Propositions with the Thesis

Electronic Cleansing for Visualization in CT Colonography

Iwo Serlie, Delft University of Technology, Tuesday October 9 2007.

1) Electronic colon cleansing is facilitated by an accurate description of the composition around material transitions (this thesis).

2) 3-D visualization for CT Colonography requires image processing (this thesis).

3) The unfolded cube combines an omnidirectional viewing frustum with an intuitive appearance (this thesis).

4) Knowledge develops from creativity and creativity from a lack of it.

5) A single workflow to make a diagnosis does not exist.

6) Collaboration between an engineering scientist and a medical practitioner is difficult because one thinks in terms of new possibilities and the other thinks in terms of existing solutions.

7) Intuitive user-interfaces of systems with a lot of functionality do not exist.

8) In modern society, the status derived from a device is more important than its functionality.

9) The choice of Lower Rhine societies to assimilate one of their most important indigenous gods, Magusanus, with the Roman god Hercules is not based solely on its martial aspect.

10) Wasting energy requires little energy.

These propositions are considered opposable and defendable and as such have been approved by the supervisor, prof.dr.ir. L.J. van Vliet.

Stellingen behorende bij het proefschrift

Electronic Cleansing for Visualization in CT Colonography

Iwo Serlie, Delft University of Technology, dinsdag 9 oktober 2007.

1) Digitale darmreiniging is mogelijk dankzij een nauwkeurige beschrijving van de samenstelling rond materiaalovergangen (dit proefschrift).

2) 3D visualisatie voor CT Colonografie kan niet zonder beeldverwerking (dit proefschrift).

3) De opengeklapte kubus combineert een omnidirectioneel blikveld met een intuïtieve manier van weergeven (dit proefschrift).

4) Kennis komt voort uit creativiteit en creativiteit uit een gebrek hieraan.

5) Een eenduidige werkwijze bij het stellen van een diagnose bestaat niet.

6) De samenwerking tussen een technisch-wetenschappelijk onderzoeker en een behandelend geneesheer is moeilijk omdat de een denkt in nieuwe mogelijkheden en de ander in bestaande oplossingen.

7) Intuïtieve user-interfaces van systemen met veel functionaliteit bestaan niet.

8) In de moderne samenleving is de status die ontleend wordt aan een apparaat belangrijker dan zijn functionaliteit.

9) De keuze van de Neder-Rijnse samenlevingen om één van hun belangrijkste goden, Magusanus, gelijk te stellen aan de Romeinse god Hercules, is niet gebaseerd op louter het martiale aspect.

10) Energieverspilling vergt weinig energie.

Deze stellingen worden opponeerbaar en verdedigbaar geacht en zijn als zodanig goedgekeurd door de promotor, prof.dr.ir. L.J. van Vliet.

Electronic Cleansing for Visualization in CT Colonography

Electronic Cleansing for Visualization in CT Colonography

Proefschrift

Ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof.dr.ir. J.T. Fokkema,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op dinsdag 9 oktober 2007 om 10 uur

door

Iwo Willem Oscar SERLIE

Informatica Ingenieur
geboren te Woerden

Dit proefschrift is goedgekeurd door de promotor:

prof.dr.ir. L.J. van Vliet

Samenstelling promotiecommissie:

Rector Magnificus		voorzitter
prof.dr.ir. L.J. van Vliet	Delft University of Technology,	promotor
dr. F.M. Vos	Delft University of Technology, Academic Medical Center,	
ir. F.H. Post	Delft University of Technology	
prof.dr. A.M. Vossepoel	Delft University of Technology, Erasmus MC - University Medical Center Rotterdam	
prof.dr.ir. C.A. Grimbergen	Delft University of Technology, Academic Medical Center	
prof.dr. J. Stoker	Academic Medical Center, University of Amsterdam	
prof.dr.ir. B.M. ter Haar Romeny	Eindhoven University of Technology	
prof.dr. I.T. Young	Delft University of Technology,	reservelid

Contents

Chapter One

Introduction

Computed (axial) tomography (CT) is a non-invasive method to generate three-dimensional (3-D) images of an object [1]. Hounsfield created the CT device in 1972 and received the Nobel Prize in medicine in 1979. The first 2-D cross-sections were composed of 80×80 pixels. Today's 3-D images are composed of 512^3 pixels and provide details typically in the range of $0.5 - 1.0$ mm. Reduction in acquisition time has also been responsible for the rise in popularity of CT imaging. The first images were acquired in four minutes. Today, multiple contiguous images are acquired in less than one second enabling breath-hold imaging and imaging of the heart.

From this ever increasing source of image based information, it is the medical specialist's challenge to effectively and efficiently extract relevant diagnostic information. Humans are well equipped to extract qualitative information, but less equipped to extract quantitative information. The main problem is the massive amount of data that needs to be searched through to extract relevant diagnostic information. In reply to this problem medical image processing aims at assisting the specialist in gaining "insight into the data". In order to provide the best assistance we need to get the information from the images in a reproducible best manner.

This thesis focuses on the application of CT colonography that can benefit substantially from advances in medical image processing. It is a method that aims at finding the latent stage of colon cancer (adenomatous polyp) or colorectal cancer itself. Traditionally, getting the diagnostic information is based upon an endoscope that is inserted into the digestive system, while CT colonography provides a non-invasive alternative.

To enable CT colonography to become the preferred primary method for colon examination, medical image processing techniques are created to enable the medical specialist to focus upon the primary task of discriminating between the normal and pathological condition.

1.1   Colon Cancer

Colorectal cancer (CRC) is an often deadly disease in the Western World. It accounts for 11% to 15% of all cases of cancer in this region (~100,000 diagnosed cases and ~50,000 deaths in the US every year) [2]. CRC (Figure 1.2c) is a malignant tumor that after a relatively long pre-malignant period of localization in the bowel wall invades the wall and metastasizes to the lymph nodes and other parts of the body [3]. The risk of CRC increases with age and occurs most frequently in the proximal colon [4].

The large bowel or colon (Figure 1.1) is part of the digestive tract and approximately two meters long. The main functions of the colon are absorption of water from digested materials and absorption of several synthesized vitamins.

Name of segments:
a   cecum (proximal colon)
b   colon ascendens
c   flexura hepatica
d   colon transversum
e   flexura linealis
f   colon descendens
g   sigmoid (distal colon)
h   rectum

mucosa

plica semilunaris
haustra

(a)   (b)
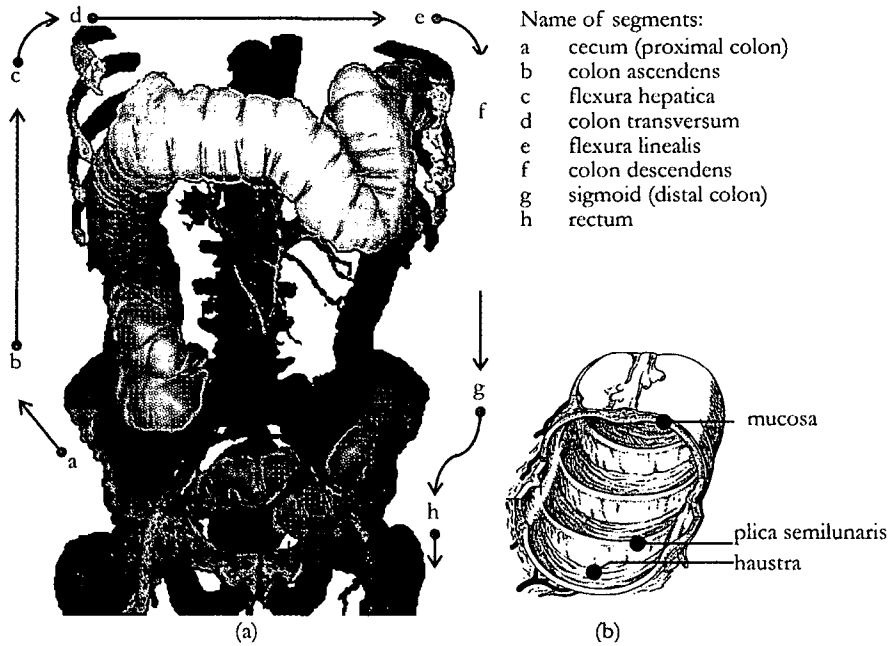
Figure 1.1   (a) Structure of the colon and (b) detailed anatomy.

(a)   (b)   (c)
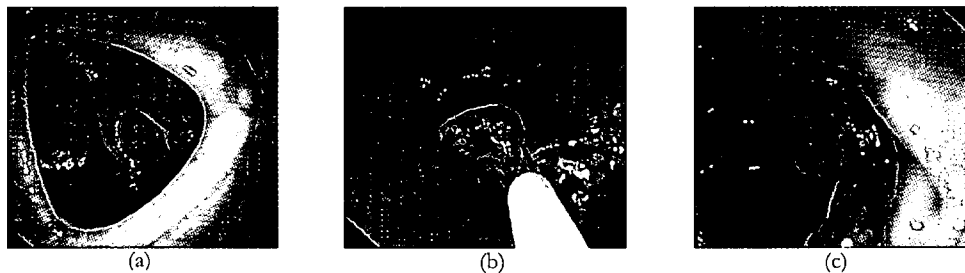
Figure 1.2   (a) Normal finding in the colon, (b) polyp just about to be removed and (c) a tumor (CRC).

Colorectal cancer (Figure 1.2c) can be prevented due to two important aspects.

First, the development from latent to diagnosed disease is understood. *Polyps* are masses of mucosa (Figure 1.2b) found in the colon that differ in clinical importance and histology (cellular formation). Genetic alterations are supposed to lead to an overgrowth of cells (adenomatous polyps), and further genetic changes are thought to be responsible for the progression of these polyps to cancer (Figure 1.2c) [5]. Approximately one-half to two-thirds of all polyps are adenomatous (adenomas), although most likely only a

minority will develop into cancer [6][7]. The probability that invasive cancer is contained within a colorectal adenomatous polyp increases with the size of the adenoma and the degree of dysplasia (abnormal tissue growth) [8]. The progression from the adenoma to CRC has a natural history of approximately 10 years [9][10]. However, the duration of the pre-symptomatic stage, has been estimated between 4.5 and 5 years [11].

Second, a suitable examination may exploit the long pre-malignant stage to detect polyps. In addition there is an accepted treatment [12]. It is a relatively simple and painless procedure to remove a polyp (Figure 1.2b): there are no nerve endings in the mucosa. Detecting and removing adenomas, at an early pre-symptomatic stage, will reduce CRC incidence [13].

1.2 Colonoscopy

The standard examination and treatment method is (optical) colonoscopy. A tube with a bright light is moved into the distal part of the colon after extensive bowel cleansing. The American Cancer Society recommends that every adult over 50 years old, with a family history of colon cancer or having pre-existing conditions, should have colonoscopy.

However, the overall benefit of the method when applied to a screening population with low probability of having a polyp does not outweigh the potential harm and discomfort from its application. A serious complication associated with this invasive colonoscopy is a perforation (1/10,000) that requires immediate surgery [14]. In addition colonoscopy requires extensive bowel preparation and demands great technical skills on the part of the examiner.

Another problem arises due to the shape of the colon (Figure 1.1a). With optical colonoscopy, due to narrow areas or the curved shape, the proximal colon ('start' of colon), is not reached in all procedures [15]. Even if the colonoscope can reach the proximal colon, surface parts are out of sight due to the unidirectional nature of the colonoscope and the folds in the colon surface [16](Figure 1.5).

1.3 CT colonography

CT colonography is a newer colorectal screening method that uses colon cleansing, inflation and non-invasive 3-D imaging to view the inside of the colon. It overcomes some important problems of colonoscopy mentioned before.

The general procedure comprises the following steps. First, the patient's colon is cleansed to remove fecal remains. Second, the colon is distended using an automated carbon dioxide inflation technique and scanned using multi-detector spiral CT imaging in supine and prone orientation. Third, the colon data are examined by a specialist. After having been scanned it takes about 15 minutes for the medical specialist to examine a colon using CT colonography [17]. If polyps with a 10 mm or larger diameter are found, the patient must undergo a conventional colonoscopy so that they can be removed.
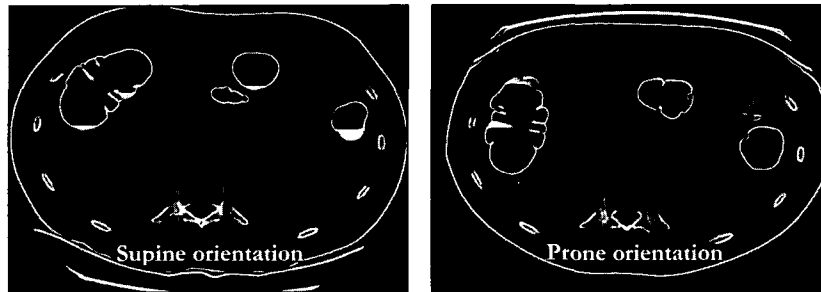
Figure 1.3   CT images of the colon. Air is black, tissue is grey and intraluminal remains are white if tagged or tissue equivalent if not tagged.

It has been demonstrated that the probability of dying from colorectal cancer by removing polyps is reduced from 3.06 to 2.30 (lifetime %) when participating in a screening program [14]. There is a very small potential risk associated with follow up colonoscopy if polyps are diagnosed (e.g. due to a perforation) as the probability of death due to colonoscopy is 0.005 (lifetime %). Because of these risks and the very low likelihood of malignancy in polyps with a 5 mm or smaller diameter, these polyps are generally not considered clinically important and rescreening should be done at more than five-year intervals.

