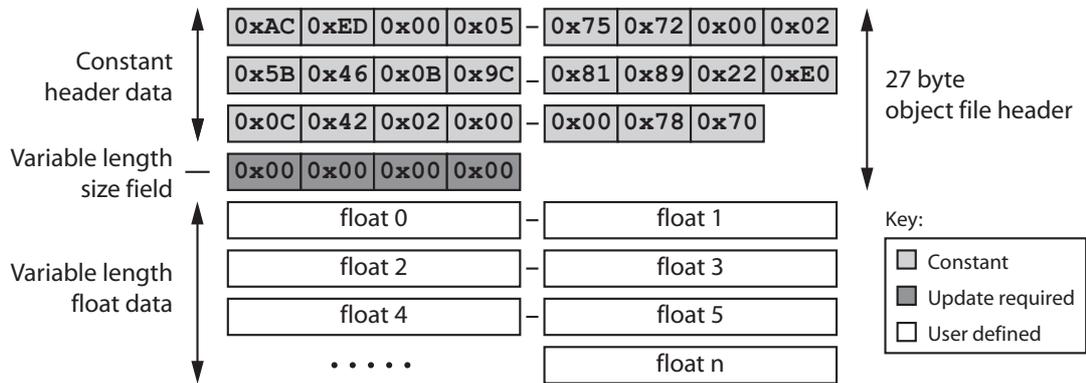
Property Name	Data Type	Description
Patient ID	String	A numerical value (stored as a string to preserve leading zeros) that uniquely identifies a particular patient.
Patient position	String	The position of the patient when scanned. The possible value of this field are prone, supine and unknown.
Original voxel size	Point3f	The voxel size in mm determined from the original DICOM data set.
Interpolated voxel size	Point3f	The voxel size in mm after interpolation. The values for $x$ , $y$ and $z$ should be equal.
Origin after cropping	Point3f	The offset in voxels from the cropped origin to the original origin.
Isosurface density	Float	The isosurface density in HU used for segmentation and isosurface extraction.

### 6.3.3 Miscellaneous Data Files

Other data files are used to store 1-D and 2-D arrays of `float` and `int` primitives as Java objects. These files are generally used in conjunction with the extracted isosurface associated with the colonic mucosa and store vertex coordinates, vertex indices and vertex normals. These files are also used to store the curvature measurements for each vertex, the colon segment index for each vertex and the density information for each vertex. A range of special classes were developed to allow an array object to be written as a series of primitive values. The resulting data file is compatible with, and can be retrieved using, a suitably constructed `ObjectInputStream` object. This significantly reduces the amount of memory required by the modified marching cubes algorithm and also improves the performance of subsequent system stages. The file format for an array of `float` primitives is illustrated in Figure 6.4.

## 6.4 User Interfaces

The majority of system modules, being autonomous, require no user interaction. Hence, user interfaces are mostly used for visualisation and navigation purposes. A limited level of user interaction is supported to enable tasks that must be performed manually i.e. colon lumen segmentation, polyp marking (in 2-D) and candidate



**Figure 6.4:** The format of a file representing an array of `float` primitives that was created using a custom `OutputStream` class. First the constant header is written, then four zero bytes are written for the length field. All of the `float` values are written to the file and upon completion the length field is overwritten with the actual number of entries.

classification (in 3-D)<sup>2</sup>. In order to support these operations, user interfaces are provided for visualisation of data representing either a series of slices or a polygonal surface mesh. Both user interfaces are multipurpose and support several modes of operation.

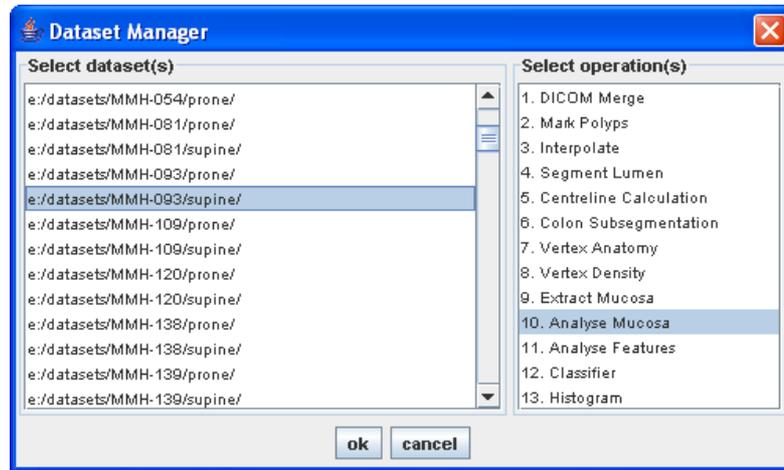
### 6.4.1 The Main System Interface

The main system user interface (see Figure 6.5) allows any of the available operations to be performed on any of the available data sets. A list of the available data sets is obtained by automatically querying the file system. Once an operation and a data set are selected, the main user interface is hidden and the selected operation is initiated. All operations require only a single data set, with the exception of the classifier where multiple data sets must be specified. The main user interface can be easily updated to include new operations, and upon loading the interface recalls the previously selected data set(s) and operation.

### 6.4.2 The Slice Viewer

The slice viewer facilitates the visualisation of 3-D voxel data as a series of 2-D slices. The slices can be orientated in any of the axial, sagittal or coronal directions (with axial being the default). The display slice is selected by means of a slider or by using the navigation buttons. Additional visualisation features, including zoom and simulated scout x-ray generation, can be accessed via the menu system. Colour

<sup>2</sup>Segmentation is the only aspect of the system that requires manual intervention. Polyp marking and candidate classification are only required for testing and evaluation purposes.



**Figure 6.5:** The main user interface for the completed system allowed selected data sets to be processed using selected operations.

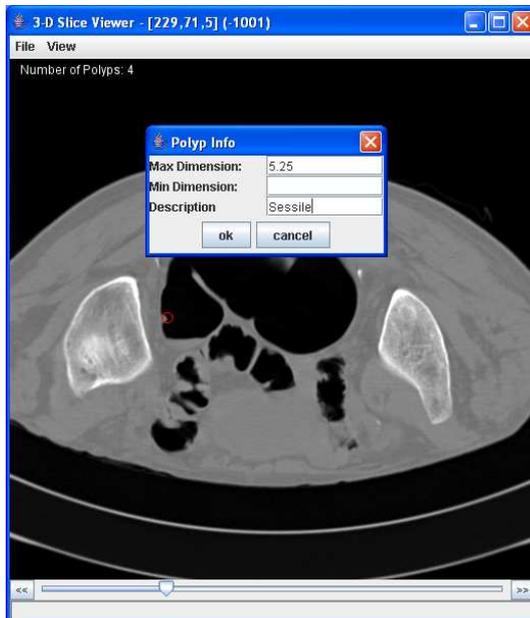
maps can be used to associate a colour with a particular intensity value, flag bit or flag value. The position and density of the current voxel are displayed in the title bar. The slice viewer and its associated visualisation modes are illustrated in Figure 6.6. The three modes that support user interaction will be outlined in the following text.

#### 6.4.2.1 Polyp Marking

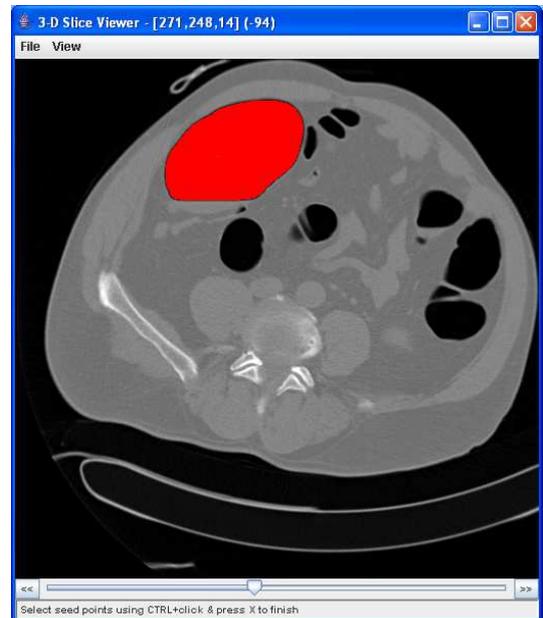
One of the first tasks that must be performed upon receipt of a data set is to mark the actual polyp locations. This is achieved using gold standard information that was determined by a Radiologist based on an unblinded review of the CTC data sets where the Radiologist has access to the conventional colonoscopy findings (see Section 7.5.1). These locations are stored in the file system and are used at a later stage for classifier training and evaluation purposes. Other polyp information including polyp size and type is also recorded at this stage. The polyp marking process is illustrated in Figure 6.6 (a).

The polyp marking process is supported by the slice viewer through a custom mouse listener interface. Two types of event are supported by this interface:

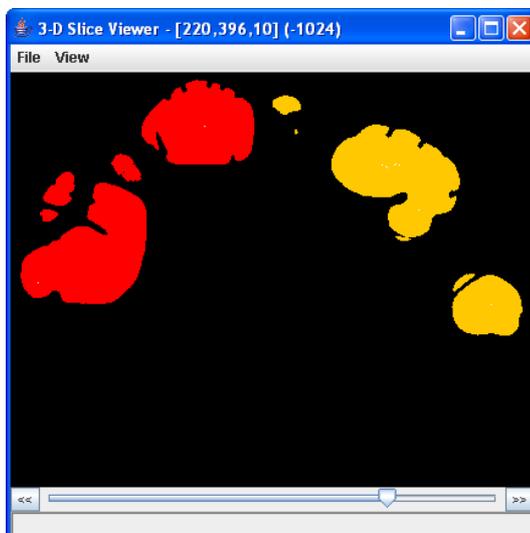
- **Add polyp:** Initiated by a mouse click on the display panel of the slice viewer, this event records the current mouse coordinates  $(x, z)$  and slice number  $y$  as a polyp.
- **Remove polyp:** Initiated by a key press (CTRL key) and a mouse click on the display panel of the slice viewer. This event removes a previously recorded



(a)



(b)



(c)



(d)

**Figure 6.6:** Four possible visualisation modes of the slice viewer. (a) Polyp marking (interactive), (b) Segmentation (interactive), (c) Centreline calculation (display only), and (d) density calculation (display only).

polyp from the list if the click location is within a certain distance of the actual polyp location.

The position of each marked polyp is indicated by a red circle with a radius of 3 mm representing the region occupied by that polyp. If the number of marked polyps is greater than zero, the marked polyp count is displayed in the top left corner of the slice viewer display. The gold standard polyp information identified using the polyp marking interface is saved to the file system in a file called `polyp.info`.

#### 6.4.2.2 Seed Point Selection

The only aspect of polyp detection requiring user input, is the task of seed point selection for colon lumen segmentation (see Section 4.4). Note that a fully automated colon lumen segmentation technique has been developed by the VSG (Chowdhury et al. 2005a). This technique will ultimately be incorporated into the CAD-CTC system described in this thesis, thus fully automating the polyp detection process.

In order to initiate the segmentation process, a seed point is selected using the slice viewer interface (see Figure 6.6 (b)). This operation is facilitated by a custom mouse listener interface. Only one event is supported by this interface:

- **Seed point select:** Initiated by a mouse click on the display panel of the slice viewer. This event initiates region growing from the current mouse coordinates  $(x, z)$  and slice number  $y$ .

Multiple seed points can be selected in cases where the colon lumen is presented as multiple segments due to collapse or blockage. The simulated scout x-ray can be used to monitor the segmentation process. Upon completion, when the slice viewer window is closed, the segmented version of the data set is saved to the file system as `dataset.iso.seg`.

#### 6.4.2.3 Viewpoint Selection

The slice viewer can be used to update the current viewpoint for the 3-D surface rendered model of the colon that is displayed in an associated surface viewer interface (see Section 6.4.3). This operation is facilitated by a custom viewpoint update listener interface. One event is supported by this interface:

- **Viewpoint update:** Occurs following a mouse drag event. The coordinates associated with the mouse release and mouse press represent the start and end

points of the view vector respectively and the viewpoint is represented by the mouse release coordinates.

The viewpoint selection tool essentially bridges the gap between the slice and surface representations of the colon lumen and greatly simplifies the viewpoint selection process.

### 6.4.3 The Surface Viewer

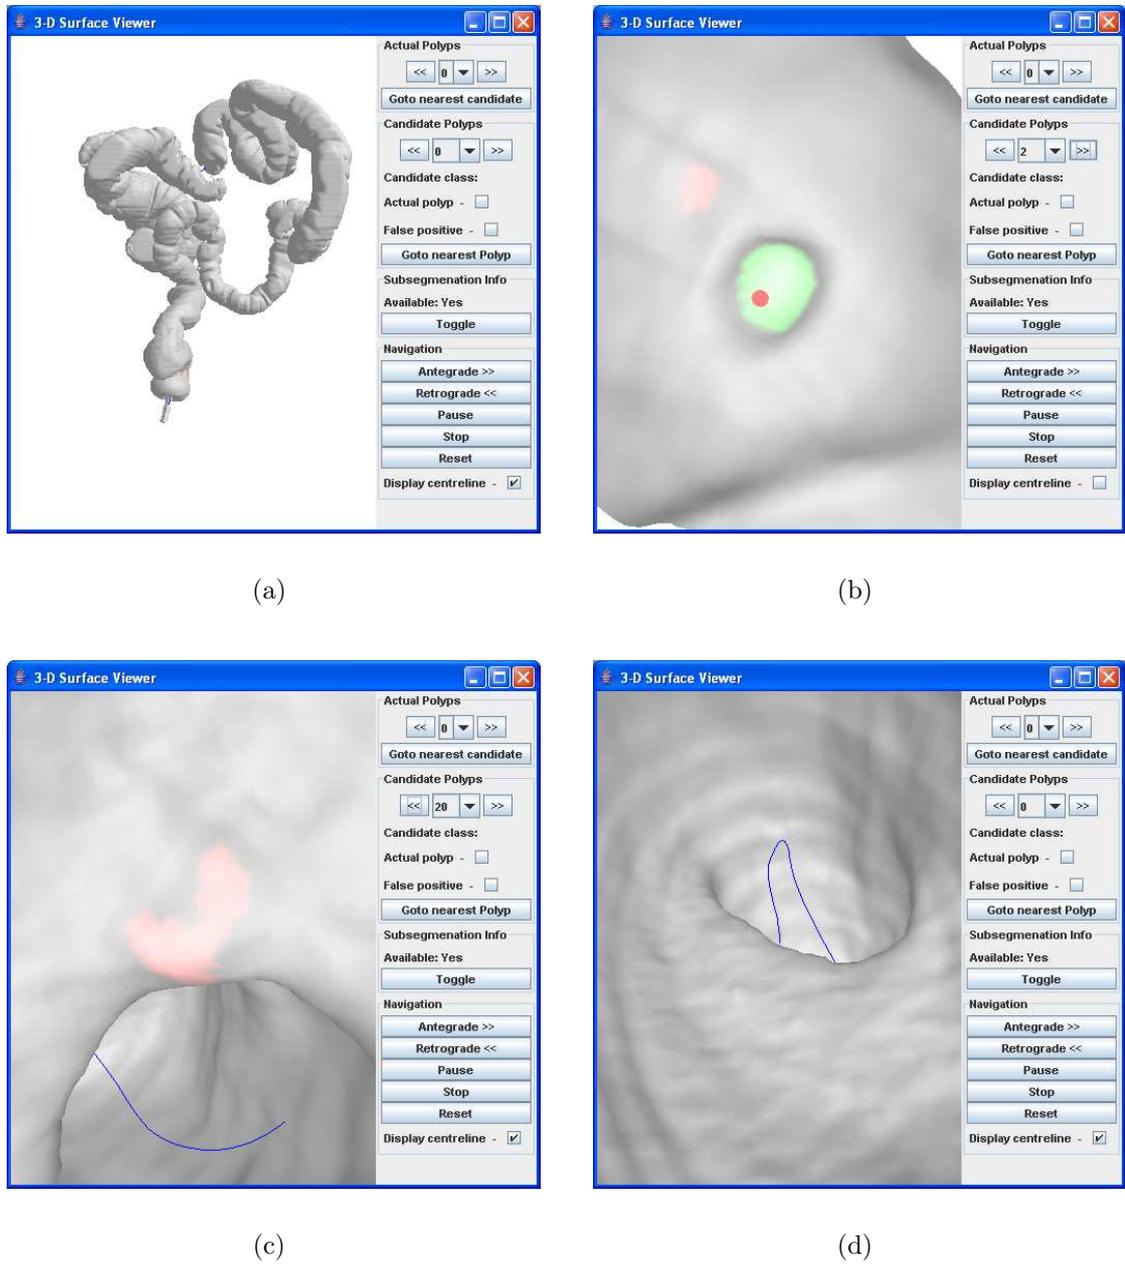
The surface viewer is a custom Java3D application that is used to visualise the 3-D triangular mesh representation of the colon generated by the modified MCA that was described in Section 5.3. The surface viewer has several mandatory inputs including vertex coordinates, normals and vertex indices. All of these inputs must be specified for the surface to be displayed. Additional optional inputs may also be specified to enable additional functionality and visualisation modes. These optional inputs include the centreline points and the vertex subsegmentation information. The initial view provided by the surface viewer is presented in Figure 6.7 (a). There are three other main modes of operation provided by the surface viewer: polyp view, candidate view and fly through.

#### 6.4.3.1 Polyp View

Each actual polyp can be viewed using backward and forward navigation buttons or by selecting the relevant polyp index from a dropdown menu. When a polyp is selected, the viewer is positioned directly above the requested polyp (in the direction of the central normal for the polyp) and moved back a fixed distance (20 mm) so that the entire polyp can be visualised. A button is provided so that the viewer can jump to the nearest candidate to enable manual candidate classification (required for classifier training). The polyp view mode is illustrated in Figure 6.7 (b). This polyp visualisation technique is based on the method described by Vining et al. (1998).

#### 6.4.3.2 Candidate View

A similar set of controls is provided to view each of the polyp candidates identified by the initial polyp candidate identification stage (see Section 5.4). Extra functionality is provided so that each candidate can be assigned to a class. This manual classification can be used in the training and automatic evaluation of the classifier



**Figure 6.7:** The different modes of operation of the surface viewer. (a) The default external view of the segmented colon lumen. (b) The polyp view where each actual polyp can be viewed. (c) The candidate view where each candidate can be visualised and classified. (d) The fly through mode where the user is automatically guided through the surface model of the colon using the set of points generated during the centreline calculation stage.

stage of the system. When a candidate is assigned to a class, its colour is changed to reflect the class to which it belongs. A save option is provided so that the classified polyp information can be saved for use in conjunction with the classifier. A button is also provided so that the viewer can jump to the nearest actual polyp. The candidate view mode is illustrated in Figure 6.7 (c).

#### 6.4.3.3 Fly Through

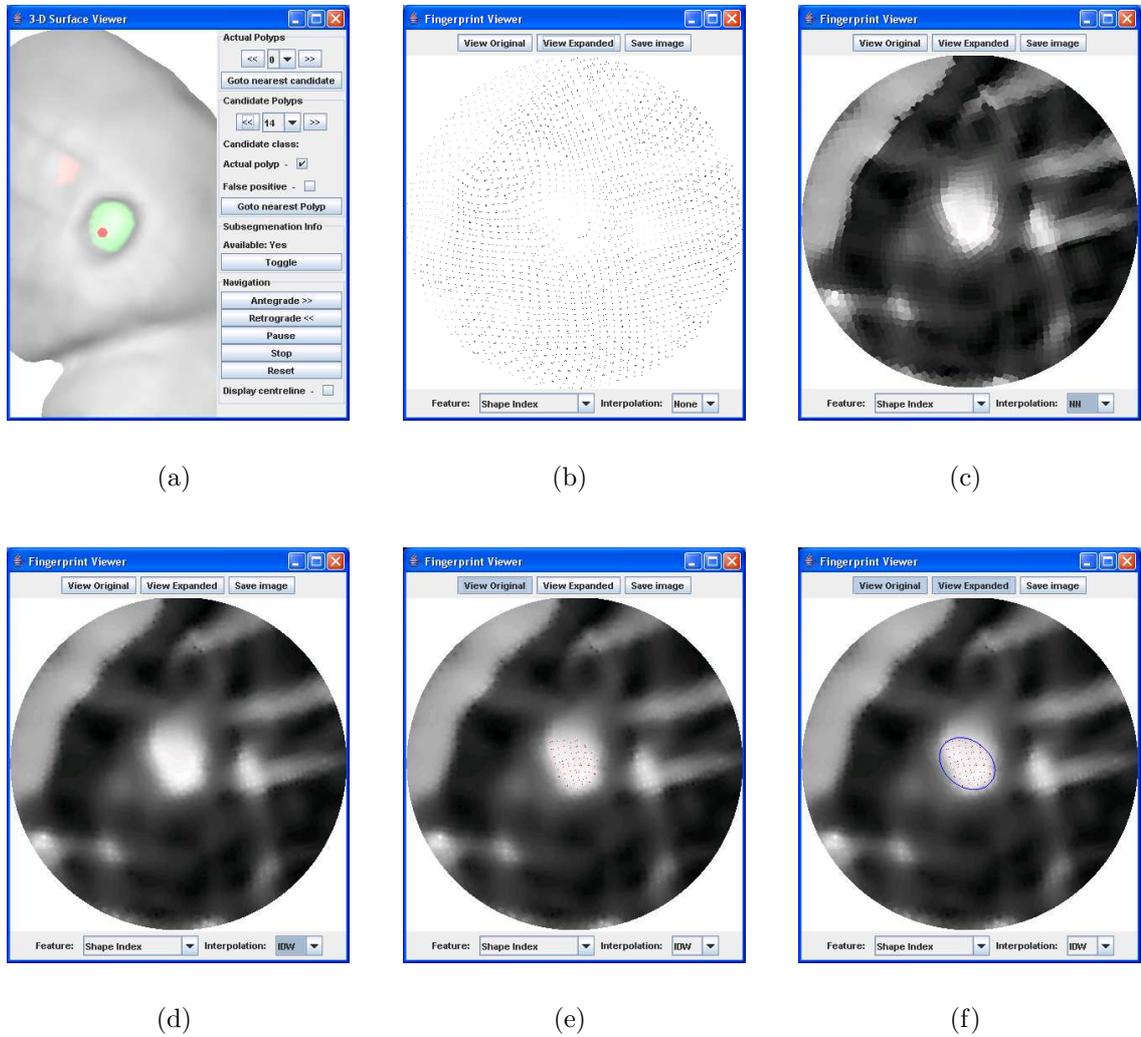
The set of points generated by the centreline calculation stage can be used to automatically guide the user through the data set between endpoints located in the rectum and the caecum. The centreline points must be lowpass filtered before they can be used. This ensures that the transition from one point to the next is smooth. Dedicated navigation controls can be used to initiate fly through, indicate direction, pause and ultimately stop the navigation process. The centreline path is displayed by default but can be turned on and off depending on user preference. The fly through mode is illustrated in Figure 6.7 (d).

#### 6.4.4 The Fingerprint Viewer

A dedicated user interface is provided for visualising the fingerprint representation that is associated with a polyp candidate. The fingerprint view automatically updates as the user sequences through the list of polyp candidates. The fingerprint view has a number of different display modes. The interpolation can be set to none, nearest neighbour or inverse distance weighted. In addition, the original and/or expanded candidate points can be displayed in conjunction with the MBE. The different modes of operation of the fingerprint viewer are illustrated in Figure 6.8.

### 6.5 Discussion

The task of system implementation is complicated by the large volumes of data that must be dealt with. As a result, an implementation strategy that enabled relatively fast data processing while requiring a minimum amount of system memory was employed. The selected approach facilitates the realisation of the system in a phased manner, and resulted in very well defined boundaries for the majority of the component system tasks. Ultimately, only a small number of user interfaces were required due to the minimal amount of user interaction that is required for system



**Figure 6.8:** The different modes of operation of the fingerprint viewer. (a) a sample polyp candidate (in this case an actual polyp) displayed using the 3-D surface viewer. (b) The uninterpolated fingerprint points associated with the candidate. (c) The fingerprint points interpolated with nearest neighbour interpolation. (d) The fingerprint points interpolated using inverse distance weighting based interpolation. (e) The superimposition of the original candidate points. (f) The superimposition of the minimum bounding ellipse and the expanded candidate points.

operation. In most cases, the user interfaces were provided for visualisation purposes and for the generation of data relating to the testing process.

# Chapter 7

## Testing and Results

A series of novel CTC techniques are introduced in this thesis. These techniques deal with different aspects of CTC including: Centreline calculation, colon lumen subsegmentation and automated polyp detection. This Chapter outlines the testing strategy that was used to evaluate these techniques, and presents the results of the evaluation. The test data included clinical CTC data sets obtained from a number of sources as well as a range of computer generated synthetic colon phantoms.

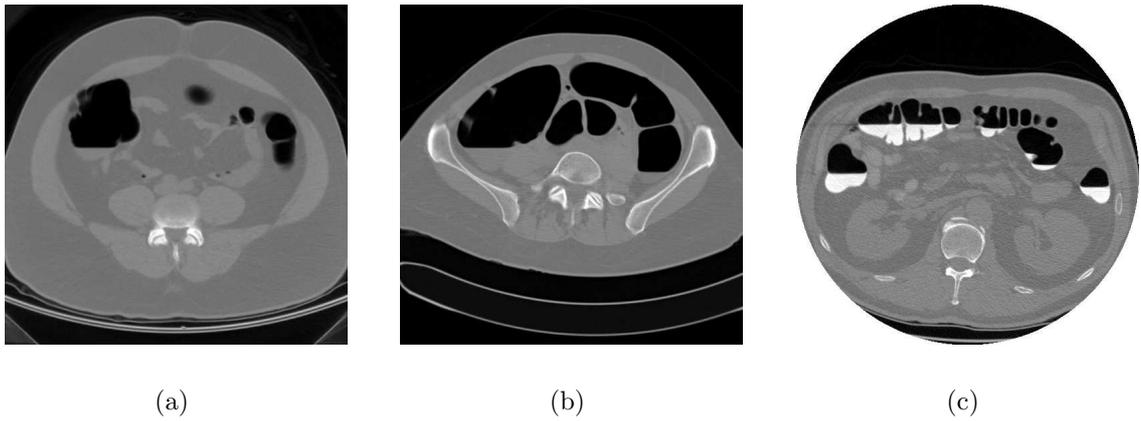
### 7.1 Test Data

A range of CTC data was used in the development, training and evaluation of the different techniques outlined in this thesis. This data was also used to illustrate numerous concepts, issues and examples throughout the thesis. The data was obtained from three sources:

- Boston University School of Medicine (BUSM)
  - Location: Boston, MA, United States.
  - Number of Patients: 1
  
- Mater Misericordiae Hospital (MMH)
  - Location: Dublin, Ireland.
  - Number of patients: 37
  
- Walter Reed Army Medical Centre(WRAMC)<sup>1</sup>

---

<sup>1</sup>This data is provided courtesy of Dr. Richard Choi, Virtual Colonoscopy Centre, Walter Reed Army Medical Centre. It is available online at: [http://nova.nlm.nih.gov/wramc/data\\_index.html](http://nova.nlm.nih.gov/wramc/data_index.html); Accessed May 26th 2006.



**Figure 7.1:** A series of sample images selected from data sets obtained from the BUSM (a) (BUSM-001/supine), the MMH (b) (MMH-640/supine) and from the WRAMC (c) (WRAMC-239/supine). There is usually a small amount of residue in the data sets obtained from the BUSM and the MMH. The samples in (a) & (b) were selected specifically to illustrate this point. This residue is approximately the same density as soft tissue. In the case of the WRAMC data, there is a significant amount of contrast enhanced residue. This is evident from the sample illustrated in (c).

- Location: Washington, DC, United States.
- Number of Patients: 52

The data provided by these three sources will be reviewed below with particular attention given to the MMH data, which represented the main data used in this study. Sample images from each of the sources are illustrated in Figure 7.1.

### 7.1.1 BUSM Data

The single patient data set received from the Department of Radiology, BUSM, was used in the early stages of the research that is outlined in this thesis, and was the only CTC data that was available while the MMH were initiating their CTC programme. This data set was previously used by Fenlon & Ferrucci (1997) in their publication outlining the CTC technique and its associated issues.

The BUSM data was used in the development of the data interpretation software that is discussed in Appendix A. It was also used to determine the implementation strategy outlined in Section 4.4, and in the early stages of the development of the centreline calculation technique that is discussed in Section 4.5. The main difference between this data and the data that was ultimately provided by the MMH was the slice spacing. The slice spacing in the BUSM data is 2.0 mm whereas the slice

spacing in the MMH data was 1.5 mm. This was not a significant problem as all data was ultimately interpolated to generate isometric voxels.

### 7.1.2 MMH Data

The data provided by the Department of Radiology and the Gastrointestinal Unit at the MMH represented the main data used in this research, specifically in relation to the development and evaluation of the CAD-CTC technique that is outlined in Chapter 5. The data provided by the MMH was generated for their comparative evaluation of conventional colonoscopy and CTC. A detailed summary of this data is provided in Appendix B. Thirty seven patient data sets were received from the MMH on a phased basis resulting in a total of 68 usable data sets (from a possible 74). A total of 155 colorectal neoplasia are visible in these data sets; two (1.3%) are flat polyps, 98 (63.2%) are polyps < 5 mm in size, 44 (28.4%) are polyps 5 mm or greater in size and the remaining 11 (7.1%) represent colorectal masses.

#### 7.1.2.1 Patient Preparation

Prior to their scheduled examination, patients were instructed to take a low-residue diet for 48 hours followed by clear fluids for 24 hours. The day before their examination, patients were instructed to take one sachet of Picolax at 08.00 hrs, a second sachet of Picolax at 12.00 hrs, a sachet of clean prep in a litre of cold water at 18.00 hrs and a Senokot tablet at 23.00. Immediately before the examination the colon was gently insufflated with room air to a maximum tolerable level and the patient was subsequently scanned in the prone and supine positions to reduce the effect of residual material in the colon.

#### 7.1.2.2 Data Acquisition

All scans were obtained using a Siemens Somatom 4, four slice multidetector spiral CT scanner. The scanning parameters were 120 kVp, 2.5 mm collimation, 3 mm slice thickness, 1.5 mm reconstruction interval, 0.6 s gantry rotation and field of view of 380°. The scanning time ranged from 20 s to 30 s, so that the acquisition is performed in a single breath-hold. The procedure was first performed in the prone position and then repeated in the supine position to reduce the effect of residual material in the colon. Each slice in the resulting data set contained  $512 \times 512$  pixels and the number of slices varied from 153 to 316 (see Appendix B) depending on

the height of the patient and the extent of the colon. The typical total size of the volumetric data is approximately 150 Mb. Patients were sent for a conventional colonoscopy examination immediately after CTC.

### 7.1.3 WRAMC Data

The WRAMC data consists of prone and supine data sets for 52 patients. Detailed virtual and optical colonoscopy reports are provided for each patient. In some cases optical colonoscopy videos and pathology reports are also provided. The 52 patients had a total of 70 colonoscopy confirmed polyps. Forty nine (70%) of these polyps were reported to be visible in the prone, in the supine or in both data sets. The main difference between the WRAMC data sets and the data sets received from the BUSM and the MMH is the type of bowel preparation that was used. The CTC protocol utilised by the WRAMC incorporated a contrast enhanced bowel preparation (this is evident in Figure 7.1 (c)). Consequently, a colon lumen segmentation technique incorporating digital subtraction (see Section 3.2.3) is required to identify the full extent of the colon.

The colon lumen segmentation technique outlined in Section 4.4 was developed for use with patient data sets where a relatively dry bowel preparation was used. This was typically the case with the BUSM and MMH data sets. A shortcoming of the segmentation technique outlined in this thesis, is that it is unable to deal with wet contrast enhanced data sets e.g. those provided by the WRAMC. In the case of the WRAMC data sets, the segmented version of the colon represents only voxels attributable to endoluminal air. The contrast enhanced regions are not included, the boundary of the contrast enhanced regions and the colon wall are not reconstructed, and the boundary artifacts between the endoluminal air and the contrast enhanced material are not removed. Consequently, limited use was made of the WRAMC data, as it was not possible to obtain a complete segmentation.

## 7.2 Test Platform

All tests were carried out using a Dell Precision 370 Workstation with a single Intel Pentium 4 CPU running at 3.2 GHz with 2 GB of RAM. The operating system was Microsoft Windows XP Professional Edition. The following additional software was installed:

- Java 2 Runtime Environment, Standard Edition 5.0 (build 1.5.0\_04-b05)
- Java 3D for Windows (OpenGL Version) Runtime for the JRE (build 1.4.0-build5-experimental)

The choice of Java for system implementation was motivated by the factors listed in Section 6.1. In addition, the uncertainty regarding the performance of Java, particularly in the case of computationally intensive operations typically found in a CAD-CTC system, was resolved in the study by Sadleir & Whelan (2005). This study compared the use of interpreted Java and natively executed C++, for the implementation of the centreline calculation technique outlined in Section 4.5, and found the two alternatives to be comparable in terms of performance.

## 7.3 Centreline Evaluation

The centreline calculation technique outlined in Section 4.5 was evaluated in relation to the accuracy of the generated centreline and the overall performance of the centreline calculation process.

### 7.3.1 Test Data

The test data that was used in the evaluation of the centreline calculation algorithm consisted of 20 real patient data sets and four computer generated synthetic phantoms that were developed specifically for centreline evaluation.

#### 7.3.1.1 Clinical Cases

The clinical data sets were selected to include a mixture of prone and supine studies with varying levels of collapse, blockage and residual material. This diverse range of test data ensured that a realistic indication of overall algorithm operation would be obtained. The clinical cases that were used are listed in Table 7.1.

#### 7.3.1.2 Phantom Data Set

Computer generated colon phantoms with known centrelines have been used for quantifying the accuracy of previously published centreline calculation techniques (Samara et al. 1999, Iordanescu & Summers 2003, Frimmel et al. 2004, Bouix et al.

**Table 7.1:** The data sets used for centreline evaluation. The average number of segments in this test set was 2.25 and the average number of lumen voxels was 5,521,687.

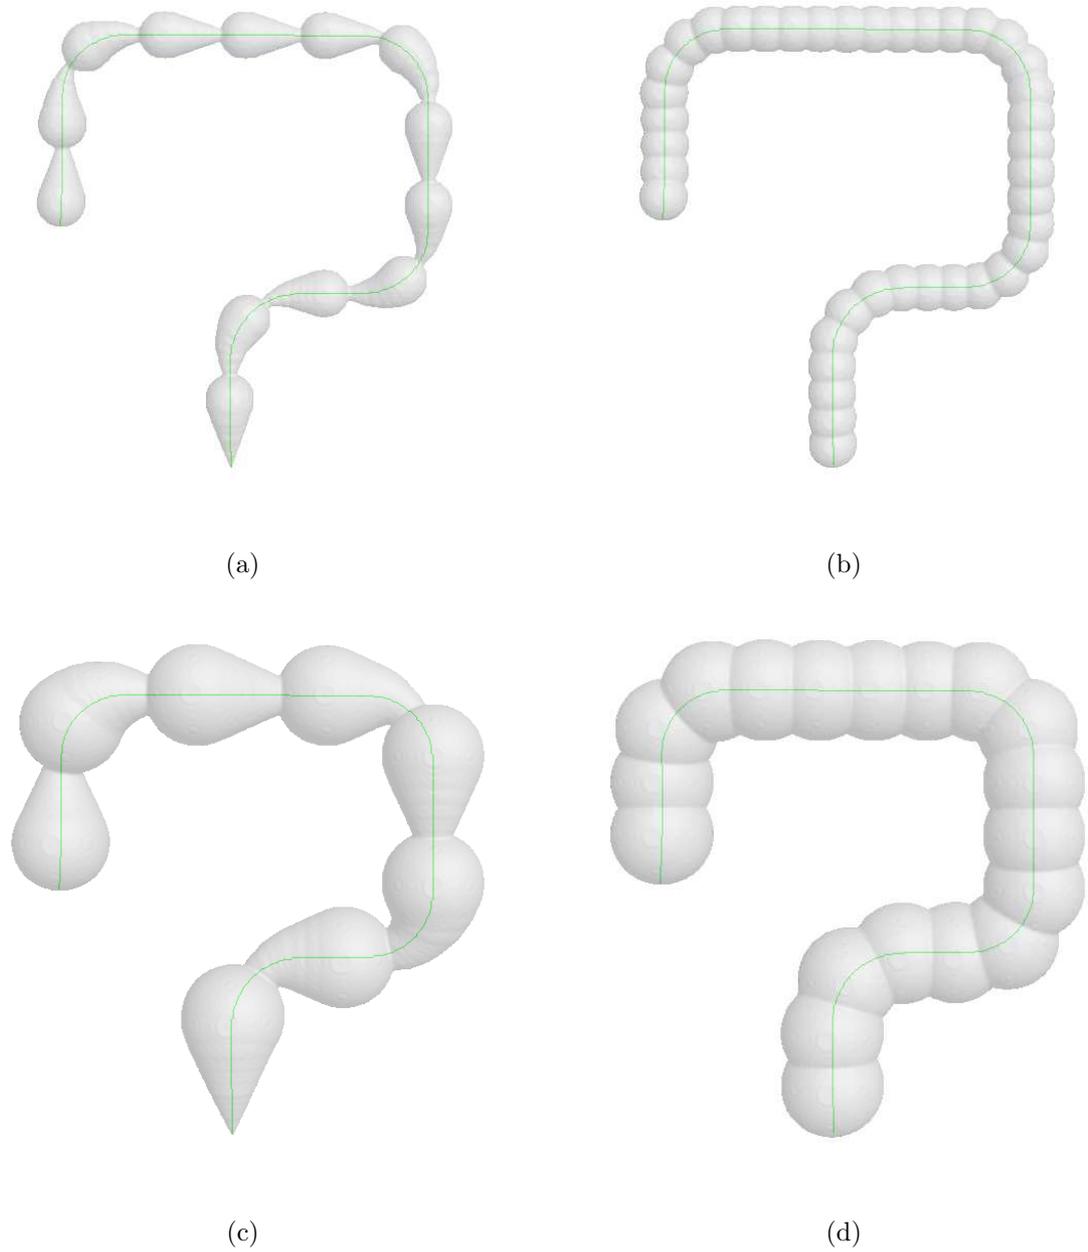
Patient ID	# Segs	# Voxels	Patient ID	# Segs	# Voxels
MMH-010S	3	3,667,269	MMH-234S	2	7,199,653
MMH-033S	4	2,258,287	MMH-291P	2	4,643,454
MMH-093P	1	3,584,789	MMH-316S	3	2,345,333
MMH-057S	2	7,009,275	MMH-368S	1	4,871,789
MMH-915S	3	1,662,761	MMH-451P	5	7,354,106
MMH-081S	1	3,479,881	MMH-526P	2	5,478,727
MMH-109S	3	2,991,453	MMH-640P	1	8,150,997
MMH-138P	2	3,905,869	MMH-672P	2	4,878,404
MMH-139S	2	4,770,474	MMH-765S	2	11,153,259
MMH-166S	2	6,404,688	MMH-909S	2	5,478,789

2005). Similarly, a phantom study was carried out to evaluate the centreline calculation technique described in this thesis. A series of phantom data sets were created specifically to establish the accuracy of the centreline calculation technique. Each phantom was generated by initially specifying a curved path in 3-D space. A sphere was generated at each point along the path. The radius of successive spheres was cycled between two extrema in order to simulate the undulating effects caused by the Haustral folds in a real patient data set. The 3-D curve, minimum and maximum sphere radii and frequency of the transitions between the minimum and maximum radii were all varied in order to generate different phantoms. All of the phantom data sets were lowpass filtered using a  $3 \times 3 \times 3$  filter in order to simulate the partial volume effect at the boundary between air and soft tissue. A total of four distinct phantoms data sets (P1 - P4) were generated for evaluation purposes, these are illustrated in Figure 7.2.

### 7.3.2 Accuracy Evaluation

The accuracy evaluation consisted of three distinct tests that involved the use of clinical and phantom data sets to:

1. Compare the location of the automatically detected endpoints with the actual end points (phantom data).
2. Establish, on a point by point basis, the correspondence between the automatically generated centreline and the actual centreline (phantom data).
3. Verify that the calculated centreline agreed with the expected centreline in



**Figure 7.2:** An illustration of the four computer generated synthetic phantom data sets that were used for centerline evaluation.

**Table 7.2:** The endpoint detection error distances for each of the four phantom studies. The mean endpoint detection error was approximately 2 mm.

	<b>P1</b>	<b>P2</b>	<b>P3</b>	<b>P4</b>	<b>Mean</b>
<b>Rectum Error (mm)</b>	0.0	1.5	0.5	3.0	1.25
<b>Caecum Error (mm)</b>	2.0	2.0	3.5	4.0	2.875
<b>Mean (mm)</b>	1.0	1.75	2.0	3.5	<b>2.063</b>

all cases, particularly in cases where the colon lumen consisted of multiple segments (clinical data).

Euclidean distance measurements were used extensively in the first two tests.

### 7.3.2.1 Endpoint Detection Accuracy

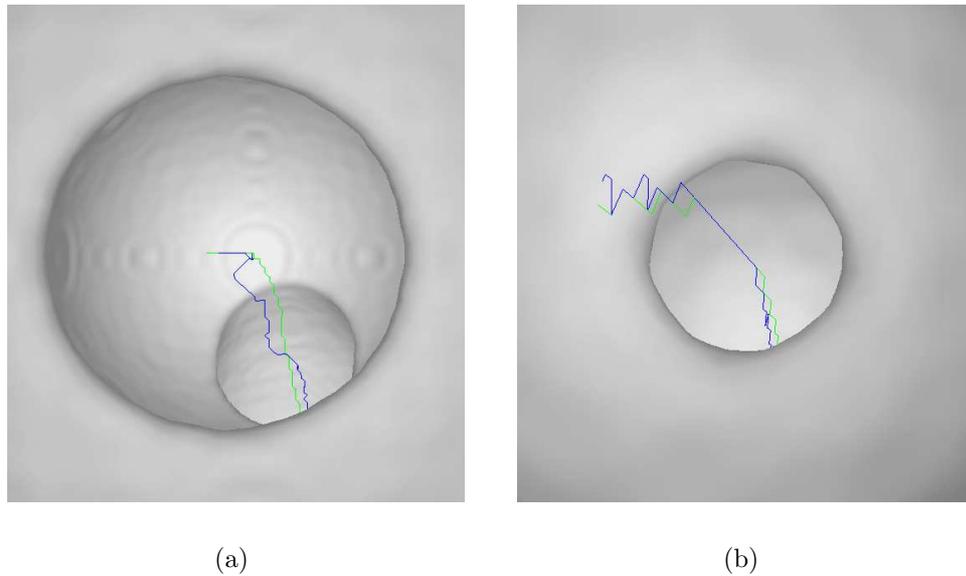
There is no precise definition for the location of the two endpoints of the colon centreline. In general these points are taken as being the lowest points in the rectum and the caecum. This was the definition that was used in the development of the centreline calculation algorithm. In many cases the colon lumen consists of multiple segments due to collapse or blockage, and, in these cases, the locations of the intermediate endpoints are even less obvious. The process of determining the accuracy of the automatically detected centreline endpoint locations involved quantifying the distance between the calculated endpoint and the known endpoints of the phantom centreline. The results of this evaluation for the four phantom studies are presented in Table 7.2. The mean error for the rectum point was found to be 1.25 mm and the mean error for the caecum point was found to be 2.875 mm, giving an overall mean endpoint detection error of approximately 2 mm.

### 7.3.2.2 Overall Centreline Accuracy

The evaluation of the overall accuracy of the automatically generated centreline was effectively an extension of the approach used to determine the accuracy of the automatically detected endpoints. The process involved measuring the distance between each voxel in the calculated centreline and its nearest neighbour in the original phantom centreline (this is similar to the mean displacement metric proposed by Frimmel et al. (2004)). If the distance was zero then a perfect match was recorded, otherwise, the measured distance value was recorded. Two centreline properties were generated during this evaluation: The standard deviation of the calculated centreline from the known centreline, and the percentage correspondence between the two centrelines.

**Table 7.3:** The results of the overall accuracy evaluations for the four phantom studies. The average standard deviation between the actual and calculated centrelines was less than 1 mm.

	P1	P2	P3	P4	Mean
<b>Correspondence (%)</b>	39	42	24	28	<b>33.25</b>
<b>Standard Deviation (mm)</b>	0.463	0.425	0.952	1.105	<b>0.736</b>
<b>Maximum Deviation (mm)</b>	2.693	2.179	3.606	5.679	<b>3.539</b>

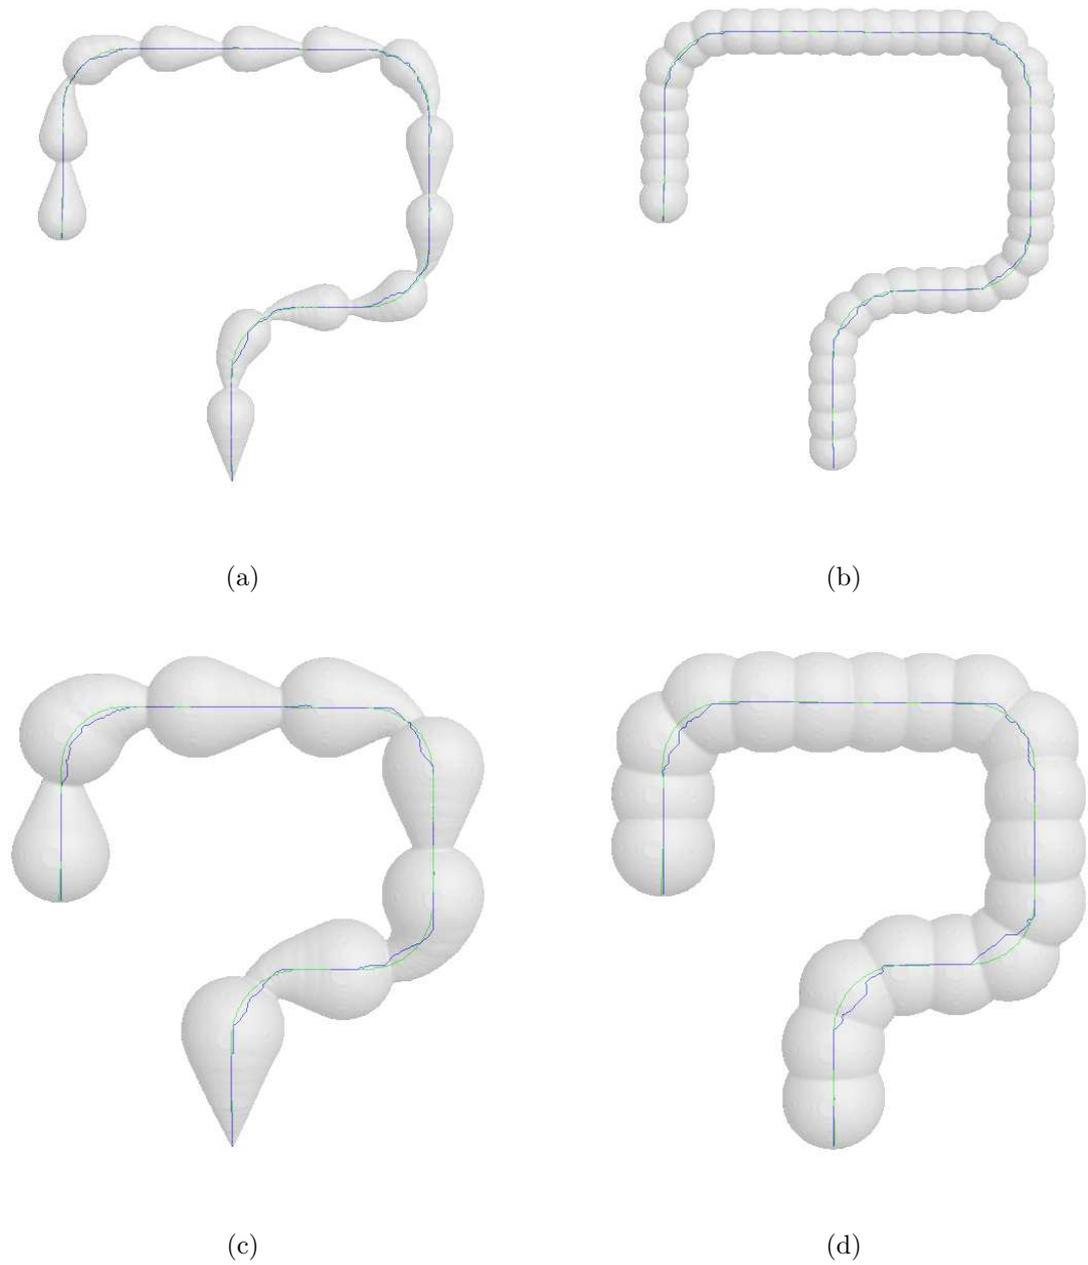


**Figure 7.3:** A comparison of the actual and calculated centrelines from an endoluminal perspective. (a) The two centrelines closely track each other but never converge. (b) The calculated centreline converges with the actual centreline. Note that in both cases the standard deviation between the actual and calculated centrelines is less than 1 mm.

The results of the accuracy evaluation are presented in Table 7.3. On average, the calculated centreline voxels perfectly matched the actual centreline voxels in approximately 33% of instances, the standard deviation between the two centrelines was less than 1 mm and the maximum deviation was 3.539 mm. This level of accuracy is comparable with levels reported by other groups (see Table 3.1). Sample renderings of the actual and calculated centrelines are illustrated in Figure 7.3 and Figure 7.4.

### 7.3.3 Performance Analysis

The performance analysis involved calculating the centrelines for the clinical cases discussed earlier. This was intended to provide an indication of overall system performance. An additional study was performed to quantify the effects the optimisations outlined in Section 4.5.4 had on system performance. Both techniques produced



**Figure 7.4:** A comparison of the actual and calculated centreline, for each of the phantoms, from an external perspective.

**Table 7.4:** Detailed results comparing the standard and optimised approaches to centreline calculation. Note: The optimisation technique only affected the skeleton generation stage of the process. The time required for all other stages remained the same as the unoptimised approach.

	Data Set A Best Case	Data Set B Worst Case	Average for all data sets
Patient position	supine	supine	-
Data set size (slices)	239	281	264
Lumen size (voxels)	3,479,881	7,199,653	5,291,415
Number of Colon Segments	1	2	2.1
Ratio of Centreline/skeleton voxels	0.573	0.523	0.515
Loading of LUT (s)	0.312	0.25	0.307
Loading of cropped Data set (s)	0.687	4.156	2.366
Endpoint Detection (s)	0.218	1.578	0.948
Caecum Point Revision (s) 0.047	0.078	0.066	
Surface detection (s)	0.282	0.453	0.329
Skeleton generation (s)	0.828	1.625	1.247
Skeleton reduction (s)	0.031	0.078	0.096
Overall with optimisation $t_1$ (s)	2.343	8.093	5.284
Overall without optimisation $t_2$ (s)	110.157	225.906	169.468
Improvement ( $t_2/t_1$ )	47.015	27.916	<b>32.739</b>

perfectly matching centrelines in all cases. The comparison was performed by calculating the distance from each voxel in the optimised centreline with the nearest voxel in the standard centreline (similar to the phantom accuracy evaluation). In all cases the sum of these distances was zero. The average time required for centreline calculation using the optimised approach was 5.284 seconds compared to 169.468 seconds for the standard approach (see Table 7.4 for a detailed breakdown of results). This represented an improvement of over 32 times. The once off task of lookup table initialisation required approximately 41 minutes. However, loading the table prior to each centreline calculation required less than one second.

### 7.3.4 Centreline Correctness

The accuracy of the colon centreline technique was established quantitatively for the ideal case (i.e. for a synthetic phantom) using the two tests outlined in Sections 7.3.2.1 and 7.3.2.2. The purpose of the final test was to perform a qualitative evaluation of the calculated centreline for the clinical cases that were used in the study. There is no accepted definition for the actual centreline of a human colon (Iordanescu & Summers 2003). In the many studies a visual inspection is performed to ensure the correctness of the calculated centreline. A similar approach to evalu-

ating centreline correctness was used in this study. The evaluation consisted of two main components:

1. To verify that the calculated centreline followed the expected path from the rectum to the caecum and that it was correctly centred in all cases.
2. To verify, in the case of multi-segment colons, that the individual colon segments were linked up correctly.

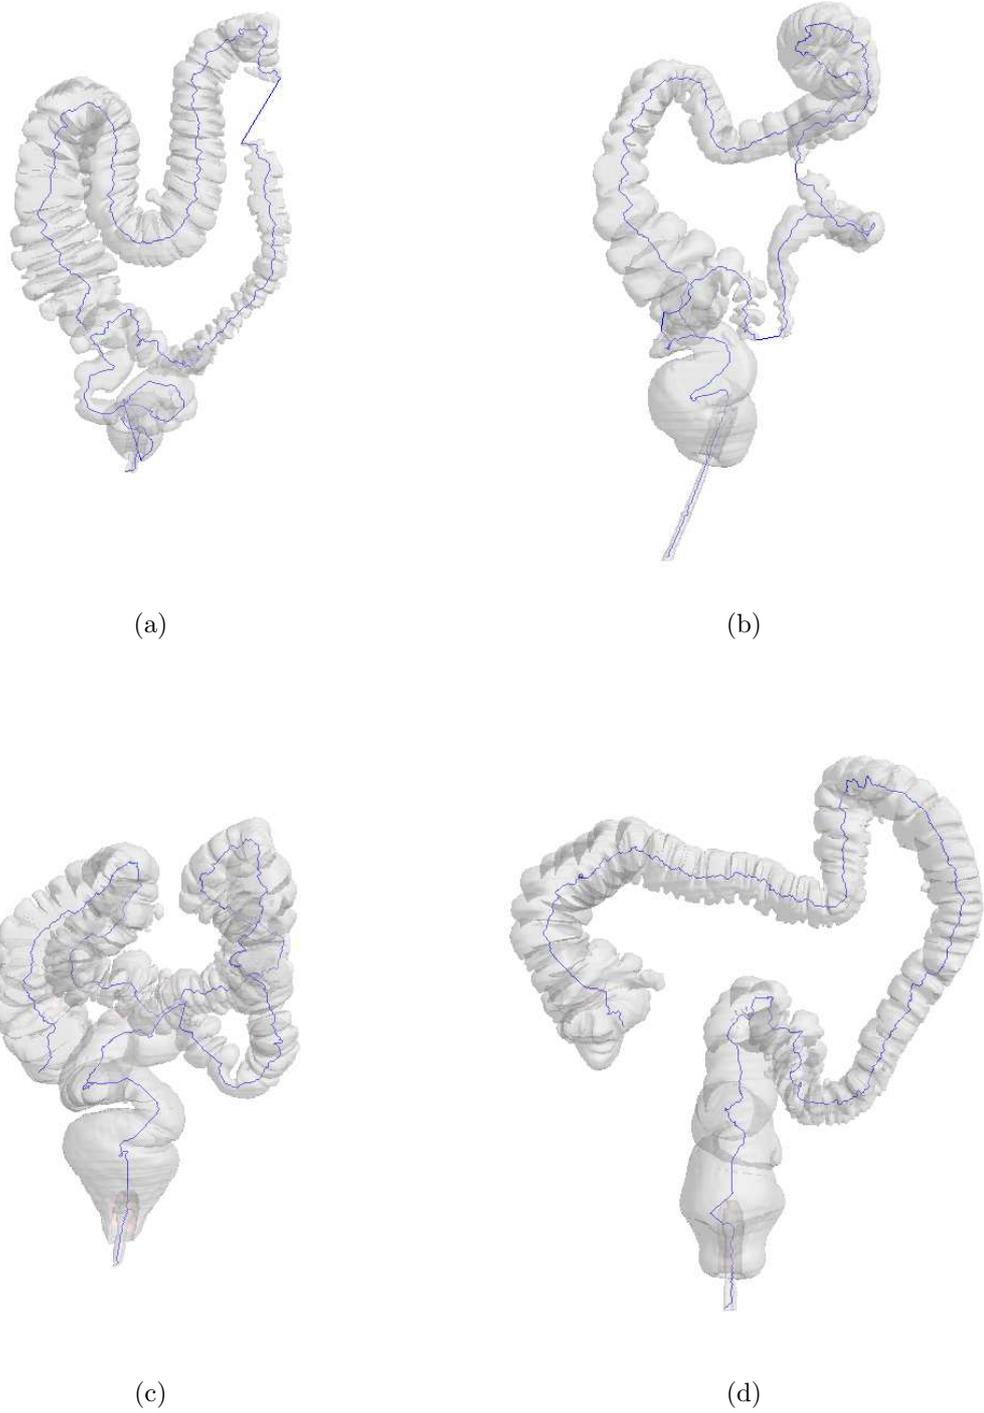
This evaluation was purely qualitative and involved a visual inspection of the centrelines calculated for the clinical cases. The correct centreline was calculated in the majority of cases, even when the colon lumen consisted of several disjointed segments (see Figure 7.5). However, the process failed in two cases. In the first case, the failure occurred due to shine through artefact that resulted in the caecum end point being detected in the mid ascending colon (see Figure 7.6 (a)). In the second case, disjointed centreline segments were incorrectly connected. This occurred due to a significant collapse in the sigmoid colon that resulted in the ascending colon/caecum being closer to the rectum than the descending colon. Thus, the rectum was incorrectly connected to the caecum, and the caecum point was incorrectly placed at the sigmoid flexure (see Figure 7.6 (b)).

## 7.4 Subsegmentation Evaluation

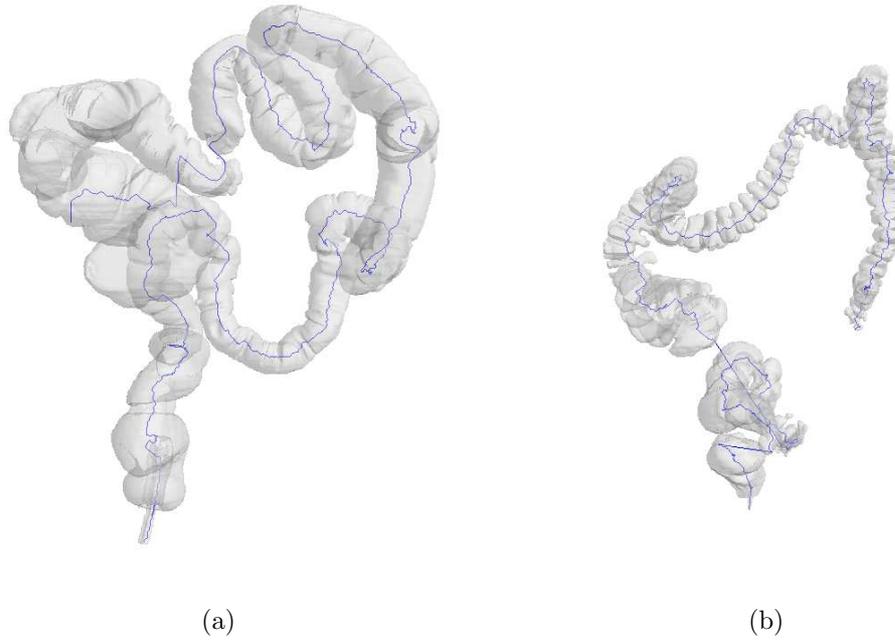
The subsegmentation process outlined in Section 4.6 was evaluated using the same clinical data set that was used in the centreline evaluation (see Table 7.1). In five of these cases, the flexure points were identified by a radiologist experienced in CTC. The subsegmentation evaluation consisted of establishing the difference between the automatically detected flexure points and those specified by the radiologist. In addition, a visual inspection of the remaining data sets was performed to ensure repeatability, and verify the correctness, of the subsegmentation process.

### 7.4.1 Flexure Detection Accuracy

In each of the five cases, the four flexures (rectosigmoid junction, sigmoid flexure, splenic flexure and hepatic flexure) were uniquely detected from each of the relevant colon centrelines in an average time of 2.55 seconds (range 0.55 - 4.04). No false positive detections were made. The position of the automatically detected and manually



**Figure 7.5:** Examples of the centreline displayed in conjunction with transparently rendered colons for colons with 3, 4, 2 and 1 segments, respectively.



**Figure 7.6:** The two cases (out of 20) where centreline calculation failed. (a) A shine through artefact that resulted in the caecum point being detected in the mid ascending colon. (b) Disjoined centreline segments were incorrectly connected due to a large collapse in the sigmoid colon.

selected flexure points  $d_p$  was calculated relative to the centreline endpoint located in the rectum.

$$d_p = \sum_{c=1}^n \sqrt{\begin{matrix} (v_s(x_c - x_{c-1}))^2 + \\ (v_s(y_c - y_{c-1}))^2 + \\ (v_s(z_c - z_{c-1}))^2 \end{matrix}} \quad (7.4.1)$$

where  $v_s$  represents the voxel size,  $(x_c, y_c, z_c)$  are points on the centreline, and  $n$  is the number of centreline points. For all data sets used in this study the voxel size ranged between 0.617 mm and 0.847 mm. The *mean error* (ME) was calculated for each of the automatically detected flexure points, this was found to be 31.03 mm. The most accurately detected point was the hepatic flexure (ME = 7.53 mm) while the least accurately detected point was the sigmoid flexure (ME = 54.30 mm). It should be noted that it is impossible to uniquely identify the flexure points and that mean errors of this magnitude are acceptable. In addition, these errors should be considered in relation to the average colon length (including rectum) which was approximately two metres<sup>2</sup>. Therefore, the mean error for the detection of all points

<sup>2</sup>Based on the average lengths of the correctly calculated centrelines from Table 7.1 (2,173.375 mm).

was 1.43% of the actual colon length.

### 7.4.2 Subsegmentation Correctness

A qualitative evaluation of the correctness of the subsegmentation process was carried out based on a visual inspection the subsegmented colons. The subsegmentation process was applied in cases where the centreline was correctly calculated. The subsegmentation process operated as expected in all cases. Samples of the output from the subsegmentation process are illustrated in Figure 7.7. Note: These are the same data sets that are illustrated in Figure 7.5.

## 7.5 CAD-CTC Evaluation

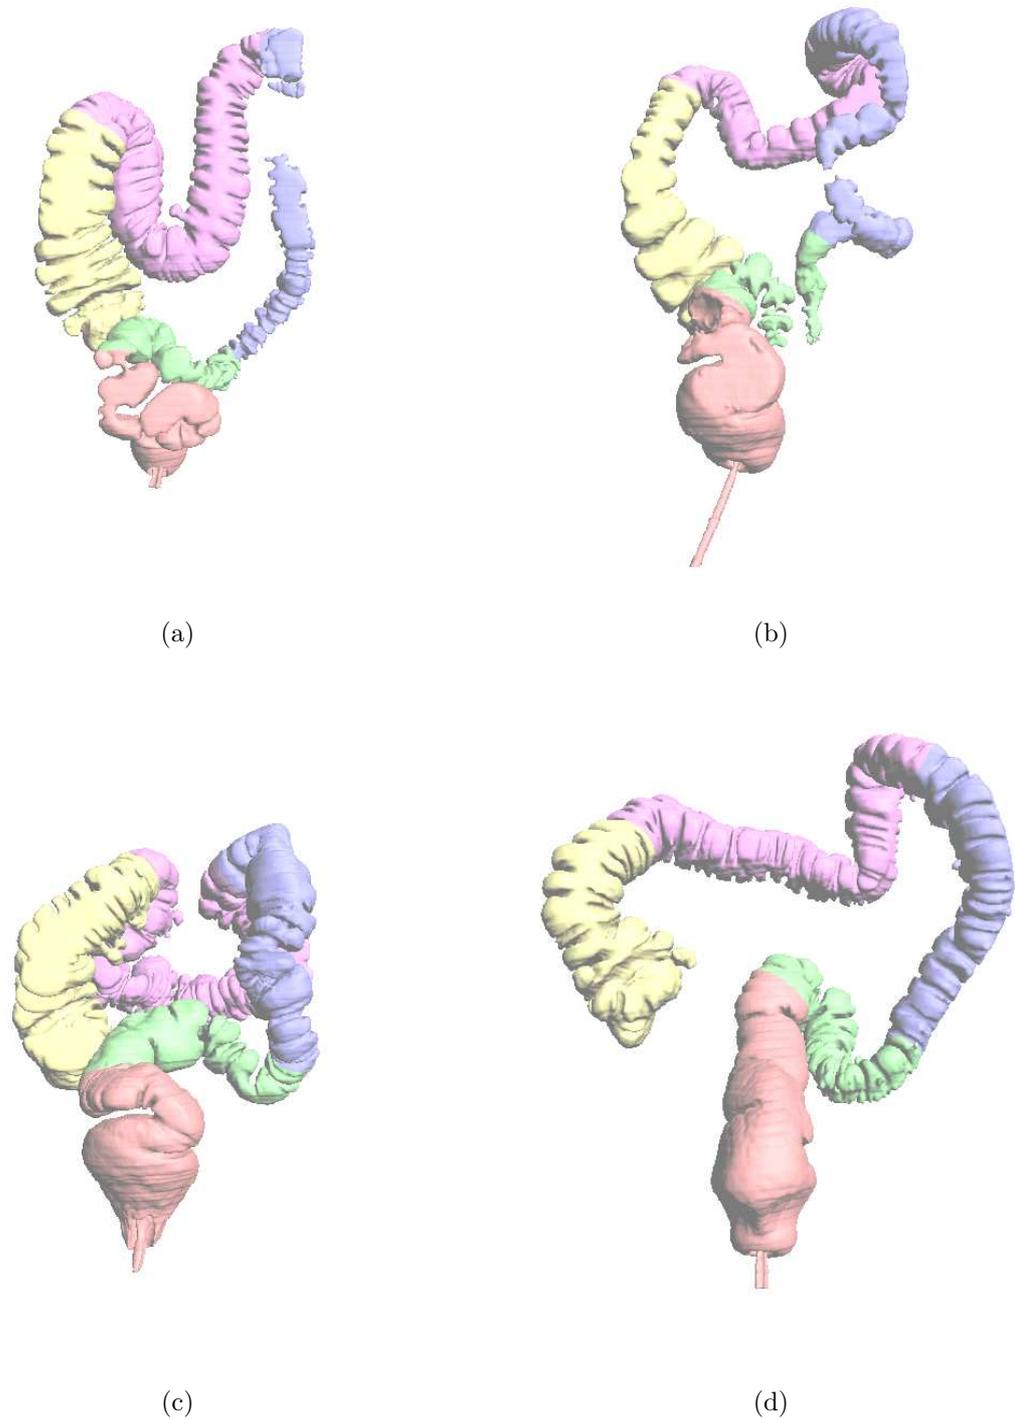
The evaluation of the CAD-CTC technique outlined in Chapter 5 consisted of three main stages:

1. A comparison of the two main surface curvature based techniques for identifying initial polyp candidates.
2. A comparative evaluation of the use of regular candidate features versus expanded candidate features in discriminating between true polyps and false positives.
3. An overall comparison of the two classes of features using a linear classifier to determine the optimum candidate representation.

All aspects of CAD-CTC training, testing and evaluation were carried out using the MMH data sets.

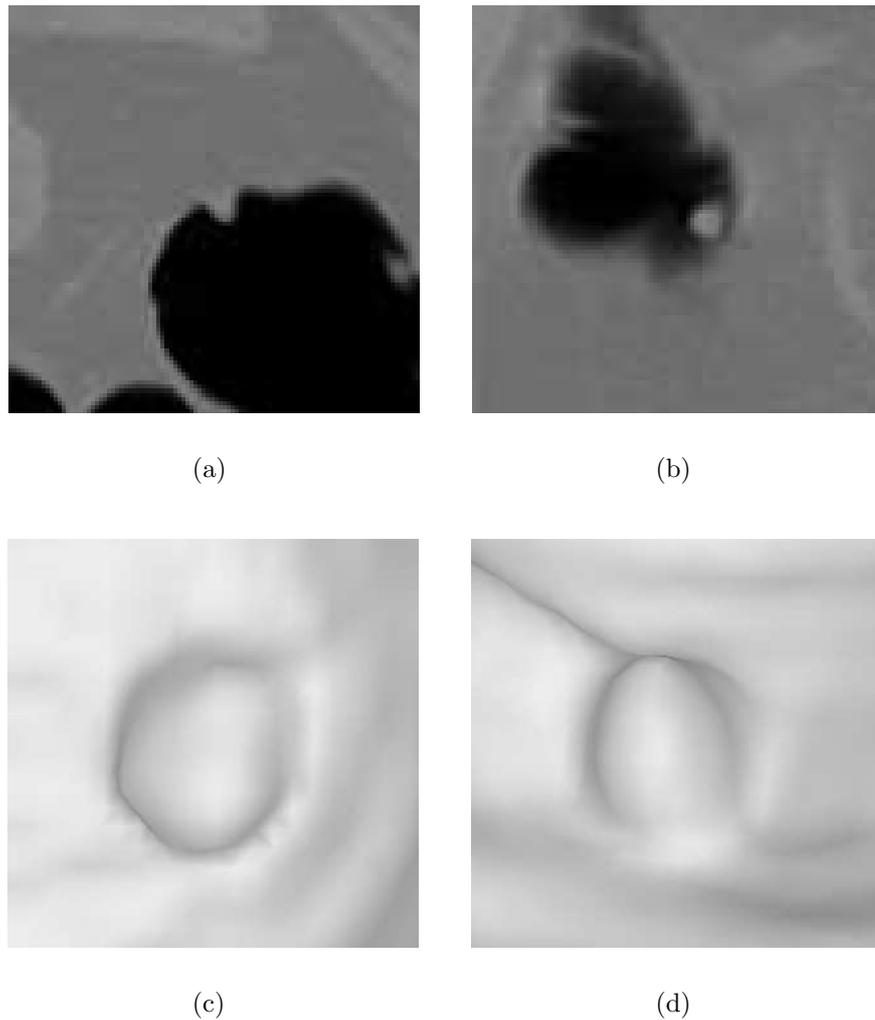
### 7.5.1 Test Data

The MMH data (summarised in Appendix B) was used exclusively in the testing and evaluation of the CAD-CTC technique. Patients in the study underwent CTC followed immediately by conventional colonoscopy. The gold standard information regarding polyp size and location was determined based on an unblinded evaluation of the CTC data sets (i.e. the radiologist had access to the colonoscopy results). This ensured that the most accurate polyp information was available for training and evaluating the CAD-CTC system.



**Figure 7.7:** Examples of the output from the subsegmentation process for colons with 3, 4, 2 and 1 segments respectively. Note: These are from the same data sets as the centrelines that are illustrated in Figure 7.5.

In order to increase the amount of test data available, each prone or supine data set was treated as a separate patient. This effectively doubled the amount of polyps available for CAD-CTC testing while ensuring a more rigorous evaluation as prone/supine polyp registration (see Section 3.4) was no longer an option to increase polyp detection sensitivity. The justification behind using the same polyp scanned in different positions as different polyps is illustrated in Figure 7.8. The gold standard information was entered manually using the dedicated user interface that was described in Section 6.4.2.1.



**Figure 7.8:** An illustration of the same polyp as it appears in the prone (a) and (c) and supine (b) and (d) data sets of a patient (MMH-393). It is evident that there are significant variations in the perceived polyp size and polyp morphology between the two positions.

## 7.5.2 Primary Search Criteria

The two main groups involved in surface curvature based CAD-CTC research are from the National Institutes of Health (NIH) and the University of Chicago (UC). Each of these groups uses a different approach for identifying initial polyp candidates.

- Candidate vertices were identified by the NIH approach as having a positive nonzero mean curvature ( $H > 0$ ) and a positive nonzero Gaussian curvature ( $K > 0$ ). Candidate regions with a maximum dimension greater than a minimum size threshold  $C_{min}$  were retained as polyp candidates.
- The UC approach involved the use of shape index (R) and curvedness (S) curvature characteristics. The approach was described in relation to voxels but will be discussed here in relation to vertices. Two classes of candidate vertices were identified:
  - **Cap** - Vertices of the polyp cap class have a shape index greater than 0.9 and a curvedness in the range  $0.08 \text{ mm}^{-1} - 0.33 \text{ mm}^{-1}$ .
  - **Collar** - Vertices of the polyp collar class have a shape index in the range 0.8 - 1.0 and a curvedness in the range  $0.05 \text{ mm}^{-1} - 0.5 \text{ mm}^{-1}$ .

Candidate regions (consisting of both cap and collar vertices) and a maximum dimension that was greater than a minimum size threshold  $C_{min}$  were retained as polyp candidates.

### 7.5.2.1 Initial Evaluation

A comparative evaluation of the two sets of features was carried out to establish which set of features provided the best overall performance. In both cases  $C_{min}$  was set to a value of 2 mm in order to guarantee the detection of all clinically significant polyps. This evaluation involved the use of 10 data sets (all with polyps) that were selected from the MMH CTC database (see Appendix B). The results of the evaluation are presented in Table 7.5 and will be summarised below.

- The  $H/K$  features retained on average 6.912% (removed on average 93.088%) of the colon surface, whereas the  $R/S$  features retained on average 1.172% (removed on average 98.828%). This difference is evident from a visual inspection

**Table 7.5:** The results of a comparative evaluation of the  $R/S$  and  $H/K$  features for initial polyp candidate detection using 10 randomly selected CTC data sets with a total of 27 polyps.