1.3.1   2-D and 3-D examination

The first method to inspect the inside of the colon is direct analysis of axial CT slices (Figure 1.3). Virtual-colonoscopy applies image processing to delineate the colon wall and surface rendering to create a 3-D view of the colon using a virtual camera (Figure 1.4b) mimicking the optical endoscope (Figure 1.4a). The field was pioneered by Vining *et al.* in 1993 on virtual bronchoscopy [18] and in 1994 on virtual colonoscopy [19][20], and by Geiger and Kikinis on virtual bronchoscopy in 1994 [21]. This frontier work rapidly matured and recent results suggest that the virtual colonoscopy procedure approaches the sensitivity achieved by colonoscopy [22].
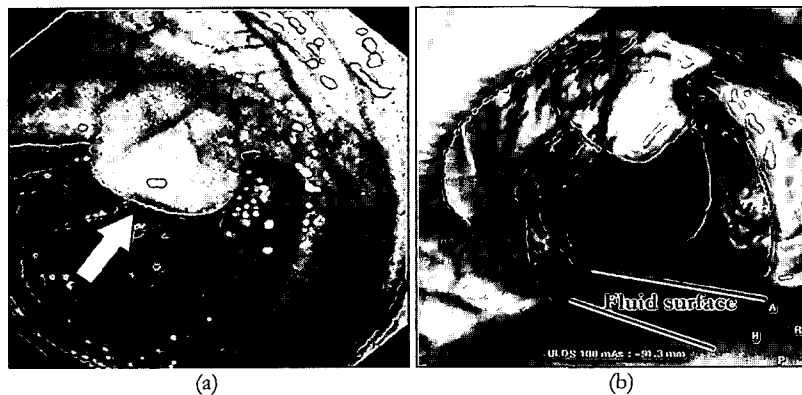
Figure 1.4   (a) Optical colonoscopy (OC) and (b) virtual colonoscopy (VC) showing the same polyp.

1.3.2   3-D examination with electronic cleansing

In virtual colonoscopy fluid layers may block the 3-D view of the colon wall (Figure 1.4b). Normally, laxatives are applied to reduce this problem. In this thesis image processing is applied to transform the medical data into electronically cleansed data enabling a clear view of the colon in spite of remaining fluids [23]. Several technologies enable 3-D examination in CT colonography:

Man machine interface. To begin with, visualization (preferably real-time) enables the interactive examination of segmented surfaces. The specialist can interact with the camera similar to interaction with the optical camera [24].

Segmentation of the colon-surface. Segmentation assists the specialist in reducing the 3-D volume into a 2-D surface. Techniques are developed that allow fully automatic determination of the surface. It enables the 3-D display of a segmented surface as visible in (Figure 1.4b).

Automatic path tracking. In addition, image processing assists in defining an optimal trajectory through the colon. This trajectory serves two purposes. First, the virtual camera navigates along the path in the 3-D volume. Second, the 2-D surface is reduced to a 1-D coordinate system with the distance to the endpoint as a variable. This allows annotation of the position of a lesion. The distance along the path is an estimate of the position in the colon that is used with an (optical) endoscope to remove the polyp.

Fecal Tagging. Fecal tagging is introduced to mark the fecal remains using orally administered contrast to be able to discriminate between tissue and fecal remains (Figure 1.3).

Supine and prone matching. Two images are acquired, one in supine and one in prone position to solve the problem of collapsed colon-segments (Figure 1.3). Image processing assists in the registration of both images. This may allows the specialist to combine information of a single polyp that is present in two images.

1.4   Research objectives

Despite clear potential and progress in visualization techniques and image processing, two main problems emerged that remain critical to the success of CT colonography. (1) The display mode was not optimal and (2) the segmentation of the colon had to be enhanced. This thesis focuses on creating image processing tools to solve these problems to allow the physicians to benefit from the advances in CT imaging.

The primary requirement of the designed image processing methods is that they allow the specialist to focus upon the principle task: discriminate between a normal and a pathological condition.

1.4.1   Unfolded cube visualization

Unfolded cube visualization is a medical imaging technique that projects six overview directions onto the six sides of the fold-up of a cube. The unfolded cube substantially reduces the problem of blinded areas caused by folds when using forward and/or backward viewing (Figure 1.5) [25].
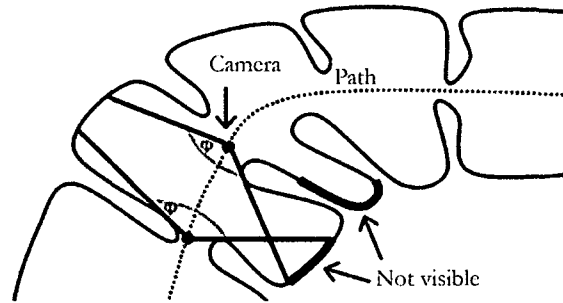
Figure 1.5   Blinded areas in virtual colonoscopy using bidirectional viewing.

It has been proven that this 3-D view of a segmented colon using six perspective projections is more efficient than browsing through a stack of 2-D image slices. The unfolded cube method and validation is described in [17] (chapter three). Other publications on the unfolded cube are [26][27][28].

(a)   (b)
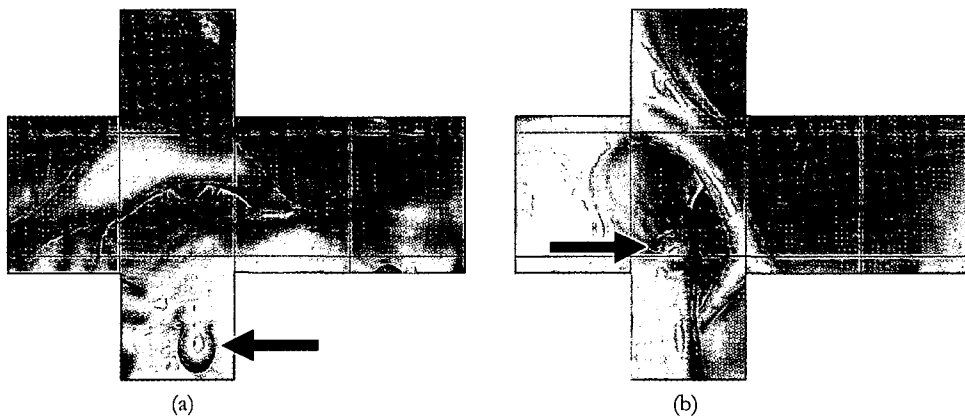
Figure 1.6   An example of the unfolded cube rendering using (a) supine and (b) prone acquisitions showing a polyp. The visualization is based upon iso material fraction rendering (see 1.4.2).

<center>(a)</center> <center>(b)</center>

Figure 1.7   (a) An artifact occurs at a junction due to using a two-material model, (b) which is suppressed by using a three-material model and iso-material fractions.

1.4.2   Volume visualization using material fractions

Volume visualization of a surface with constant tissue fraction is introduced in order to better approximate the surface of an object than with standard techniques [29]. It allows medical imaging to assist the specialist in automatic segmentation of objects even if they connect to more than one other material. Figure 1.6 and Figure 1.7 show colon imaging based upon iso-material fraction rendering.

Traditional volume visualization uses a surface through constant datavalue to approximate the object-surface. In 3-D colon imaging this may lead to fluid blocking the view of the colon surface. Occlusion of polyps will lower the sensitivity (true-positive fraction) and the detection of polyp-like structures will lower the specificity (1 − false-positive fraction).

Two advantages of using iso-material fractions when compared to traditional iso-surface imaging are:

1.   Artifacts that typically occur at locations where three materials meet are suppressed.

2.   Traditionally, the user has to manually define an opacity function. Surfaces through constant material fraction simplify automatic surface segmentation.

In general, the method determines the relative contributions of the materials to the voxel-value based on a modeled relation of differential invariants across a transition. A patent is described in [30].

The method to determine relative contributions across two-material transitions is described in chapter four. Three-material transitions are addressed in chapter five. An evaluation using optical colonoscopic data is conducted and described in chapter six.

1.5 References

[1]   G.N. Hounsfield, A method of an apparatus for examination of a body by radiation such as x-ray or gamma radiation, *The Patent Office. UK*, 1972.

[2]   National Cancer Institute, Cancer Statistics, *in: U.S. National Institute of Health*.

[3]   Australian Health Technology Advisory Committee, Colorectal cancer screening, cat. no. 9801293, *Australian Government Printing Service*, 1997.

[4]   National Cancer Institute of Canada, Canadian Cancer Statistics 2001, *in: Annual demographic statistics* 2001.

[5]   S.J. Winawer, R.H. Fletcher, L. Miller, F. Godlee, *et al.*, Colorectal cancer screening: clinical guidelines and rationale, *in: Gastroenterology*, vol. 112, 1997, pp. 594–642.

[6]   M.J. O'Brien, S.J. Winawer, A.G. Zauber, L.S. Gottlieb, S.S. Sternberg, B. Diaz, G.R. Dickerson,   The National Polyp Study: Patient and polyp characteristics associated with high-grade dysplasia in colorectal adenomas, *in Gastroenterology*, vol. 98, 1990, pp. 371–379.

[7]   M.H. Vatn, H. Stalsberg, The prevalence of polyps of the large intestine in Oslo: an autopsy study, *in: Cancer*, vol. 49, 1981, pp. 819–825.

[8]   J.J. Decosse, G.J. Tsioulias, J.S. Jacobson. Colorectal Cancer: Detection, Treatment, and Rehabilitation, *in: CA Cancer J Clin*, vol. 44, 1994, pp. 27–42.

[9]   P.J. Prichard, J.J. Tjandra, Colorectal cancer, *Med J*, vol. 169, 1998, pp. 493–498.

[10]   S.J. Winawer, E.T. Stewart, A.G. Zauber, J.H. Bond, *et al.*, A comparison of colonoscopy and doublecontrast barium enema for surveillance after polypectomy, *in: N Engl J Med*, vol. 342, 2000, pp. 1766–1772.

[11]   G. Launoy, T.C. Smith, S.W. Duffy, V. Bouvier. Colorectal cancer mass-screening: estimation of faecal occult blood test sensitivity, taking into account cancer mean sojourn time, *in: Int J Cancer*, vol. 73, 1997, pp. 220–224.

[12]   S. Shapiro, Goals of screening. *in: Cancer*, vol. 70, 1992, pp. 1252–1258.

[13]   L. Gordis. *Epidemiology. Philadelphia: W.B. Saunders, 1416025308*, 1996, p. 229.

[14]   S.A. Lieberman, D.G. Weiss, J.H. Bond, D.J. Ahnen, *et al.*, Use of colonoscopy to screen asymptomatic adults for colorectal cancer, *in: N Engl J Med*, vol. 343, no. 3, 2000, pp. 162–168.

[15]   S. Cherian, P. Singh, Is routine ileoscopy a useful? An observational study of procedure times, diagnostic yield, and learning curve, *Am J Gastroenterology*, vol. 99, 2004, pp. 2324–2329.

[16]   P.J. Pickhardt, P.A. Nugent, P.A. Mysliwiec, J.R.Choi, W.R. Schindler, Location of Adonomas Missed by Optical Colonoscopy, *in: Ann Intern Med*, vol 141(5), 2004, pp. 352–359.

[17]   F.M. Vos, R.E. van Gelder, I.W.O. Serlie, J. Florie, C.Y. Nio, A.S. Glas, F.H. Post, R. Truyen, F.A. Gerritsen, J. Stoker, Three-dimensional display modes for CT colonography: conventional 3D virtual colonoscopy versus unfolded cube projection, *Radiology*, vol. 228 no. 3, 2003, pp. 878–885.

[18]   D.J. Vining, A.R. Padhani, S.Wood, *et al.*, Virtual Bronchoscopy, A new perspective for viewing the tracheobronchial tree, *in: Radiology*, vol. 189(P), 1993, p. 438.

[19]   D.J. Vining, D.W. Gelfand, Noninvasive colonoscopy using helical CT scanning, 3D reconstruction and virtual reality, *in: SGR Scientific program*, 1994.

[20] D.J. Vining, R. Shifrin, E. Girshaw, K. Liu, R. Choplin, Virtual Colonoscopy, *in Radiology*, vol. 193, 1994, p. 446.

[21] B. Geiger, R. Kikinis, Simulation of Endoscopy, *in: AAAI Spring Symposium Series: Application of Computer Vision in Medical Image Processing*, 1994, pp. 138–140.

[22] S. Halligan, H.M. Fenlon, Virtual colonoscopy, *BMJ*, vol. 319, 1999, pp. 1249–1252.

[23] P.J. Pickhardt, J.R. Choi, I. Hwang, J.A. Butler, M.L. Puckett, H.A. Hildebrandt, R.K. Wong, P.A. Nugent, P.A. Mysliwiec, W.R. Schindler, Computed tomographic virtual colonoscopy to screen for colorectal neoplasia in asymptomatic adults, *in: N Engl J Med.*, vol. 349, 2003, pp. 2191–2200.

[24] L. Wei, A. Kaufman, K. Kreeger, Real-Time Volume Rendering for Virtual Colonoscopy, *in: proc. of 2nd Workshop on Volume Graphics*, 2001, pp. 27–33.

[25] I.W.O. Serlie, F.M. Vos, R. van Gelder, J. Stoker, R. Truyen, F. Gerritsen, Y. Nio, F. H. Post, "Improved Visualization in Virtual Colonoscopy Using Image-Based Rendering", *in: Data Visualization 2001*, proc. of the Joint Eurographics and IEEE TCVG Symposium on Visualization, 2001, pp. 137–146.

[26] F.M. Vos, I.W.O. Serlie, R.E. van Gelder, H. Vrooman, F. H. Post, A review of technical advances in virtual colonoscopy, *in: Proceedings of MedInfo*, 2001, pp. 938-942.

[27] F.M. Vos, I.W.O. Serlie, R.E. van Gelder, F.H. Post, R. Truyen, F.A. Gerritsen, J. Stoker, A.M. Vossepoel, A New Visualization Method for Virtual Colonoscopy", *in: LNCS*, vol. 2208, MICCAI, 2001, pp. 645–654.

[28] F.A. Gerritsen, R. Truyen, I.W.O. Serlie, Method, system and computer program for producing a medical report, *in: US2002190980, EP1393259*.

[29] I.W.O. Serlie, R. Truyen, J. Florie, F.H. Post, L.J. van Vliet, F.M. Vos, Computed cleansing for virtual colonoscopy using a three-material transition model, *in: LNCS*, vol. 2879, MICCAI, 2003, pp. 175–183.

[30] R. Truyen, I.W.O. Serlie, Volume visualization using tissue mix, US2006143215.

Chapter Two

CT Imaging Characteristics

This chapter tests several basic assumptions on which the proposed methods in this thesis rely. It explores basic aspects of CT imaging for Colonography: the noise characteristics, the spatial resolution and the model that describes the CT value across a material transition [1]. It is important to notice that the image processing discussed in this chapter is restricted to the edges of the colon. When applied to other areas or sources of data, the assumptions that are made have to be tested again.

2.1   Introduction

Computed Tomography (CT) is based upon measuring x-ray attenuation (absorption and scattering) characteristics of an object. First, x-ray projections, i.e. transmission measurements along lines - actually thin stripes - through the object, are acquired as the x-ray source rotates around the object. Second, the projections are used as input data in a tomographic reconstruction algorithm to calculate a 3-D CT image consisting of CT values on a grid of voxels.

For mono-energetic radiation, the measured intensity at an x-ray detector is described by

$$I_\theta = I_0 \exp\left(-\int \mu(z)\,dz\right) = I_0 \exp(-Z) \tag{2.1}$$

where I_θ is the measured intensity, I_0 the intensity of the unattenuated x-ray beam and $\mu(z)$ the attenuation coefficient at position z. (2.2) The integrated beam attenuation Z along the path of the ray is given by

$$Z = \ln\left(\frac{I_0}{I_\theta}\right) \tag{2.2}$$

However, the imaging process is not ideal and the detectors have finite sizes. Consequently, each measured signal value emanates from a compound of materials instead of pure materials. This generally yields an underestimation of the estimated density [2]. The larger the attenuation gradients perpendicular to the beam, the larger the underestimation will be.

Because the x-ray beam is not mono-energetic, and the attenuation coefficients are dependent on the energy of radiation, equation (2.1) is only an approximation. Current CT reconstruction algorithms use a nonlinear function to map the Z-values determined with (2.2) on values expected for mono-energetic radiation. In this mapping procedure, it is usually assumed that the attenuating object consists of water. Small errors in the reconstructed images are present because of the presence of other materials in the attenuating object, for example bone.

Nearly all CT scans are made nowadays in spiral mode. In this mode the object is not stationary, but moves slowly through the imaging plane (z-direction) while the x-ray tube rotates continuously. For the reconstruction of a cross section, line integrals have to be available in the plane of that cross section. Because of the movement of the object, these line integrals are not all available right away. Therefore, in spiral CT, an added interpolation step is necessary before the images can be reconstructed. This interpolation causes that noise is correlated at an approximately constant rate in the z-direction (paragraph 2.2.2).

The CT value of each voxel is linearly related to the linear attenuation coefficient of the material in the volume element corresponding to that voxel μ_v, with the linear attenuation coefficient of water μ_{water} serving as a reference:

$$\text{CT Number} = 1000 \frac{\mu_v - \mu_{water}}{\mu_{water}} \tag{2.3}$$

Although CT values are dimensionless, they are conventionally expressed as Hounsfield Units (HU). This convention is used throughout this thesis. Hence, in all CT scans, the CT value of water should be zero, the CT value of soft tissue that attenuates the x-rays a few percent more, in the order of 30 HU. The CT value of air, with virtually no attenuation, should be -1000 HU. Slight deviations from these expected values may occur due to small calibration errors. Tagging of water with iodine enhances the CT value, depending on the concentration of iodine. In this chapter, a mixture was used with a CT value of approximately 400 HU.

2.2   Noise and CT value across image transitions for CT colonography

The main sources of image degradation are taken to be loss of spatial resolution and noise. We will propose models that describe the noise and the CT value across the material transitions to deal with these deficiencies in order to obtain reliable measurements for CT colonography.

A primary source of image degradation is noise. Especially in CT colonography imaging where the radiation dose is kept to a minimum, excessive noise may be present. The 3-D image reconstruction causes that the noise *is* correlated. We assert, however, that that correlation only occurs at a very small scale without hampering the detection of polyps.

Furthermore, it is assumed that the amplitude of the noise may be modeled by a zero-mean Gaussian.

The *edge-spread-function* (ESF) is modeled as a 1-D Gaussian with a scale that depends on the orientation of the edge. It is demonstrated in this chapter that the erf-function, i.e. a step edge convolved with a Gaussian, serves as an excellent approximation to model the observed CT value across a transition between tissue/air and tissue/tagging. Note that we do not try to recover the specific point-spread-function (PSF), the image produced by a delta function, from CT data. The real PSF of a spiral multi-detector CT scanner is a complicated, space variant function that reminds of spiral pasta [3]. Hence, the model that we apply does not generally hold. The real PSF neither is a separable Gaussian, nor may be independent of material and position, thus slightly violating the linear-time-invariant (LTI) assumptions underneath a convolution model. Since the analysis is limited to approximate fixed value transitions, we may assume the ESF to be modeled by an anisotropic Gaussian function.

The approach that we will take is not robust in the presence of extreme artifacts that for example occur in the presence of metal prostheses. Strategies to reduce such artifacts are outside the scope of this research. Reconstruction errors that occur due to structures in patient-data are considered small and are not accounted for.

2.2.1   Test image

The validity of the assumptions was checked using CT images of a phantom (Figure 2.1a-d) and real CT data. The phantom consisted of a Lucite tube (CT number in the order of 100 HU) inside a cylindrical barrel with a diameter of 34 cm filled with water. Lucite has a CT number that is only slightly higher than that of soft tissue. The tube had an internal diameter of 50 mm. The tube was partly filled with a mix of iodine and water (CT number in the order of 400 HU); above the liquid was air.

The phantom was imaged on a 4-slice spiral CT scanner (Mx-8000, Philips Medical Systems Nederland B.V.). The scan parameters were: 120 kV, 4 x 1.3 mm collimation, pitch 1.25, a 180° interpolation algorithm, a field of view of 256 mm and a slice increment of 0.5 mm.

The image reconstruction kernels available on the Mx-8000 CT scanner are named 'A', 'B', 'C' and 'D', with the resulting image sharpness varying from smooth to sharp, and the noise from low to high. The choice for a specific kernel depends on the application; it is a compromise between the resolution and the noise that may be tolerated. The AMC hospital in Amsterdam has chosen kernel 'C' for its CT colonography examinations.

The raw CT projection data were used to reconstruct the volume of interest containing the phantom in an ordinary way, in the same volume of interest as if it was scanned at lower doses as well. For that purpose, the raw data were modified in a controlled way. The raw transmission measurements were replaced by a realization of the Poisson

process given a scaled version of the high-dose datavalue as the expected value. Note that the variance and the mean values of the Poisson distribution are identical. The CT scan was made at the highest dose level that was available on the CT scanner (400 mAs) given the scan parameters described above. The images from this scan were used to locate the surface of the tube serving for reference. Simulated images were made with average dose (50 mAs) and low-dose (20 mAs) (Figure 2.2).

It is expected that the scale of the PSF would be different along the z-axis than in the xy-plane: the image reconstruction kernels operate in the xy-plane. Accordingly, the tube was oriented such that edges along the tube would be affected by both components of the PSF (Figure 2.1a,b).
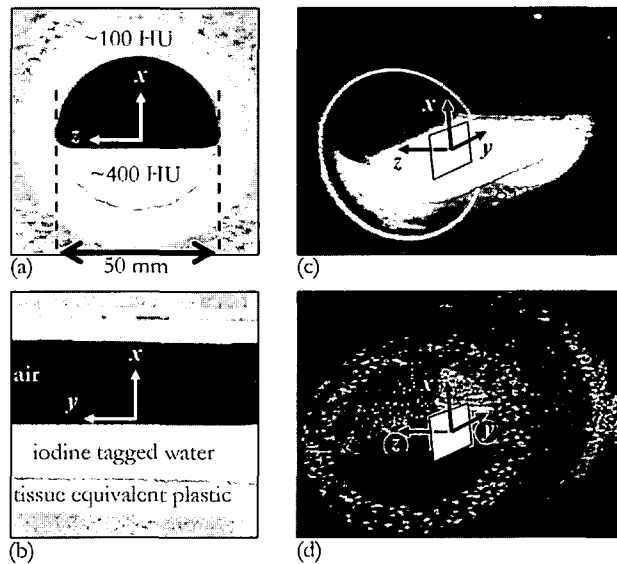
Figure 2.1 ·   (a),(b) Multi-planar-reformats of the phantom showing in-plane (x,y) and out-of-plane (z) orientation. (c) Max-intensity projection of the phantom. (d) Volume rendering of the phantom.
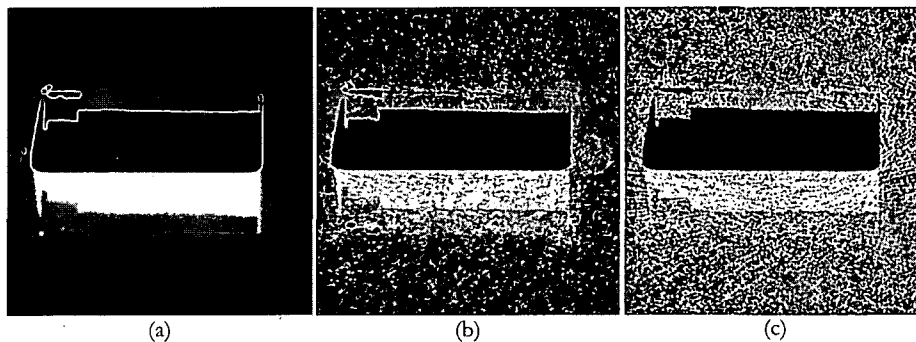
Figure 2.2   (a) High-dose, (b) average-dose and (c) low-dose images of the tube phantom.

2.2.2   Noise correlation

CT detectors measure the x-ray transmissions through a patient. The measurements are subject to noise due to photon counting statistics. When combining these projection-data to reconstruct a 3-D CT image, correlation of noise, the degree to which noise quantities are linearly associated, is introduced between a voxel and its neighboring voxels. It is tested if the correlation of noise responses is rotation invariant in the xy-plane and occurs at a small scale. CT slices of the phantom that only contain water are used. This analysis does not consider reconstruction errors due to structure in the data (e.g. edges of the colon).

The noise correlation was analyzed in three dimensions using the autocorrelation function (ACF), i.e. the cross-correlation of the signal with itself. Let $I = I(x, y, z)$ be the datavalue at a voxel-position (x, y, z). Then, the sample autocorrelation function for a selected area is

$$ACF(k, l, m) = \frac{1}{S_N^2} \sum_x \sum_y \sum_z \left(\{ I(x, y, z) - E(I) \} \{ I(x + k, y + l, z + m) - E(I) \} \right) \quad (2.4)$$

in which $E(I)$ is the mean datavalue and S_N^2 is the sample variance of I over the entire volume. The function has the property of being in the range $[-1, 1]$ with 1 indicating perfect correlation and -1 indicating perfect anti-correlation. The ACF-image displays the correlation-coefficient as a function of the voxel-shift in grey-value.