Patient ID	# Polyps	Detections		Candidates		Coverage (%)	
		$H/K$	$R/S$	$H/K$	$R/S$	$H/K$	$R/S$
MMH-017S	4	4	4	977	178	6.808	1.224
MMH-120S	3	3	3	1050	227	7.030	1.363
MMH-161P	1	1	1	657	134	6.054	1.004
MMH-208P	3	3	3	629	103	6.500	0.831
MMH-234S	2	2	2	1393	206	7.300	1.134
MMH-368P	1	1	1	779	55	5.487	0.631
MMH-379P	2	2	2	1962	438	7.850	1.396
MMH-451S	3	3	3	1437	310	8.177	1.748
MMH-483S	5	5	5	643	66	4.786	0.830
MMH-708S	3	3	2	853	177	8.133	1.526
mean	2.7	2.7	2.6	1038	135.5	6.912	1.172

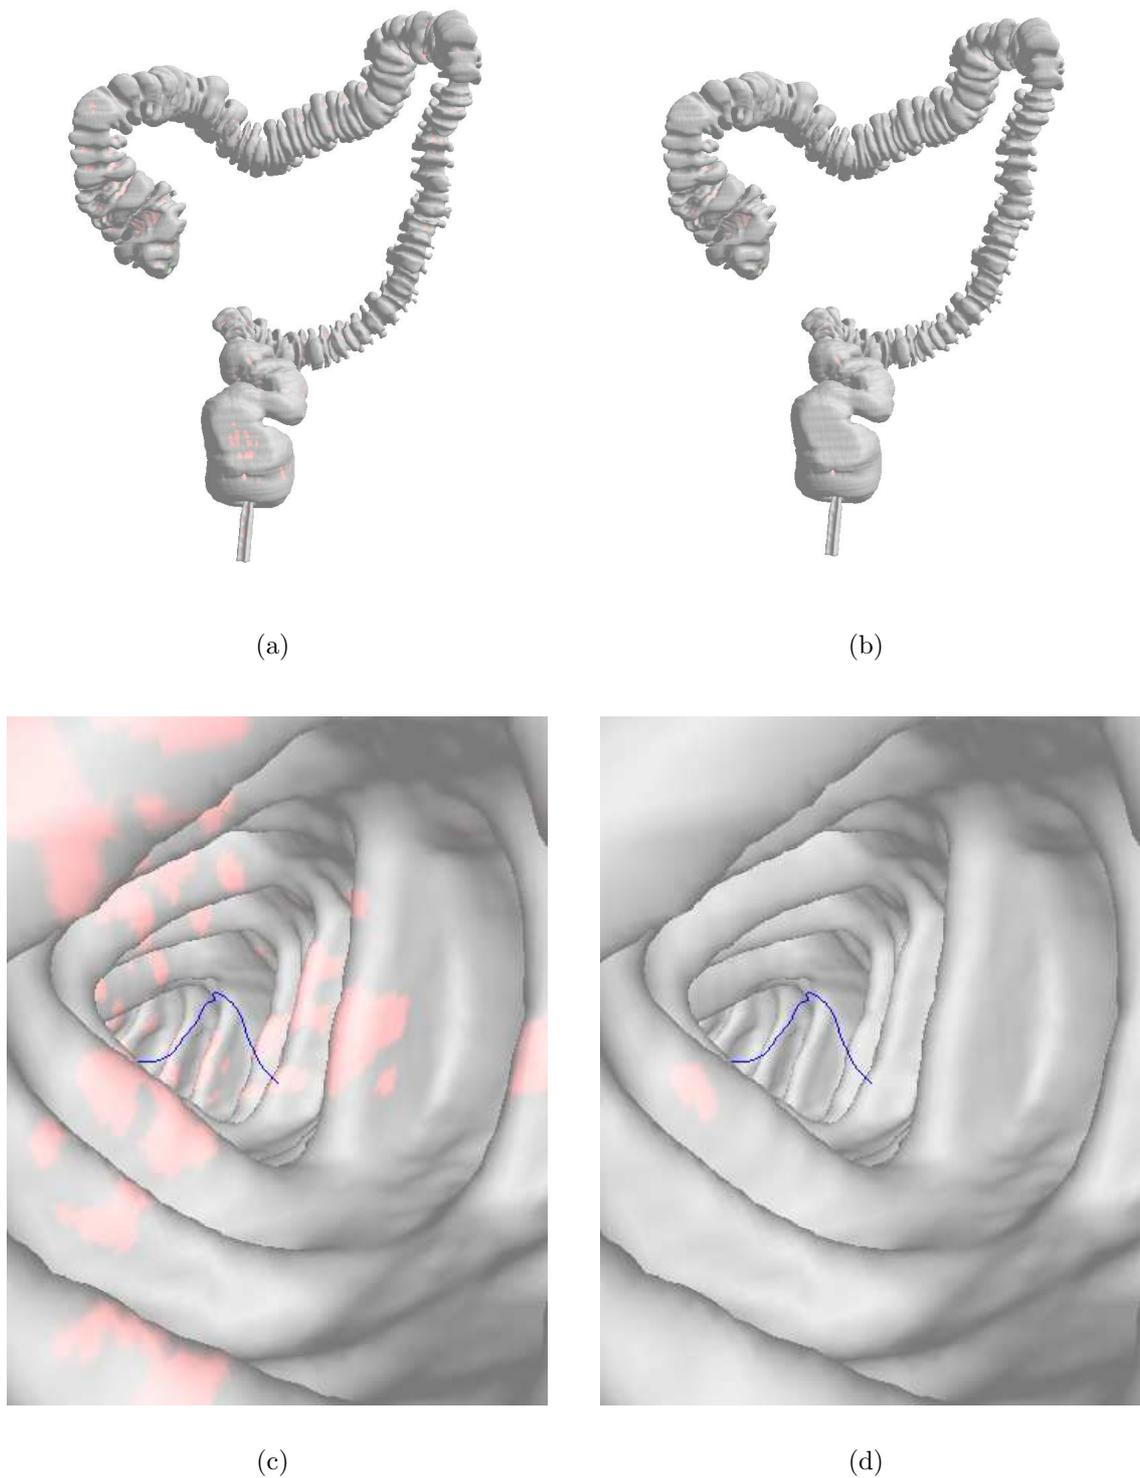
of the same colon processed using the  $H/K$  features (see Figure 7.9 (a) & (c)) and the  $R/S$  features (see Figure 7.9 (b) & (d))

- The mean number of candidates generated by the  $H/K$  features was 1,038 whereas the  $R/S$  features generated an average of just 135.5 candidates<sup>3</sup> (i.e. the  $H/K$  features generated of seven times more candidates than the  $R/S$  features).
- The  $H/K$  features identified all 27 polyps that were present in the 10 data sets (100% sensitivity). However, the  $R/S$  features failed to detect a single 3 mm polyp in one of the data sets (96% sensitivity).

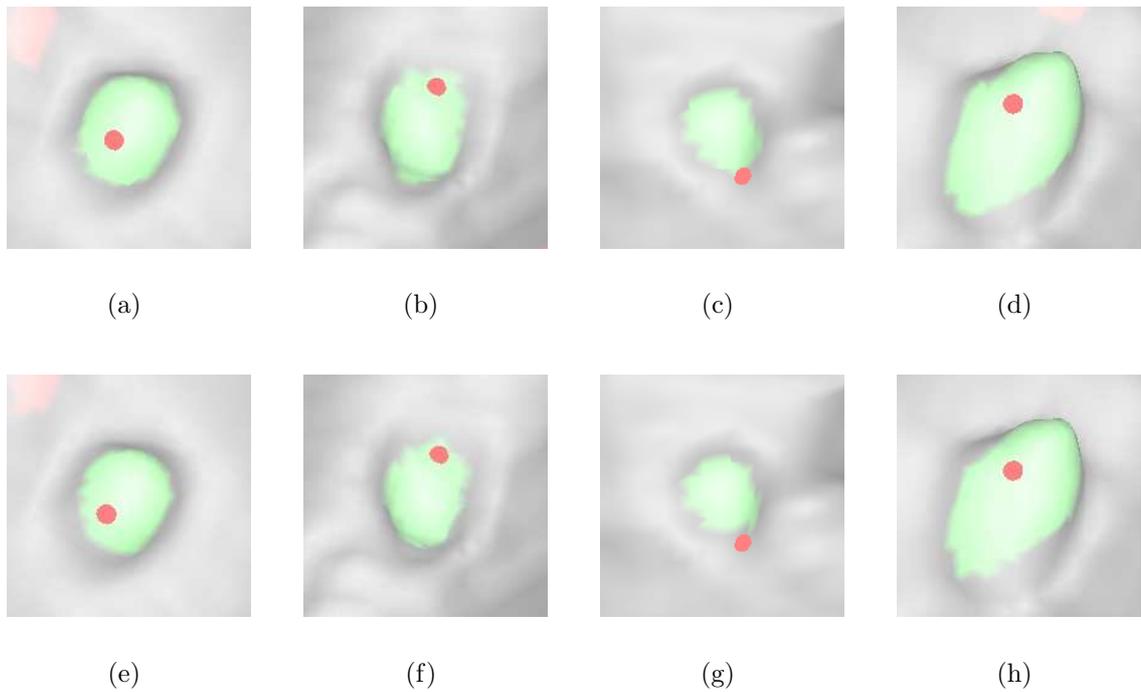
Both initial polyp candidate detection techniques provided similar sensitivity levels. However, the  $H/K$  features generated over seven times more candidates than the  $R/S$  features. The single polyp that was not detected by the  $R/S$  features was 3 mm in size and was not clinically significant. A qualitative evaluation of the differences in candidate patch size for actual polyps found no significant differences between the two approaches, see Figure 7.10.

The  $R/S$  features detected all clinically significant polyps in the test set and provided adequate polyp surface coverage with the relatively low number of false

<sup>3</sup>Other group that used shape index for initial polyp candidate identification reported a similar number of detections, e.g. Wang et al. (2005) (SUNY) reported that their approach identified over 100 initial polyp candidates.



**Figure 7.9:** Internal and external renderings of the same data set (MMH-708/supine) where initial candidate detection has performed using the  $H/K$  features (a) & (c) and  $R/S$  features (b) & (d). It is clear from a visual inspection of these renderings that the  $R/S$  features generate far fewer polyp candidate regions.



**Figure 7.10:** A comparison of polyp coverage of the same set of polyps selected from MMH-017/supine using  $H/K$  features (a) - (d) and  $R/S$  features (e) - (h) initial candidate identification techniques. Note: The polyp sizes from left to right are 5.25, 3.0, 3.0 and 5.5 mm respectively.

positive detections per data set (when compared with the use of the  $H/K$  features). As a result, the  $R/S$  features used for initial polyp candidate detection.

### 7.5.2.2 Overall Findings

When the  $R/S$  features were used for initial candidate detection in the full test set, all 13 large polyps were identified, all 31 medium polyps were identified and of the 98 small polyps, a total of 90 (92%) were identified. The average number of false positives per data set was 183.

### 7.5.3 Automated Scoring

Each of the candidate regions detected by the primary stage must be manually classified as either a true polyp or false positive in order to automate subsequent tasks including feature evaluation and classifier evaluation. The number of candidates generated by the primary stage was quite high (on average 183 per data set). As a result, an automated approach to scoring based on the technique described by Paik et al. (2004) was employed. Using this approach a polyp candidate was defined as a true polyp if it was within a fixed distance of the gold standard coordinates of an actual

polyp. This approach fully automated the task of candidate definition and greatly reduced the potential for errors that would be associated with manually defining the class of so many candidates. Following this process, candidate patches that were associated with polyps were shaded green, whereas those that were associated with false positives were shaded red.

#### 7.5.4 Original versus Expanded Polyp Candidates

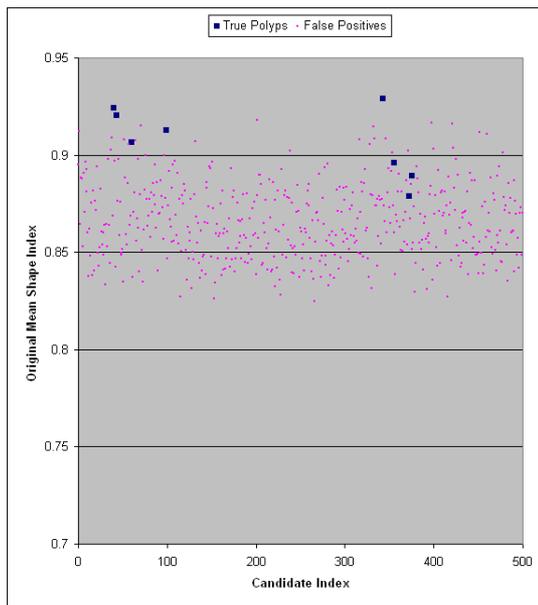
A series of tests were performed to compare the use of original candidates and expanded candidates for polyp detection. It was only possible to carry out an individual evaluation of three features: Shape index, sphericity and density. The remaining features: Mean curvature, Gaussian curvature and curvedness all needed to be combined with at least one other feature in order to effectively discriminate between true polyps and false positives. The capability of each feature to discriminate between true polyps and false positives was evaluated using scattergrams and ROC curves. Scattergrams enabled a visual comparison of the two candidate types, and ROC curves enabled a quantitative comparison of the sensitivities and false positive rates for the two candidate types over all three features.

##### 7.5.4.1 Initial (Limited) Individual Feature Evaluation

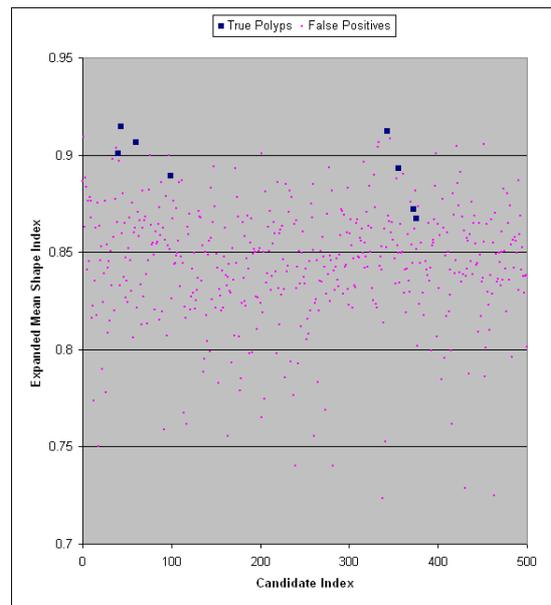
A proof of concept evaluation was carried out using the prone and supine data sets for a single patient. The patient in question (MMH-017), had four polyps, all of which were visible in both the prone and the supine views. Following initial candidate identification, the prone study yielded 329 candidates and the supine study yielded 178 candidates (i.e. a total of 507 candidates). The scattergrams and ROC curves for each of the three candidate features are presented in Figures 7.11 to 7.13. Note: In the expanded feature scattergrams, some of the outliers have been truncated. With the exception of the density feature, the use of fingerprinting and candidate expansion reduced the number of false positive detections by at least 25%, thus, demonstrating that the use of candidate expansion has the potential to markedly reduce the number of false positive detections in CAD-CTC.

##### 7.5.4.2 Extended Individual Feature Evaluation

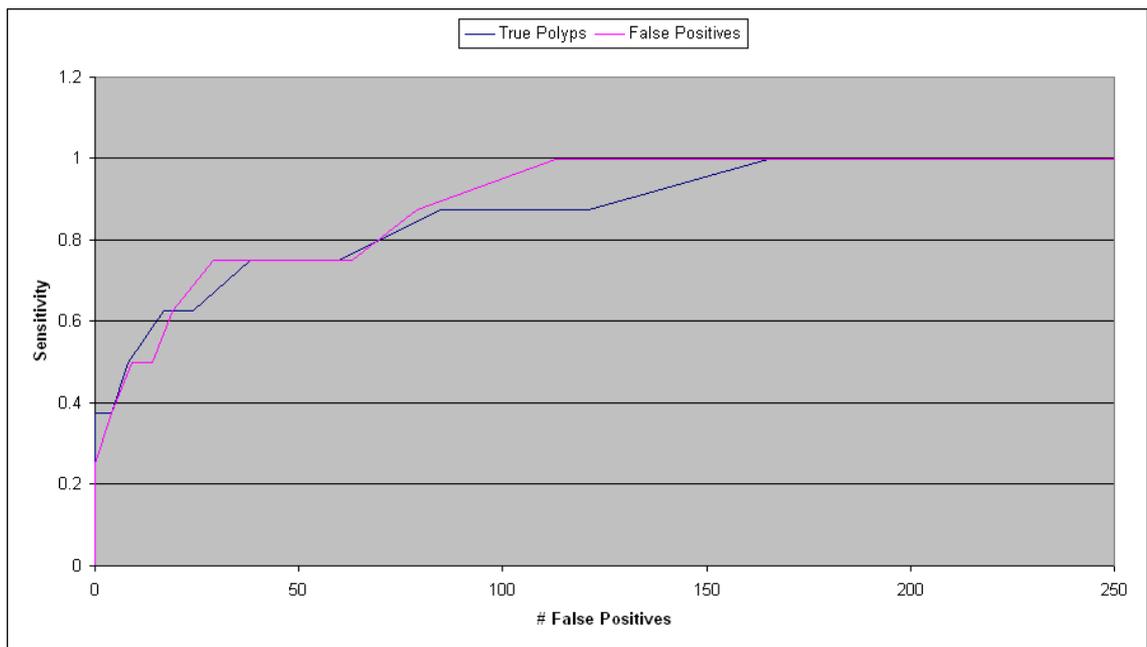
A more rigorous evaluation of the use of expanded candidates was carried out using a larger test set consisting of prone and supine data sets from five patients. These



(a)

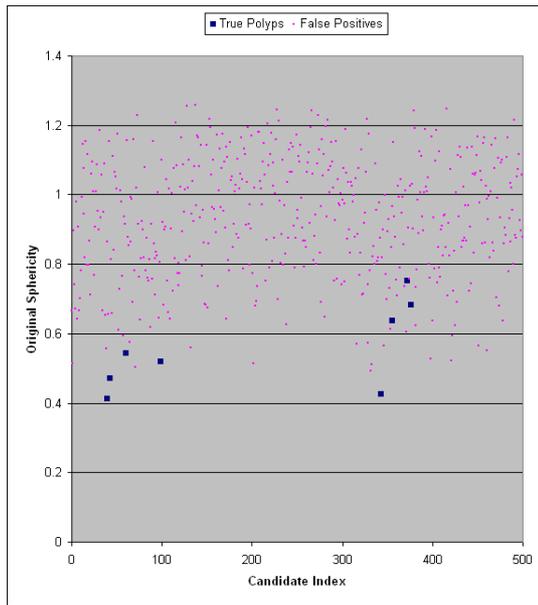


(b)

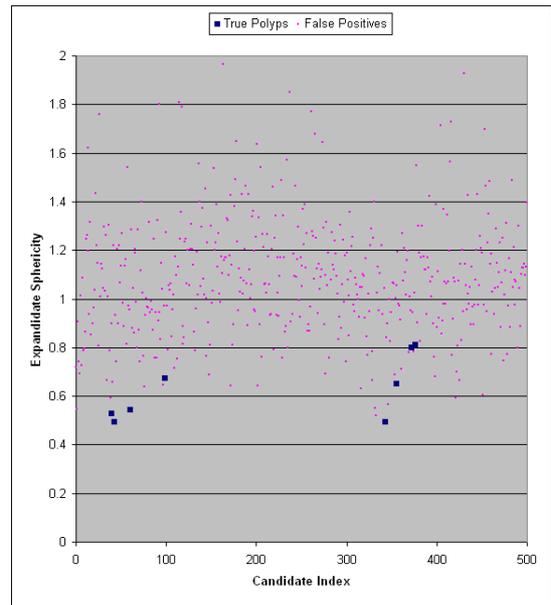


(c)

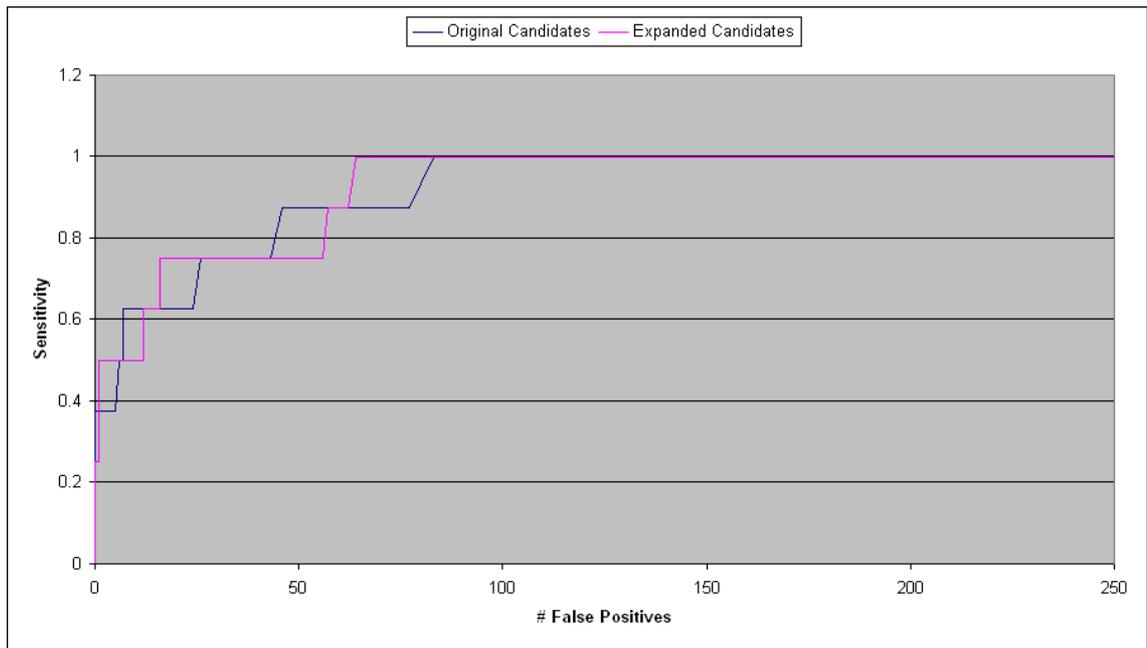
**Figure 7.11:** A comparative evaluation of the original versus expanded shape index in a single patient study: (a) A scattergram representing the distribution of the original candidates. (b) A scattergram representing the distribution of the expanded candidates. (c) A ROC analysis comparing the performance of the original and expanded candidates. Note: At 100% sensitivity the expanded candidates provide a 32% reduction in false positive detections.



(a)

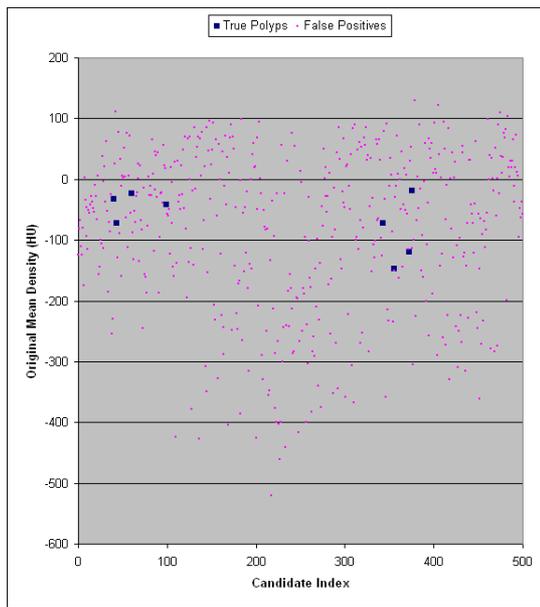


(b)

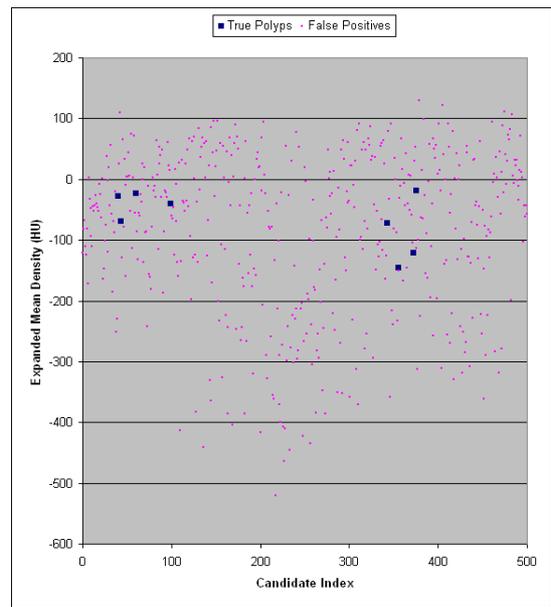


(c)

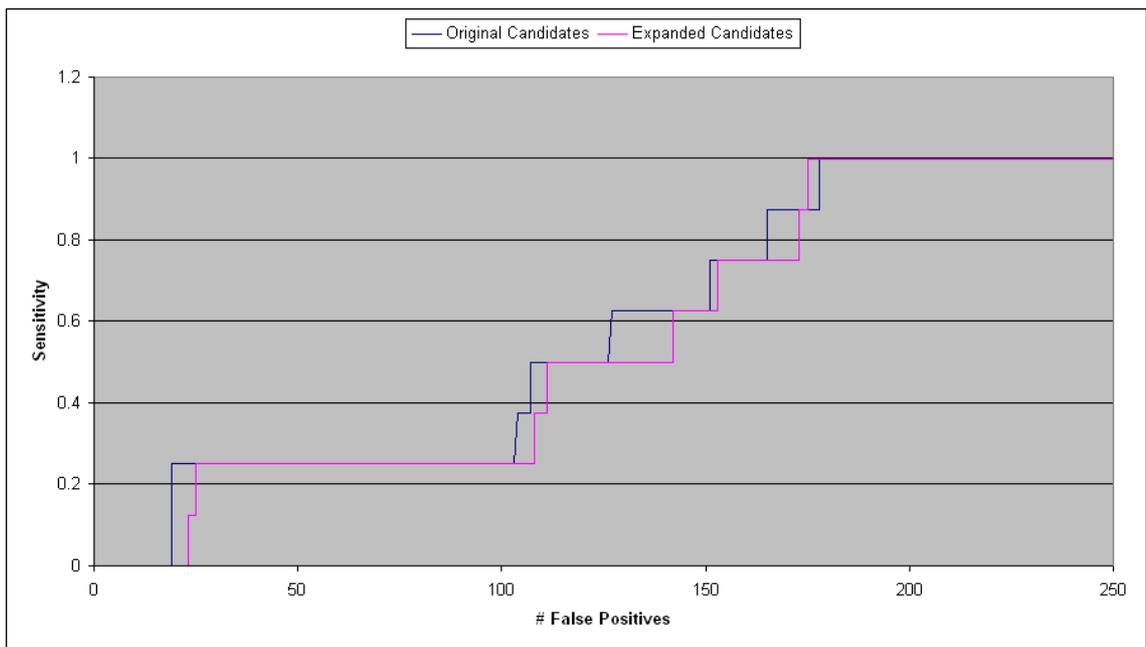
**Figure 7.12:** A comparative evaluation of the original versus expanded sphericity in a single patient study: (a) A scattergram representing the distribution of the original candidates. (b) A scattergram representing the distribution of the expanded candidates. (c) A ROC analysis comparing the performance of the original and expanded candidates. Note: At 100% sensitivity the expanded candidates provide a 30% reduction in false positive detections.



(a)

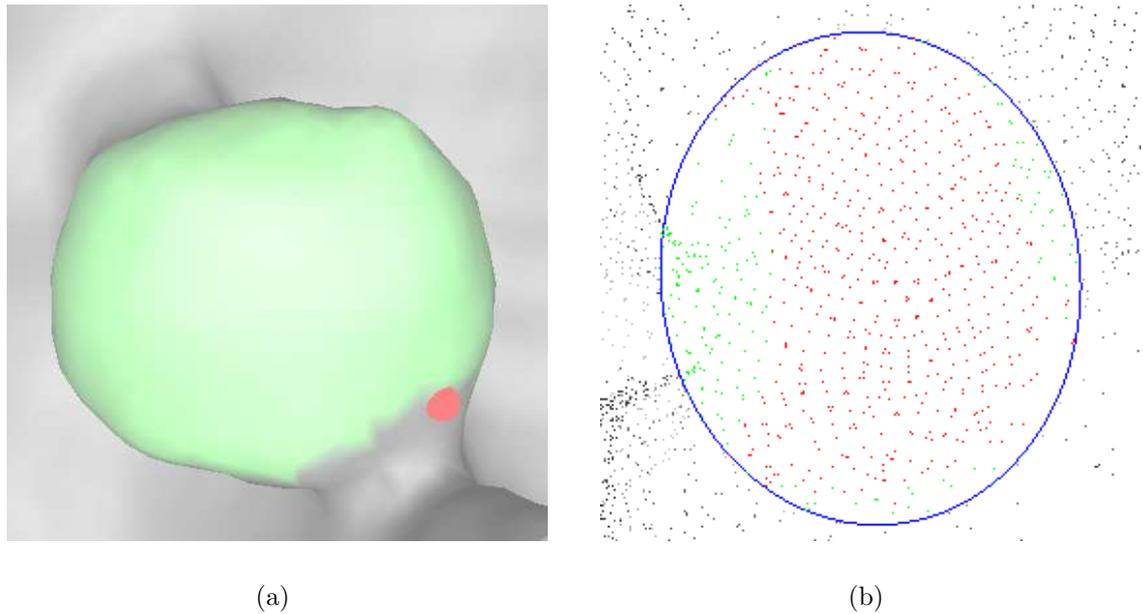


(b)



(c)

**Figure 7.13:** A comparative evaluation of the original versus expanded density in a single patient study: (a) A scattergram representing the distribution of the original candidates. (b) A scattergram representing the distribution of the expanded candidates. (c) A ROC analysis comparing the performance of the original and expanded candidates. Note: The density feature does not provide a good discrimination between true and false positives.

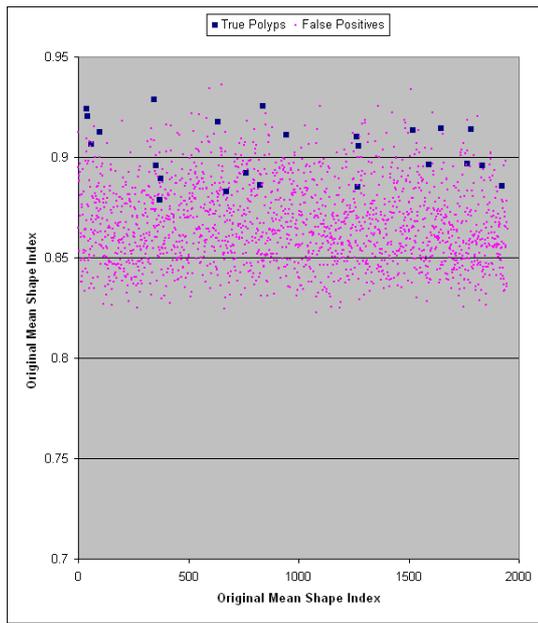


**Figure 7.14:** A large polyp located on a fold (a) and its associated fingerprint (b) (from MMH-483/supine). The irregular boundary of the polyp causes the MBE to include points that belong to the fold. This dilutes the polyp’s features e.g. the mean shape index drops from 0.92 in the original candidate to 0.81 in the expanded candidate.

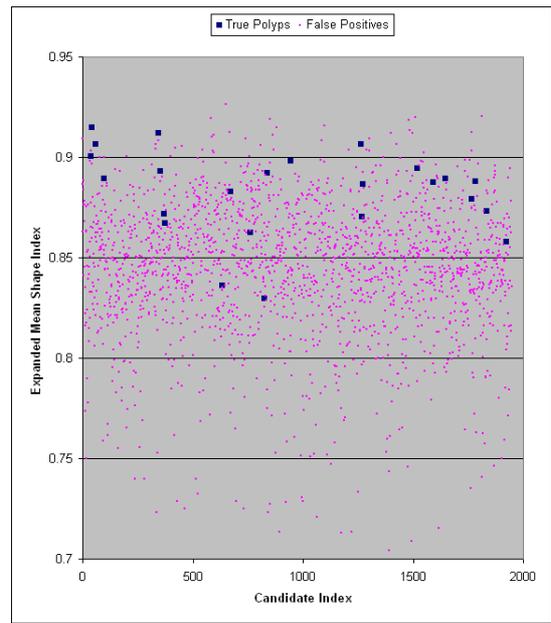
data sets contained a total of 24 polyps and initial polyp candidate detection yielded 1,948 candidates. The scattergrams and ROC curves for each of the three candidate features are illustrated in Figures 7.15 to 7.17. It is clear from an inspection of these plots that the use of expanded candidates had a detrimental effect on polyp detection sensitivity in the larger test set. Upon inspection of the missed polyps, it was found that candidate expansion diluted the properties of some polyps to such an extent that they were considered to be false positives. This problem occurred in the case of large polyps, or polyps that included sections of the folds that they were attached to. In both cases the major axis of the associated MBE was large, i.e. this problem did not occur in smaller candidate regions. An example of this problem is illustrated in Figure 7.14. Consequently, this study showed that candidate expansion can fail in certain cases, particularly in the case of large polyp candidates.

#### 7.5.4.3 Candidate Density as a Discriminating Feature

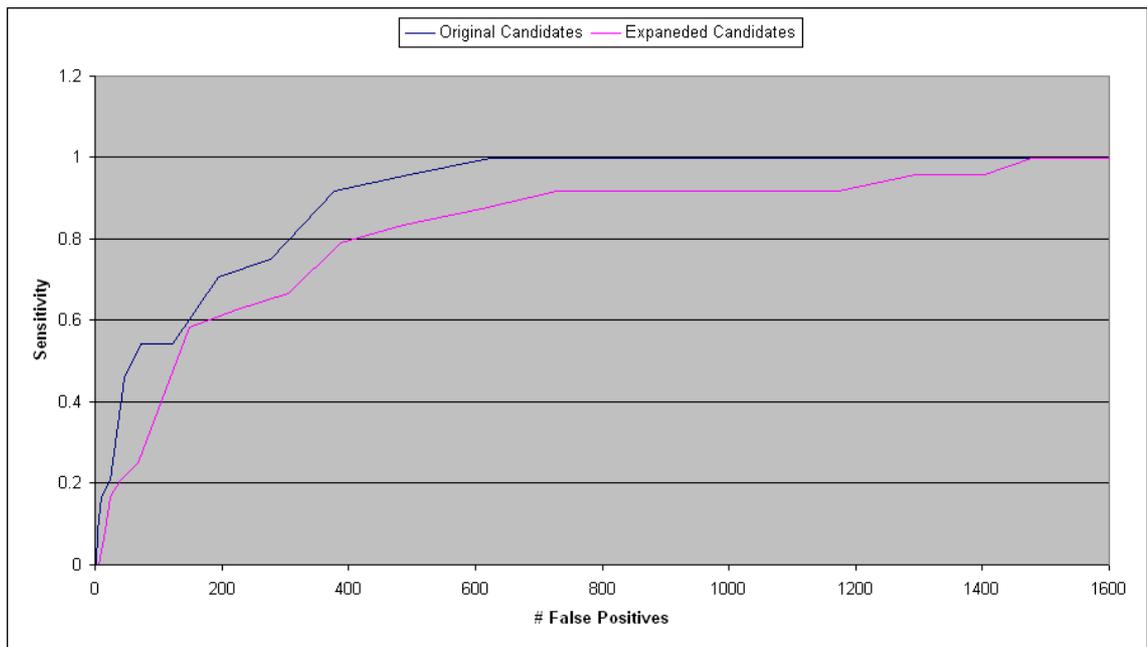
The density feature was found, in all cases, to provide very bad discrimination between true polyp and false positives (i.e. the ROC curve approximated a straight line). As a result, the density feature was not included in the final polyp candidate feature set.