For clarity reasons, 2-D cross-sectional ACF images are displayed that were derived from the 3-D ACF images. Four images obtained along the xy-plane are depicted in Figure 2.3 (left) and four images along the xz-plane are depicted in Figure 2.3 (right). The autocorrelation function was determined from the central 32 by 32 pixel area in a slice containing water to only measure the correlation of noise. The aforementioned filters 'A', 'B', 'C' and 'D' were used (paragraph 2.2.1). The ACF-images are displayed using nearest neighbor interpolation (pixel-replication) to show the autocorrelation with discrete voxel-shifts. To further demonstrate the pattern of noise correlation surface-plots are added using cubic-spline interpolation (Figure 2.3). In addition, vertical profiles were made through the ACF images to visualize the scale of the autocorrelation function (Figure 2.4).
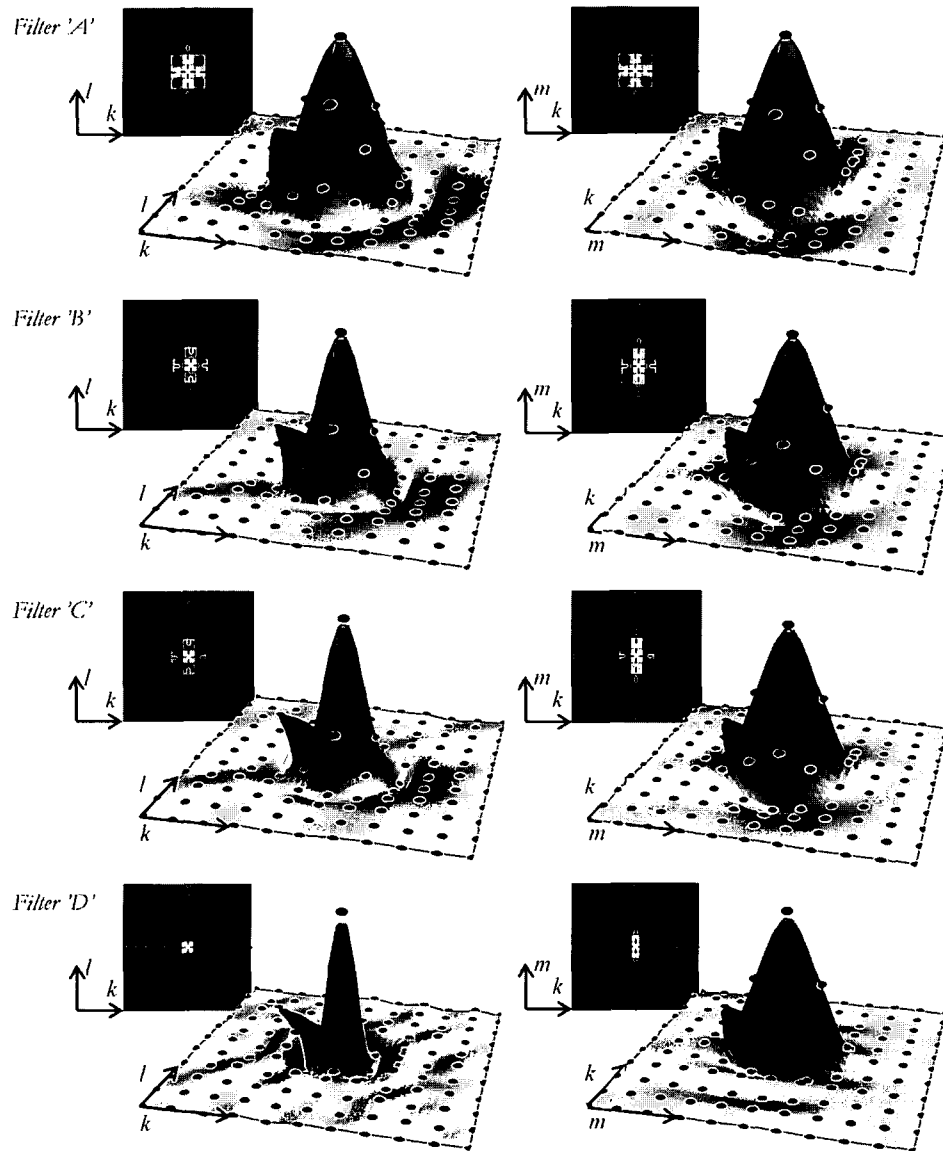
Figure 2.3   Autocorrelation functions generated from the central area of CT images in which the phantom was not present. The images were reconstructed using the available kernels; k is the shift in x-direction, l in y-direction and m in z-direction. The voxel-positions are displayed as dots.
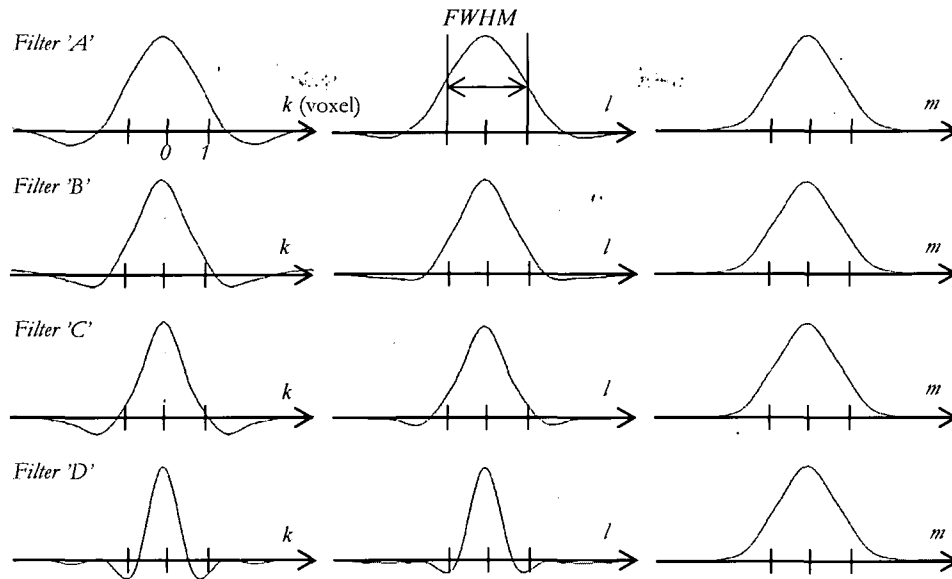
Figure 2.4   Cross-sections of the ACF along the k- l- and m-axis for the four reconstruction-kernels of the CT scanner.

For the reconstruction filters of the CT scanner the Full Width at Half Maximum (FWHM) value of the ACF along the main orientations are given in Table 2.1.

Table 2.1   Scale measurements on the autocorrelation function in unit voxel.

FILTER	FWHM X	FWHM Y	FWHM Z
A	2.40	2.70	2.40
B	1.70	1.94	2.40
C	1.40	1.56	2.40
D	0.98	1.06	2.40

Autocorrelation results in the cross-correlation of the reconstruction-filter with itself resulting in a FWHM for the reconstruction filter 'C' of $\sim 1.48/\sqrt{2} \approx 1.05$ voxel. The measured scale of the in-plane ESF (paragraph 2.2.7) is ~0.8 voxel with a FWHM of ~1.88 voxel. From this we see that the scale of the ESF across edges is dominated by the scale of the detectors: $\sqrt{1.88^2 - 1.05^2} \approx 1.56$.

To test the positional variance, the autocorrelation function was determined for an area in the center and in the periphery of the image that was created using filter type 'D' (selected for its small scale) (Figure 2.5).
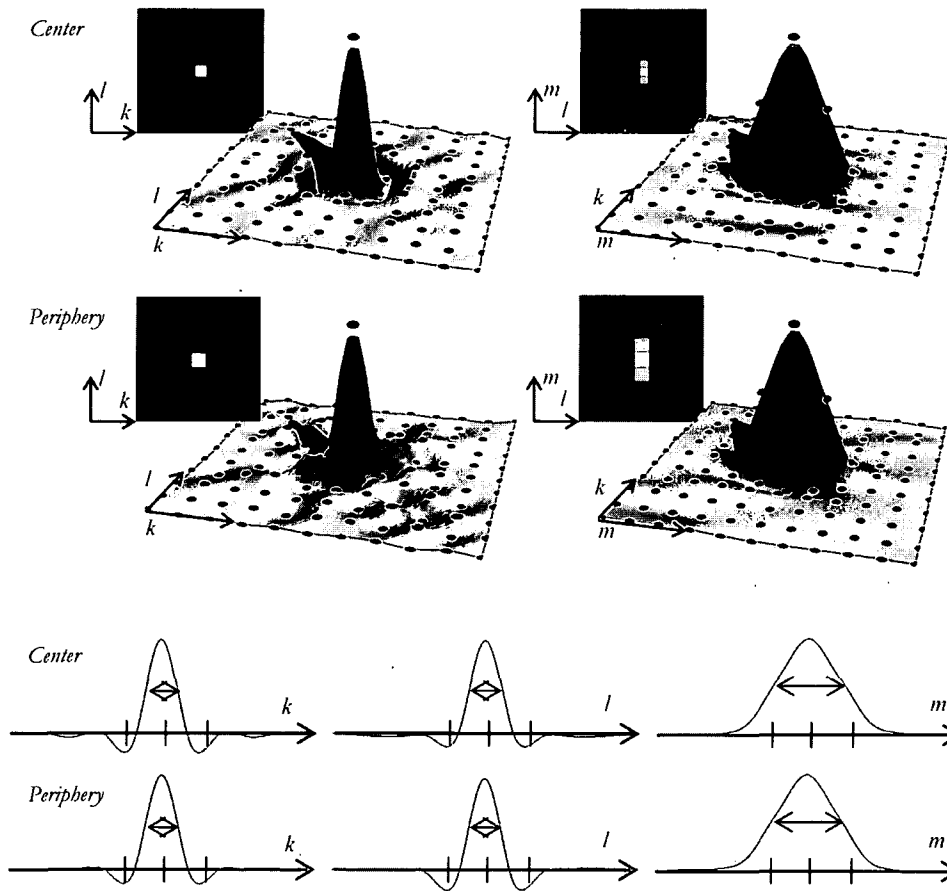
Figure 2.5   Autocorrelation function of areas in the center and in the periphery of an image created using reconstruction filter 'D'.

The previous results show that the in-plane pattern of noise correlation is small and approximately rotation invariant. Apparently, the ACF-image shows only a slight angular dependency in the periphery of the image.

2.2.3   Noise variance

It was tested to what extent the noise variance $\text{Var}(n_s)$ varies over the CT image plane. Slices of the phantom-data that only contain water are used. The local noise variance was estimated from the 'sample variance' S_N^2 of N datapoints. The sample variance S_N^2 for x_i, $i = 1..N$ is defined by

$$S_N^2 = \frac{1}{N-1}\sum_{i=1}^{N}\left(x_i - \overline{x}\right)^2 = \frac{1}{N-1}\left[\sum x_i^2 - \frac{1}{N}\left(\sum x_i\right)^2\right] \qquad (2.5)$$

The raw transmission measurements were used to reconstruct four CT images using reconstruction kernels 'A'-'D'. For every voxel, N datapoints were collected in a 5×5 local neighborhood of that voxel in the xy-plane. Four images display the $\log(S_N)$ in grey-value (Figure 2.6) as an indicator of the noise variance across the image-plane. In addition to the four images, 1-D vertical profiles of S_N -values are depicted in Figure 2.6b using a log-scale along the S_N -axis.
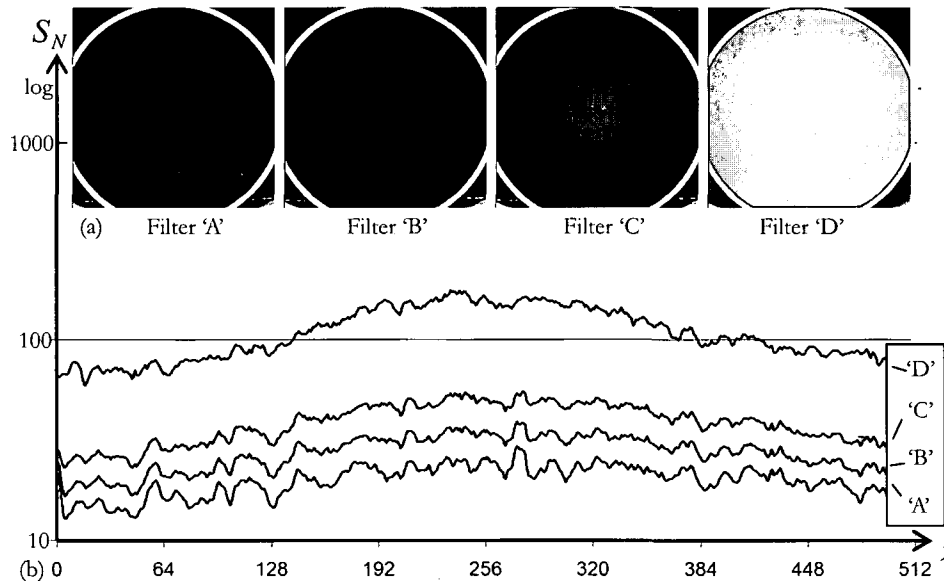
Figure 2.6   (a) S_N -images derived from four CT volumes: CT projection-data are reconstructed using four CT filters: 'A', 'B', 'C' and 'D'. (b) Vertical profiles across these images.

From these results we conclude that the noise variance changes over the image-plane. Its center/perimeter ratio is ~2, such that it certainly may not be considered a constant value. In a clinical CT image the noise is expected to vary even more due to for example bone and is difficult to measure due to absence of homogeneous structures. In this thesis, we do not attempt to locally estimate the local absolute noise variance. Our solution to this problem is using the relationship between the variance of the noise on measured

derivatives and the measured data-value. This relationship only depends on the scale of the operator and order of the derivative.

2.2.4   Noise isotropy

The gradient magnitude L_w and the second derivative in gradient direction L_{ww} are obtained using Gaussian operators with identical scales. However, they may display different variances of the noise. Isotropic noise in $(L, v_1 L_w, v_2 L_{ww})$ facilitates an isotropic metric to be used in fitting functions to the data (chapter 4). The problem is that v_1 and v_2 has to be known in advance. The relation between the noise variance before σ_{ni}^2 and after convolution σ_{no}^2 with a D-order Gaussian derivative of scale σ_{op} in N-dimensional space is given by [5]

$$\frac{\sigma_{no}^2}{\sigma_{ni}^2} = \frac{(2D)!}{\sigma_{op}^{N+2D} \pi^{N/2} D! 2^D 2^{N+D}} \qquad (2.6)$$

Let $\sigma^2{}_l$ be the variance of the noise on the datavalue after filtering with a Gaussian filter of scale σ_{op}. Let $\sigma^2{}_{l_w}$ be the variance of the noise on the gradient magnitude and $\sigma^2{}_{l_{ww}}$ the variance of the noise on the second derivative in gradient direction. Then v_1 and v_2 are obtained from $v_1 = \sigma_l / \sigma_{l_w}$ and $v_2 = \sigma_l / \sigma_{l_{ww}}$. Figure 2.7 displays v_1 and v_2 as a function of the scale of the operator σ_{op} using solid grey lines.
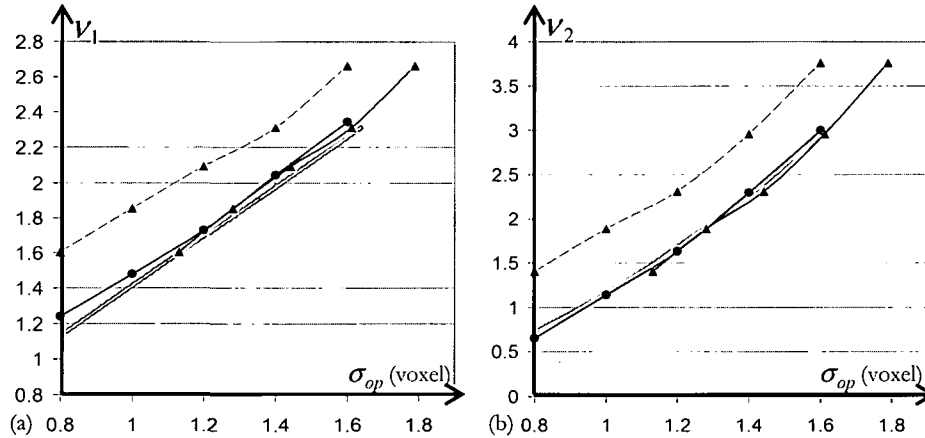
Figure 2.7   Grey-solid lines (a) v_1 and (b) v_2 as a function of the scale of the Gaussian operator σ_{op} using equation 2.6. Black solid lines with circles (a) v_1 and (b) v_2 measured from noisy data as a function of σ_{op}. Dashed lines with triangles (a) v_1 and (b) v_2 as a function of σ_{op} measured from noisy data that has been pre-filtered using a Gaussian with $\sigma_g = 0.8$ (voxel). Solid lined with triangles (a) v_1 and (b) v_2 as a function of $\sigma_{op} * \sigma_g$ measured from noisy data that has been filtered using a Gaussian with $\sigma = 0.8$ (voxel).

Synthetic 3-D data are used to test (2.6). To begin with, Gaussian noise is added to the voxels and V_1 and V_2 are measured as a function of the scale of the operator σ_{op}. These measurements are depicted in Figure 2.7 using black solid lines connected with dots. They exactly follow V_1 and V_2 as described by the model. Secondly, the noisy data are filtered using a Gaussian with $\sigma_g = 0.8$ voxel, to simulate correlated noise. Again V_1 and V_2 are measured as a function of the scale of the operator σ_{op} and depicted in Figure 2.7 using dashed lines with triangles. However, when measuring V_1 and V_2 as a function of the effective scale of σ_{op} convolved with σ_g, then (2.6) is an excellent model for V_1 and V_2, depicted in Figure 2.7 using dashed lines with triangles. From this we conclude that scale parameter used in equation 2.6 should include the scale of noise correlation as well.

3-D CT phantom data are used that have been reconstructed using CT filter 'A-D' to test (2.6). Both V_1 and V_2 are obtained as a function of the scale of the operator σ_{op} and depicted in Figure 2.8a. In addition V_1 and V_2 are obtained taking into account the degree of noise correlation (Figure 2.8b).
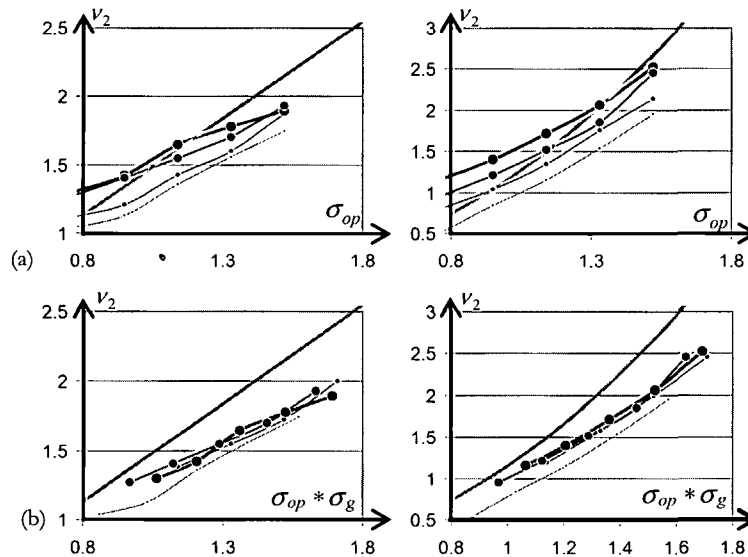
Figure 2.8   (a) V_1 and (b) V_2 measured from CT data as a function of the scale of the Gaussian operator σ_{op}. (b) V_1 and (b) V_2 measured from CT data as a function of the scale of the operator convolved with the degree of noise correlation σ_g: $\sigma_{op} * \sigma_g$.

From Figure 2.7 we conclude that equation (2.6) is a good model to predict the variance of noise on derivatives after convolution with a Gaussian of scale σ_{op}. From Figure 2.8 obtained 3-D CT phantom data we conclude that using the knowledge about correlation of noise improves the estimates of V_1 and V_2. However, there remains as small bias when comparing the model to the measured values of V_1 and V_2 using real CT data. A model that describes this bias is part of future research.

2.2.5   Noise distribution

The assumption of Gaussian distributed noise was tested by creating a histogram using measurements obtained in one material at sufficient distance from edges in a local neighborhood (see Figure 2.9). Subsequently, a Gaussian was fitted to the histogram by minimizing the squared error. The data that were used were reconstructed using CT kernel 'C' as in our CT colonography images.
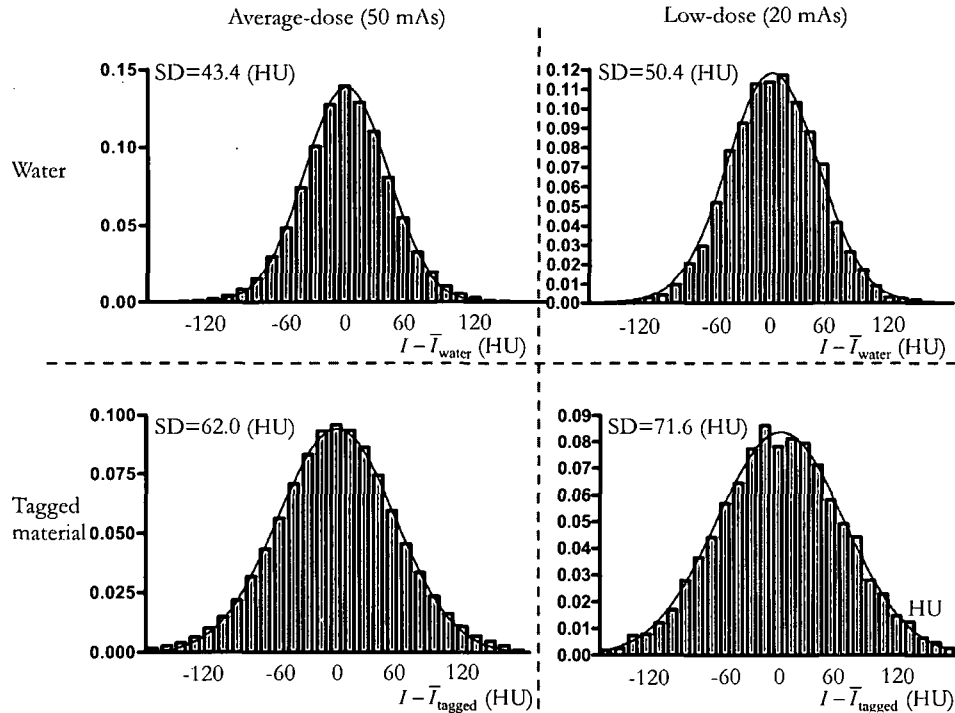
Figure 2.9   The frequency distributions of datavalue measured in water and tagged material from average-dose and low-dose CT scans. The fitted Gaussian function is superimposed on the histogram.

From these results we conclude that the noise distribution resembles a Gaussian. The contrast to noise ratio was determined using

$$CNR = 20\log\left(\frac{H-L}{\sigma_s}\right) \qquad (2.7)$$

Along tissue to tagged material the $CNR = {\sim}16\text{dB}$, using $L = 100$, $H = 400$ and $\sigma_s = 50$. Along tissue to air the $CNR = {\sim}27\text{dB}$, using $L = -1000$ HU, $H = 100$ and $\sigma_s = 50$.

2.2.6   Edge model

We assume that the CT values across transitions between two materials can be modeled by the erf-function (2.8): the edge-spread-function. The scale of the underlying Gaussian may be anisotropic (i.e. varying with the angle), but it is verified that for each angle it remains approximately constant over the image-plane. CT colonography operates only on the edges of the colon, so that is what we focus on.

Six parameters are involved to describe an edge in 3-D by the erf model: (1) L, the expected datavalue of low-intensity material; (2) H, the expected datavalue of high-intensity material; (3) μ, the shortest distance from the model edge to the actual edge-position; (4) σ, the apparent scale of the underlying Gaussian and (5,6) two angles α_1, α_2:

$$G(w; L, H, \sigma, \mu, \alpha_1, \alpha_2) = L + (H - L)\left(\frac{1}{2} + \frac{1}{2}\mathrm{erf}\left(\frac{w - \mu}{\sigma\sqrt{2}}\right)\right) \qquad (2.8)$$

in which the projection of a voxel $v(x, y, z)$ onto w (Figure 2.10) is described by

$$w = x\left(\cos(\alpha_1)\cos(\alpha_2)\right) - y\left(\sin(\alpha_1)\cos(\alpha_2)\right) + z\sin(\alpha_2) \qquad (2.9)$$

where α_1 is the angle in the x, y-plane between the projected gradient and x-axis and α_2 is the angle between the gradient and the positive z-axis.
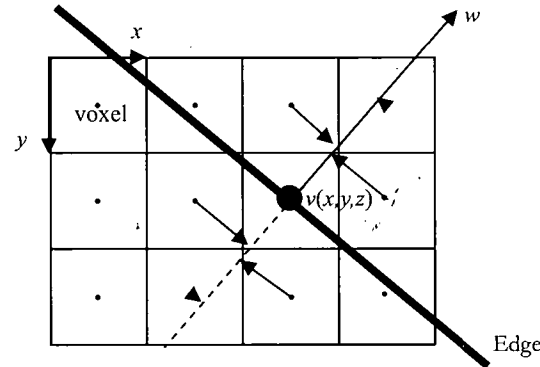
Figure 2.10   The projection of voxels onto the gradient-vector.

The function was fitted to the data and the six parameters were estimated by minimizing ε^2, which is a function of the sum of the squared errors $r_i = y_i - G(w_i; L, H, \sigma, \mu, \alpha_1, \alpha_2)$ [4]

$$\varepsilon^2 = \sum_{i=1}^{n} \sigma_e^2 \left(1 - \exp\left(\frac{-r_i^2}{2\sigma_e^2}\right)\right) \qquad (2.10)$$

in which σ_e is a weight factor that should be set larger than the SD of the noise. We set σ_e safely to the largest standard deviation encountered (Figure 2.6). This procedure is used to control outliers.

Three types of edges were manually selected in the high-dose CT data reconstructed using CT filter 'C' (Figure 2.11 to Figure 2.13, left). The air-tissue transition is considered in Figure 2.11, the tagged-tissue transition in Figure 2.12 and the air/tagged material transition in Figure 2.13. The function-values (black lines) and the original voxel-values (black dots) are plotted in the top rows. The residuals are plotted in bottom rows to indicate the quality of the fit.

Voxel-positions were projected upon the gradient-vector to increase the number of sample points. However, only a very small environment is used to avoid problems due to curvedness of edges: fitting a model to the data using voxel-values in a small neighborhood assumes local flatness of the surface. Results are given for a small environment of 2 voxels (left column) and a larger neighborhood of 4 voxels (right column).