(a)

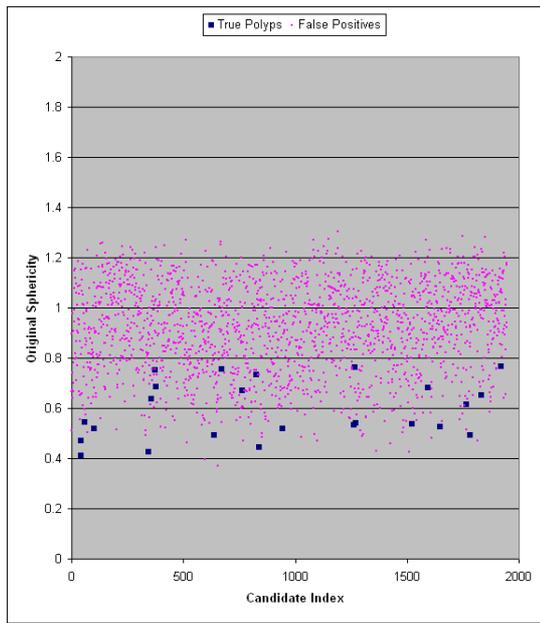


(b)

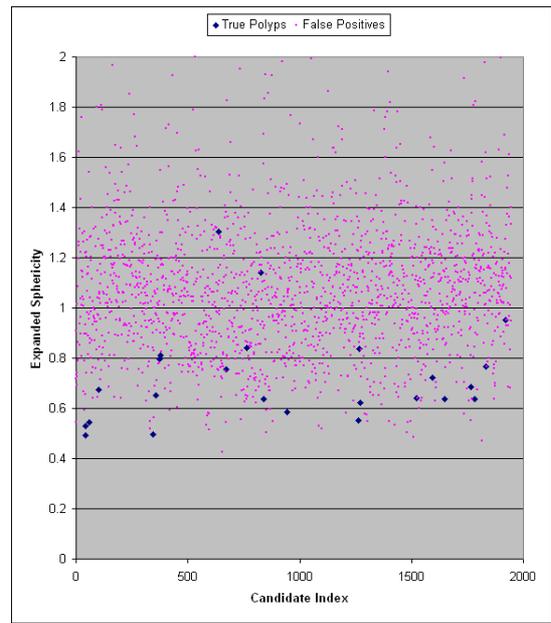


(c)

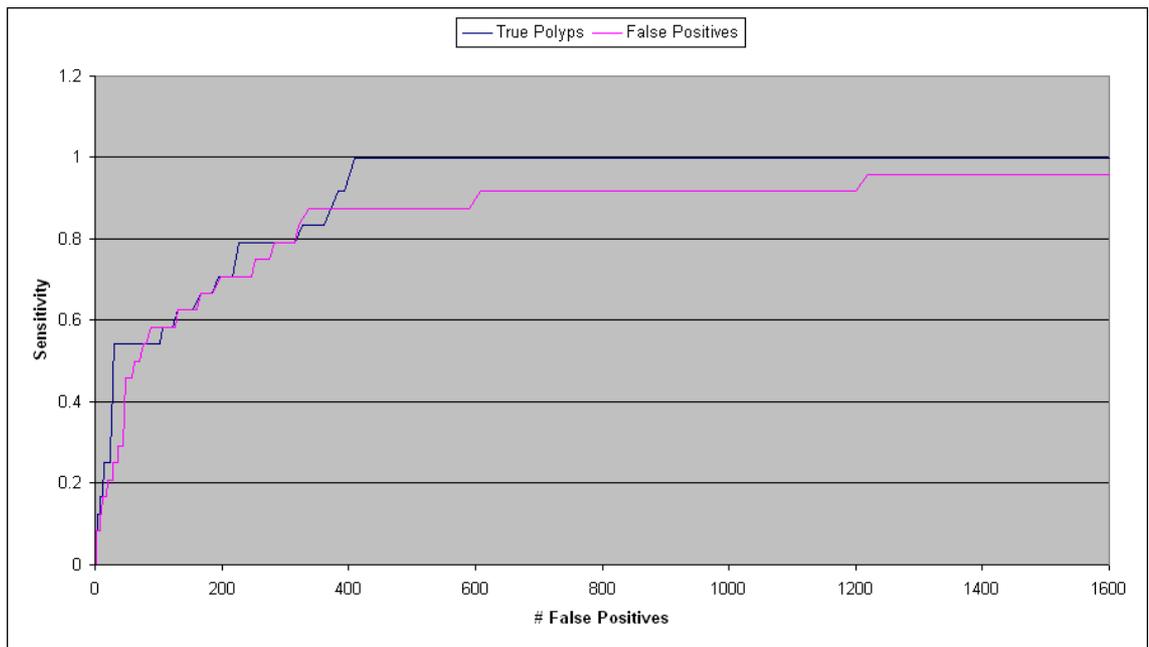
**Figure 7.15:** A comparative evaluation of the original versus expanded shape index in a five patient study: (a) A scattergram representing the distribution of the original candidates. (b) A scattergram representing the distribution of the expanded candidates. (c) A ROC analysis comparing the performance of the original and expanded candidates. Note: At 100% sensitivity the expanded candidates caused a 137% increase in false positive detections.



(a)

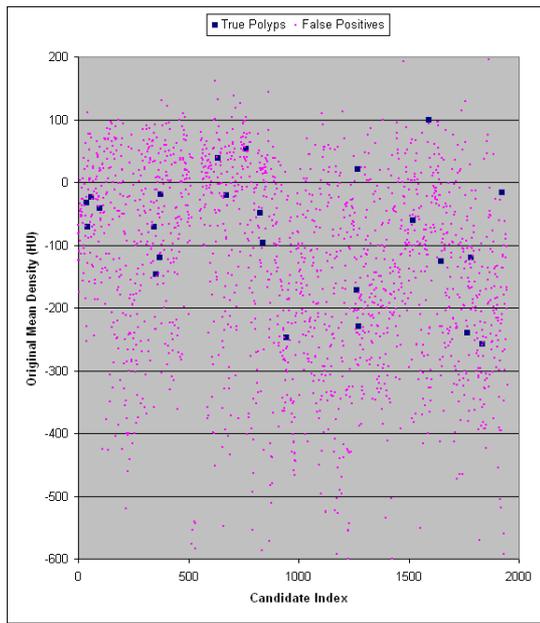


(b)

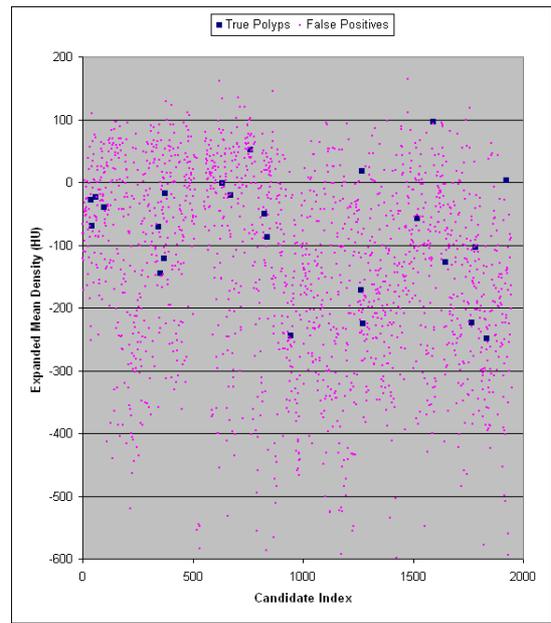


(c)

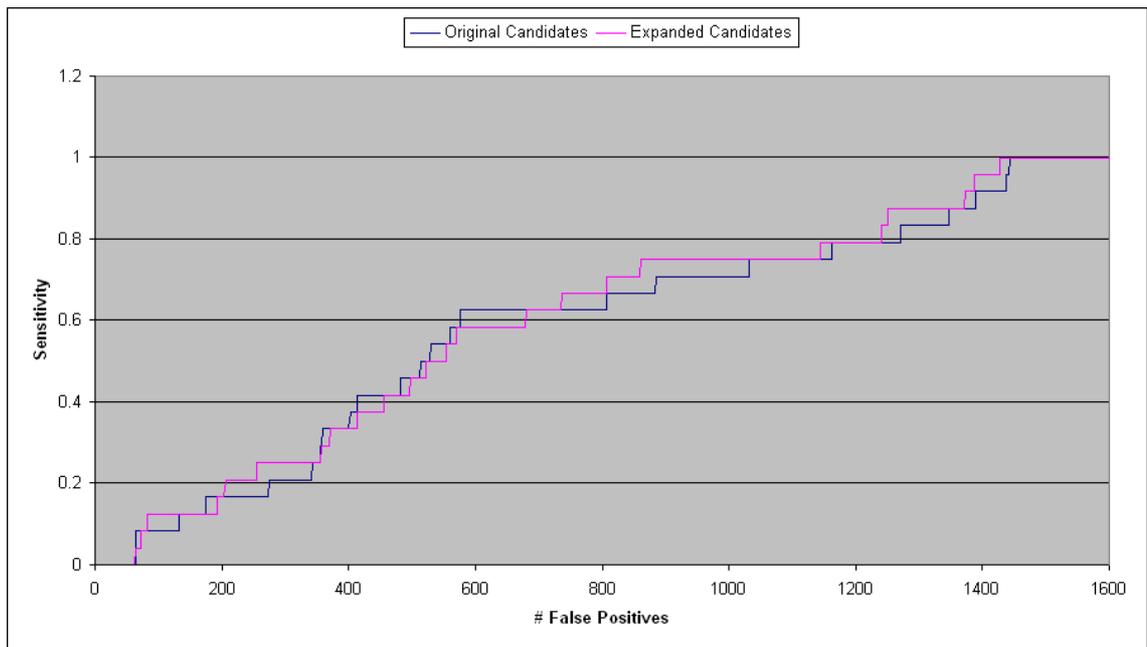
**Figure 7.16:** A comparative evaluation of the original versus expanded sphericity in a five patient study: (a) A scattergram representing the distribution of the original candidates. (b) A scattergram representing the distribution of the expanded candidates. (c) A ROC analysis comparing the performance of the original and expanded candidates. Note: A sensitivity of 100% was not achievable in this experiment with expanded candidates.



(a)



(b)



(c)

**Figure 7.17:** A comparative evaluation of the original versus expanded density in a five patient study: (a) A scattergram representing the distribution of the original candidates. (b) A scattergram representing the distribution of the expanded candidates. (c) A ROC analysis comparing the performance of the original and expanded candidates. Note: The density feature does not provide a good discrimination between true and false positives.

### 7.5.5 Combined Feature Evaluation

A classifier study was ultimately used to quantify the difference between using standard and expanded candidates for polyp detection. All of the candidate features, with the exception of candidate density, were used in this evaluation.

### 7.5.6 Classification of Standard Candidates

The classification of standard polyp candidates was performed using a MDC (described in Section 5.5.5.1). Each of the individual candidate features were normalised in the range 0 - 1 prior to classifier training. The leave one out approach was used for evaluating the classifier<sup>4</sup>. The decision boundary was initially set to be midway between the mean true polyp vector and the mean false positive vector. At this setting, the use of standard candidates was found to give a sensitivity of 97% for polyps in the range 5 - 9 mm, and 92% for polyps that were 10 mm or larger. There were on average 17.5 false positive detections per data set when standard candidates were used.

### 7.5.7 Classification of Expanded Candidates

The classification of expanded polyp candidates was performed in a similar manner with one minor modification. The classifier was trained using standard polyp candidates as before, however, unknown candidates were expanded prior to classification. Polyp candidate expansion was not applied to all candidates due to the limitations discussed in Section 7.5.4.2. Only polyps below a certain size were expanded, the remainder retained their original features. A study was carried out to determine the maximum size threshold for reliable candidate expansion. This study had two objectives:

1. To determine the effect on polyp detection sensitivity as the size threshold was increased, and hence, determine the optimal operating threshold for the candidate expansion process.
2. To determine the effect on false positive detections as the size threshold was increased, and consequently, determine the level of false positive detections at the optimal operating threshold.

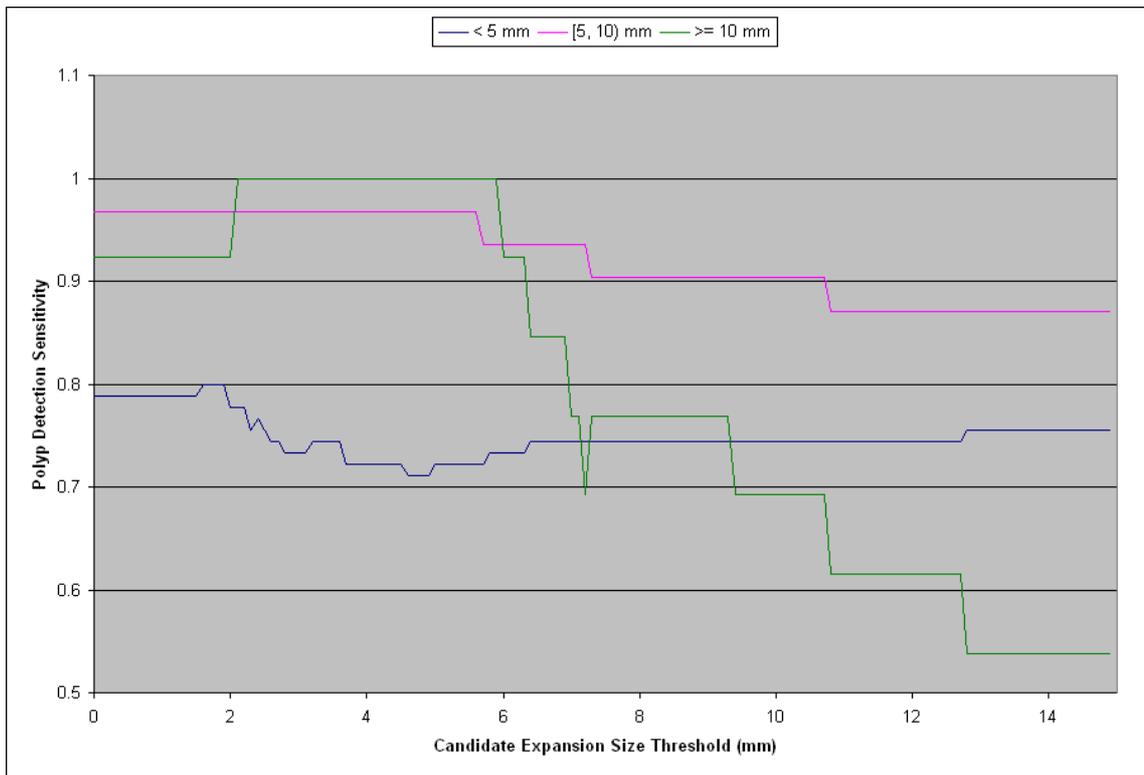
---

<sup>4</sup>This approach to classifier evaluation has been used extensively by other groups involved in CAD-CTC research (see Section 3.5).

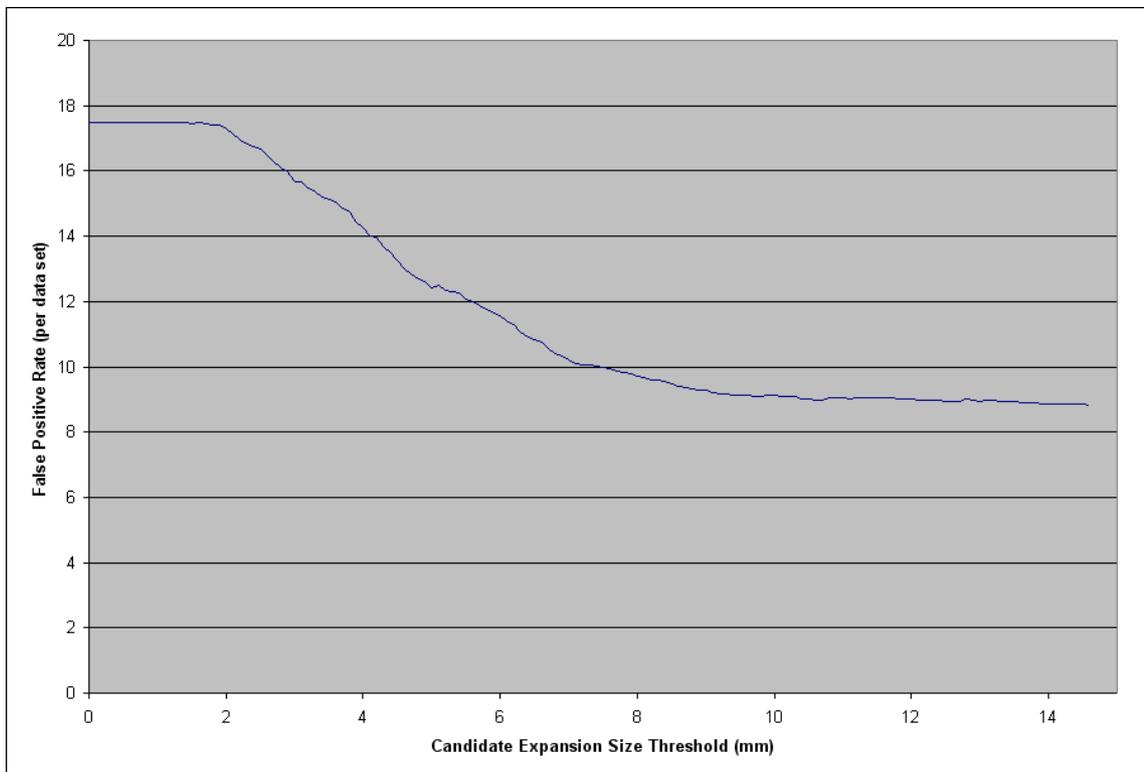
The required performance characteristics (i.e. sensitivity and false positive rate) were obtained for a range of candidate expansion thresholds. The range extended from zero to the size of the largest polyp candidate and was sampled at regular intervals. The size of a candidate was taken to be the length of the major axis of its MBE.

A plot of the polyp detection sensitivity versus candidate expansion size threshold is illustrated in Figure 7.18. It is evident that the sensitivity for the detection of large polyps ( $\geq 10$  mm) improves when the candidate expansion size threshold is low. However, as the threshold is increased the sensitivity for the detection of large polyps decreases the most of all the size ranges that are plotted. The impact on the sensitivity for medium polyps is also negative but less pronounced and the impact on the sensitivity for the detection of small polyps is slightly negative initially before stabilising and increasing. These results are as expected based on the observations from Section 7.5.4.2. It is clear from this plot, that there are a range of candidate expansion sizes where polyp detection sensitivity for clinically significant polyps is improved or unaffected. This range extends from zero to a maximum value of 5.6 mm. This is the optimal operating threshold for candidate expansion. It should be noted at this point that a candidate region of this size may be indicative of a much larger polyp.

A second plot, representing the false positive rate per patient versus the candidate expansion size threshold is illustrated in Figure 7.19. The false positive rate associated with standard candidates, i.e. where the candidate expansion size threshold is set to zero, is 17.5 per data set. As the size threshold increases the false positive rate decreases. The false positive rate at the maximum candidate expansion size threshold is 12.0 per data set. This represents a 31.4% reduction in false positive detections, compared with the use of standard candidates. The reduction in false positive detections is achieved without adversely affecting polyp detection sensitivity. In fact, the candidate expansion process slightly increases the sensitivity for the detection of large polyps. At the optimum threshold it was found that 10,900 (87%) of the 12,250 candidates were expanded and 114 (85%) of the 134 actual polyps were expanded. Therefore, candidate expansion successfully meets its original objectives in the large majority of cases.



**Figure 7.18:** A plot of polyp detection sensitivity versus increasing candidate expansion size threshold. Note that the system performance remains largely unaffected up to a size value of 5.6 mm.



**Figure 7.19:** A plot of false positive detection rate versus increasing candidate expansion size threshold. Note that the false positive detection rate at the optimum candidate expansion size threshold (5.6 mm) is 12.0 per data set.

### 7.5.8 Classifier Performance Characterisation

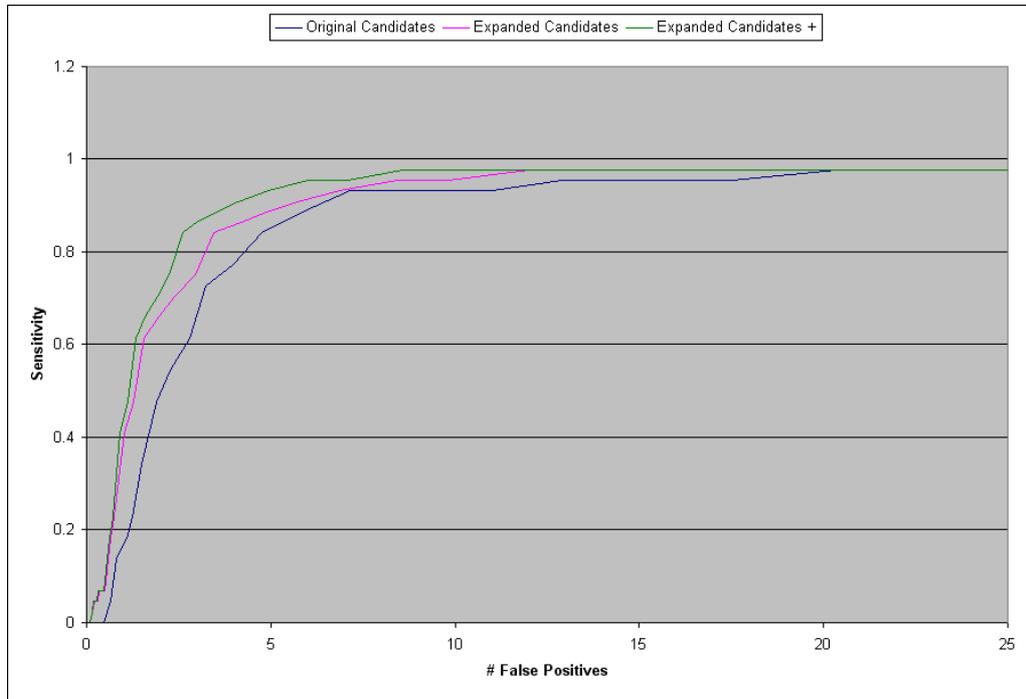
The final part of this evaluation involved characterising the performance of the classifier for both modes of operation. This enabled the optimal classifier performance to be determined and thus, enabled the most accurate comparison between the use of regular and expanded polyp candidates. The performance characterisation was performed using ROC analysis. A series of ROC curves were generated for clinically significant polyps (i.e. those  $\geq 5$  mm in size). This was achieved by moving the decision boundary of the classifier between the mean true polyp feature vector and the mean false positive feature vector. The ROC curve for medium sized polyps is illustrated in Figure 7.20, at 97% sensitivity, original candidates generated 20.4 false positives per data set whereas, expanded candidates generated 12.0 (a 41.2% reduction).

### 7.5.9 Additional Candidate Features

Recall that three additional features were calculated during the fingerprinting and candidate expansion process: Eccentricity, protrusiveness and expansion (see Section 5.5.3). An additional experiment was carried out to determine the maximum reduction in false positives achievable through the use of these features. In each case, the maximum feature value was calculated for the set of actual polyps that were used in this study. In the case of the protrusiveness feature, the minimum feature value was also calculated. The combination of minimum and maximum protrusiveness values enabled the identification of flat as well as highly curved false positive regions. The calculated values were used as upper and lower thresholds to filter out false positive detections prior to candidate classification. The use of these three features in addition to candidate expansion resulted in a false positive rate of 8.6 per data set at a sensitivity of 97% (see Figure 7.20). This represents a 57.8% reduction in false positives when compared with the use of standard polyp candidates.

## 7.6 Discussion

The three main contributions outlined in this thesis deal with colon centreline calculation, colon lumen subsegmentation and the computer assisted detection of colorectal polyps. This Chapter outlined the testing strategy that was employed to evaluate each of these contributions, and presented the results of this evaluation.



**Figure 7.20:** A ROC curve comparing polyp detection sensitivity and false positive detection rate for clinically significant polyps  $\geq 5$  mm in size. Note that at 97% sensitivity, original candidates generated 20.4 false positives per data set whereas the use of expanded candidates generated 12.0 (a 41.2% reduction). The use of the additional features calculated during the candidate expansion process enabled a further reduction in the false positive rate to a minimum of 8.6 per data set (a 57.8% reduction).

The centreline calculation technique has been demonstrated to be fast, accurate and robust. On average, the optimised approach to 3DTT was found to be over 32 times faster than the standard approach. It reduced the average time required for centreline calculation from 169 seconds to 5.284 seconds. This included the time required to load the data, detect the relevant endpoints, obtain the skeleton for each colonic segment and finally, to reduce the skeleton in order to yield the single path of voxels connecting the endpoints located in the rectum and the caecum. The accuracy of the technique has been confirmed using a series of computer generated synthetic colon phantoms with known centrelines. Endpoints were detected within approximately 2 mm of their actual locations and the standard deviation between the calculated centreline points and the actual centreline point was found to be less than 1 mm, with perfect matches in 33% of cases. This level of accuracy is comparable to, and in the majority of cases better than, previously published centreline calculation algorithms (see Table 3.1<sup>5</sup>). The centreline calculation technique has also

<sup>5</sup>Note that the performance difference between the previously documented centreline calculation algorithm (Sadleir & Whelan 2005), and the approach outlined in this thesis, is due to the inclusion of support for multisegment colons in the later.

been demonstrated to be robust using wide range of clinical test cases, including cases consisting of multiple colon segments due to instances of collapse or blockage. The centreline calculation process is fully automated requiring no user input.

The subsegmentation algorithm has been demonstrated to accurately identify the flexures and junctions of the colon based on an analysis of the centreline. On average, the automatically detected flexure and junction points were found to be within approximately 30 mm of points identified by a radiologist. This error is equivalent to approximately 1.5% of the length of the colon centreline. The flexure points were ultimately used to associate each colon surface vertex with a particular colon segment. A visual inspection of the resulting subsegmented colon surface confirmed the correctness of the outcome, even in cases where the colon consisted of multiple disconnected segments due to blockages or collapse. The subsegmentation information can ultimately be used in conjunction with CAD, or manually read CTC, to automatically assign a polyp to a particular colon segment. Thus it can be used to enable automated reporting or to automatically register polyps between prone and supine data sets (see Section 3.4).

The evaluation of the CAD-CTC contributions outlined in this thesis consisted of two main components. An initial study was carried out to determine the optimal approach to initial polyp candidate detection. This study involved the comparative evaluation of the features used by two of the main groups involved in surface curvature based CAD-CTC research: shape index/curvedness (UC) and mean/Gaussian curvature (NIH). Initial polyp candidate detection, using these two sets of features, was performed on ten sample data sets containing a total of 27 polyps. In both cases all of the clinically significant polyps were detected. However, the use of the shape/index curvedness features resulted in far fewer false positive detections. When the shape index/curvedness features were used to identify initial polyp candidate in the full test set, all of the clinically significant polyps, and 92% of the polyps < 5 mm in size, were detected. The number of false positive detections was 182. The initial polyp candidate detection stage of the CAD-CTC system outlined in this thesis is essentially a hybrid approach that combines the colon surface representation suggested by the NIH group with the candidate features suggested by the UC group.

A subsequent study was carried out to establish the effect of candidate expansion on polyp detection sensitivity. A total of 68 data sets were used in this evaluation. Fifty four of these data sets contained 42 clinically significant polyps (i.e.  $\geq 5$  mm in

size) and the remaining 14 data sets did not contain any polyps. Prone and supine data sets were treated as different patients to increase the number of individual polyps that were available for evaluation purposes. The justification behind this decision is illustrated in Figure 7.8. An initial evaluation of the candidate expansion technique found that it could not be applied to all candidates, particularly in the case of large candidates. A subsequent, more comprehensive evaluation, found that candidates up to 5.6 mm in size could be reliably expanded. At this setting, 87% of all candidates, and 85% of all actual polyps, could be expanded without adversely affecting polyp detection sensitivity. A comparison between standard and expanded candidates was carried out using ROC analysis. At a maximum sensitivity of 97%, the use of standard candidates resulted in 20.4 false positives per data set. At the same sensitivity expanded candidates resulted in 12.0 false positives per data set (a 41.2% reduction). When the additional features obtained during the candidate expansion and fingerprinting processes were used in conjunction with the expanded candidates the number of false positives per data set reduced to 8.6 (a 57.8% reduction when compared with standard candidates).

It should be noted that individual per-polyp sensitivities were reported in this study. This contrasts with other studies documented in the literature where combined per-polyp, or even combined per-patient sensitivities, are reported. In the case of combined per-polyp sensitivity, a polyp is deemed to be detected if it is identified in either the prone or the supine data sets. For example, if all polyps were detected in either the prone scan or the supine scan of a particular patient (but not in both scans), it would be acceptable to report a 100% combined per-polyp sensitivity. However, this would correspond to an individual per-polyp sensitivity of just 50%. It should also be noted that at an individual per-polyp sensitivity of 50%, the number of false positives generated using expanded candidates in conjunction with additional features is 1.32 per data set. This is comparable with the lowest false positive rates that have been reported in the literature (see Table 3.2).

# Chapter 8

## Conclusions and Future Work

This thesis introduced a novel methodology for automating various aspects of CTC, particularly the tasks of polyp detection and reporting. The formulation of the proposed methodology was based on an extensive review of the related literature and gave rise to a series of research contributions. The theoretical contributions involved:

- improving the performance of 3DTT based colon centreline calculation by using a novel LUT based optimisation approach. An evaluation of the optimised approach demonstrated that a significant improvement in performance could be achieved without compromising accuracy.
- accurately identifying the individual flexure points of the colon based on an analysis of the colon centreline. The detected flexure points were ultimately used to automatically divide the colon into its main anatomical regions, thus enabling automated reporting of CAD findings.
- determining the optimal approach for the detection of initial polyp candidates. This involved a comparative evaluation of different surface curvature characteristics associated with a polygonal mesh representation of the colonic mucosa.
- comparing the use of regular and expanded polyp candidates for the automatic detection of colorectal polyps. The main objective of this comparative study was to establish what effect the use of expanded candidates had on the rate of false positive detections.

The practical contributions related to the realisation of a comprehensive CAD-CTC system that incorporated all of the techniques and enhancements outlined in the

preceding Chapters of this thesis. The completed CAD-CTC system consisted of the following components:

1. An extensible framework developed specifically to enable rapid prototyping and evaluation of CTC algorithms.
2. A suite of tools, utilities and algorithms that were developed during the course this research programme to deal with various aspects of CTC, specifically:
  - (a) **Image interpretation** - accesses the DICOM encoded CTC data stored in the file system, extracts the pertinent patient details and acquisition information and reformats the relevant voxel data so that it is suitable for subsequent analysis.
  - (b) **Colon lumen segmentation** - identifies each region in a CTC data set that is attributable to the colon lumen. This operation essentially creates a binary volumetric representation of the colon lumen.
  - (c) **Centreline calculation** - reduces the volumetric representation of the colon lumen to a single path of centralised voxels spanning the entire length of the colon and terminating at endpoints located in the rectum and the caecum.
  - (d) **Colon subsegmentation** - identifies the flexure points based on an analysis of the colon centreline. The flexure points are then used to divide the colon lumen into five distinct anatomical segments.
  - (e) **Isosurface extraction** - converts the binary volumetric representation of the colon lumen resulting from segmentation into a polygonal mesh representation that is suitable for both processing and realtime visualisation.
  - (f) **Vertex classification** - calculates a series of properties for each vertex in the mesh representation of the colon surface. These properties include shape characteristics, density information and information regarding anatomical location.
  - (g) **Initial polyp candidate identification** - highlights patches of the colon surface that have polyplike properties. This process involves analysing the surface curvature characteristics of the entire colon surface.

- (h) **Polyp candidate expansion** - reformats polyp candidates and restructures them so that their shape conforms more closely to that of an actual polyp. Actual polyps retain their original polyplike characteristics, whereas, false positives expand to include more false positive characteristics.
- (i) **Polyp candidate classification** - determines whether an unknown polyp candidate is either a true polyp or a false positive and thus represents the final stage of the CAD-CTC system.

The completed system operates with minimal user interaction and has been validated using a range of clinical data.

## 8.1 Research Contributions

### 8.1.1 Major Contributions

The main contribution of this research is the novel approach to the computer aided detection of colorectal polyps at CTC. The primary objective of the CAD-CTC technique was to address the issue of false positive reduction. This was achieved by incorporating a novel intermediate stage into the conventional two stage model for polyp detection. The purpose of the intermediate stage was to expand polyp candidate regions so that their shape was more representative of actual polyps. The expansion process was intended to amplify the false positive characteristics of nonpolyps, while retaining the true positive characteristics of actual polyps. An evaluation of the expansion process demonstrated that it could be successfully applied to 87% of candidates without adversely affecting polyp detection sensitivity. Ultimately, the use of expanded candidates and their associated features was demonstrated to reduce false positive detections by 57.8%, from 20.4 per data set to 8.6 per data set, while maintaining a sensitivity level of 97% for the detection of clinically significant polyps. Other significant contributions include:

- A Java based medical image interpretation toolkit that supports images encoded using either DICOM or ANALYZE formats. This toolkit has been made available to the medical imaging community and is currently being used by a number of different groups involved in various aspects of medical imaging research.

- An efficient centreline calculation algorithm that identifies the central path through the colon, even in cases where the colon consists of multiple disconnected segments due to collapse or blockage. The accuracy of the algorithm has been verified using a series of computer generated colon phantoms with known centrelines. An evaluation of the technique using clinical cases demonstrated that centreline calculation typically requires approximately five seconds.
- A novel approach to identifying the individual flexure points along the centreline of the colon and thus enabling the identification of the main anatomical regions of the colon. The resulting information can be used to automatically assign a polyp to a particular colonic region. Thus enabling automated reporting in accordance with the recommendations made by Zalis et al. (2005b).
- A detailed comparative evaluation of the two main surface curvature based techniques for identifying initial polyp candidates. This evaluation demonstrated that the use of shape index and curvedness features generated fewer false positive detections than the use of mean and Gaussian curvature, while maintaining the same polyp detection sensitivity for clinically significant polyps.

### 8.1.2 Minor Contributions

Although the research outlined in this thesis dealt specifically with CTC, many of the techniques discussed have potential applications in other areas of medical imaging research. Consequently, other contributions of this research include:

- A demonstration of the utility of the Java programming language and its associated APIs for large scale medical imaging application development.
- A series of development techniques that can be used to streamline the implementation of a large scale medical imaging system, particularly in the case of a collaborative research environment.
- A range of medical image processing and visualisation utilities (e.g. the modified marching cubes algorithm) that have potential uses in other areas of medical imaging.

## 8.2 Future Work

It is clear from the reviews in Chapters 2 and 3 that CTC is a multifaceted and extremely active research area. This thesis addressed a number of different aspects of CTC research. However, numerous opportunities remain to be investigated.

### 8.2.1 Improved Colon Lumen Segmentation

The segmentation of the colon lumen is currently carried out manually. The user must select a seed point for each gas filled region that belongs to the colon. It is important to note that this is the only aspect of the entire system that requires user interaction. The reason for this, is to ensure that a perfect segmentation of the colon is available for analysis in the CAD stage of the system.

Automation of the segmentation stage would completely remove the need for user interaction. The VSG has developed a fully automated colon lumen segmentation technique that is capable of identifying 100% of colon lumen voxels where the colon is well distended, and 95% of colon lumen voxels where there is collapse or blockage, with only 1% of detected voxels attributable to extracolonic structures (Chowdhury et al. 2005a) (reviewed in Section 3.2.1). The use of this colon segmentation technique would fully automate the task of computer aided polyp detection.

An additional improvement would involve using an orally ingested contrast agent to opacify residual faecal material prior to a CTC examination. This makes the bowel preparation much more patient friendly and may remove the need for scanning in dual positions, thus significantly reducing the amount of radiation that the patient is exposed to. Zalis et al. have published a series of studies dealing with incremental improvements to the digital subtraction colonography technique. However, there are still areas of this technique that warrant further attention, specifically the removal of PVE artifacts at the air/contrast interface and the accurate reconstruction of the air/tissue boundary at the subtraction interface. The WRAMC protocol incorporates an orally ingested contrast enhanced bowel preparation and thus the WRAMC data sets could be used to enable research in this area.

### 8.2.2 Enhanced False Positive Reduction

The candidate expansion process was found to be applicable to 87% of polyp candidates. The reason it could not be applied to the remaining 13% was because in

certain cases polyp candidate patches extended beyond the boundary of the actual polyp. This was particularly true in the case of large candidate regions. A refinement of the initial candidate detection process, so that the candidate regions more accurately cover actual polyps, would enable candidate expansion to be applied to all candidates and enable a further reduction in false positive detections.

In addition, there are a number of consistent sources of false positive detection that commonly occur in CTC data sets. These include the rectal tube, the ileocecal valve and residual fecal material. The issue of residual fecal material has been addressed by the recent introduction of an oral contrast agent to opacify fecal material prior to CTC. The other sources, i.e. the rectal tube and the ileocecal valve, have been addressed by Iordanescu & Summers (2004) and Summers et al. (2004) respectively. Inclusion of these false positive reduction techniques in the CAD-CTC system outlined in this thesis would further reduce the number of false positive detections.

### 8.2.3 Further Evaluation

The evaluation of the techniques outlined in this thesis was limited by the number of suitable data sets provided by the Mater Misericordiae Hospital. Only 37 of the 250 patient data sets received contained polyps. Further evaluation would provide a more accurate indication of the performance of these techniques. The WRAMC has made their series of data sets available via the Internet for general usage. These data sets are all contrast enhanced. Once the updated segmentation scheme has been incorporated into the system, these data sets can be used for evaluation purposes. In addition, the American College of Radiology Imaging Network (ACRIN) CTC database will be archived and made available for academic and commercial groups to evaluate their CAD systems (Aberle et al. 2005). The inclusion of the WRAMC and ACRIN studies in the test database will facilitate a more detailed evaluation of all aspects of the system.

### 8.2.4 CTC Education

The need for reader training has been highlighted in the literature (Soto et al. 2005). A research project that is addressing this issue is currently underway at the VSG. This project involves the development of a remote access CTC training system that can be accessed over the Internet. The trainee marks and annotates suspected polyps and then submits their work for automatic evaluation (Luauté et al. 2006a). Many of

the techniques outlined in this thesis are compatible with, and can be incorporated into, this remote access CTC training system.

# Bibliography

- Aberle, D. R., Chiles, C., Gatsonis, C., Hillman, B. J., Johnson, C. D., McClennan, B. L., Mitchell, D. G., Pisano, E. D., Schnall, M. D. & Sorensen, A. G. (2005), ‘Imaging and cancer: Research strategy of the American College of Radiology Imaging Network’, *Radiology* **235**(3), 741–751.
- Acar, B., Beaulieu, C. F., Gokturk, S. B., Tomasi, C., Paik, D. S., Jeffrey Jr, R. B., Yee, J. & Napel, S. (2002), ‘Edge displacement field-based classification for improved detection of polyps in ct colonography’, *IEEE Transactions on Medical Imaging* **21**(12), 1461–1467.
- Acar, B., Napel, S., Paik, D. S., Li, P., Yee, J., Jeffrey Jr, R. B. & Beaulieu, C. F. (2001), Medial axis registration of supine and prone CT colonography data, *in* ‘Proceeding of the 23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society’, Istanbul, Turkey, 25-28 October 2001, pp. 2433–2436.
- ACR-NEMA (2004), *Digital imaging and communications in medicine*, National Electrical Manufacturers Association.
- Ahlquist, D. A., Skoletsky, J. E., Boynton, K. A., Harrington, J. J., Mahoney, D. W., Pierceall, W. E., Thibodeau, S. N. & Shuber, A. P. (2000), ‘Colorectal cancer screening by detection of altered human DNA in stool: Feasibility of a multitarget assay panel’, *Gastroenterology* **119**(5), 1219–1227.
- Bernarding, J., Thiel, A., Decker, I. & Tolxdorff, T. (2001), ‘Implementation of a dynamic platform-independent DICOM-server’, *Computer Methods and Programs in Biomedicine* **65**(1), 71–78.
- Bitter, I., Kaufman, A. E. & Sato, M. (2001), ‘Penalized-distance volumetric skeleton algorithm’, *IEEE Transactions on Visualization and Computer Graphics* **7**(3), 195–206.

## Bibliography

- Bitter, I., Sato, M., Bender, M., McDonnell, K. T., Kaufman, A. & Wan, M. (2000), CEASAR: A smooth, accurate and robust centerline extraction algorithm, *in* 'Proceedings of IEEE Visualization 2000', Salt Lake City, UT, USA, 8-13 October 2000, pp. 45–52.
- Bouix, S., Siddiqi, K. & Tannenbaum, A. (2005), 'Flux driven automatic centreline extraction', *Medical Image Analysis* **9**, 209–221.
- Burkitt, D. P. (1971), 'Epidemiology of cancer of the colon and rectum', *Cancer* **28**, 3–13.
- Callstrom, M. R., Johnson, C. D., Fletcher, J. G., Reed, J. E., Ahlquist, D. A., Harmesen, W. S., Tait, K., Wilson, L. A. & Corcoran, K. E. (2001), 'CT colonography without cathartic preparation: feasibility study', *Radiology* **219**(3), 693–698.
- Campbell, S. R. & Summers, R. M. (2004), Analysis of kernal method for surface curvature estimation, *in* 'Computer Assisted Radiology and Surgery (CARS), 18th International Congress and Exhibition', Chicago, USA, June 23rd - 26th, 2004. (available in: *International Congress Series* **1268**, 999-1003).
- Chen, D., Li, B., Liang, Z., Wan, M., Kaufman, A. E. & Wax, M. (2000*b*), Tree-branch searching multiresolution approach to skeletonization for virtual endoscopy, *in* K. M. Hanson, ed., 'Proc. SPIE Medical Imaging 2000: Image Processing', Vol. 3979, pp. 726–734.
- Chen, D. Q., Liang, Z. R., Wax, M. R., Li, L. H. & Kaufman, A. E. (2000*a*), 'A novel approach to extract the colon lumen from CT images for virtual colonoscopy', *IEEE Transactions on Medical Imaging* **19**(12), 1220–1226.
- Chen, S. C., Lu, D. S. K., Hecht, J. R. & Kadell, B. M. (1999), 'CT colonography: value of scanning in both the supine and prone positions', *American Journal of Roentgenology* **172**(3), 595–599.
- Chiou, R. C. H., Kaufman, A. E., Liang, Z. R., Hong, L. C. & Achniotou, M. (1999), 'An interactive fly-path planning using potential fields and cell decomposition for virtual endoscopy', *IEEE Transactions on Nuclear Science* **46**(4), 1045–1049.
- Chiou, R., Kaufman, A., Liang, Z., Hong, L. & Achniotou, M. (1998), Interactive

## Bibliography

- Path Planning for Virtual Colonoscopy, *in* 'Proceedings of IEEE Nuclear Science Symposium and Medical Image Conference', Vol. 3, pp. 2069–2072.
- Chowdhury, T. A., Ghita, O. & Whelan, P. F. (2005*b*), A statistical approach for robust polyp detection in CT colonography, *in* 'Proceedings of the 27th Annual International Conference of the IEEE Engineering in Medicine and Biology Society', Shanghai, China, 1-4 September 2005.
- Chowdhury, T. A., Sadleir, R. J. T., Whelan, P. F., Moss, A., Varden, J., Short, M., Fenlon, H. & MacMathuna, P. (2004), The impact of radiation dose on imaged polyp characteristics: Experiments with a synthetic phantom, *in* 'Proceedings of the Association of Physical Scientists in Medicine, 2004 Annual Scientific Meeting', Dublin, June 11th, 2005.
- Chowdhury, T. A., Whelan, P. F., Fenlon, H. & MacMathuna, P. (2005*c*), Evaluation of radiation dose on automatic polyp detection at CT colonography: Experiments with a synthetic phantom, *in* 'Proceedings of the Association of Physical Scientists in Medicine, 2005 Annual Scientific Meeting', Galway, Ireland, February 2005.
- Chowdhury, T. A., Whelan, P. F. & Ghita, O. (2005*a*), A Method for automatic segmentation of collapsed colons at CT colonography, *in* '2nd Indian International Conference on Artificial Intelligence (IICAI-05)', Pune, India, December 20-22 2005, pp. 3517–3532.
- Coady-Fariborzian, L., Angel, L. P. & Procaccino, J. A. (2004), 'Perforated colon secondary to virtual colonoscopy: Report of a case', *Diseases of the Colon & Rectum* **47**(7), 1247–1249.
- Coin, C. G., Wollett, F. C., Coin, J. T., Rowland, M., DeRamos, R. K. & Dadrea, R. (1983), 'Computerized radiology of the colon: a potential screening technique', *Computerized Radiology* **7**(4), 215–221.
- Cotton, P. B., Durkalski, V. L., Pineau, B. C., Palesch, Y. Y., Mauldin, P. D., Hoffman, B., Vining, D. J., Small, W. C., Affronti, J., Rex, D., Kopecky, K. K., Ackerman, S., Burdick, J. S., Brewington, C., Turner, M. A., Zfass, A. & Wright, A. R. (2004), 'Computer tomographic colonography (virtual colonoscopy): A multicenter comparison with standard colonoscopy for detection of colorectal neoplasia', *The Journal of the American Medical Association* **291**(14), 1713–1719.

## Bibliography

- Dachman, A., Nappi, J., Frimmel, H. & Yoshida, H. (2002), ‘Sources of false positives in computerized detection of polyps in CT colonography’, *Radiology* **225**, 303–304. (abstract).
- Deriche, R. (1987), ‘Using Canny’s criteria to derive a recursively implemented optimal edge detector’, *International Journal of Computer Vision* **1**(2), 167 – 187.
- Deschamps, T. & Cohen, L. (2000), Minimal path in 3D images and application to virtual endoscopy, *in* ‘European Conference on Computer Vision, (ECCV’00)’, Dublin, June 2000.
- Deschamps, T. & Cohen, L. (2001), ‘Fast extraction of minimal paths in 3D images and applications to virtual endoscopy’, *Medical Image Analysis* **5**(4), 281–299.
- Dijkstra, E. W. (1959), ‘A note on two problems in connexion with graphs’, *Numerische Mathematik* **1**, 269–271.
- Erberich, S. G., Song, K., Arakawa, H., Huang, K., Richard, W., Hoo, K. S. & Loo, D. W. (1997), ‘Knowledge-based lung nodule detection for helical CT’, *Radiology* **205P**, 617. (abstract).
- Fenlon, H. & Ferrucci, J. (1997), ‘Virtual Colonoscopy: What will the issues be?’, *American Journal of Roentgenology* **169**(2), 453–458.
- Fenlon, H. M., Nunes, D. P., Schroy, P. C., Barish, M. A., Clarke, P. D. & Ferrucci, J. T. (1999), ‘A comparison of virtual and conventional colonoscopy for the detection of colorectal polyps’, *New England Journal of Medicine* **341**(20), 1496–1503. Erratum, *New England Journal of Medicine* 2000 (342) 524.
- Fernandez-Bayo, J., Barbero, O., Rubies, C., Sentis, M. & Donso, L. (2000), ‘Distributing medical imaging with internet technologies: a DICOM web server and DICOM java viewer’, *Radiographics* **20**, 581–591.
- Frimmel, H., Näppi, J. & Yoshida, H. (2004), ‘Fast and robust computation of colon centerline in CT colonography’, *Medical Physics* **31**(11), 3046–3056.
- Frimmel, H., Näppi, J. & Yoshida, H. (2005), ‘Centerline-based colon segmentation for CT colonography’, *Medical Physics* **32**(8), 2665–2672.

## Bibliography

- Ge, Y., Stelts, D. R., Wang, J. & Vining, D. J. (1999), 'Computing the centerline of a colon: a robust and efficient method based on 3D skeletons', *Journal of Computer Assisted Tomography* **23**(5), 786–794.
- Ge, Y., Stelts, D. & Vining, D. (1996), 3D skeleton for virtual colonoscopy, *in* K. Höhne & R. Kikinis, eds, 'Lecture Notes in Computer Science, VBC '96 - 4th International Conference on Visualization in Biomedical Computing', Vol. 1131, September 22nd - 25th, Hamburg, Germany, pp. 449–454.
- Gluecker, T. M., Johnson, C. D., Harmsen, W. S., Offord, K. P., Miller, A. M., Wilson, L. A. & Alquist, D. A. (2003), 'Colorectal cancer screening with CT colonography, colonoscopy, and double-contrast barium enema examination: Prospective assessment of patient perceptions and preferences', *Radiology* **227**(2), 378–381.
- Göktürk, S. B., Tomasi, C., Acar, B., Beaulieu, C. F., Paik, D. S., Jeffrey, R. B., Yee, J. & Napel, S. (2001), 'A statistical 3D pattern processing method for computer-aided detection of polyps in CT colonography', *IEEE Transactions on Medical Imaging* **20**(12), 1251–1260.
- Gonzalez, R. C. & Woods, R. E. (1992), *Digital image processing*, Addison Wesley Longman.
- Hara, A., Johnson, C., MacCarty, R. & Welch, T. (2000), 'Incidental extracolonic findings at CT colonography', *Radiology* **215**(2), 353–357.
- Hara, A., Johnson, C., Reed, J., Ehman, R. & Ilstrup, D. (1996), 'Colorectal polyp detection with CT colography: two- versus three-dimensional techniques. Work in progress', *Radiology* **200**(1), 49–54.
- Hecht, J., Chen, S., Lu, D. & Gralnek, I. (1997), 'A patient satisfaction study of CT colonography (virtual colonoscopy) and colonoscopy', *Gastrointestinal Endoscopy* **45**(4), Suppl:333.
- Hong, L., Kaufman, A., Wei, Y.-C., Viswambharan, A., Wax, M., Liang, Z., Loew, M. & Gershon, N. (1995), 3D virtual colonoscopy, *in* 'Proceedings 1995 Biomedical Visualisation', pp. 26–32.
- Hong, L., Muraki, S., Kaufman, A., Bartz, D. & Taosong, H. (1997), Virtual Voyage:

## Bibliography

- Interactive Navigation in the Human Colon, *in* ‘Computer Graphics Proceedings, SIGGRAPH97’, pp. 27–34.
- Hounsfield, G. (1973), ‘Computerized transverse axial scanning (tomography): Part I. Description of system’, *British Journal of Radiology* **46**, 1016–1022.
- Hunt, G. W., Hemler, P. F. & Vining, D. J. (1997), Automated virtual colonoscopy, *in* Y. Kim, ed., ‘Proc. SPIE Medical Imaging 1997: Image Display’, Vol. 3031, pp. 535–541.
- Iordanescu, G., Pickhardt, P. J., Choi, J. R. & Summers, R. M. (2005), ‘Automated seed placement for colon segmentation in computed tomography colonography’, *Academic Radiology* **12**(2), 182–190.
- Iordanescu, G. & Summers, R. (2004), ‘Reduction of false positives on the rectal tube in computer-aided detection for CT colonography’, *Medical Physics* **31**(10), 2855–2862.
- Iordanescu, G. & Summers, R. M. (2003), ‘Automated centreline for computed tomography’, *Academic Radiology* **10**(11), 1291–1301.
- Jensch, S., Gelder, R. V., Venema, H., Laméris, J. & Bossuyt, P. (2004), Effective radiation doses in CT colonography: Results of an inventory among research institutions, *in* ‘Annual meeting of the Radiological Society of North America’, Chicago, Nov 28th - Dec 3rd.
- Jerebko, A. K., Malley, J. D., Franaszek, M. & Summers, R. M. (2003*b*), ‘Multiple neural network classification scheme for detection of colonic polyps in CT colonography data sets’, *Academic Radiology* **10**(2), 154–160.
- Jerebko, A. K., Malley, J. D., Franaszek, M. & Summers, R. M. (2005), ‘Support vector machines committee classification method for computer-aided polyp detection in CT colonography’, *Academic Radiology* **12**(4), 479–486.
- Jerebko, A. K., Summers, R. M., Malley, J. D., Franaszek, M. & Johnson, C. D. (2003*a*), ‘Computer-assisted detection of colonic polyps with CT colonography using neural networks and binary classification trees’, *Medical Physics* **30**(1), 52–60.

## Bibliography

- Johnson, C., Hara, A. & Reed, J. (1998), ‘Virtual endoscopy: what’s in a name?’, *American Journal of Roentgenology* **171**(5), 1201–2.
- Kang, D.-G. & Ra, J. (2005), ‘A new path planning algorithm for maximizing visibility in computed tomography colonography’, *IEEE Transactions of Medical Imaging* **24**(8), 957–968.
- Kiss, G., Van Cleynenbreugel, J., Thomeer, M., Suetens, P. & Marchal, G. (2002), ‘Computer-aided diagnosis in virtual colonography via combination of surface normal and sphere fitting methods’, *European Radiology* **12**(1), 77–81.
- Koenderink, J. J. (1990), *Solid Shape*, MIT Press.
- Levin, B., Brooks, D., Smith, R. A. & Stone, A. (2003), ‘Emerging technologies in screening for colorectal cancer: CT colonography, Immunochemical fecal occult blood tests, and stool screening using molecular markers’, *CA* **53**(2), 44–55.
- Li, H., Pineau, B. & Santago, P. (2005), ‘Efficient computerized polyp detection for CT colonography’, *Journal of Digital Imaging* **18**(1), 55–65.
- Li, H. & Santago, P. (2005), ‘Automatic colon segmentation with dual scan CT colonography’, *Journal of Digital Imaging* **18**(1), 42–54.
- Li, P., Napel, S., Acar, B., Paik, D. S., Jeffrey, R. B. & Beaulieu, C. F. (2004), ‘Registration of central paths and colonic polyps between supine and prone scans in computed tomography colonography: Pilot study’, *Medical Physics* **31**(10), 2912 – 2923.
- Ling, S., Summers, R., Loew, M., NcCollough, C. & Johnson, C. (2002), ‘Computer-aided detection of polyps in a colon phantom: effect of scan orientation, polyp size, collimation, and dose’, *Journal of Computer Assisted Tomography* **26**(6), 1013–8.
- Lohmann, G. (1998), *Volumetric Image Analysis*, Wiley.
- Lorensen, W. E. & Cline, H. E. (1987), Marching cubes: A high resolution 3D surface construction algorithm, *in* ‘Proceedings of the 14th annual conference on computer graphics’, pp. 163 – 169.
- Luauté, V., Sadleir, R. J. T. & Whelan, P. F. (2006*a*), A remote access training system for the detection of colorectal polyps at computed tomography colonography, *in* ‘European Congress of Radiology (ECR)’, Vienna, Austria, March 3rd - 7th.

## Bibliography

- Luauté, V., Sadleir, R. J. T. & Whelan, P. F. (2006b), An automatic evaluation strategy for a remote access CT colonography training system, *in* 'Biosignal 2006 - The 18th International EURASIP (European Association for Signal, Speech and Image Processing) Conference', Brno, Czech Republic, June 28th - 30th.
- Luboldt, W., Bauerfeind, P., Steiner, P., Fried, M., Krestin, G. & Debatin, J. (1997), 'Preliminary assessment of three-dimensional magnetic resonance imaging for various colonic disorders', *The Lancet* **349**(9061), 1288–1291.
- Luboldt, W., Mann, C., Tryon, C., Vonthein, R., Stueker, D., Kroll, M., Luz, O., Claussen, C. & Vogl, T. (2002), 'Computer-aided diagnosis in contrast-enhanced CT colonography: an approach based on contrast', *European Radiology* **12**(9), 2236–2241.
- Luboldt, W., Tryon, C., Kroll, M., Toussaint, T., Holzer, K., Hoepffner, N. & Vogl, T. (2005), 'Automated mass detection in contrast-enhanced CT colonography: an approach based on contrast and volume', *European Radiology* **15**(2), 247–253.
- Macari, M., Berman, P., Dicker, M., Milano, A. & Megibow, A. (1999), 'Usefulness of CT colonography in patients with incomplete colonoscopy', *American Journal of Roentgenology* **173**(3), 561–564.
- Macari, M., Bini, E., Xue, X., Milano, A., Katz, S., Resnick, D., Chandarana, H., Krinsky, G., Klingensbeck, K., Marshall, C. & Megibow, A. (2002), 'Colorectal neoplasms: Prospective comparison of thin-section low-dose multi-detector row CT colonography and conventional colonoscopy for detection', *Radiology* **224**(2), 383–392.
- Macari, M., Lavelle, M., Pedrosa, I., Milano, A., Dicker, M., Megibow, A. & Xue, X. (2001), 'Effect of different bowel preparations on residual fluid at CT colonography', *Radiology* **218**(1), 561–564.
- Malone, K. M., Molloy, D., Satck, J., Conroy, D., Egan, M., Murphy, D., Sadleir, R., Whelan, P., Saykin, A. & Wishart, H. (2003), A functional brain magnetic resonance imaging (fMRI) study of strategic thinking/problem solving and hope in healthy volunteers, *in* 'Proceedings of the 43th Annual New Clinical Drug Evaluation Unit (NCDEU) Meeting', Boca Raton, Florida, USA, May 26th - 28th.

## Bibliography

- Mani, A., Napel, S., Paik, D. S., Jeffrey, R. B., Yee, J., Olcott, E. W., Proesch, R., Davila, M., Schraedley-Desmond, P. & Beaulieu, C. F. (2004), 'Computed tomography colonography - Feasibility of computer-aided polyp detection in a "first reader" paradigm', *Journal of Computer Assisted Tomography* **28**(3), 318–326.
- Masutani, Y., Yoshida, H., MacEneaney, P. M. & Dachman, A. H. (2001), 'Automated segmentation of colonic walls for computerized detection of polyps in CT colonography', *Journal of Computer Assisted Tomography* **25**(4), 629–638.
- McDermott, R., McFarland, E., Brink, J., Risvedt, S., Manias, C. & Littenberg, B. (2001), 'Prospective comparison of air and CO<sub>2</sub> insufflation techniques at CT colonography: evaluation of image quality and patient reactions', *Radiology* **221**(Suppl), 578–579. (abstract).
- McFarland, E. G., Wang, G., Brink, J. A., Balfe, D. M., Heiken, J. P. & Vannier, M. W. (1997), 'Spiral computer tomographic colonography: determination of the central axis and digital unraveling of the colon', *Academic Radiology* **4**(5), 367–373.
- Mikolajczyk, K., Szabatin, M., Rudnicki, P., Grodzki, M. & Burger, C. (1998), 'A Java environment for medical image data analysis: Initial application for brain PET quantitation', *Medical Informatics* **23**(3), 207–214.
- Mildenberger, P., Eichelberg, M. & Martin, E. (2002), 'Introduction to the DICOM standard', *European Radiology* **12**(4), 920–927.
- Monga, O. & Benayoun, S. (1995), 'Using partial derivatives of 3D images to extract typical surface features', *Computer Vision and Image Understanding* **61**(2), 171–189.
- Morrin, M., Farrell, R., Kruskal, J. & LaMont, J. (1999), 'Virtual Colonoscopy: a kinder, gentler colorectal cancer screening test?', *Lancet* **354**(9184), 1048–1049.
- Morrin, M., Farrell, R., Kruskal, J., Reynolds, K., McGee, J. & Raptopoulos, V. (2000), 'Utility of intravenously administered contrast material at CT colonography', *Radiology* **217**(3), 765–71.

## Bibliography

- Näppi, J., Akihiko, A., Frimmel, H., Dachman, A. & Yoshida, H. (2005a), ‘Region-based supine-prone correspondence for the reduction of false-positive CAD polyp candidates in CT colonography’, *Academic Radiology* **12**(6), 695–707.
- Näppi, J., Dachman, A. H., MacEneaney, P. & Yoshida, H. (2002), ‘Automated knowledge-guided segmentation of colonic walls for computerized detection of polyps in CT colonography’, *Journal of Computer Assisted Tomography* **26**(4), 493–504.
- Näppi, J., Frimmel, H., Okamura, A., Dachman, A. H. & Yoshida, H. (2004a), Region-based prone-supine correspondence for reduction of false positives in CAD for CT colonography, in ‘Computer Assisted Radiology and Surgery (CARS), 18th International Congress and Exhibition’, Chicago, USA, June 23rd - 26th, 2004. (available in: *International Congress Series* **1268**, 993-998).
- Näppi, J., Frimmel, H. & Yoshida, H. (2005b), ‘Virtual endoscopic visualization of the colon by shape-scale signatures’, *IEEE Transactions on Information Technology in Biomedicine* **9**(1), 120–131.
- Näppi, J. J., Frimmel, H., Dachman, A. H. & Yoshida, H. (2004b), ‘Computerized detection of colorectal masses in CT colonography based on fuzzy merging and wall-thickening analysis’, *Medical Physics* **31**(4), 860–872.
- Näppi, J. & Yoshida, H. (2002), ‘Automated detection of polyps with CT colonography: Evaluation of volumetric features for reduction of false-positive findings’, *Academic Radiology* **9**(4), 386–397.
- Näppi, J. & Yoshida, H. (2003), ‘Feature-guided analysis for reduction of false positives in CAD of polyps for computed tomographic colonography’, *Medical Physics* **30**(7), 1592–1601.
- NCRI (2005), *Cancer in Ireland, 1994-2001: Incidence, mortality and treatment*.
- Nielson, G. M. (2003), ‘On marching cubes’, *IEEE Transactions on Visualisation and Computer Graphics* **9**(3), 283–297.
- Okamura, A., Dachman, A., Parsad, N., Näppi, J. & Yoshida, H. (2004), Evaluation of the effect of CAD on observers’ performance in detection of polyps in CT colonography, in ‘Computer Assisted Radiology and Surgery (CARS), 18th

## Bibliography

- International Congress and Exhibition', Chicago, USA, June 23rd - 26th, 2004. (available in: *International Congress Series* **1268**, 989-992).
- Paik, D., Beaulieu, C., Jeffrey, R., Yee, J., Steinaur, A. & Napel, S. (2000), 'Computer aided detection of polyps in CT colonography: method and free-response ROC evaluation of performance', *Radiology* **217**(Suppl), 370 (704). (abstract).
- Paik, D., Beaulieu, C., Rubin, G., Acar, B., Jeffrey, R., Yee, J., Dey, J. & Napel, S. (2004), 'Surface Normal Overlap: A Computer-Aided Detection Algorithm With Application to Colonic Polyps and Lung Nodules in Helical CT', *IEEE Transactions on Medical Imaging* **23**(6), 661-675.
- Paik, D. S., Beaulieu, C. F., Jeffrey, R. B., Rubin, G. D. & Napel, S. (1998), 'Automated flight path planning for virtual endoscopy', *Medical Physics* **25**(5), 629-637.
- Patel, K. & Hoffman, N. (2001), 'The anatomical distribution of colorectal polyps at colonoscopy', *Journal of Clinical Gastroenterology* **33**(3), 222-225.
- Pennebaker, W. B. & Mitchell, J. L. (1993), *JPEG still image data compression standard*, Chapman & Hall, New York (1st edition).
- Pickhardt, P. (2003), 'Three-dimensional endoluminal CT colonography (virtual colonoscopy): Comparison of three commercially available systems', *American Journal of Roentgenology* **181**(6), 1599-1606.
- Pickhardt, P., Choi, J. R., Hwang, I., Butler, J., Puckett, M., Hildebrandt, H., Wong, R., Nugent, P., Mysliwiec, P. & Schindler, W. (2003), 'Computed Tomographic Virtual Colonoscopy to Screen for Colorectal Neoplasia in Asymptomatic Adults', *The New England Journal of Medicine* **349**(23), 2191-2200.
- Pickhardt, P. J., Lee, A. D., McFarland, E. G. & Taylor, A. J. (2005), 'Linear polyp Measurement at CT colonography: In vitro and in vivo comparison of two-dimensional and three-dimensional displays', *Radiology* **236**(3), 872-878.
- Pickhardt, P., Nugent, P., Choi, J. H. & Schindler, W. (2004), CT colonoscopy for primary colorectal cancer screening: Effect of patient age, patient gender, and reader experience on reader performance, *in* 'Annual meeting of the Radiological Society of North America', Chicago, Nov 28th - Dec 3rd.

## Bibliography

- Pochaczewsky, R. (2002), 'Digital subtraction bowel cleansing in CT colonography', *American Journal of Roentgenology* **178**(1), 241. (correspondence).
- Post, M. J. (1994), Minimum spanning ellipsoids, *in* 'Proc 16th Annual ACM Symposium on Theory of Computing', pp. 108–116.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T. & Flannery, B. P. (1992), *Numerical recipes in C: The art of scientific computing*, 2nd edn, Cambridge University Press.
- Rembacken, B., Fujii, T., Cairns, A., Dixon, M., Yoshida, S., Chalmers, D. & Dixon, A. (2000), 'Flat and depressed colonic neoplasms: a prospective study of 1000 colonoscopies in the UK', *Lancet* **355**(9211), 1211–1214.
- Robinson, K., Whelan, P. F. & Stack, J. (2002), Segmentation of the biliary tree in MRCP data, *in* 'Proceedings of OPTO-Ireland: SPIE's Regional Meeting on Optoelectronics, Photonics and Imaging', Galway, Ireland, Sept 5th - 6th, pp. 192–200.
- Rockey, D., Kock, J., Cello, J., Sanders, L. & McQuaid, K. (1998), 'Relative frequency of upper gastrointestinal and colonic lesions in patients with positive fecal occult-blood tests', *New England Journal of Medicine* **339**(3), 153–159.
- Sadleir, R. J. T. & Whelan, P. F. (2002), Colon centerline calculation using optimised 3D topological thinning, *in* 'Proc 1st International conference on 3D data processing, visualisation and transmission', Padova, June 19th - 21st, pp. 800–803.
- Sadleir, R. J. T. & Whelan, P. F. (2005), 'Fast colon centreline calculation using optimised 3D topological thinning', *Computerized Medical Imaging and Graphics* **29**(4), 251–258.
- Sadleir, R. J. T., Whelan, P. F., Bruzzi, J. F., Moss, A. C., Fenlon, H. M. & MacMathuna, P. (2002), A novel technique for identifying the individual regions of the human colon at CT colonography, *in* 'IEE - Medical Applications of Signal Processing', London, October 7th, 2002, pp. 8/1–8/5.
- Sadleir, R. J. T., Whelan, P. F., Bruzzi, J. F., Moss, A. C., MacMathuna, P. & Fenlon, H. M. (2002a), 'A portable toolkit for providing straightforward access to medical image data', *Radiology* **225**, Supplement P 762.

## Bibliography

- Sadleir, R. J. T., Whelan, P. F., MacMathuna, P. & Fenlon, H. M. (2004a), ‘A portable toolkit for providing straightforward access to medical image data’, *Radiographics* **24**(4), 1193–1202.
- Sadleir, R. J. T., Whelan, P. F., Sezille, N., Chowdhury, T. A., Bruzzi, J., Moss, A., MacMathuna, P. & Fenlon, H. (2004b), Automatic detection of colorectal polyps at CT colonography using shape information, *in* ‘Association of Physical Scientists in Medicine Annual Scientific Meeting’, Dublin, June 11th.
- Samara, Y., Fiebich, M., Dachman, A. H., Doi, K. & Hoffmann, K. R. (1998), Automated centerline tracking of the human colon, *in* K. M. Hanson, ed., ‘Proc. SPIE Medical Imaging 1998: Image Processing’, Vol. 3338, pp. 740–746.
- Samara, Y., Fiebich, M., Dachman, A. H., Kuniyoshi, J. K., Doi, K. & Hoffmann, K. R. (1999), ‘Automated calculation of the centerline of the human colon on CT images’, *Academic Radiology* **6**(6), 352–359.
- Sato, M., Bitter, I., Bender, M., Kaufman, A. & Nakajima, M. (2000), TEASAR: Tree-Structure Extraction Algorithm for Accurate and Robust Skeletons, *in* ‘Pacific Graphics 2000’, pp. 281–289.
- Sato, M., Lakare, S., Wan, M., Kaufman, A. & Wax, M. (2001), ‘An automatic colon segmentation for 3D virtual colonoscopy’, *IEICE Transactions on Information and Systems* **E84D**(1), 201–208.
- Schroeder, W. J., Avila, L. S. & Hoffman, W. (2000), ‘Visualizing with VTK: A tutorial’, *IEEE Computer Graphics and Applications* **20**(5), 20–27.
- Seidel, P. (1990), Linear programming and convex hulls made easy, *in* ‘Proc. 6th Annual ACM Symposium on Computational Geometry’, pp. 211–215.
- Shapero, T., Alexander, P., Hoover, J., Burgis, E. & Schabas, R. (2001), ‘Colorectal cancer screening: Video reviewer flexibly sigmoidoscopy by nurse endoscopists - A Canadian community-based perspective’, *Canadian Journal of Gastroenterology* **15**(7), 441–445.
- Shepard, D. (1968), A two-dimensional interpolation function for irregularly-spaced data, *in* ‘Proceedings of the 1968 23rd ACM national conference’, pp. 517 – 524.

## Bibliography

- Sheppard, D., Iyer, R., Herrin, D. & Charnsangavej, C. (1999), 'Subtraction CT colonography: feasibility in an animal model', *Clinical Radiology* **54**(2), 126–132.
- Soto, J. A., Barish, M. A. & Yee, J. (2005), 'Reader Training in CT colonography: How much is enough?', *Radiology* **237**(1), 26–27. (editorial).
- Soto, J. A., Barish, M. & Ferrucci, J. (2004), CT colonography interpretation: Guidelines for interpretation, *in* 'Annual meeting of the Radiological Society of North America', Chicago, Nov 28th - Dec 3rd.
- Summers, R., Beaulieu, C., Pusanik, L., Malley, J., Jeffrey, R., Glazer, D. & Napel, S. (2000), 'Automated polyp detector for CT colonography: feasibility study', *Radiology* **216**(1), 284–290.
- Summers, R., Jerebko, A., Franaszek, M., Malley, J. & Johnson, C. (2002), 'Colonic polyps: complementary role of computer-aided detection in CT colonography', *Radiology* **225**(2), 391–399.
- Summers, R., Johnson, C., Pusanik, L., Malley, J., Youssef, A. & Reed, J. (2001), 'Automated polyp detection at CT colonography: feasibility assessment in a human population', *Radiology* **219**(1), 51–59.
- Summers, R. M. (2000), *Handbook of medical imaging processing and analysis*, Academic Press, chapter 45 - Morphometric methods for virtual endoscopy, pp. 747–755.
- Summers, R. M. (2003), Current concepts and future directions in CT colonography computer-aided detection, *in* 'Annual meeting of the Radiological Society of North America', Chicago, Nov 28th - Dec 3rd.
- Summers, R. M., Franaszek, M., Miller, M. T., Pickhardt, P. J., Choi, J. R. & Schindler, W. R. (2005), 'Computer-aided detection of polyps on oral contrast-enhanced CT colonography', *American Journal of Roentgenology* **184**(1), 105–108.
- Summers, R. M., Selbie, W. S., Malley, J. D., Ousanik, L. M., Dwyer, A. J., Courcoutsakis, N. A., Shaw, D. J., Kleiner, D. E., Sneller, M. C., Langford, C. A., Holland, S. M. & Shelhamer, J. H. (1998), 'Polypoid lesions of airways: early experience with computer-assisted detection using virtual bronchoscopy and surface curvature', *Radiology* **208**(2), 331–337.

## Bibliography

- Summers, R., Yao, J. & Johnson, C. (2004), ‘CT colonography with computer-aided detection: automated recognition of ileocecal valve to reduce number of false-positive detections’, *Radiology* **233**(1), 266–272.
- Svensson, M., Svensson, E. & Hellström, M. (2000), Patient acceptance of CT colonography (virtual colonoscopy) as compared with conventional colonoscopy, *in* ‘Proceedings of the Second International Symposium on Virtual Colonoscopy’, Boston, October 16th-17th, 2000, p. 109. (abstract).
- Taylor, S. A., Halligan, S., Burling, D., Roddie, M. E., Honeyfield, L., McQuillian, J., Amin, H. & Dehmeshki, J. (2006), ‘Computer-assisted reader software versus expert reviewers for polyp detection on CT colonography’, *American Journal of Roentgenology* **186**(3), 696–702.
- Thirion, J.-P. & Gourdon, A. (1993), The marching lines algorithm: new results and proofs, Technical Report 1881-1, INRIA.
- Thirion, J.-P. & Gourdon, A. (1995), ‘Computing the differential characteristics of iso-intensity surfaces’, *Computer Vision and Image Understanding* **61**(2), 190–202.
- Tran, D. Q., Rosen, L., Kim, R., Riether, R. D., Stasik, J. J. & Khubchandani, I. T. (2001), ‘Actual colonoscopy: what are the risks of perforation?’, *American Surgeon* **67**(9), 845–847.
- Truyen, R., Deschamps, T. & Cohen, L. (2001), Clinical evaluation of an automatic path tracker for virtual colonoscopy, *in* ‘Medical Image Computing and Computer Assisted Intervention - MICCAI 2001’, pp. 85–93.
- Tsao, Y. F. & Fu, K. S. (1981), ‘A Parallel thinning algorithm for 3-D pictures’, *Computer Graphics and Image Processing* **17**, 315–331.
- Vining, D. J., Ge, Y., Ahn, D. K. & Stelts, D. R. (1998), Virtual Colonoscopy with computer assisted polyp detection, *in* ‘Proceedings of the first international workshop on computer aided diagnosis’, Elsevier Science B.V., pp. 445–452.
- Vining, D. J., Gelfand, D. W., Bechtold, R. E., Scharling, E. S., Grishaw, E. K. & Shifrin, R. Y. (1994), ‘Technical feasibility of colon imaging with helical CT and virtual reality’, *American Journal of Roentgenology* **162**, Suppl:104.

## Bibliography

- Vining, D. J., Stelts, D. R., Ahn, D. K., Hemler, P., Ge, Y., Hunt, G., Siege, C., McCorquodale, D., Sarojak, M. & Ferretti, G. (1997), FreeFlight: a virtual endoscopy system, *in* ‘First Joint Conference, Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics’, Vol. LNCS 1205, pp. 413–416.
- Wallace, G. K. (1991), ‘The JPEG still image compression standard’, *Communications of the ACM*, 1991 **34**(4), 30–44.
- Wan, M., Dachille, F. & Kaufman, A. (2001), Distance-field based skeletons for virtual navigation, *in* ‘12th IEEE Visualization 2001 (VIS’01)’, San Diego, CA.
- Wan, M., Liang, Z., Ke, Q., Hong, L., Bitter, I. & Kaufman, A. (2002), ‘Automatic centerline extraction for virtual colonoscopy’, *IEEE Transactions on Medical Imaging* **21**(12), 1450–1460.
- Wan, M., Tang, Q., Kaufman, A., Liang, Z. & Wax, M. (1999), Volume rendering based interactive navigation within the human colon, *in* ‘Proc of the IEEE visualisation conference 1999’.
- Wang, Z., Liang, Z., Li, L., Li, X., Li, B., Anderson, J. & Harrington, D. (2005), ‘Reduction of false positives by internal features for polyp detection in CT-based virtual colonoscopy’, *Medical Physics* **32**(12), 3602–3616.
- Welzl, E. (1991), ‘Smallest enclosing disks (balls and ellipoids)’, *Lecture Notes in Computer Science* **555**, 359–370.
- Whelan, P. F. & Molloy, D. (2000), *Machine Vision Algorithms in Java: Techniques and Implementation*, Springer-Verlag London.
- Whelan, P. F. & Sadleir, R. J. T. (2004), ‘A visual programming environment for machine vision engineers’, *Sensor Review* **24**(3), 265–270.
- Whelan, P. F., Sadleir, R. J. T. & Ghita, O. (2004), ‘NeatVision: Visual programming for computer-aided diagnostic applications’, *Radiographics* **24**(6), 1779–1789.
- Wyatt, C., Ge, Y. & Vining, D. (2000), ‘Automatic segmentation of the colon for virtual colonoscopy’, *Computerized Medical Imaging and Graphics* **24**(1), 1–9.