Figure 2.14 was based on publicly available data from the Walter Reed Army Medical Center (patient number 38).
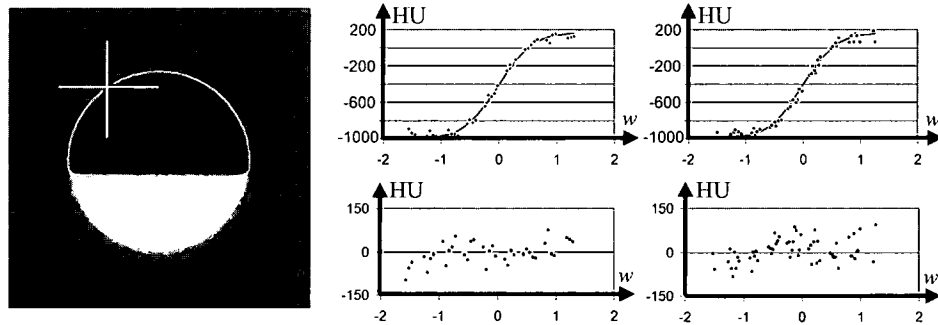
Figure 2.11   The datavalue along the indicated edge in the average-dose phantom was modeled by the cumulative Gaussian. The parameters of the fitted function were: $L = -994$ HU, $H = 162$ HU, $\mu = -0.01$, $\sigma = 0.476$, $\alpha_1 = 2.28$ Rad, $\alpha_2 = 0.01$ Rad, correlation > 0.99. Residuals are plotted in the bottom images to indicate the quality of the fit. In the middle column a small neighborhood and in the right column a broader neighborhood was used. For the middle column the S_N of the residuals is 38 HU.
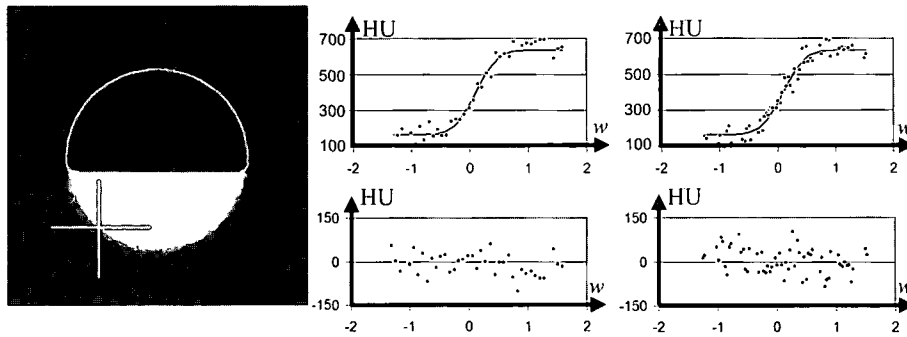
Figure 2.12 As in Figure 2.11. The parameters of the fitted function were: $L = 167$ HU, $H = 635$ HU, $\mu = 0.1$, $\sigma = 0.35$, $\alpha_1 = 0.892$ Rad, $\alpha_2 = 0.01$ Rad, correlation > 0.98. For the middle column the S_N of the residuals is 38 HU.
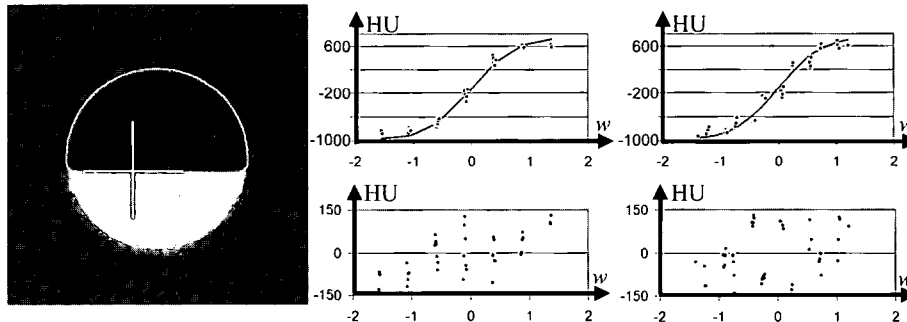
Figure 2.13 As in Figure 2.11. The parameters of the fitted function were: $L = 34$ HU, $H = 1727$HU, $\mu = 0.0$, $\sigma = 0.58$, $\alpha_1 = 0.892$ Rad, $\alpha_2 = 0.01$ Rad, correlation > 0.99. For the middle column the S_N of the residuals is 76 HU.
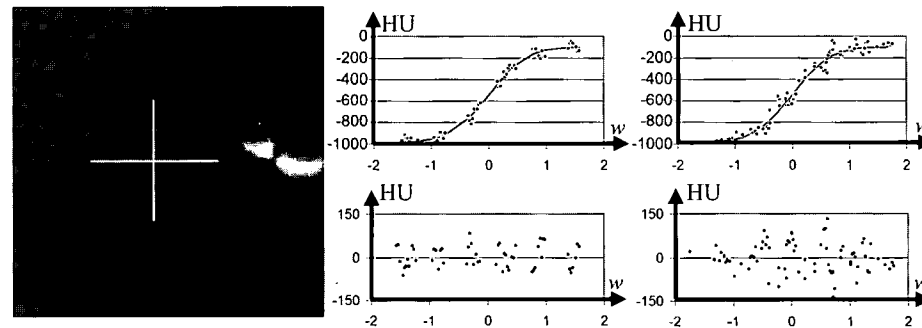
Figure 2.14 As in Figure 2.11, now using patient data. The parameters of the fitted function were $L = 17$ HU, $H = 900$ HU, $\mu = 0$, $\sigma = 0.53$, $\alpha_1 = 1.53$ Rad, $\alpha_2 = -0.51$ Rad, correlation > 0.99. For the middle column the S_N of the residuals is 37 HU.

Considering the values of the standard deviation of the residuals along tissue transitions of 38 HU (bottom-middle column of the Figures 2.11, 2.12 and 2.14) it can be concluded that they represent the noise. Hence, the erf-function, the step-edge model convolved with the 1-D edge-spread-function, is an excellent model for the observed CT value across a two-material transition.

2.2.7   Edge scale

In the previous section it was demonstrated that the erf-function is an excellent model to describe the edge-spread-function (ESF).

The 'apparent' size of the underlying Gaussian may depend, though, on the orientation of an edge with respect to the edge-spread-function. The latter is expected to have a smaller scale in xy-direction (in plane) than in z-direction (out of plane). The effective scale of edges in the gradient-direction is modeled as a function of the angle α between the gradient vector and the positive z-axis:

$$\sigma_w(\alpha) = \sqrt{\left(\sin(\alpha)\sigma_{\perp z}\right)^2 + \left(\cos(\alpha)\sigma_z\right)^2} \qquad (2.11)$$

where σ_z is the scale of the out-of-plane underlying Gaussian edge-spread-function and $\sigma_{\perp z}$ the scale of the in-plane underlying Gaussian edge-spread-function.

In this thesis, both σ_z and $\sigma_{\perp z}$ were estimated from the data by fitting equation (2.11) to all scale-estimates obtained along edges of the colon. The scale in the gradient direction $\sigma_w(\alpha)$ was used as well for noise-isotropy (Chapter 4). The assumption that this approach is valid was tested using the tube phantom and a high-dose reconstruction using CT filter 'C'. Notice that the orientation of the tube surface was such that edges were affected both by the in-plane and out-of-plane scale.

First, σ_z and $\sigma_{\perp z}$ were estimated by fitting the erf-function to edges with gradients that are in xy-orientation or z-orientation: bottom part of Figure 2.15. Second, σ_z and $\sigma_{\perp z}$ were obtained. Third, σ_z and $\sigma_{\perp z}$ were computed by fitting (2.11) to $\sigma_w(\alpha)$ estimates that were obtained from fitting the erf-function to edges with gradients, which were not in xy-orientation or z-orientation top of Figure 2.15.

From Figure 2.15 we conclude that the measured scale is identical to the scale obtained from equation (2.11). The graph illustrates the validity of the equation. Fitting equation (2.11) to the data results in an in-plane scale of ~0.40 mm and an out-of-plane scale of ~0.82 mm (0.8 unit voxel). This agrees with the scale measured from only using in-plane and out-of-plane gradients (Figure 2.15 b,c).
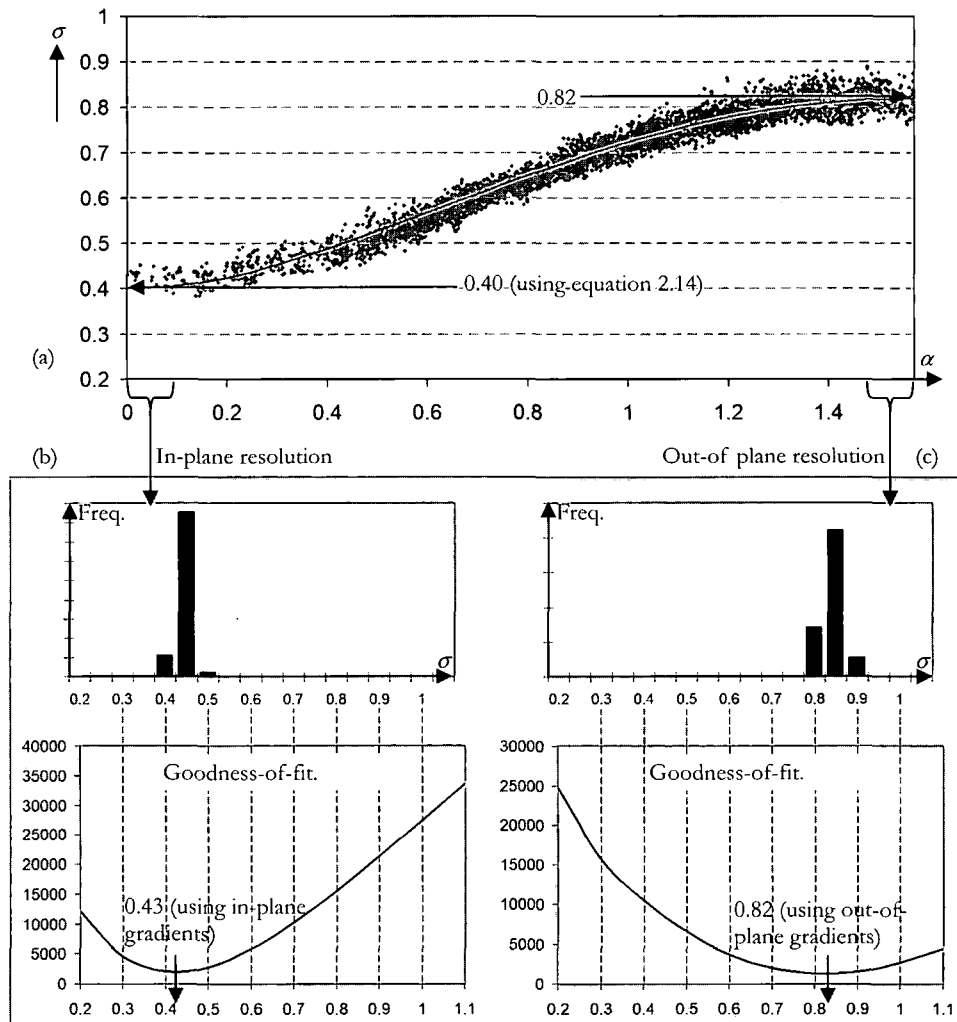
Figure 2.15   (a) The apparent scale of the erf-function is plotted as a function of orientation. The line
              results from fitting equation 2.14. From the indicated part of the measurements, both the
              histogram of found scales and per scale the average goodness-of-fit are displayed,
              (b) using approximate in-plane orientations and (c) using out-of-plane orientations.

It might be expected that the scale of the ESF is slightly larger towards the outside of
CT image slices compared to the center. This will be tested in true patient data.

The main problems with measuring the scale using patient data are that the colon surface
is curved and that structures close to the surface may slightly violate the assumption of a
Gaussian-filtered step edge model. To measure the true scale of the ESF along edges
proper candidates will have to be selected.

The erf-function is fitted at each air-tissue transition as described above. Two groups of voxels are obtained in the neighborhood of an edge voxel. The first group has a small distance to the gradient vector of the edge-voxel (i.e. smaller than 2 voxels). The second group has a larger projected distance to the gradient vector (i.e. larger than 2 and smaller than 4 voxels). If both groups lead to approximately identical fits measured by the standard deviation of the errors, an approximate flat surface with homogeneous contrast is assumed.

The test only focuses on scales that were measured using approximately in-plane image gradients. To test the dependency on position in a slice, the estimated scale is plotted as a function of the distance to the center of the xy-plane in unit voxel (Figure 2.16).
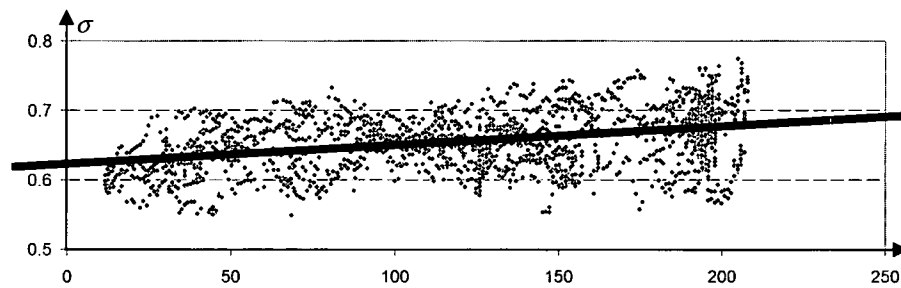
Figure 2.16   The scale is plotted as a function of the distance v to the center of the xy-plane. A line is fitted to the data.

From the measured scales we conclude that the scale increases slightly towards the perimeter of the image.

2.3   Conclusion

It is concluded that for the purpose of CT colonography it may safely be assumed that the noise distribution resembles a Gaussian and that the edge-spread-function (ESF) may be modeled by a 3-D Gaussian.

It has been shown in this chapter that the erf-function, the step-edge model convolved with the 1-D edge-spread-function, is an excellent model to describe the CT value across a two-material transition. This was demonstrated by the residuals showing a random pattern and standard deviation that was of the same magnitude as the noise. Additionally, it was observed that for a given direction, the edge spread function is approximately constant over the image.

However, the variance of noise varies over the image-plane and may not be assumed constant. The variance of noise in CT data with water is smaller in the periphery than in the center of the image.

The measurements show that the in-plane pattern of noise correlation is approximately rotation invariant. The four reconstruction filters at the Mx-8000 have this property, but it was shown that filter 'A' is the most rotationally symmetric. The ACF-image shows only a slight angular dependency in the periphery of the image.

2.4   References

[1]   P. Razifar, M. Sandström, H. Schnieder, B. Langström, E. Maripuu, E. Bengtsson, M. Bergström, Noise correlation in PET, CT, SPECT, and PET/CT data evaluated using autocorrelation function: a phantom study on data reconstructed using FBP and OSEM, *in BMC Medical Imaging*, vol. 5, 2005.

[2]   G. Dougherty, D. Newman, Measurement of thickness and density of thin structures by computed tomography: A simulation study, *in: Medical Physics*, vol. 26–7, 1999, pp. 1341–1348.

[3]   G. Schwarzband, N. Kyriati, The Point Spread Function of Spiral CT, *in: Phys. Med. Biol.*, vol. 50, 2005, pp. 5307–5322.

[4]   T.Q. Pham, Spational adaptivity in Super-Resolution of undersampled image sequences, *in: PhD thesis, Delft University of Technology, Delft, The Netherlands*, www.ist.tudelft. nl/qi: group pub.: PhD, 2006.