## Bibliography

- Yao, J., Miller, M., Franaszek, M. & Summers, R. (2004), ‘Colonic polyp segmentation in CT colonography-based on fuzzy clustering and deformable models’, *IEEE Transactions on Medical Imaging* **23**(11), 1344–1352.
- Yee, J., Akerkar, G., Hung, R., Steinauer-Gebauer, A., Wall, S. & McQuaid, K. (2001), ‘Colorectal neoplasia: performance characteristics of CT colonography for detection in 300 patients’, *Radiology* **219**(3), 685–92.
- Yoshida, H. & Dachman, A. H. (2005), ‘CAD techniques, challenges, and controversies in computed tomographic colonography’, *Abdominal Imaging* **30**(1), 26–41.
- Yoshida, H., Masutani, Y., MacEneaney, P., Rubin, D. & Dachman, A. (2002a), ‘Computerized detection of colonic polyps at CT colonography on the basis of volumetric features: pilot study’, *Radiology* **222**(2), 327–336.
- Yoshida, H. & Näppi, J. (2001), ‘Three-dimensional computer-aided diagnosis scheme for detection of colonic polyps’, *IEEE Transactions on Medical Imaging* **20**(12), 1261–1274.
- Yoshida, H., Nappi, J., MacEneaney, P., Rubin, D. & Dachman, A. (2002b), ‘Computer-aided diagnosis scheme for detection of polyps at CT colonography’, *Radiographics* **22**(4), 963–979.
- Zalis, M. E., Barish, M. A., Choi, J. R., Dachman, A. H., Fenlon, H. M., Ferrucci, J. T., Glick, S. N., Laghi, A., Macari, M., McFarland, E. G., Morrin, M. M., Pickhardt, P. J., Soto, J. & Yee, J. (for the working group on virtual colonoscopy) (2005b), ‘CT colonography reporting and data system: A consensus proposal’, *Radiology* **236**(1), 3–9.
- Zalis, M. E., Perumpillichira, J., Del Frate, C. & Hahn, P. F. (2003), ‘CT colonography: digital subtraction bowel cleansing with mucosal reconstruction? Initial observations’, *Radiology* **226**(3), 911–917.
- Zalis, M. E., Perumpillichira, J. & Hahn, P. F. (2004), ‘Digital subtraction bowel cleansing for CT colonography using morphological and linear filtration methods’, *IEEE Transactions on Medical Imaging* **23**(11), 1335–1343.
- Zalis, M. E., Perumpillichira, J. J., Kim, J. Y., Del Frate, C., Magee, C. & Hahn, P. F.

## Bibliography

- (2005a), ‘Polyp size at CT colonography after electronic subtraction cleansing in anthropomorphic colon phantom’, *Radiology* **236**(1), 118–124.
- Zalis, M. E., Perumpillichira, J. J., Magee, C., Kohlberg, G. & Hahn, P. F. (2006), ‘Tagging-based, electronically cleansed CT colonography: evaluation of patient comfort and image readability’, *Radiology* **239**(1), 149–159.
- Zalis, M. & Hahn, P. (2001), ‘Digital subtraction bowel cleansing in CT colonography’, *American Journal of Roentgenology* **176**(3), 646–648.
- Zhou, Y., Kaufman, A. & Toga, A. W. (1998), ‘Three-dimensional skeleton and centerline generation based on an approximate minimum distance field’, *Visual Computer* **14**(7), 303–314.
- Zhou, Y. & Toga, A. W. (1999), ‘Efficient skeletonisation of volumetric objects’, *IEEE Transactions on Visualisation and Computer Graphics* **5**(3).
- Zucker, S. W. & Hummel, R. A. (1981), ‘A three-dimensional edge operator’, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **3**(3), 324–331.

# Appendix A

## DICOM Interpretation

Custom medical imaging applications are becoming more common place. This is primarily driven by the increasing role of computer aided image processing and analysis in radiology research. An important and common stage in the development of such applications is the interpretation of medical image data. This data is generally stored in accordance with the Digital Imaging and Communications in Medicine (DICOM) standard (ACR-NEMA 2004) (summarised by Mildenberger et al. (2002)). Interpretation of medical images involves decoding the relevant DICOM data and making it readily available to the application developer for analysis and display. Our group has developed a Java based medical imaging toolkit to facilitate the rapid development and deployment of medical imaging applications in a research environment (see Figure A.1). Over the past number of years this toolkit has been used internally as a platform for the development of a number of research applications in the areas of CT colonography (Sadleir et al. 2002, Sadleir & Whelan 2005), MRCP (Robinson et al. 2002) and also functional MRI (Malone et al. 2003). In the course of this work the toolkit has been exposed to a wide range of medical images obtained from various imaging modalities.

This medical imaging toolkit, which will be referred to as the NeatMed application developers interface (API), was developed using the Java programming language (Sun Microsystems, Mountain View, Calif.). The use of Java has previously been demonstrated in the development of custom medical imaging applications. Fernandez-Bayo et al. (2000), describe a web based Java viewer developed for use with a custom web server, Bernarding et al. (2001), describe a framework for the development of DICOM server applications implemented using Java and Mikolajczyk et al. (1998), describe a Java environment for the analysis of positron emission tomography (PET) image data. The choice of Java for implementing the



overhead.

- **Core functionality:** The core Java libraries encapsulate an extensive range of functionality that can be easily reused to create reliable, diverse and powerful applications. Some of the main capabilities supported by the core libraries include: networking, file IO, image processing, database access and graphical user interface development.
- **Extension APIs and toolkits:** There are a large number of extension APIs and toolkits that can be used in conjunction with the core Java libraries. These extensions usually deal with a particular speciality such as advanced image processing, 3D graphics or speech recognition.

The concept of a Java based medical imaging toolkit is not entirely new. In fact there are a number of commercial toolkits already available: the DICOM Image I/O Plugin (Apteryx, Issy-les-Moulineaux, France), the LEADTOOLS Medical Imaging suite (LEAD Technologies, Charlotte, N.C.) and the Java DICOM toolkit (Softlink, Halle Zoersel, Belgium). NeatMed was developed as a freely available alternative to these commercial toolkits. It is primarily intended for use in a research environment, specifically for dealing with offline medical image data. A major benefit associated with NeatMed is that it was designed with simplicity and ease of use in mind. This is especially important as it makes NeatMed accessible to all those involved in medical imaging research including radiologists, computer scientists and engineers. Other freely available toolkits e.g. ImageJ<sup>1</sup> and jViewbox<sup>2</sup> facilitate the development of applications that deal with DICOM images. However, they lack the low level functionality required to create applications that deal with all aspects of the DICOM file format.

This appendix is intended to provide an introduction to the NeatMed API and deals with implementation issues, toolkit structure, sample applications, support material and current API status. The reader is encouraged to visit the official NeatMed website<sup>3</sup> for further information about the NeatMed API and associated resources. The NeatMed API is distributed in accordance with the terms and conditions laid

---

<sup>1</sup>ImageJ, Image Processing and Analysis in Java. Available at <http://rsb.info.nih.gov/ij/>; Accessed May 26th 2006.

<sup>2</sup>jViewbox, Laboratory of Neuro Imaging, UCLA. Available at [http://www.loni.ucla.edu/Software/Software\\_Detail.jsp?software\\_id=1](http://www.loni.ucla.edu/Software/Software_Detail.jsp?software_id=1); Accessed May 26th 2006.

<sup>3</sup>NeatMed, Medical Imaging Application Developer Interface. Available at: <http://www.eeng.dcu.ie/~vsl/DICOM/>; Accessed May 26th 2006.

out in the GNU Lesser General Public License. This licence was selected to ensure that the NeatMed API is accessible to all potential users.

## A.1 API Implementation

NeatMed was developed using the Java programming language. Java was created by Sun Microsystems in the early 1990s. It was initially intended for the development of software for use in the consumer electronics industry (e.g. set-top boxes). Unfortunately, the Java development team were unable to find a target market for their new programming language. Forced to reconsider, it was decided to redeploy Java as an Internet technology. Java programs known as Applets first appeared on the Internet in 1994. These interactive Applets were embedded in standard hypertext mark-up language (HTML) web pages and greatly enhanced the previously static Internet. The Java programming language can also be used for the development of stand-alone client side applications which can run independently of a web browser and are not subjected to the same security restrictions as Applets. The core Java libraries maintained by Sun Microsystems are used as the foundation for the development of any Java application. These libraries can be used in conjunction with extension APIs in order to develop specialised applications. See Whelan & Molloy (2000) for further information about the Java programming language, particularly in relation to image processing and analysis.

An extension API is a set of classes that can be instantiated by a programmer to create a particular type of application, thus facilitating software reuse. NeatMed is an example of an extension API that can be used for the development of applications that deal with offline medical image data. A large number of Java APIs exist, these deal with a broad range of applications ranging from communicating with the serial and parallel ports to advanced image processing. The set of classes representing the API is deployed in some type of library. Java provides a packaging tool that can be used to package a set of class files and associated resources into a Java archive or JAR file. In order to be useful, an API must be well documented. Java provides a documentation tool called Javadoc that allows an API developer to document software as it is being written. The resulting documentation provides detailed information about each class, method and variable that is defined in the associated API. The structure of Javadoc documentation is more or less the same

for every API. This makes it very easy for programmers to familiarise themselves with a new API once they are comfortable with the basic Javadoc documentation structure. The API documentation is generated in HTML and can be viewed using any standard web browser.

Java has a wide range of benefits associated with it. However, there are also some limitations. One example is performance: Java is a multiplatform programming language, the byte code (i.e. binary form) that represents a java program is interpreted and not executed directly. This reduces the performance of a Java program compared to a natively executed program. Overall, however, the benefits associated with Java (listed previously) far outweigh the drawbacks, hence its selection for the development of NeatMed.

## A.2 API Overview

The NeatMed API is a group of core and support classes that can be used to interpret, represent and manipulate images and related data that are stored in DICOM compliant files. The central class in the API is the `DICOMImage` class. A `DICOMImage` object can be instantiated by specifying a reference to a suitable data source in the constructor. The constructor will accept data from a number of sources (e.g. local file, data stream and remote URL). Once constructed a `DICOMImage` object provides direct access to all of the data elements stored within the specified DICOM source. Other classes in the API are used to represent individual components within a DICOM file such as `DICOMTag`: a data element tag (i.e. group number and element number combination) or `DICOMDataElement`: a special wrapper class created for storing data associated with a particular value representation. There is a subclass of `DICOMDataElement` for each of the 26 value representations, examples include: `DICOMAgeString`, `DICOMPersonName` and `DICOMUniqueIdentifier`. The API also provides utilities classes such as `DICOMDataDictionary`: a data dictionary containing value representation definitions for all of the data elements defined in the DICOM 3.0 standard (required when the value representation is not specified explicitly). This section describes how the `DICOMImage` class operates and how it interacts with other classes in the API to provide access to DICOM encoded data.

### A.2.1 Data Interpretation

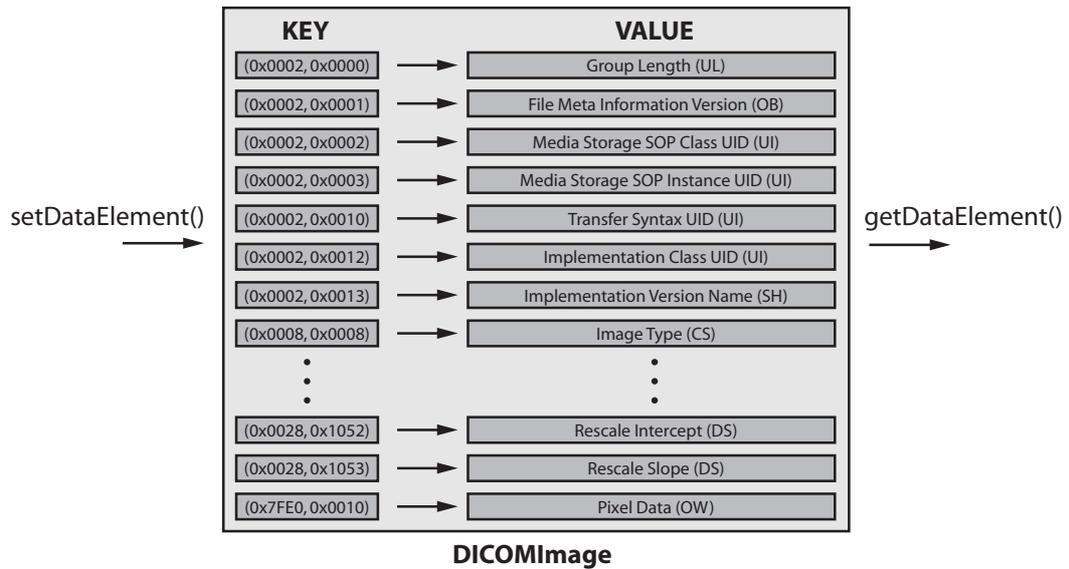
When a `DICOMImage` object is constructed, it reads and decodes all the data from the specified source. The decoder automatically determines the type of data and the transfer syntax to be used. If the value representation is not specified explicitly, then the required value representation for a particular data element is obtained from the pre-programmed data dictionary (i.e. `DICOMDataDictionary`). Each data element in the DICOM file is subsequently decoded and stored using the relevant wrapper class. Data stored within a wrapper class can easily be accessed for further processing. The image data (0x7fe0, 0x0010) is a special type of data element that can be stored using one of two possible value representations: Other byte string or other word string. The image data is always packed and sometimes compressed using either *joint photographic experts group* (JPEG) (Wallace 1991, Pennebaker & Mitchell 1993) or *run length encoding* (RLE) compression. Information required to unpack and decompress the image data is usually stored within group 0x0028 data elements. The DICOM decoder automatically unpacks and decodes the image data and stores it internally as a one dimensional array of signed 32-bit integer primitives.

### A.2.2 Data Management

Once decoded all data elements are stored internally using a modified hash table structure. A hash table allows information to be stored as key/value pairs. In this case, the key is the data element tag that consists of a group number and an element number and the value is the data element associated with the relevant tag. A tag is represented by a `DICOMTag` class and all data elements are represented by dedicated wrapper classes, which are subclasses of `DICOMDataElement`. A specific data element value can be retrieved from the hash table by specifying the relevant group number and element number combination. The method employed for storing and querying decoded data elements is illustrated in Figure A.2.

### A.2.3 Data Access/Manipulation

Data can be accessed at several levels. At the most basic level, a data element can be retrieved from the hash table using the `getDataElement()` method of the `DICOMImage` class. There are two versions of the `getDataElement()` method. The first version takes a single argument of type `DICOMTag`, which encapsulates the group



**Figure A.2:** A representation of how data is stored and queried at the most basic level of the `DICOMImage` class.

number and element number of the desired data element and the second version takes two integer arguments, one for the group number and the second for the element number. In either case, the `getDataElement()` method will return an object of type `DICOMDataElement` representing the requested data element. This object must be cast (i.e. converted) into the relevant subclass of `DICOMDataElement` before the actual data element value can be accessed. An existing data element can be modified, or a new data element can be added to a `DICOMImage` object by calling the `setDataElement()` method. This method takes arguments that represent a tag and the new data element value to be associated with that tag. The basic level representation does not provide direct access to the actual image data. This data is packed and in some cases compressed and further decoding is required in order to facilitate pixel level operations.

There are an extremely large number of data elements defined by the DICOM standard. It is possible for any of these data elements to be present in a DICOM file. Some data elements occur more frequently than others, and have particular significance to the application developer. In order to facilitate straightforward access to important, and frequently used data elements, a number of special accessor methods are defined by the `DICOMImage` class. This set of methods represents a higher level of access than that provided by the `getDataElement()` and `setDataElement()` methods, and simplifies access to particular frequently used or important data elements.

Internally each accessor method calls the `getDataElement()` method with the relevant `DICOMKey` argument and casts the returned data element value into the relevant Java primitive (or object). Some examples of the accessor methods provided by the `DICOMImage` class are as follows:

- `String getPatientID()` retrieves the data element value associated with the key (0x0010, 0x0020) from the hash table. The `DICOMLongString` object at this location is converted into a Java `String` object that is then returned.
- `int getSeriesNumber()` retrieves the data element value associated with the key (0x0020, 0x0011) from the hash table. The `DICOMIntegerString` object at this location is converted into a signed 32-bit integer primitive and returned.
- `int getBitsAllocated()` retrieves the data element value associated with the key (0x0028, 0x0100) from the hash table. The `DICOMUnsignedShort` object at this location is converted into a signed 32-bit integer primitive and returned.

Accessor methods are also provided for adding new data element values or modifying existing data element values. These methods take a single argument that represents the new value of the relevant data element. Internally, these methods call `setDataElement()` with the relevant `DICOMKey` and new value arguments. Examples include `setPatientID(String ID)`, `setSeriesNumber(int seriesNumber)` and `setBitsAllocated(int bitsAllocated)`.

The final level of abstraction is used to provide direct access to image data. As mentioned previously the image data is stored in data element (0x7fe0, 0x0010). After the initial decoding/interpretation stage this data is still packed and sometimes compressed. The final stage of decoding performed by the `DICOMImage` class, automatically unpacks and decompresses the image data that is ultimately stored within the `DICOMImage` class as a one dimensional array of signed integer values. Individual pixel values can be accessed/modified using special accessor methods: `getSample()` and `setSample()`. In either case the (x, y) coordinates of the relevant pixel must be specified in the argument list of the accessor method. In the case of multi-frame image data (i.e. where the number of frames is greater than one) a time coordinate must also be specified, and for multi-plane image data (e.g. ARGB or CYMK) a colour plane index must also be specified. Some examples of the pixel level accessor methods provided by the `DICOMImage` class are as follows:

- `getSample(int x, int y)` returns an integer primitive that represents the pixel value at the specified (x, y) coordinates.
- `getSample(int x, int y, int t)` returns an integer primitive that represents the pixel value at the specified (x, y, t) coordinates.
- `setSample(int x, int y, int value)` sets the value of the pixel at the specified (x, y) coordinates to be the same as that indicated by the value argument.

### A.2.4 Data Storage

In many cases an application developer will only need to read from a DICOM file. However, there are some cases where it may be useful to modify the contents of a DICOM file. In order to facilitate this operation the `DICOMImage` class provides a `save()` method. This method saves the current state of the associated `DICOMImage` object, and takes a single argument that represents the output stream where the data will be written. The saving operation involves writing all of the data elements represented by the `DICOMImage` object (including the packed pixel information) to the specified output stream. The saved data will be stored according to the specified transfer syntax, i.e. data element (0x0002, 0x0010). If this data element is not present, the default transfer syntax will be used. The data storage feature completes the range of functionality provided by the NeatMed API.

## A.3 Sample Applications

The NeatMed API can be used to develop a wide variety of medical imaging applications. This section describes a number of sample applications that demonstrate the power, flexibility and ease of use of the NeatMed API. In some cases NeatMed is used in conjunction with extension APIs and toolkits in order to demonstrate its development potential. The source code and associated instructions for the compilation and execution of each sample application can be downloaded directly from the NeatMed website.

### A.3.1 Simple DICOM viewer

A simple DICOM image viewer application can be created quite easily. The source code for the application is extremely compact (see Listing A.1). The application

```

1.  import java.awt.*;
2.  import javax.swing.*;
3.  import neatmed.DICOM.*;
4.
5.  public class DICOMViewer
6.  {
7.      public static void main(String args[])
8.      {
9.          new DICOMViewer(args[0]);
10.     }
11.
12.     public DICOMViewer(String filename)
13.     {
14.         try
15.         {
16.             // Load the requested DICOM image...
17.             DICOMFile file = new DICOMFile(filename);
18.             DICOMImage dimage = file.getImage();
19.             BufferedImage bimage = dimage.getAsBufferedImage();
20.
21.             // Add the image to the frame...
22.             JLabel label = new JLabel(new ImageIcon(bimage));
23.             getContentPane().setLayout(new BorderLayout());
24.             getContentPane().add(label, BorderLayout.CENTER);
25.         }
26.         catch(Exception e)
27.         {
28.             System.out.println(e.toString())
29.         }
30.     }
31. }

```

**Listing A.1:** The complete source code for a DICOM image viewer application. The application takes one argument, the name of the DICOM compliant image that is to be loaded.

takes a single argument which represents the name of the image to be displayed. This filename is subsequently used to construct a `DICOMImage` object. A `BufferedImage` object representation of the `DICOMImage` object is then obtained. This `BufferedImage` object is ultimately displayed using Swing graphical user interface (GUI) components (see `DICOMViewer.java`).

### A.3.2 Sequence Viewer

The simple DICOM viewer application can be extended to deal with multi-frame DICOM images. This extension is facilitated by adding two buttons to the application GUI. These buttons allow the user to sequence backward and forward through the available frames. Internally the buttons update the value of a counter, the backward button decrements the counter and the forward button increments the counter. After either button is pressed the frame with the index that corresponds to the counter

value is displayed. This application demonstrates how user interaction is handled by Java. (see `SequenceViewer.java`).

### A.3.3 Volume Rendering

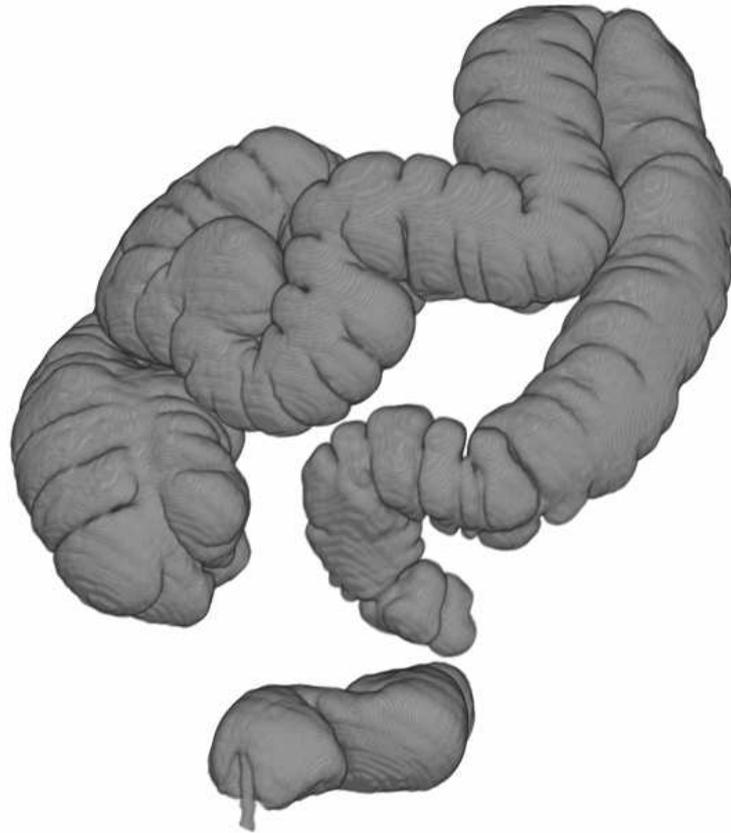
The NeatMed API can be used in conjunction with the visualisation toolkit (VTK) (Schroeder et al. 2000) to provide 3-D volume or surface renderings of medical image data. NeatMed is used to read in axial slices from a volumetric data set. The image data from the slices is used to populate an array of scalar values that represents the volumetric data. Two transfer functions (opacity and colour) are specified to indicate how the volume should be displayed. The opacity transfer function indicates the opacity values associated with particular voxel intensity. The colour transfer function indicates the RGB colour value associated with a particular voxel intensity, this can be used for pseudo colouring the volume data. The rendering of the volume is handled by the VTK, and the user can interact with the resulting model using the mouse. Supported operations include rotate, zoom and translate. A sample volume rendering using this application is illustrated in Figure A.3 (see `VolumeRendering.java`).

### A.3.4 Anonymising Data

A DICOM file usually contains a number of data elements that hold sensitive patient information. It is important to be able to anonymise this information in certain cases in order to protect the patient's identity. The anonymisation process involves constructing a `DICOMImage` object from a suitable source, overwriting all of the sensitive data element values (e.g. patient's name, patient's birth date and patient's address) and saving the modified `DICOMImage` object to an output stream using the `save()` method. The saved file represents the anonymized data. Each of the sensitive data element values can be overwritten using the `setDataElement()` method or special accessor methods (where available). Note that this application is a command line only application, there is no GUI (see `Anonymise.java`).

### A.3.5 Image Processing

The NeatMed API can be used for the development of medical image analysis applications. This can be achieved using any of the well documented image processing operations that are described in the literature, see Whelan & Molloy (2000) for examples of image processing algorithms in Java. The threshold operation is a simple



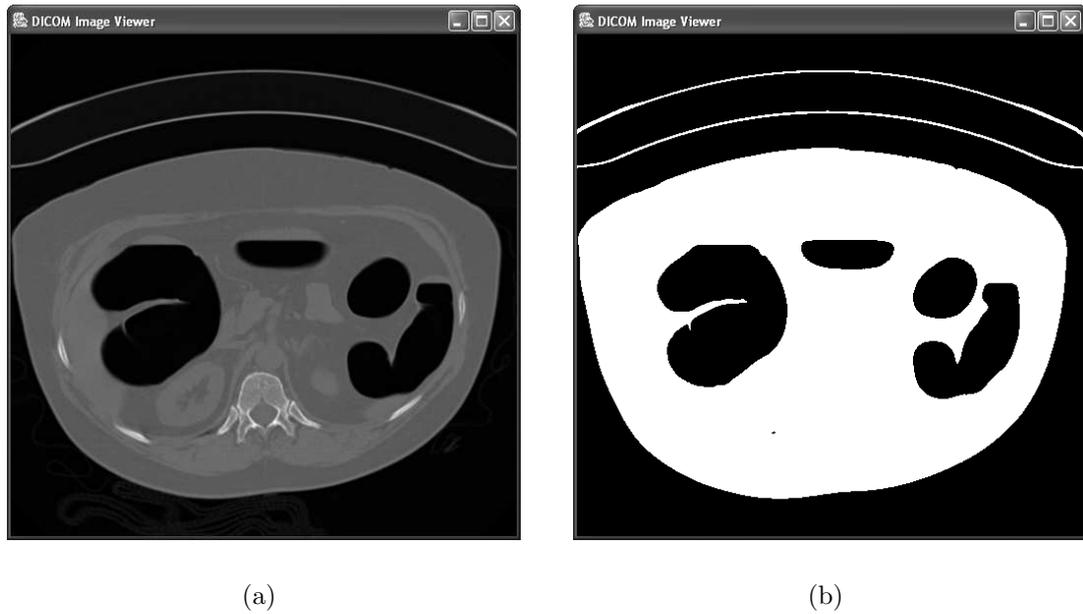
**Figure A.3:** A volume rendering of the human colon generated from an abdominal CTC data set using a combination of the NeatMed API and the VTK.

example of a point operation that is based on a raster scan of the input image. Thresholding converts a grey scale image into a binary (black and white) image based on a user defined threshold value. Any pixel that is greater than or equal to the threshold is set to white, while any pixel less than the threshold is set to black. The threshold operation performs a rough segmentation of the image and demonstrates both looping and conditional data processing. The threshold operation is illustrated in Figure A.4 (see `Threshold.java`).

As mentioned earlier, NeatMed has been used internally over the past number of years as a platform for the development of a number of medical imaging research applications. Examples of this research are illustrated in Figure A.5.

## A.4 Support Material

All support material for the NeatMed API is accessible through the official NeatMed website. In addition to providing access to the latest version of the NeatMed API,



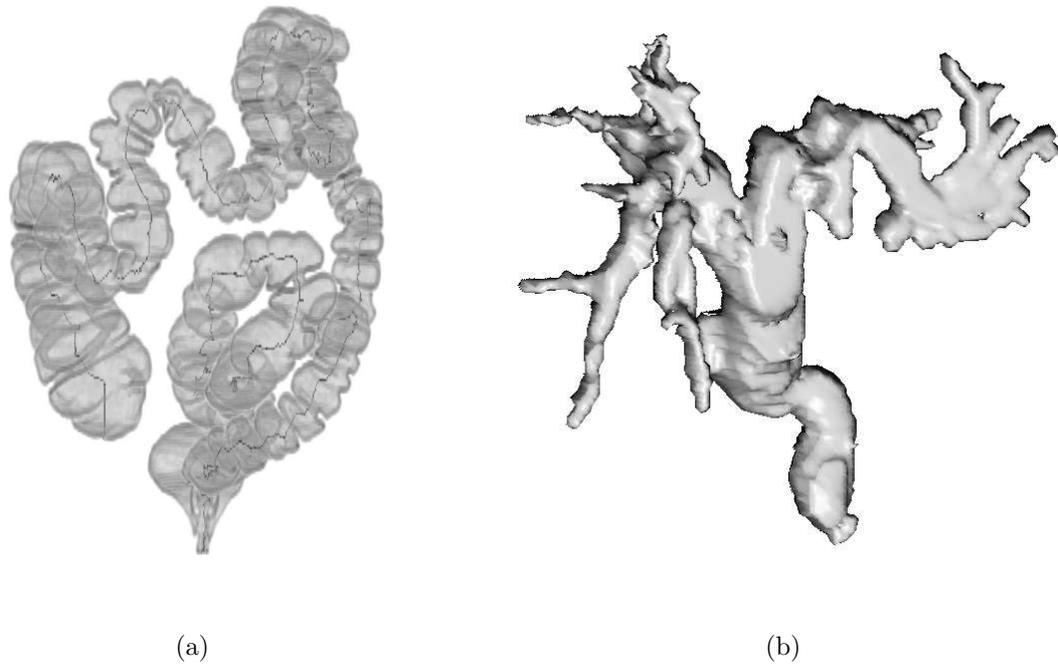
**Figure A.4:** The threshold image processing operation. A sample axial CTC slice (a) before and (b) after application of the threshold.

the website also provides access to documentation, sample source code and contact information for the NeatMed development team.

The NeatMed API documentation is generated using Javadoc. The NeatMed API documentation (see Figure A.6) can be browsed online or downloaded as a ZIP file for subsequent offline access. In either case, the documentation can be viewed using a standard web browser. The documentation itself provides an overview of the entire API. Each class and its associated methods and variables are described in detail. The relationship between classes is indicated and hyperlinks are provided to facilitate easy navigation through the documentation.

Fully commented source code for each of the sample applications described earlier in this paper can be downloaded from the NeatMed website. These applications can be compiled and executed using the Java 2 Platform, Standard Edition Development Kit (JDK) from Sun Microsystems. The latest version of the JDK (currently v5.0 update 6) can be downloaded at no cost from official Java website<sup>4</sup>. All sample applications require the NeatMed API to be included in the Java class path and some applications also require the installation of additional APIs or toolkits such as the VTK. The NeatMed website should be consulted for detailed information about

<sup>4</sup>Sun Microsystems Inc. Java 2 Platform Standard Edition 5.0 Downloads. Available at <http://java.sun.com/j2se/1.5.0/download.jsp>; Accessed May 26th 2006.



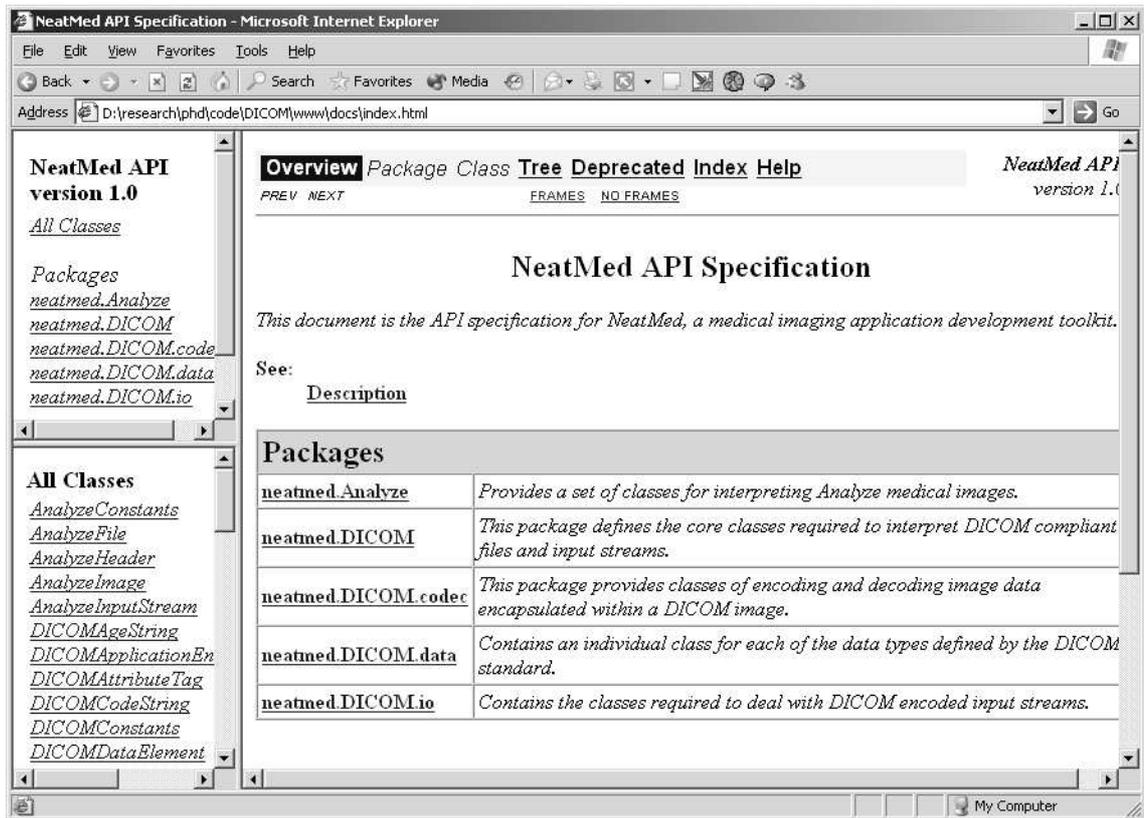
**Figure A.5:** Advanced image processing facilitated by NeatMed (a) Centreline calculation at CT colonography. (b) Segmentation of the biliary tree from MRCP data (Image courtesy of K. Robinson (VSG)).

compiling and executing the sample applications. Links to sites providing DICOM images that can be used in conjunction with the sample applications can also be accessed via the NeatMed website.

Finally the NeatMed team can be contacted by email either for general queries or to suggest additional functionality that may be useful to other NeatMed users. All queries regarding the NeatMed API are dealt with via email ([neatmed@eeng.dcu.ie](mailto:neatmed@eeng.dcu.ie)).

## A.5 Current Status

The NeatMed API is constantly evolving. Its ongoing development is driven by internal requirements and external requests for additional features. NeatMed supports data that conforms to version 3 of the DICOM standard as well as the older American College of Radiology - National Electrical Manufacturers Association (ACR-NEMA) standard. NeatMed currently supports all of the uncompressed DICOM transfer syntaxes as well as the lossless RLE transfer syntax (1.2.840.10008.1.2.5). Support for JPEG compressed image data is currently being incorporated into the API. Dedicated wrapper classes are provided for storing each of the 26 possible data element value representations defined in the DICOM standard. The pixel data decoder is



**Figure A.6:** The HTML Javadoc documentation for the NeatMed API. This documentation can be accessed online at <http://www.eeng.dcu.ie/~vsl/DICOM/>

extremely flexible and supports any valid pixel encoding (packing) scheme, the majority of commonly used photometric interpretations are also supported and the recent addition of the basic data write feature completes the range of functionality available for dealing with DICOM data.

NeatMed has recently been updated to support version 7.5 of the ANALYZE file format. As with DICOM data, the ANALYZE file pair (image and header) is automatically interpreted, and all encoded information is made readily available to the programmer. The ANALZE file format has a finite number of header fields and direct access is provided to each of these. Individual image sample values can be obtained by simply specifying the relevant 2-D, 3-D or 4-D coordinates. The ANALYZE section of the API can also be used in conjunction with other toolkits and APIs to create powerful medical imaging applications.

The NeatMed API is intended for use by programmers interested in developing medical imaging applications, particularly for computer aided analysis. There is no reason why nonprogrammers should be excluded from developing such applications. In order to facilitate nonprogrammers, NeatMed has recently been incorporated

into the NeatVision visual programming environment (Whelan et al. 2004, Whelan & Sadleir 2004). NeatVision is a free, Java based software development environment designed specifically for use in image analysis applications. NeatVision is both user-friendly and powerful, providing high-level access to over 300 image manipulation, processing and analysis algorithms. A visual program can be created by simply instantiating the required algorithms (blocks) and creating data paths (interconnections) between the inputs and outputs to indicate the data flow within the program. This combination of NeatVision and NeatMed makes a vast library of image processing operators easily accessible to the medical image analysis community.

## A.6 NeatMed Users

The NeatMed API was released in 2002 at the 88th scientific assembly and annual meeting of the Radiological Society of North America. Since then its user base has grown steadily. A selection of groups who have downloaded and are using NeatMed include:

- Montreal Neurological Institute and Hospital (a teaching and research institute of McGill University), Montreal, Canada.  
<http://www.mni.mcgill.ca/>
- GE Healthcare.  
<http://www.gehealthcare.com/>
- DeVry University, Columbus, Ohio, United States.  
<http://www.devrycols.edu/>
- Medical Micrographics, Three Rivers, Michigan, United States.  
<http://medicalmicrographics.com/>
- New Rochelle Radiology, New Rochelle, New York, United States.  
<http://www.newrorad.com/>
- United States Department of Veterans Affairs.  
<http://www.med.va.gov>
- University Hospital of Lund, Lund, Sweden.  
(no website available)

- Medical School of Ribeirao Preto, Brazil.  
<http://www.fmrp.usp.br/>
- Ninewells Hospital and Medical School, Dundee, Scotland.  
<http://www.show.scot.nhs.uk/>
- Kybele Group, Universidad Rey Juan Carlos, Madrid, Spain.  
<http://kybele.escet.urjc.es/>
- Institute of Diagnostic Radiology and Neuroradiology, Ernst-Moritz-Arndt-Universitt, Greifswald, Germany.  
<http://www.uni-greifswald.de/>
- University of North Texas, Denton, Texas, United States.  
<http://www.unt.edu>
- Northwestern Radiology, Chicago, Illinois, United States.  
<http://www.radiology.northwestern.edu/>

Note: All URLs accessed May 26th 2006.

## A.7 Summary

The interpretation of medical image data is an important and common stage in the development of any medical imaging application. NeatMed removes the need to deal directly with encoded medical image data, thus increasing productivity and allowing the developer to concentrate on other aspects of application development. NeatMed is written in Java, a multiplatform programming language with a large amount of freely available support material that is straightforward to learn and use. These and other features of Java make NeatMed accessible to a large group of potential users. NeatMed is well supported with documentation and sample code available through the NeatMed website. Most importantly, NeatMed is a freely available research tool whose ongoing development is driven by the needs and requirements of its users.

# Appendix B

## Summary of Available Data

This Appendix provides a detailed summary of the CTC data and gold standard polyp information that was obtained from the Department of Radiology and the Gastrointestinal Unit at the Mater Misericordiae University Hospital, Eccles Street, Dublin 7, Ireland. All patients were scanned in the prone and supine positions and underwent conventional colonoscopy immediately following the CTC examination. The gold standard was obtained from an inspection of the data sets where the reviewer was unblinded to the results from colonoscopy examination. This ensured that the most accurate information regarding colorectal neoplasia was available for CAD-CTC training and evaluation. A comprehensive overview of the test data is presented in Table B.2 and the more important information is summarised in Table B.1. Surface renderings of all available data sets are presented in Figures B.1, B.2 and B.3.

There were a total of 37 individual patients studies yielding 68 data sets. There were four instances where only one data set was provided for a patient and two instances where only partial data sets were received. In the majority of cases the tube current was 250 mAs. The average number of slices per study was 263 and, thus, the average data set consisted of just under 70 million voxels. In general, the slice thickness was 1.5 mm and the pixel size ranged from 0.486 mm to 0.848 mm. In 25 data sets (37%), the colon was found to be well distended and the segmentation process resulted a single segment that represented the entire colon lumen. Prone positioning resulted in the best overall distention. The remaining 43 data sets consisted of multiple colon segments. The maximum number of segments was seven and the mean number of segments per data set was 2.162. These observations were comparable with those made by Li & Santago (2005), who reported that a fully segmented colon can have between 1 - 10 individual segments with an average of 2.4

**Table B.1:** A brief summary of the data that was used in this study.

<b>Data set information</b>			
Number of patients	37		
Number of usable data sets	68		
<b>Data set statistics</b>	<b>min</b>	<b>max</b>	<b>mean</b>
Radiation dose (mAs)	167	250	241
Number of Slices	153	316	263
Pixel size (mm)	0.486	0.848	0.699
Slice depth (mm)	1.5	3.0	1.522
Slice Dimensions	512 × 512		
<b>Distension information</b>	<b>prone</b>	<b>supine</b>	<b>all</b>
Single segment colons	16	9	25
Mean number of segments	1.882	2.441	2.162
Maximum number of segments	5	7	7
<b>Neoplasia information</b>			
Number of neoplasia	155		
Number of flat polyps	2		
Number of [0-5) polyps	98		
Number of [5-10) polyps	31		
Number of ≥ 10 polyps	13		
Number of masses	11		

segments.

There were a total of 155 colorectal neoplasia, 44 of these were clinically significant polyps (i.e.  $\geq 5$  mm), 11 were colorectal masses, two were flat polyps (both  $< 5$  mm) and the remaining 98 were  $< 5$  mm in size. Each data set was treated as a separate patient in order to increase the number of polyps available for testing and evaluation purposes.

Table B.2: A detailed summary of the patient data, provided by the MMH, that used for system training and evaluation.

Patient ID	Patient Position	Full	Dose (mAs)	Slices	$V_{xz}$ (mm)	$V_y$ (mm)	Segs	Polyp/Mass Information								
								#	Type			Size (mm)		Location (x, y, z)	Size (mm)	Type
									flat	< 5	[5, 9)	≥ 10	mass			
MMH-010	prone	Yes	250	260	0.652	1.5	3	0	0	0	0	0	(x, y, z)	-	-	
	supine	Yes	250	249	0.637	1.5	3	0	0	0	0	0	(x, y, z)	-	-	
MMH-017	prone	Yes	250	266	0.596	1.5	5	4	0	2	2	0	0	(253, 70, 291)	3	P
														(154, 71, 254)	5.25	P
	supine	Yes	250	249	0.656	1.5	3	4	0	2	2	0	0	(113, 80, 354)	4	P
														(102, 101, 287)	5	P
MMH-033	prone	Yes	250	281	0.742	1.5	1	1	0	0	0	0	1	(277, 99, 196)	> 20	M
	supine	Yes	250	294	0.785	1.5	4	1	0	0	0	0	1	(218, 104, 327)	> 20	M
MMH-054	prone	Yes	250	223	0.691	1.5	4	1	0	0	1	0	0	(340, 143, 156)	6	P
	supine	No	-	-	-	-	-	-	-	-	-	-	-	-	-	-
MMH-081	prone	Part	-	-	-	-	-	-	-	-	-	-	-	-	-	-
	supine	Yes	250	239	0.665	1.5	1	2	0	2	0	0	0	(221, 38, 315)	3	P
MMH-093	prone	Yes	250	293	0.781	1.5	1	1	0	0	1	0	0	(215, 59, 215)	> 5	P
	supine	Yes	250	267	0.813	1.5	1	1	0	0	1	0	0	(279, 61, 259)	> 5	P
MMH-109	prone	Part	-	-	-	-	-	-	-	-	-	-	-	-	-	-
	supine	Yes	250	251	0.742	1.5	3	3	0	1	0	1	1	(199, 141, 133)	10	P
														(374, 65, 159)	> 20	M

Continued on next page...

Table B.2: Continued from previous page.

Patient ID	Position	Full	Dose	Slices	$V_{xz}$	$V_y$	Segs	#	flat	< 5	[5, 9)	$\geq 10$	mass	Location	Size	Type
MMH-120	prone	Yes	167	257	0.486	1.5	1	1	0	0	1	0	0	(104, 168, 174)	4	P
	supine	Yes	167	226	0.551	1.5	1	3	0	2	1	0	0	(174, 27, 277) (265, 20, 334) (247, 25, 288) (264, 26, 380)	5.24 5.24 3 3	P P P P
MMH-138	prone	Yes	250	267	0.742	1.5	2	0	0	0	0	0	0	-	-	-
	supine	Yes	250	231	0.742	1.5	4	0	0	0	0	0	0	-	-	-
MMH-139	prone	Yes	250	233	0.695	1.5	2	0	0	0	0	0	0	-	-	-
	supine	Yes	250	241	0.617	1.5	2	2	0	1	0	1	0	(407, 158, 229) (355, 52, 145)	7 2	P P
MMH-142	prone	Yes	250	290	0.742	1.5	2	0	0	0	0	0	0	-	-	-
	supine	Yes	250	153	0.814	3	4	3	0	0	1	0	2	(371, 54, 207) (134, 71, 239) (135, 58, 217)	8 > 20 > 20	P M M
MMH-161	prone	Yes	250	238	0.660	1.5	2	1	0	1	0	0	0	(305, 57, 301)	4	P
	supine	Yes	250	233	0.652	1.5	2	0	0	0	0	0	0	-	-	-
MMH-166	prone	Yes	167	313	0.789	1.5	2	1	0	0	1	0	0	(240, 113, 262)	7	P
	supine	Yes	250	301	0.813	1.5	2	1	0	0	1	0	0	(287, 97, 188)	6	P
MMH-175	prone	No	-	-	-	-	-	-	-	-	-	-	-	-	-	-
	supine	Yes	250	295	0.848	1.5	1	2	1	1	0	0	0	(368, 225, 185) (225, 42, 273)	4 2	P F
MMH-208	prone	Yes	250	235	0.742	1.5	2	3	0	3	0	0	0	(361, 170, 309) (258, 29, 152)	3 3	P P
	supine	Yes	250	220	0.742	1.5	2	2	0	2	2	0	0	(372, 101, 231) (121, 136, 106)	3 3	P P

Continued on next page...

Table B.2: Continued from previous page.

Patient ID	Position	Full	Dose	Slices	$V_{xz}$	$V_y$	Segs	#	flat	< 5	[5, 9)	$\geq 10$	mass	Location	Size	Type
MMH-234	prone	Yes	250	273	0.742	1.5	1	4	0	1	1	0	2	(169, 161, 99)	3	P
	supine	Yes	250	281	0.684	1.5	2	4	0	1	1	0	2	(268, 105, 303) (362, 175, 299) (401, 168, 190) (373, 103, 287) (225, 185, 134) (433, 185, 174) (80, 151, 194) (161, 81, 163)	7 3.5 > 20 > 20 5.5 3.5 > 20 > 20	P P M M P P M M
MMH-245	prone	Yes	167	273	0.779	1.5	1	2	0	1	1	0	0	(438, 141, 290)	4	P
	supine	Yes	167	258	0.648	1.5	3	1	0	1	0	0	0	(395, 128, 262) (89, 150, 201)	6 3	P P
MMH-263	prone	Yes	250	303	0.688	1.5	1	1	0	0	0	1	0	(240, 29, 209)	$\geq 10$	P
	supine	Yes	250	287	0.715	1.5	1	1	0	0	0	1	0	(259, 35, 290)	$\geq 10$	P
MMH-284	prone	Yes	250	282	0.742	1.5	1	1	0	1	0	0	0	(386, 171, 291)	2	P
	supine	No	-	-	-	-	-	-	-	-	-	-	-	-	-	-
MMH-291	prone	Yes	250	275	0.742	1.5	2	1	0	1	0	0	0	(326, 199, 300)	3	P
	supine	Yes	250	272	0.742	1.5	3	1	0	1	0	0	0	(223, 198, 215)	3	P
MMH-316	prone	Yes	250	267	0.742	1.5	1	1	0	1	0	0	0	(114, 91, 303)	3	P
	supine	Yes	250	240	0.668	1.5	3	0	0	0	0	0	0	-	-	-
MMH-357	prone	Yes	250	267	0.672	1.5	3	2	0	1	1	0	0	(347, 179, 161) (104, 203, 261)	7.5 2	P P
	supine	Yes	250	267	0.668	1.5	2	3	1	1	1	0	0	(158, 204, 190) (350, 190, 139) (122, 88, 186)	7.5 4 3.5	P P F

Continued on next page...

Table B.2: Continued from previous page.

Patient ID	Position	Full	Dose	Slices	$V_{xz}$	$V_y$	Segs	#	flat	< 5	[5, 9)	$\geq 10$	mass	Location	Size	Type
MMH-368	prone	Yes	250	316	0.742	1.5	1	1	0	1	0	0	0	(187, 234, 288)	3	P
	supine	Yes	250	286	0.738	1.5	1	0	0	0	0	0	0	-	-	-
MMH-378	prone	Yes	250	263	0.617	1.5	2	3	0	2	0	0	1	(228, 62, 291)	> 20	M
	supine	Yes	250	243	0.617	1.5	3	3	0	2	0	0	1	(363, 75, 369) (423, 94, 372)	4 4	P P
MMH-393	prone	Yes	250	267	0.775	1.5	1	2	0	1	1	0	0	(386, 140, 173)	6	P
	supine	Yes	250	267	0.785	1.5	1	1	0	0	1	0	0	(102, 212, 250) (142, 129, 168)	3 6	P P
MMH-451	prone	Yes	250	249	0.605	1.5	5	0	0	0	0	0	0	-	-	-
	supine	Yes	250	261	0.662	1.5	7	3	0	2	1	0	0	(377, 182, 235) (213, 168, 207) (215, 167, 294)	5.5 3.5 4.5	P P P
MMH-483	prone	Yes	250	287	0.742	1.5	4	7	0	6	0	1	0	(274, 51, 251) (404, 185, 241) (416, 109, 239)	12 4 3	P P P
	supine	Yes	250	243	0.617	1.5	5	5	0	2	1	0	1	(244, 228, 203) (274, 256, 158) (134, 97, 209) (392, 90, 275) (287, 71, 129) (243, 54, 212) (149, 192, 100) (292, 239, 158)	4 3 3 4.5 12 12 5 3	P P P P P P P P

Continued on next page...

Table B.2: Continued from previous page.

Patient ID	Position	Full	Dose	Slices	$V_{xz}$	$V_y$	Segs	#	flat	< 5	[5, 9)	$\geq 10$	mass	Location	Size	Type
MMH-523	prone	Yes	250	266	0.688	1.5	1	2	0	2	0	0	0	(394, 156, 272) (290, 96, 295)	5	P
	supine	Yes	250	261	0.688	1.5	1	0	0	0	0	0	0	-	-	-
MMH-526	prone	Yes	250	245	0.598	1.5	2	2	0	2	0	0	0	(280, 89, 297) (203, 61, 215)	4	P
	supine	Yes	250	235	0.623	1.5	3	1	0	1	0	0	0	(229, 86, 240)	4	P
MMH-576	prone	Yes	250	254	0.602	1.5	1	1	0	1	0	0	0	(145, 109, 147)	3.5	P
	supine	Yes	250	248	0.598	1.5	2	1	0	1	0	0	0	(100, 167, 299)	4.5	P
MMH-640	prone	Yes	250	267	0.703	1.5	1	0	0	0	0	0	0	-	-	-
	supine	Yes	250	252	0.672	1.5	1	1	0	1	0	0	0	(331, 225, 112)	2	P
MMH-672	prone	Yes	250	267	0.742	1.5	2	0	0	0	0	0	0	-	-	-
	supine	Yes	250	273	0.729	1.5	4	1	0	1	0	0	0	(203, 216, 181)	4	P
MMH-708	prone	Yes	250	265	0.672	1.5	1	1	0	1	0	0	0	(414, 158, 201) (141, 128, 164)	3	P
	supine	Yes	250	267	0.691	1.5	2	3	0	3	0	0	0	(72, 150, 202) (67, 151, 214)	3	P
MMH-735	prone	Yes	250	274	0.723	1.5	1	1	0	1	0	0	0	(251, 179, 257)	4	P
	supine	No	-	-	-	-	-	-	-	-	-	-	-	-	-	-
MMH-765	prone	Yes	167	284	0.582	1.5	1	21	0	17	2	2	0	(265, 90, 187)	< 5	P
														(245, 67, 198)	$\geq 10$	P
														(235, 81, 178)	< 5	P
														(251, 98, 251)	$\geq 10$	P
														(211, 71, 185)	5 - 9	P
														(191, 51, 292)	< 5	P

Continued on next page...

Table B.2: Continued from previous page.

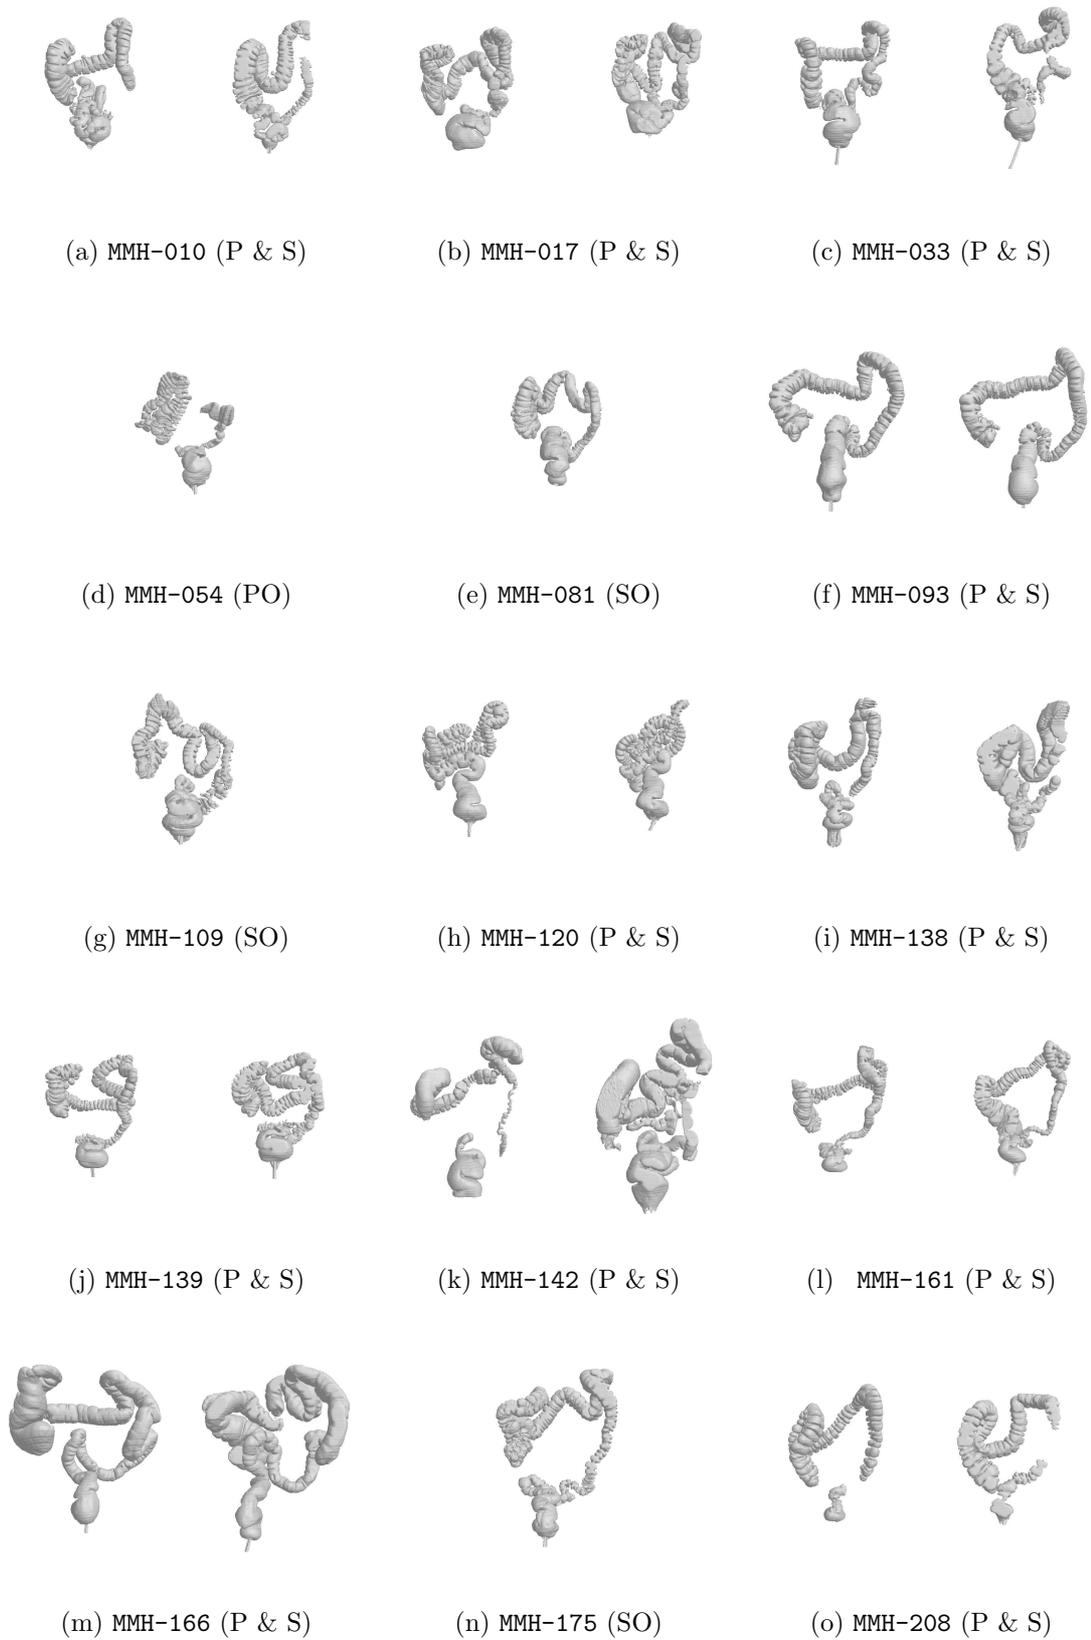
Patient ID	Position	Full	Dose	Slices	$V_{xz}$	$V_y$	Segs	#	flat	< 5	[5, 9)	$\geq 10$	mass	Location	Size	Type
														(199, 69, 210)	5 - 9	P
														(185, 61, 368)	< 5	P
														(249, 81, 191)	< 5	P
														(189, 104, 283)	< 5	P
														(202, 60, 141)	< 5	P
														(191, 90, 191)	< 5	P
														(222, 107, 244)	< 5	P
														(244, 127, 236)	< 5	P
														(178, 82, 207)	< 5	P
														(269, 107, 254)	< 5	P
														(272, 118, 253)	< 5	P
														(223, 78, 170)	< 5	P
														(182, 74, 159)	< 5	P
														(229, 58, 133)	< 5	P
														(169, 121, 270)	< 5	P
														(246, 71, 310)	$\geq 10$	P
														(281, 75, 322)	5 - 9	P
														(236, 93, 264)	$\geq 10$	P
														(372, 68, 181)	3.5	P
														(247, 88, 279)	5 - 9	P
														(305, 125, 230)	2	P
														(261, 121, 207)	2	P
														(217, 100, 240)	4	P
													0	(215, 92, 295)	4	P
	supine	Yes	167	287	0.535	1.5	2	15	0	11	2	2		(231, 106, 221)	4	P
														(291, 90, 321)	3	P

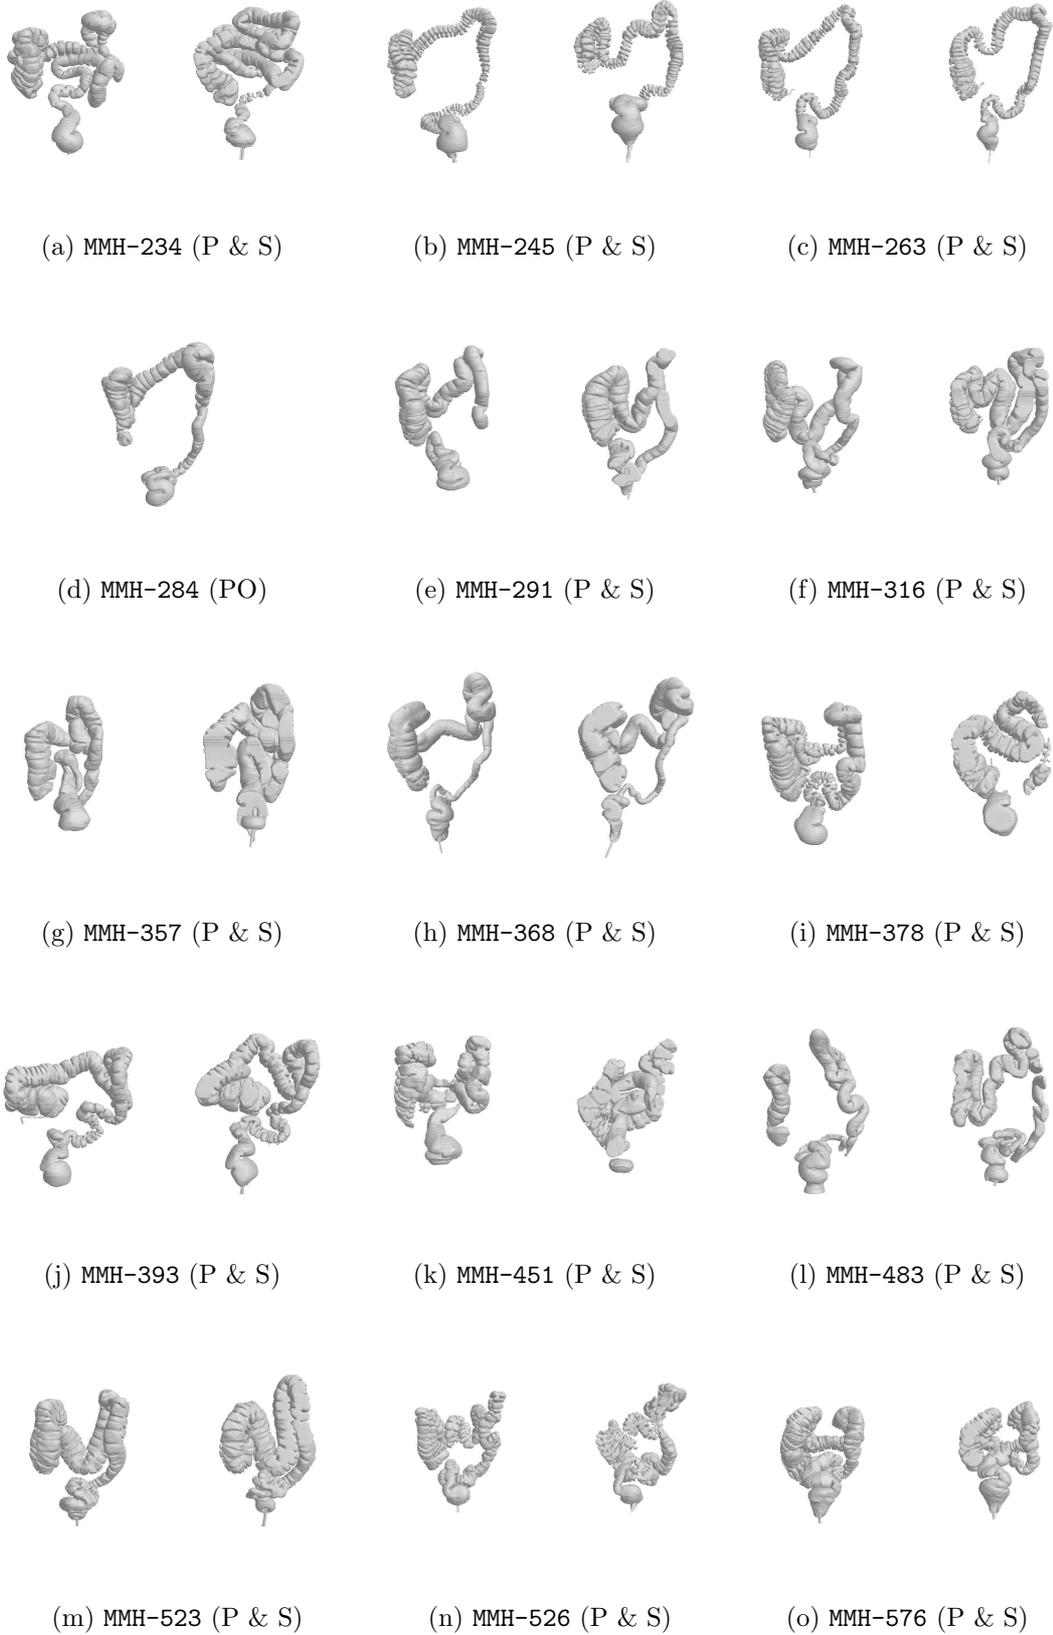
Continued on next page...

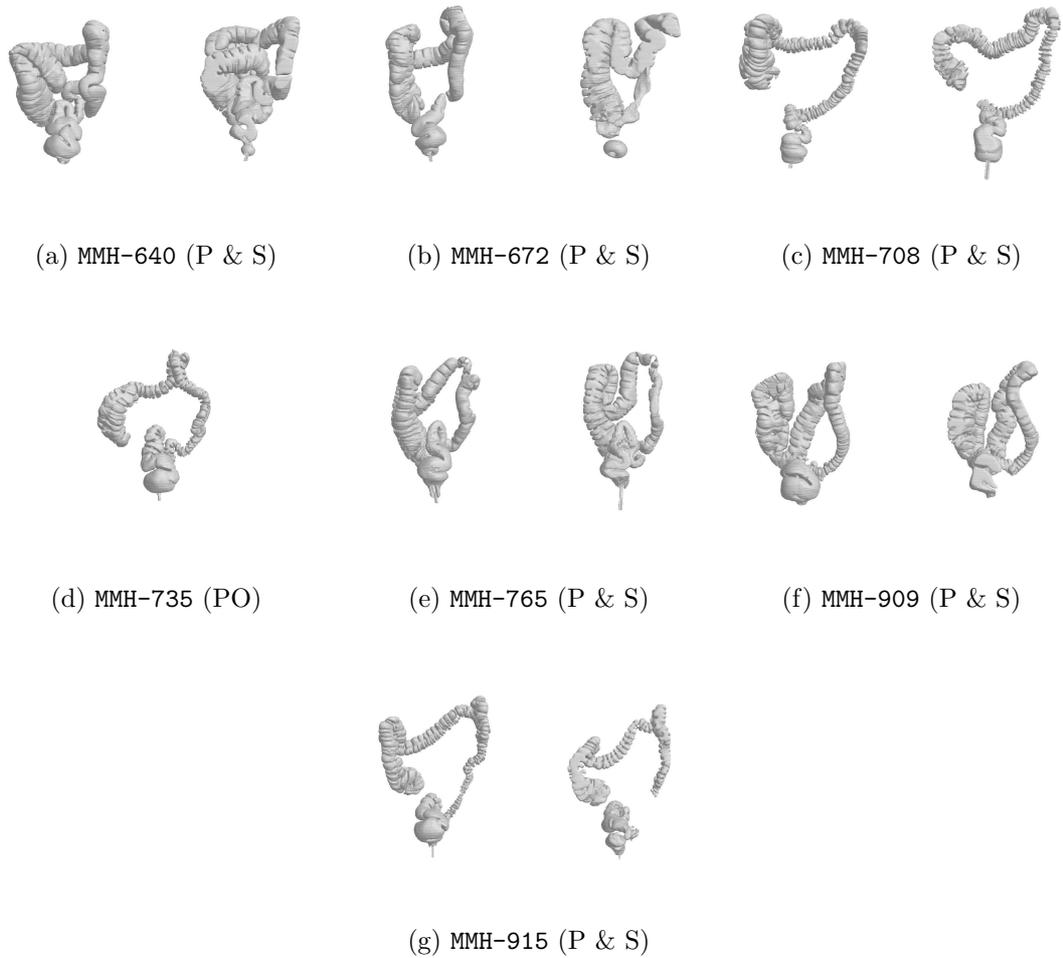
Table B.2: Continued from previous page.

Patient ID	Position	Full	Dose	Slices	$V_{xz}$	$V_y$	Segs	#	flat	< 5	[5, 9)	$\geq 10$	mass	Location	Size	Type
MMH-909														(261, 121, 208)	2	P
														(292, 126, 188)	3.5	P
	prone	Yes	250	253	0.660	1.5	2	6	0	3	2	1	0	(254, 133, 314)	5	P
														(286, 161, 329)	2	P
														(334, 135, 329)	3	P
														(347, 145, 287)	4	P
														(293, 145, 188)	4	P
														(259, 52, 174)	8	P
														(272, 42, 163)	7	P
		supine	Yes	250	250	0.707	1.5	2	11	0	8	2	1	0	(294, 65, 171)	10
MMH-915														(239, 69, 183)	3	P
														(173, 81, 153)	3	P
														(274, 132, 161)	3	P
														(341, 132, 181)	2	P
	prone	Yes	250	299	0.824	1.5	2	1	0	1	0	0	0	(273, 84, 176)	4	P
	supine	Yes	250	266	0.777	1.5	3	0	0	0	0	0	0	(316, 77, 146)	3	P
														(300, 66, 127)	2	P
														(227, 64, 151)	2	P
														-	-	-

Abbreviations: *pixel size* ( $V_{xz}$ ), *slice spacing* ( $V_y$ ), *sessile/pedunculated polyp* (P), *flat polyp* (F), *colorectal mass* (M)

**Figure B.1:** MMH data sets.

**Figure B.2:** MMH data sets (continued).



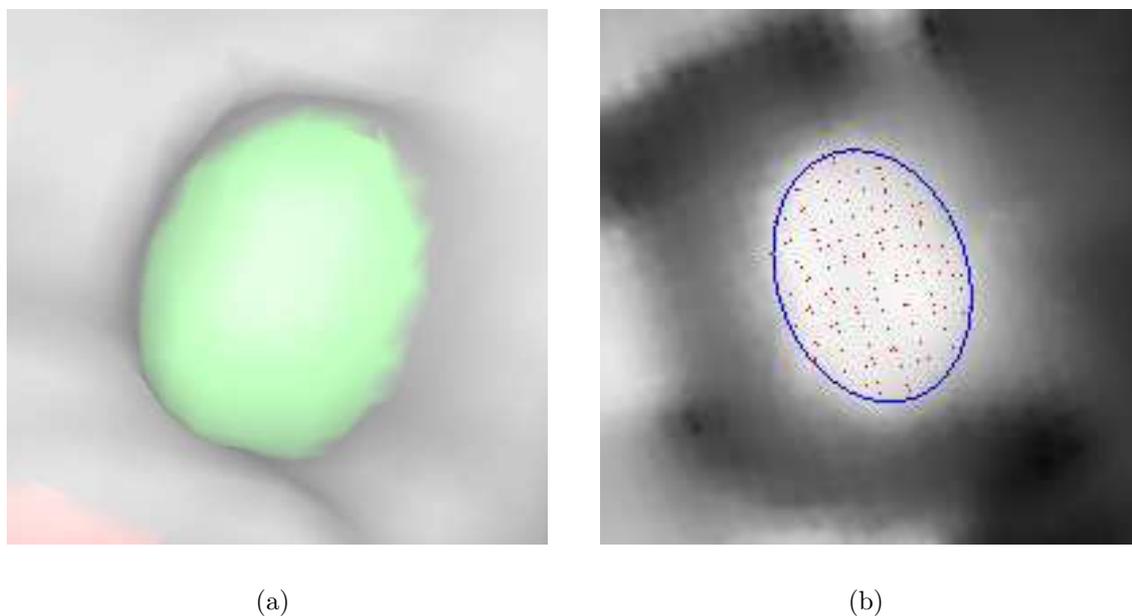
**Figure B.3:** MMH data sets (continued).

# Appendix C

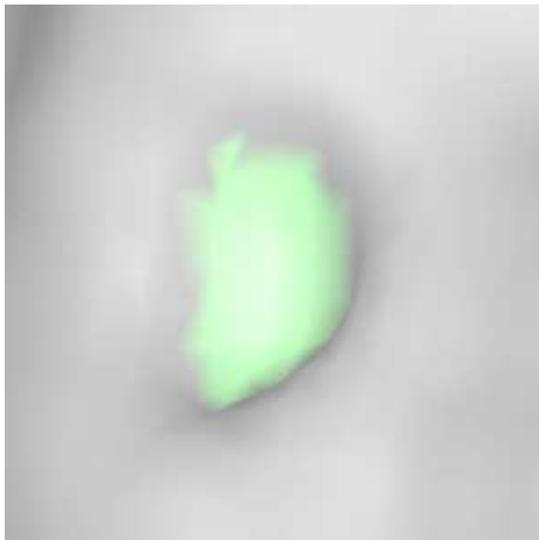
## Sample Candidates

This Appendix illustrates a range of sample polyp candidates that were detected using the initial candidate detection technique outlined in Section 5.4. The associated fingerprint and expanded candidate region are also illustrated in each case. The shape index version of the fingerprint is used as this provides the best indication of overall candidate morphology.

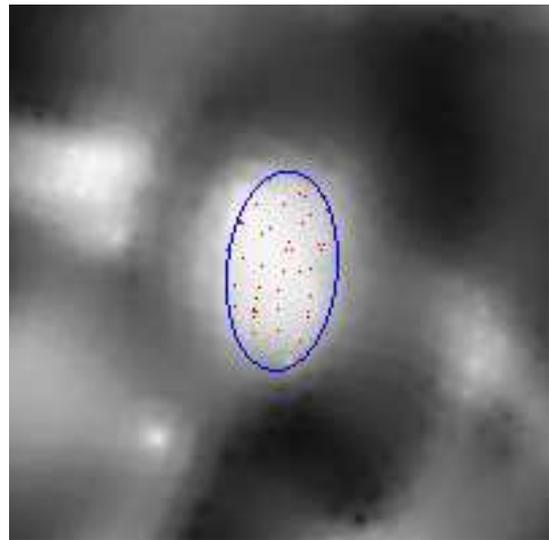
### C.1 Colorectal Polyps



**Figure C.1:** A 6 mm polyp from MMH-054/prone. The mean shape index of the original candidate is 0.924 and the mean shape index of the expanded candidate is 0.920 (a 0.4% reduction).