[5]   H. Bouma, A. Vilanova, L.J. van Vliet, and F.A. Gerritsen, Correction for the dislocation of curved surfaces caused by the PSF in 2D and 3D CT images, *IEEE Trans Pat Anal Mach Intell* 27, 2005, pp. 1501–1507.

Chapter Three

Three-dimensional Display Modes for CT Colonography: Conventional 3-D Virtual Colonoscopy versus Unfolded Cube Projection

This chapter was based upon a publication in Radiology [1]. For a technical description of the visualization method, we refer to [2].

The authors compared a conventional two-directional three-dimensional (3-D) display for computed tomography (CT) colonography with an alternative method they developed on the basis of time efficiency and surface visibility. With the conventional technique, 3-D ante- and retrograde cine loops were obtained (hereafter, conventional 3-D). With the alternative method, six projections were obtained at 90° viewing angles (unfolded cube display). Mean evaluation time per patient with the conventional 3-D display was significantly longer than that with the unfolded cube display. With the conventional 3-D method, 93.8% of the colon surface came into view; with the unfolded cube method, 99.5% of the colon surface came into view. Sensitivity and specificity were not significantly different between the two methods. Agreements between observers were $\kappa = 0.605$ for conventional 3-D display and $\kappa = 0.692$ for unfolded cube display. Consequently, the latter method enhances the 3-D endoluminal display with improved time efficiency and higher surface visibility.

3.1   Introduction

Computed tomography (CT) colonography is a widely studied technique for surveillance and screening for colorectal cancer. With typical methods for evaluating the data, transverse source images, multiplanar reformatted images obtained along the central colon axis, or a virtual three-dimensional (3-D) endoscopic display are applied [3][4]. There appears to be no consensus in the literature, however, regarding the appropriate method. Most investigators primarily use transverse source images in combination with multiplanar reformatted images and/or a 3-D display for problem solving. Findings in several studies, however, suggest that primary use of 3-D views results in higher sensitivity [5][6]. Beaulieu et al [5] demonstrate a significantly better outcome with 3-D modes after correction for lesion visibility.

The conventional 3-D method is similar to colonoscopy. Anterograde and retrograde cine images are generated off line and display forward and backward viewing planes (conventional 3-D) [1]. However, 3-D display methods are time-consuming [7]. Although hypotonic agents and adequate distention tend to minimize the problem, haustral folds may occlude the wall, thereby reducing sensitivity (Figure 3.1). Other 3-D concepts have emerged [8]-[11]. Such displays are hampered by distortions that could lead to misinterpretation [9][12]. Drawbacks prohibit large-scale use of 3-D methods.

To overcome current limitations of 3-D imaging, we developed an alternative 3-D display method that renders six planar projections (unfolded cube display) at 90° viewing angles from points on the central path [2]. The unfolding of such a cube shows the complete field of view at a path position. The aim of the image sequence of unfolded cubes is to facilitate rapid exploration of the entire colon wall.

The purpose of our study was to compare a conventional two-directional 3-D display for CT colonography with an alternative method we developed on the basis of time efficiency and surface visibility.

3.2   Materials and methods

3.2.1   Data acquisition

Patient population. Thirty patients (13 men, 17 women; mean age, 58 years; age range, 35–82 years) were included in this study. The patients were selected from a population at increased risk for colorectal cancer (a medical or family history of colorectal cancer or polyps) who were referred for colonoscopy. The number of patients with polyps was representative of the prevalence in this surveillance population (54%, 27 of 50, was reported earlier [13]). Lieberman et al. [14] reported a comparable prevalence (50%) in another screening population. Selection was based on the presence of polyps, irrespective of the location. The sample size (power) was calculated to facilitate our

Invisible surface
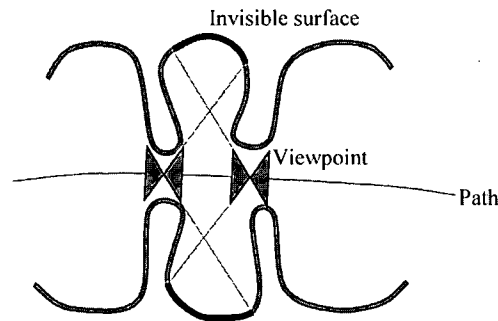
Figure 3.1   Schematic shows areas in black that are missed in conventional 3-D view.

primary aim (i.e., to compare the display methods for time efficiency and surface visibility). At colonoscopy, 16 of 30 patients had at least one polyp of any size, with a total of 78 lesions. Eight of 30 patients had polyps that were 5 mm or more in diameter, with a total of nine such lesions. The CT colonography study was approved by the medical ethics committee of the hospital. The patients were informed a priori by letter, as well as verbally, about the purpose of the study, and they gave written informed consent.

CT colonography. On the day before the examination, each patient drank 4–6 L of macrogol solution (KleanPrep; Helsinn Birex Pharmaceuticals, Dublin, Ireland) for bowel preparation. Directly before image acquisition, 20 mg of butyl scopalamine (Buscopan; Boehringer, Ingelheim, Germany) was administered intravenously to reduce peristalsis. The colon was distended with approximately 2 L of CO_2-enriched air (13.4% CO_2) to enhance resorption and reduce patient discomfort. The air was introduced by means of manual insufflation with a balloon-tipped rectal enema tube. The balloon was inflated with water. The endpoint of CO_2 administration was at maximum patient tolerance or when colon filling was considered adequate. The adequacy of distention was gauged with a scout view. Multisection spiral CT scans (Mx8000; Philips Medical Systems, Best, the Netherlands) were acquired with the patient in prone and supine positions. Scanning parameters were 120 kV, 167 mA (100 mAs: {mA × rotation time}/pitch), collimation of 4.0 × 2.5 mm, pitch of 1.25, standard reconstruction filter, and a 180° interpolation algorithm. Effective section thickness was 3.2 mm with an overlap of 1.6 mm. The insufflation procedure (including scout imaging) was repeated after the patient was repositioned. The delay time between insufflation and the start of scanning was at most 30 seconds. All CT scans were obtained within 2 hours before colonoscopy.

Colonoscopy. Colonoscopy (model CF-140I; Olympus, Tokyo, Japan) was performed by an experienced gastroenterologist. Sedation (5 mg of midazolam, Dormicum; Roche, Basel, Switzerland) and analgesics (0.05 mg of fentanyl, Fentanyl-Janssen; Janssen Pharmaceuticals, Beerse, Belgium) were administered on request. A research fellow, who was not involved in the CT colonography evaluation, was present and recorded the study

on videotape. Lesion size was estimated with the aid of an opened biopsy forceps. Lesion shape was characterized as flat, sessile, or pedunculated. The location of a polyp was classified in one of the following segments: cecum, ascending colon, hepatic flexure, transverse colon, splenic flexure, descending colon, sigmoid, and rectum. The endoscopist was unaware of the CT colonography findings.

3.2.2   Conventional 3-D display

Ante- and retrograde cine loops were evaluated for both prone and supine positions of each patient; this represents a conventional 3-D display method (EasyVision; Philips Medical Systems). The images were volume rendered at positions that were 4 mm apart on the central colon path (the sampling rate will be clarified later). The method developed by Truyen et al [15] was applied for path generation. With the transfer function, an opacity value of 0 was assigned to voxel values less than -650 HU (making such voxels completely transparent) and an opacity value of 1 to voxel values higher than -650 HU. The viewing angle was set to $120°$ to compromise between image distortion and surface visibility. Frame rate was fixed at one image per second. Multiplanar reformatted images in any direction were available from the original CT study to verify potential lesions (e.g., to check for gas). Typical lung ($-1,250$ HU) and level (-500 HU) windows were used. Lesion size was measured on the two-dimensional (2-D) reformatted images.

3.2.3   Unfolded cube projection

To avoid extreme deformations while showing the full visible field around a position, we introduced a series of unfolded cube renderings [2]. The unfolded cubes were also generated by means of sampling on the central path through the colon. A cube was virtually placed in each such point, and on the cube faces, $90°$ views from the center were projected (Figure 3.2. illustrates the principle). By folding out the six images onto a single plane (unfolded cube display), the complete field of view was rendered.

The sequence of unfolded cubes was shown as cine images. Reformatted views of the original CT data became accessible by clicking in the displayed images. Then the cine images were stopped, and the physician could manipulate the reformatted images for closer inspection of suspicious areas. Lesion size was measured on the 2D reformatted images. Real-time rendering of the corresponding wall part also facilitated local inspection of the surface by means of adaptation of the virtual camera position and orientation. Parameters of the volume-rendering transfer function, sampling and frame rate, were identical to those used with conventional 3-D inspection. Figure 3.3 shows the interface.
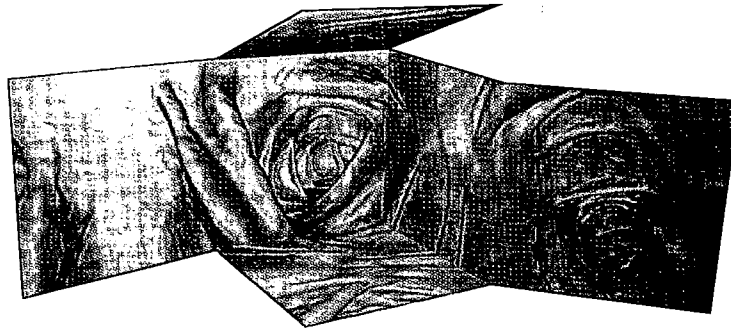
Figure 3.2   Unfolded cube display shows the projection of six overview directions onto the six sides of a cube.
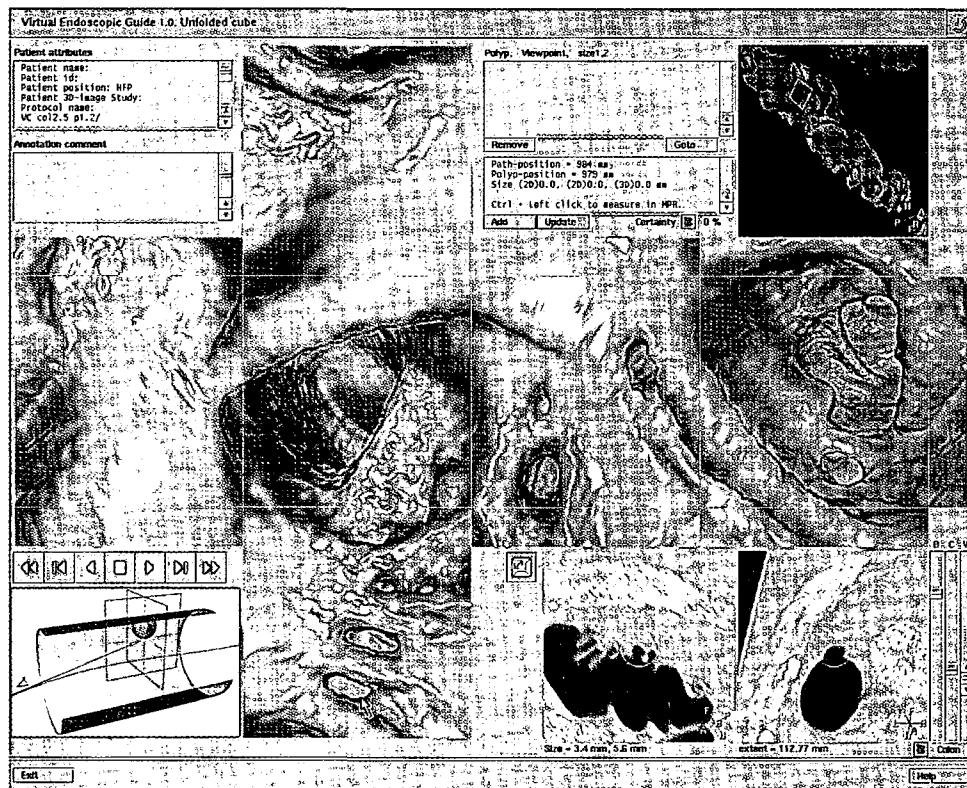
Figure 3.3   Graphical user interface shows the sequence of unfolded cubes.

3.2.4   Evaluation

An abdominal radiologist (C.Y.N.) (observer 1) and a research fellow (J.F.) (observer 2) independently evaluated the data with both display methods. The observers had previous experience with more than 50 CT colonography examinations at the start of the study. Both were blinded to findings during colonoscopy, including those for themselves for the same patient (conventional 3-D display vs. unfolded cube display) and those for each other. Additionally, they were unaware of the prevalence of polyps. We implemented the display methods with a proprietary experimental version of the workstation [2]. Currently, the unfolded cube view is also commercially available. The median interval between two evaluations of the same patient for observer 1 was 21 weeks (interquartile range, 12–28 weeks) and that for observer 2 was 11 weeks (interquartile range, 11–13 weeks). Evaluations with an unfolded cube display preceded those with conventional 3-D display. The cases were presented in random order. Each suspected polyp was scored with respect to size, morphology, and location. Figure 3.4 shows the same polyp in conventional 3-D and unfolded cube display.

3.2.5   Outcome parameters and statistical analysis.

Evaluation time. Evaluation time per patient was recorded for both observers and display modes. Evaluation time included measuring a lesion, imaging it, and generating the report. Outcome was stratified according to the presence of colonoscopically proven polyps. The classifications were compared by means of a Student paired t test. A P value of less than .05 was considered to indicate a statistically significant difference. Initialization time for generating the conventional 3-D and unfolded cube display sequences was disregarded in the analysis.