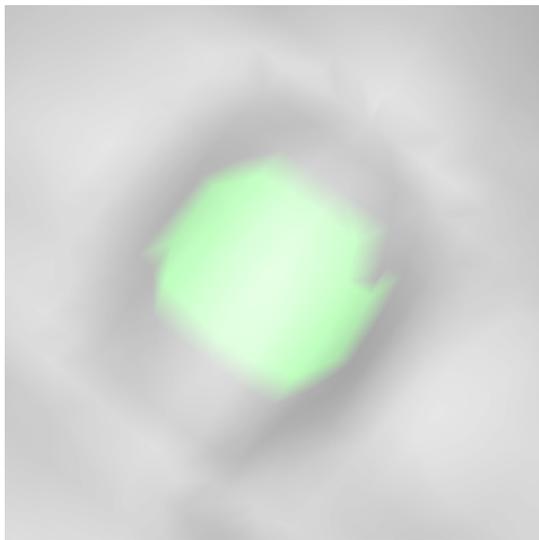


(a)

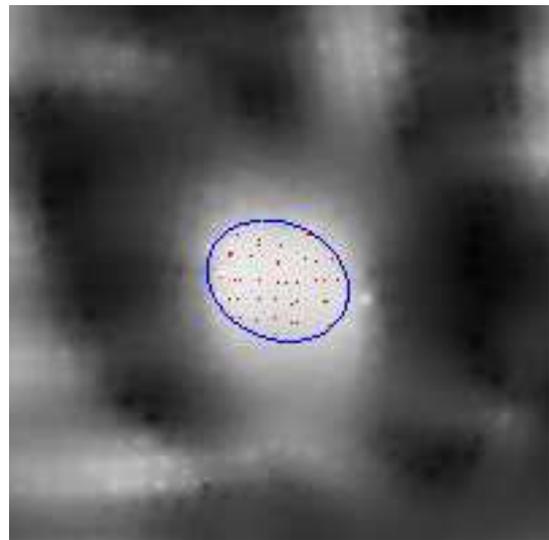


(b)

**Figure C.2:** A 3 mm polyp from MMH-208/prone. The mean shape index of the original candidate is 0.896 and the mean shape index of the expanded candidate is 0.879 (a 1.9% reduction).

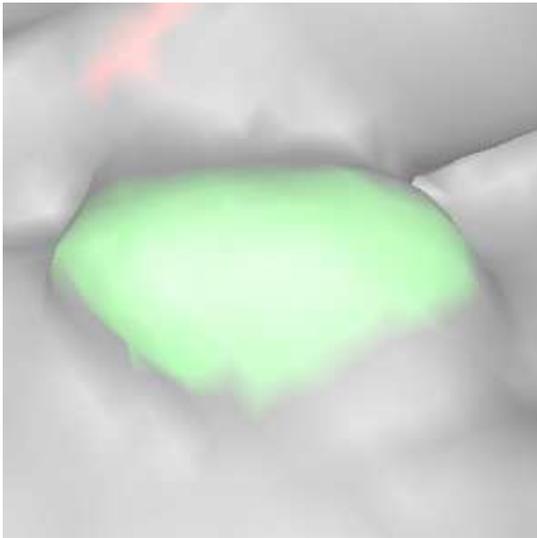


(a)

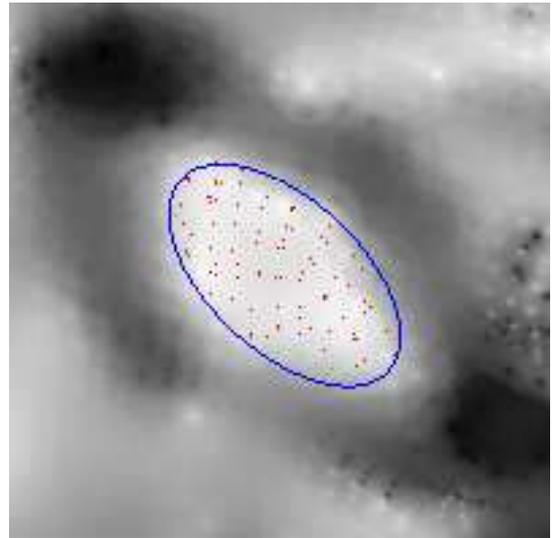


(b)

**Figure C.3:** A 4 mm polyp from MMH-161/prone. The mean shape index of the original candidate is 0.859 and the mean shape index of the expanded candidate is 0.856 (a 0.3% reduction).

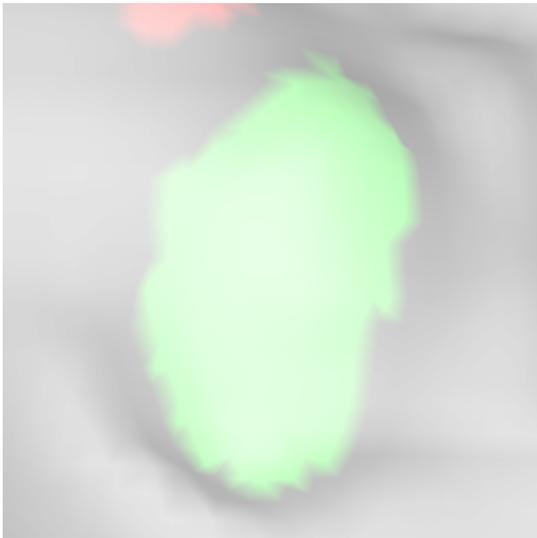


(a)

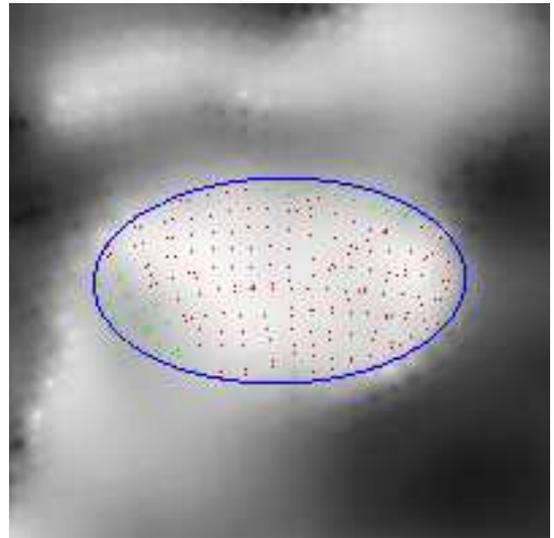


(b)

**Figure C.4:** A 6 mm polyp from MMH-245/prone. The mean shape index of the original candidate is 0.899 and the mean shape index of the expanded candidate is 0.886 (a 1.4% reduction).

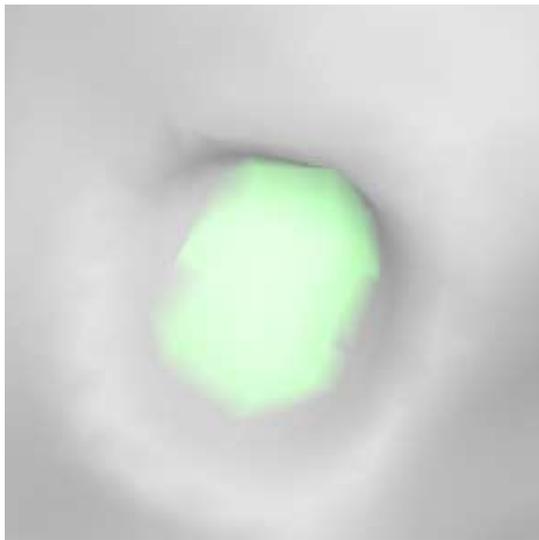


(a)

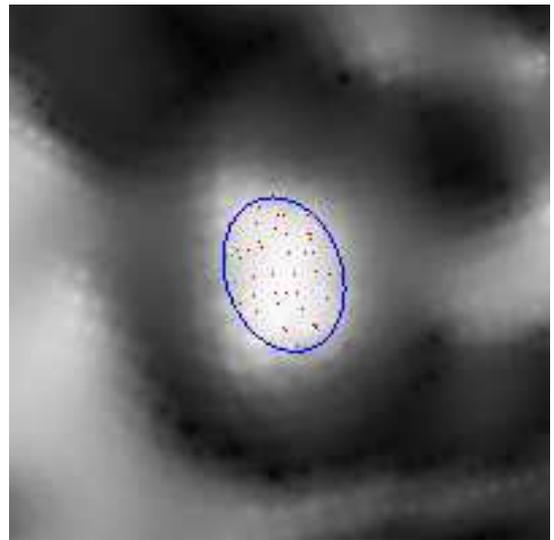


(b)

**Figure C.5:** A 7 mm polyp from MMH-357/supine. The mean shape index of the original candidate is 0.892 and the mean shape index of the expanded candidate is 0.862 (a 3.3% reduction).

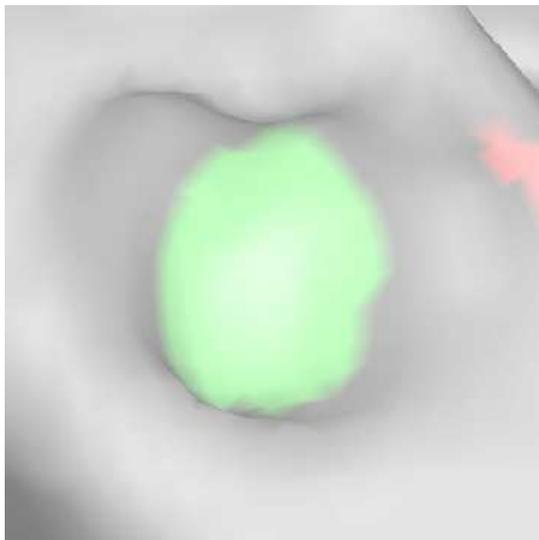


(a)

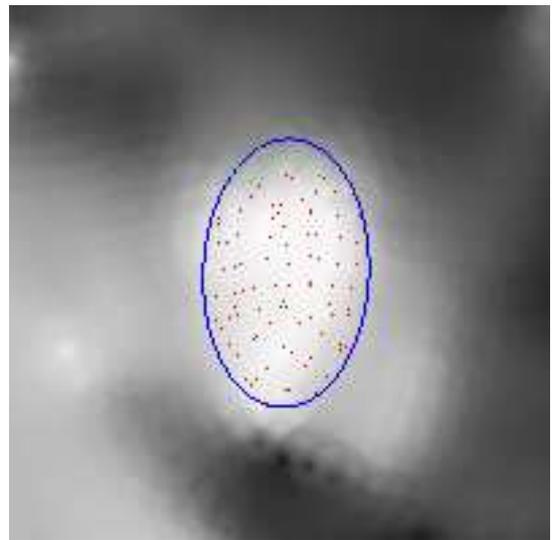


(b)

**Figure C.6:** A 3 mm polyp from MMH-708/supine. The mean shape index of the original candidate is 0.923 and the mean shape index of the expanded candidate is 0.902 (a 2.2% reduction).



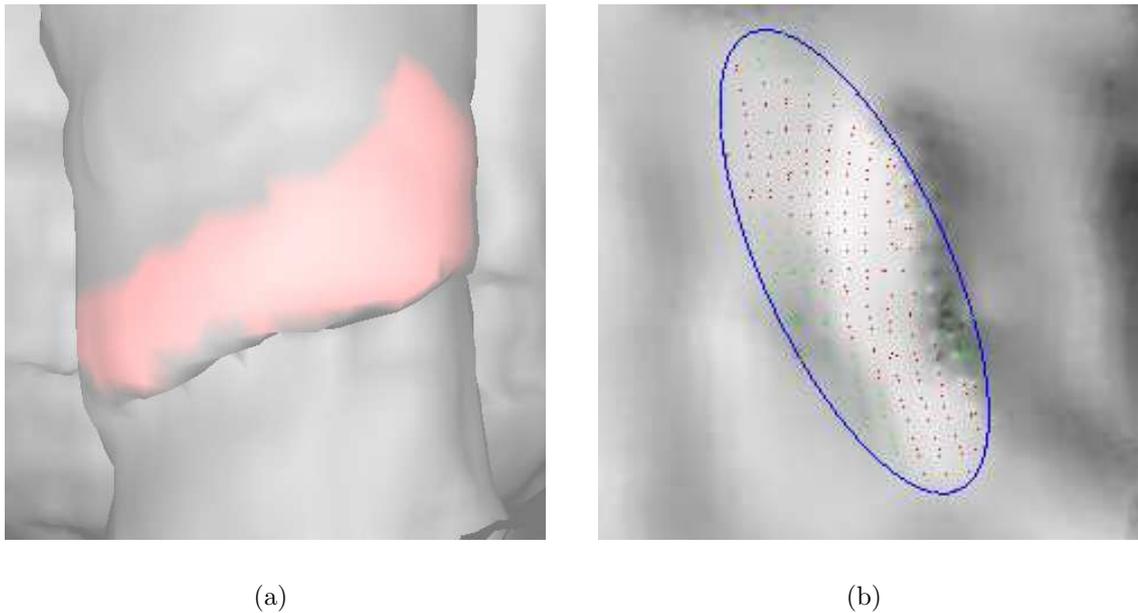
(a)



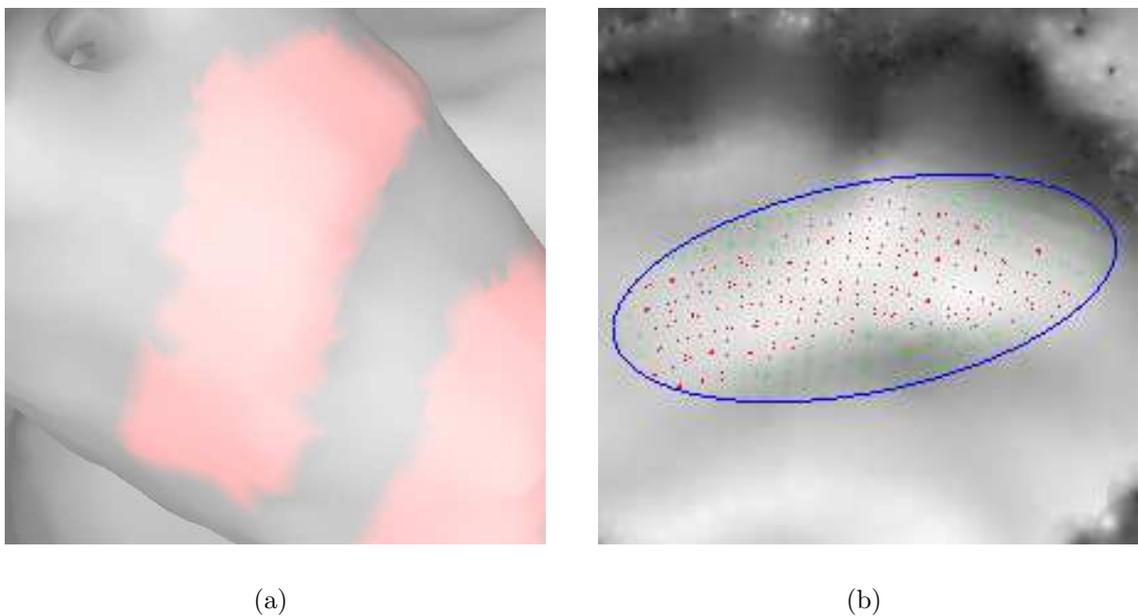
(b)

**Figure C.7:** A 4 mm polyp from MMH-735/prone. The mean shape index of the original candidate is 0.919 and the mean shape index of the expanded candidate is 0.900 (a 2.0% reduction).

## C.2 Rectal Tube

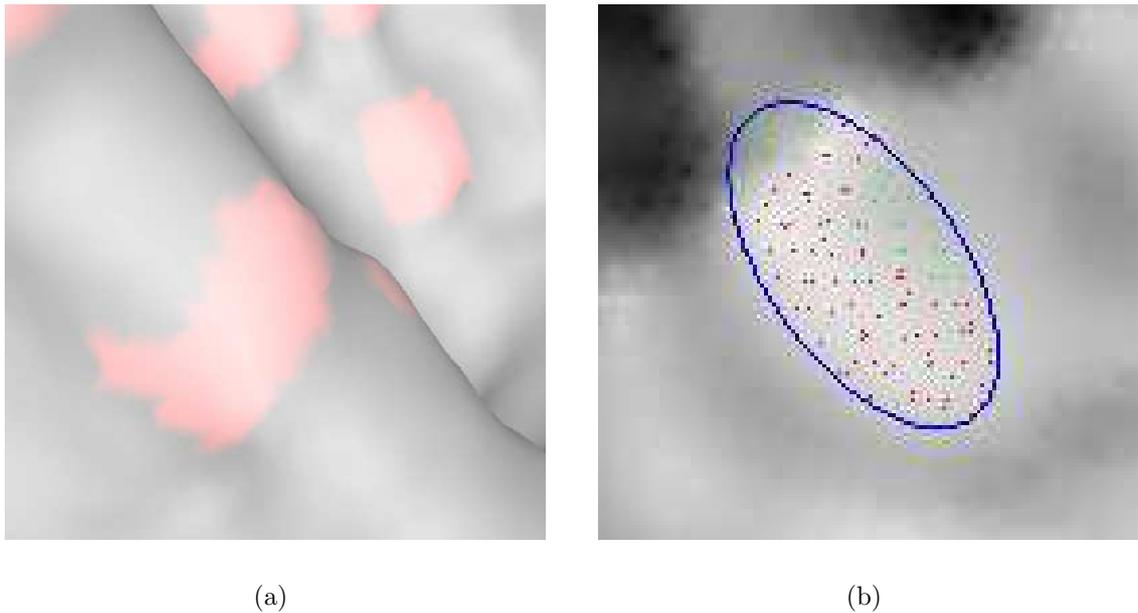


**Figure C.8:** A polyp candidate from MMH-166/prone that is attributable to the rectal tube. The mean shape index of the original candidate is 0.866 and the mean shape index of the expanded candidate is 0.784 (a 9.4% reduction).

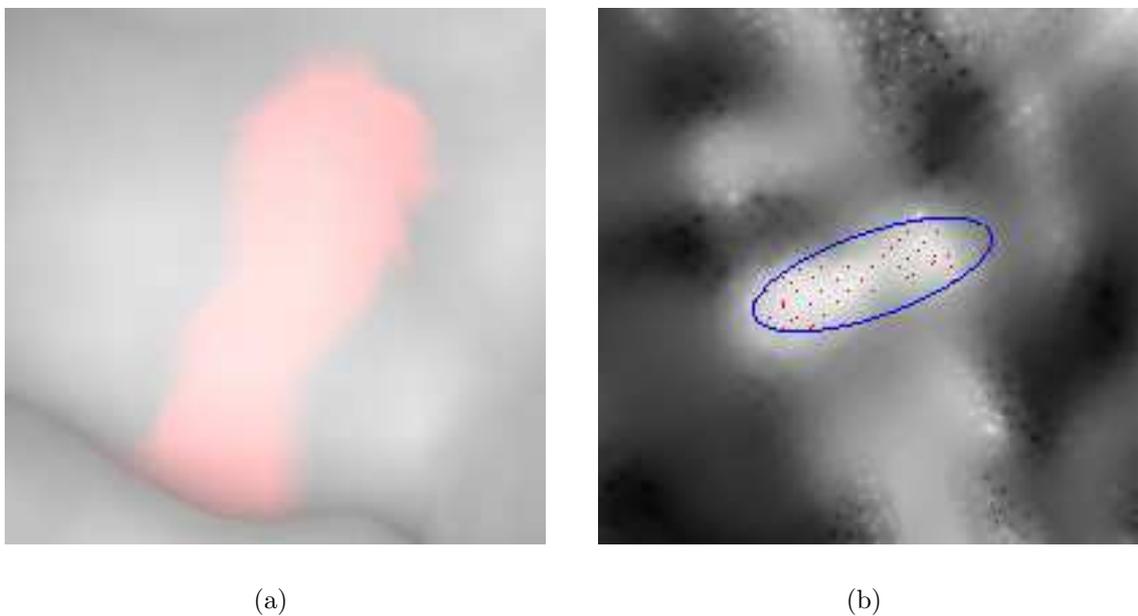


**Figure C.9:** A polyp candidate from MMH-263/prone that is attributable to the rectal tube. The mean shape index of the original candidate is 0.882 and the mean shape index of the expanded candidate is 0.806 (a 8.6% reduction).

### C.3 Haustral Folds

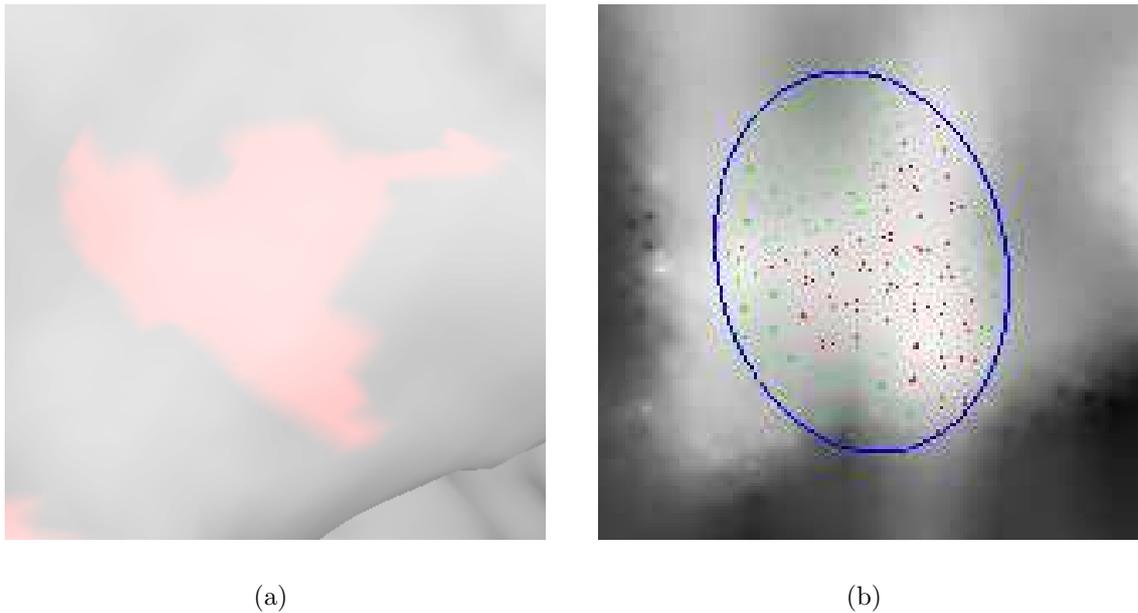


**Figure C.10:** A polyp candidate from MMH-081/supine that is attributable a haustral fold. The mean shape index of the original candidate is 0.854 and the mean shape index of the expanded candidate is 0.816 (a 4.4% reduction).

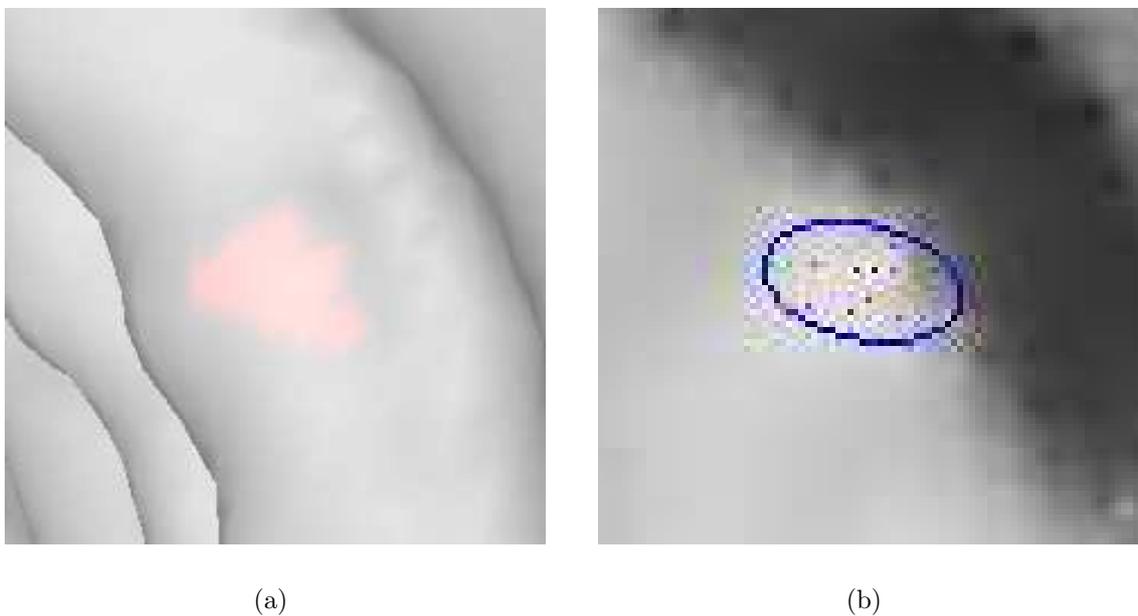


**Figure C.11:** A polyp candidate from MMH-139/supine that is attributable a haustral fold. The mean shape index of the original candidate is 0.882 and the mean shape index of the expanded candidate is 0.819 (a 7.5% reduction).

## C.4 Miscellaneous



**Figure C.12:** A polyp candidate from MMH-708/supine that is attributable to a patch of the colon surface. The mean shape index of the original candidate is 0.863 and the mean shape index of the expanded candidate is 0.778 (a 9.8% reduction).



**Figure C.13:** A polyp candidate from MMH-526/Prone that is attributable to a patch of the colon surface. The mean shape index of the original candidate is 0.844 and the mean shape index of the expanded candidate is 0.762 (a 9.7% reduction).

# Publications arising from this Work

## Refereed Journal Papers

Brennan, D., Whelan, P.F., Robinson, K., Ghita, O., **Sadleir, R.**, O'Brien, J., and S. Eustace (2005), 'Rapid automated measurement of body fat tissue distribution from whole body MRI', *American Journal of Roentgenology* **185** (2), 418-423.

**Sadleir, R. J. T.** & Whelan, P. F. (2005), 'Fast colon centreline calculation using optimised 3D topological thinning', *Computerized Medical Imaging and Graphics* **29**, 251-258.

**Sadleir, R. J. T.**, Whelan, P. F., MacMathuna, P. & Fenlon, H. M. (2004), 'A portable toolkit for providing straightforward access to medical image data', *Radiographics* **24**, 1193-1202.

## Invited Conference Papers / Presentations

**Sadleir, R. J. T.** & Whelan, P.F. (2003), 'A comparative study of selected computer aided diagnosis techniques utilised in CT colonography' The 4th International Workshop on Multislice CT, 3D Imaging and Virtual Endoscopy, Rome, November 3rd - 4th.

**Sadleir, R. J. T.**, Whelan, P. F., Sezille, N., Bruzzi, J. F., Fenlon, H. M., Moss, A. C. & MacMathuna, P. (2002) 'Automated detection and flagging of potential colorectal polyps at CT colonography' The 3rd International Workshop on Multislice CT, 3D Imaging and Virtual Endoscopy, Rome, June 6th - 8th.

## Publications Arising from this Work

### Refereed Conference Papers

Luauté, V., **Sadleir, R.J.T.**, & Whelan P.F. (2006) 'An automatic evaluation strategy for a remote access CT colonography training system' Biosignal 2006 - The 18th International EURASIP (European Association for Signal, Speech and Image Processing) Conference , Brno, Czech Republic, June 28th to 30th.

Ilea, D., Ghita, O., Robinson, K., **Sadleir, R.**, Lynch, M., Brennan, D. & Whelan, P. F. (2004) 'Identification of Body Fat Tissues in MRI Data', OPTIM 2004 9th International Conference On Optimization of Electrical and Electronic Equipment, Brasov, Romania, May 20-21, 2004.

Sezille, N., **Sadleir, R.J.T.** & Whelan, P.F. (2003) 'Fast Extraction Of Planes Normal To The Centreline From CT Colonography Datasets' VIE 2003 - IEE Visual Information Engineering Conference, Guildford, UK, July 7th - 9th

**Sadleir, R. J. T.**, Whelan, P. F., Bruzzi, J. F., Moss, A. C., Fenlon, H. M. & MacMathuna P. (2002), A novel technique for identifying the individual regions of the human colon at CT colonography, IEE - Medical Applications of Signal Processing, London, October 7th.

Sezille, N., **Sadleir, R.J.** & Whelan, P.F. (2002) 'Automated synthesis, insertion and detection of polyps for CT colonography' Opto-Ireland - SPIE's Irish meeting on optoelectronics, photonics and imaging, Galway, September 5th - 6th.

**Sadleir, R. J. T.** & Whelan, P. F. (2002), Colon centreline calculation for CT colonography using optimised 3D topological thinning, The 1st International Symposium on 3D data processing, Visualisation and Transmission, Padova, June 19th - 21st.

### Refereed Conference Abstracts (in Journals)

Luauté, V., **Sadleir, R.J.T.**, & Whelan, P.F. (2006) 'A Remote Access Training System for the Detection of Colorectal Polyps at Computed Tomography Colonography' ECR 2006 - European Congress of Radiology, Vienna, Austria, March 3rd - 7th. (European Radiology Supplement 1 to Volume 16)

## Publications Arising from this Work

**Sadleir, R. J. T.**, Whelan, P. F., Sezille, N., Chowdhury, T. A., Moss, A., Bruzzi, J., Fenlon, H. & MacMathuna, P. (2003), 'Computer-aided detection of colorectal polyps at CT colonography', *Endoscopy* **36** (presented at ISG 2003 - the Winter meeting of the Irish Society of Gastroenterology, Dublin, November 28th - 29th.)

**Sadleir, R. J. T.**, Whelan, P. F., Bruzzi, J. F., Moss, A. C., MacMathuna, P. & Fenlon, H. M. (2002), 'A portable toolkit for providing straightforward access to medical image data', *Radiology* **225**, Supplement p 762. (presented at RSNA 2002, the 88th Scientific Assembly and Annual Meeting of the Radiological Society of North America, Chicago, December 1st - 6th.)

**Sadleir, R. J.**, Moss, A., Whelan, P. F., Bruzzi, J., Fenlon, H. M. & MacMathuna, P. (2002), 'Rapid automated neoplasia detection in CT colonography using a novel technique', *Gastroenterology* **122**, Supplement 1 P M1689. (presented at DDW 2002, Digestive Diseases Week, San Francisco, May 19th - 22nd.)

**Sadleir, R. J. T.**, Moss, A. C., Whelan, P. F., Fenlon, H. M. & MacMathuna, P. (2002), 'Use of a novel technique for automated polyp detection in CT colonography', *Gut* **50**, Supplement 2, p A89. (presented at BSG 2002, the Annual Scientific Meeting of the British Society of Gastroenterology, Birmingham, March 17th - 20th.)

**Sadleir, R. J. T.**, Whelan, P.F., MacMathuna, P. & Fenlon, H. M. (2001), 'Automated detection and flagging of colorectal polyps for virtual colonoscopy: A preliminary study', *European Radiology* **11** C6. (presented at ESGAR 2001, the 12th Annual Meeting of the European Society for Gastrointestinal and Abdominal Radiology, Dublin, June 20th - 23rd.)

## Refereed Conference Abstracts (in Proceedings)

Chowdhury, T.A., **Sadleir, R.J.T.**, Whelan, P.F., Sezille, N., Moss, A., O'Hare, A., Foley, S., Fenlon, H., & MacMathuna, P. (2004) 'Automatic detection of colon at CT colonography' ISG 2004 - Winter meeting of the Irish Society of Gastroenterology, Dublin, Ireland, November 26th - 27th.

**Sadleir, R. J. T.**, Whelan, P. F., Sezille, N., Chowdhury, T. A., Bruzzi, J., Moss, A., MacMathuna, P. & Fenlon, M. (2004), 'Automatic detection of colorectal polyps

## **Publications Arising from this Work**

at CT colonography using shape information', Association of Physical Scientists in Medicine Annual Scientific Meeting, Dublin, June 11th.

Chowdhury, T.A., **Sadleir, R.J.T.**, Whelan, P.F. Moss, A. Varden, J., Short, M., Fenlon, H. & MacMathuna P. (2004), 'The impact of radiation dose on imaged polyp characteristics at CT colonography:Experiments with a synthetic phantom', Association of Physical Scientists in Medicine, 2004 Annual Scientific Meeting, Dublin, June 11th.

Malone, K.M., Molloy, D., Stack, J., Conroy, D., Egan, M., Murphy, D., **Sadleir, R.**, Whelan, P., Saykin, A. & Wishart H. (2003) 'Functional Brain Magnetic Resonance Imaging (fMRI) Study of Problem-Solving and Hope in Healthy Volunteers', 43rd Annual New Clinical Drug Evaluation Unit (NCDEU) Meeting, Boca Raton, FL, USA, May 26th - 28th.

Molloy, D., Stack, J., Conroy, D., Egan, M., Murphy, D., **Sadleir, R.**, Whelan, P., Saykin, A., Wishart, H. & Malone, K. (2003) 'A Functional Brain Magnetic Resonance Imaging (fMRI) Study of Strategic Thinking / Problem-Solving and Hope in Healthy Volunteers' Winter meeting of the Royal Academy of Medicine in Ireland, University College Dublin, Ireland, January 10th.

Whelan, P.F., **Sadleir, R.J.T.**, Robinson, K., Fenlon, H., Stack, J., & MacMathuna, P. (2002) 'Computer aided diagnosis of medical image data in gastrointestinal procedures' Healthcare Informatics Society of Ireland, 7th Annual Conference and Scientific Symposium, Dublin, Ireland, November 14th.

Moss, A., Bruzzi, J., **Sadleir, R.**, Fenlon, H. & Mac Mathuna, P. (2002) 'CT colonography for routine surveillance of patients at high-risk of colorectal carcinoma' UEGW 2002 - 10th United European Gastroenterology Week, Geneva, Switzerland, October 19th - 24th.

**Sadleir, R. J. T.**, Bruzzi, J. F., Moss, A. C., Fenlon, H. M., MacMathuna, P. & Whelan, P. F. (2002) 'Colon lumen subsegmentation: A new technique for CT colonography', the 3rd International Symposium on Virtual Colonoscopy, Boston, April 8th - 10th.

**Sadleir, R. J. T.**, Moss, A., Whelan, P. F., Bruzzi, J., Fenlon, H. & MacMathuna, P. (2001) 'Automated detection and flagging of colorectal polyps at CT (virtual)

## **Publications Arising from this Work**

colonoscopy', ISG 2001, the Summer Meeting of the Irish Society of Gastroenterology, Cork, June 8th - 9th.

Whelan, P., & **Sadleir, R.** (2000) '3-D virtual colonoscopy for the detection of colorectal cancer' IMVIP 2000 - Irish Machine Vision and Image Processing Conference, Queen's University Belfast, Northern Ireland, August 30th - September 2nd.

## **Technical Reports**

**Sadleir, R.** and Whelan, P. F. 'Virtual Colonoscopy: Proposal for an Enhanced Virtual Colonoscopy System', Dublin City University, School of Electronic Engineering, *Internal Technical Report*, 24th October 2001.

**Sadleir, R.** and Whelan, P. F. 'Virtual Colonoscopy', Dublin City University, School of Electronic Engineering, *Internal Technical Report*, 26th February 2001.

## **Publications Under Review**

### **Conference Abstracts**

Luauté, V., **Sadleir, R. J. T.**, Whelan, P. F. 'A remote access web based tutor system for reader training at CT colonography'  
(submitted to RSNA 2006 - April 13th 2006)

# Awards & Grants

## RSNA *infoRAD* Exhibit Subsidy

The paper entitled ‘A portable toolkit for providing straightforward access to medical image data’ (Sadleir et al. 2002a) received a subsidy from the RSNA. The value of the subsidy was \$582.12 and its purpose was to fund the exhibition of the paper at RSNA 2002. The exhibit was recommended for publication in RadioGraphics by a RadioGraphics. Consequently, a manuscript based on the exhibit was submitted, accepted and subsequently published in RadioGraphics (Sadleir et al. 2004a).

## EE School Research Fellowship

A grant of € 35,000 was received from the School of Electronic Engineering, Dublin City University to fund research into remote access training for CTC. This research project is ongoing and has resulted in a number of publications (Luauté et al. 2006*a,b*).

## APSM Best Presentation Prize

The paper entitled ‘Automatic detection of colorectal polyps at CT colonography using shape information’ (Sadleir et al. 2004b) received the best presentation prize at the 2004 Annual Scientific Meeting of the Association of Physical Scientist in Medicine. The award, which was in the form of a travel bursary (valued at € 500), was used in conjunction with other funding to attend RSNA 2004. A review of the papers presented at RSNA 2004, that dealt with CTC research, was subsequently submitted to the APSM for circulation to its membership:

**Sadleir, R. J. T.** & Whelan, P. F. (2005), ‘CT colonography at RSNA 2004’  
Circulated to the members of the APSM (submitted 02/02/2005).