Surface visibility. Surface visibility was tested with a method described previously [12]; it was defined as the percentage of colon surface voxels coming into view [5]. To measure the visibility, the interior volume of the colon was obtained by means of thresholding of the CT volume at − 650 HU. Subsequently, the surface voxels were identified on the basis of their adjacency to the interior. A surface voxel was defined as visible if a line could be drawn to a position on the central path that was not obstructed by another surface voxel. The Wilcoxon signed rank test was used to compare differences in surface visibility for the unfolded cube display and the conventional 3-D display.

Sensitivity, specificity, and user agreement. Lesions detected during CT colonography were matched with the findings at colonoscopy by a research fellow (R.E.v.G.) on the basis of the colonoscopic findings and the video registration. The fellow was not involved in reading the CT colonography studies.

A polyp detected at CT colonography was considered true-positive if it matched a finding at colonoscopy on the basis of size, shape, location (i.e., colonic segment), and anatomic interrelation to haustral folds. A deviation in size of no more than 5 mm was accepted to accommodate the inherent inaccuracy of colonoscopic size measurement.

(a)   (b)
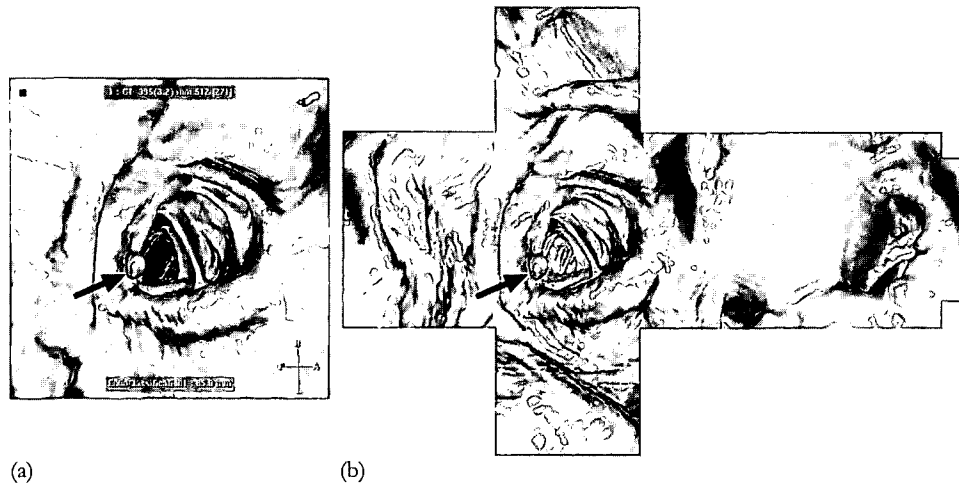
Figure 3.4   Typical example of polyp (arrow) in (a) conventional 3-D display and (b) unfolded cube
display.

A false-negative finding was defined as a polyp that was detected at colonoscopy but was
not found at CT colonography. A patient was identified correctly as having polyps if
there was at least one true-positive finding. The absence of lesions at CT colonography
was considered true-negative if lesions were also absent at colonoscopy.

Sensitivity of both display methods was determined per patient and per polyp. Specificity
of both display methods was determined per patient. Lesions detected at colonoscopy
were divided into two categories: *medium and large*, polyps 5 mm or more in diameter; and
small, polyps less than 5 mm in diameter. Analysis was stratified according to polyp
diameter.

Differences in sensitivity per patient between the display methods were tested with the
McNemar test. Differences in sensitivity per polyp between the methods were assessed
with logistic regression with random effects (i.e., a generalized version of the McNemar
test for clustered data). Logistic regression was used to adjust for clustering of polyps in
patients. A McNemar test was also applied to compare the specificities of the display
methods. Agreement between the observers was measured with the κ statistic on a per
segment basis for all detected lesions. The κ values were interpreted according to the
next qualification: $\kappa < 0.20$, poor agreement; $\kappa = 0.21$–0.40, fair; $\kappa = 0.41$–0.60,
moderate; $\kappa = 0.61$–0.80, good; and $\kappa = 0.81$–1.00, very good.

(a)   (b)
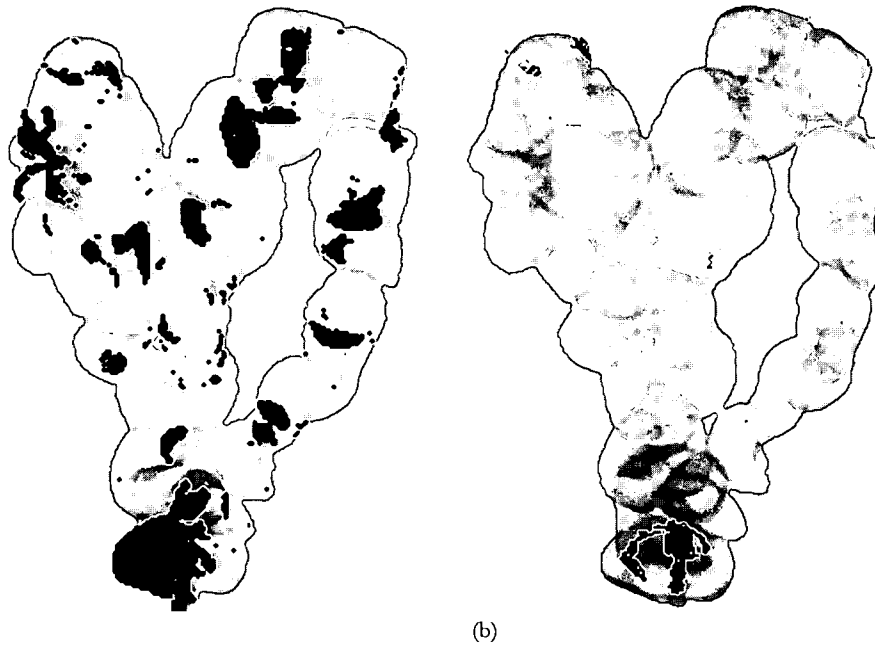
Figure 3.5   (a) Conventional 3-D display. (b) Unfolded cube display.
                Areas in black were missed.
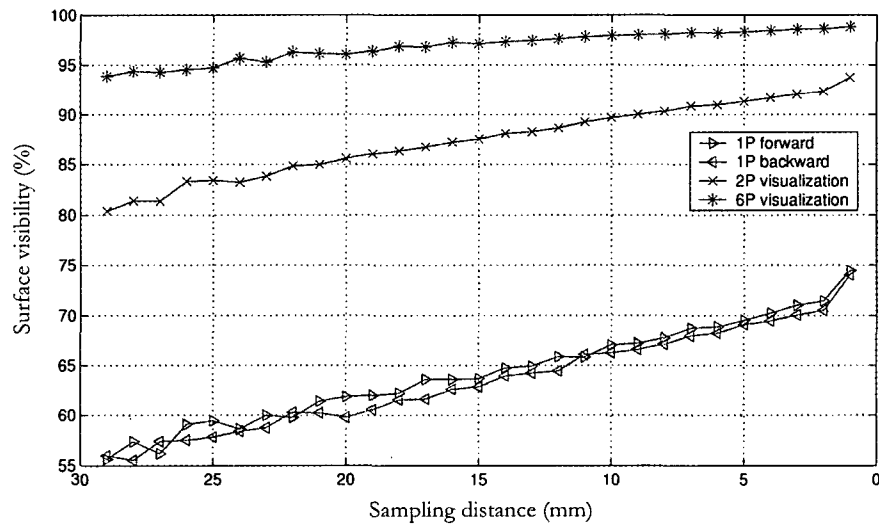
Figure 3.6   Line graph shows relationship between surface visibility and sampling rate along the
                central path for one data set. *1P* = forward or reverse viewing, *2P* = conventional
                3-D display, *6P* = unfolded cube display. Surface area coming into view asymptotically
                approximates 99.8% for the unfolded cube display.

3.3   Results

3.3.1   Evaluation Time

Mean evaluation time with conventional 3-D display (observer 1, 36 minutes 49 seconds; observer 2, 35 minutes 5 seconds) was significantly slower than that with unfolded cube display (observer 1, 19 minutes 33 seconds; observer 2, 20 minutes 9 seconds). The presence of polyps did not lead to significantly longer evaluation time with either technique. Also, the differences in evaluation time between the observers were not significant. Table 3.1 lists the mean evaluation time and SD per observer for the conventional 3-D and unfolded cube displays stratified on the basis of the presence of polyps.

Table 3.1   Evaluation time per patient

Findings at Colonoscopy	Conventional 3D Display		Unfolded Cube Display		Paired Difference	
	Observer 1	Observer 2	Observer 1	Observer 2	Observer 1	Observer 2
Proven polyps	38:22 ± 9:21	38:55 ± 17:38	21:32 ± 7:42	21:55 ± 7:34	16:50 ± 9:36	17:0 ± 13:36
No proven polyps	35:1 ± 10:30	30:41 ± 10:10	17:6 ± 4:30	18:8 ± 4:18	17:44 ± 10:46	12:33 ± 9:2
With and without polyps	36:49 ± 9:51	35:5 ± 14:59	19:33 ± 6:40	20:9 ± 6:27	17:15 ± 9:58	14:55 ± 11:43

Note.—Data are the mean evaluation time (minutes:seconds) ± SD.

3.3.2   Surface visibility

With an anterograde view only, on average, 73% of the colon surface came into view. The conventional 3-D display (antero- and retrograde views) resulted in 93.8% visibility. The unfolded cube display depicted 99.5% of the bowel surface. All differences were significant (P < .05).

A typical outcome that shows the invisible voxels is given in Figure 3.5 (the unseen parts are marked in black). The example bowel wall consisted of 387,225 voxels. With the conventional 3-D display, 25,833 voxels remained invisible, the largest cluster of which contained 10,685 voxels (the next largest cluster contained 2,070 voxels; mean cluster size, 190 voxels). This is considerably more than the surface of a polyp, which is approximately 500 voxels (on the basis of modeling with a sphere that was 5 mm in diameter, with voxel size of $0.6^2 \times 1.6$ mm). The unfolded cube display shows one relatively large area (which consisted of 3,077 voxels) in the rectum, where the path had an open end and the tube with the balloon occluded the wall. Except for these elements, the next largest cluster contained 316 voxels (mean cluster size, 69 voxels).

For one data set, Figure 3.6 shows the relationship between surface visibility and sampling rate along the central path. The surface area coming into view asymptotically approximates 99.8% for unfolded cube display. The graph illustrates that a sampling rate of once every 4 mm yields optimal visibility with the unfolded cube display. Graphs were similar for other data sets.

Sensitivity and specificity

Sensitivity on a per patient basis for medium and large polyps was not significantly different between the two methods for each observer: observer 1 with both displays, eight of eight polyps; observer 2 with conventional 3-D display, seven of eight polyps, and with unfolded cube display, eight of eight polyps. The difference in sensitivity per patient for small polyps between the conventional 3-D and unfolded cube displays was not significant for each observer. Table 3.2 contains sensitivities per patient.

Sensitivities for medium and large polyps were not significantly different between conventional 3-D and unfolded cube displays for both observers: observer 1 with both displays, nine of nine polyps; observer 2 with both displays, eight of nine polyps. Sensitivity for small polyps for observer 1 was significantly worse with conventional 3-D display (16 of 69 polyps) than with unfolded cube display (27 of 69 polyps) (logistic regression, $P < .003$). Sensitivity per small polyp for observer 2 was not significantly different with the two displays (conventional 3-D display, 26 of 69 polyps; unfolded cube display, 24 of 69 polyps). Table 3.3 lists sensitivities per polyp.

Specificity regarding the correct identification of patients without medium and large polyps for observer 1 was 16 of 22 patients with conventional 3-D display and 15 of 22 patients with unfolded cube display. For observer 2, these findings were 16 of 22 patients with conventional 3-D display and 19 of 22 patients with unfolded cube display. Differences between display methods were not significant for each observer. Also, differences in specificity for patients without any polyp were not significant between the methods for each observer. Table 3.4 summarizes the outcome regarding specificity per patient.

On a per segment basis, agreement between the observers with conventional 3-D was $\kappa = 0.605$ (standard error $= 0.052$) and with unfolded cube display was $\kappa = 0.692$ (standard error $= 0.048$). Both κ values indicate good agreement. The difference in κ values between the methods was not significant.

Table 3.2   Sensitivity per patient

Polyp Size (mm)	Conventional 3D Display		Unfolded Cube Display		No. of Patients with Polyps at Colonoscopy
	Observer 1	Observer 2	Observer 1	Observer 2	
Any size	14 (87)	12 (75)	13 (81)	13 (81)	16
<5	6 (75)	5 (62)	5 (62)	5 (62)	8
≥5	8 (100)	7 (87)	8 (100)	8 (100)	8

Note.—Data are the number of patients with polyps. Numbers in parentheses are percentages of findings at colonoscopy.

Table 3.3   Sensitivity per polyp

Polyp Size (mm)	Conventional 3D Display		Unfolded Cube Display		No. of Polyps at Colonoscopy
	Observer 1	Observer 2	Observer 1	Observer 2	
Any size	25 (32)	34 (44)	36 (46)	32 (41)	78
<5	16 (23)	26 (38)	27 (39)	24 (35)	69
≥5	9 (100)	8 (89)	9 (100)	8 (89)	9

Note.—Data are the number of polyps. Numbers in parentheses are percentages of findings at colonoscopy.

Table 3.4   Specificity per patient

Polyp Size in Patients without Polyps (mm)	Conventional 3D Display		Unfolded Cube Display		No. of Patients without Polyps at Colonoscopy
	Observer 1	Observer 2	Observer 1	Observer 2	
Any size	7 (50)	4 (29)	3 (21)	5 (36)	14
≥5	16 (73)	16 (73)	15 (68)	19 (86)	22

Note.—Data are the number of patients. Numbers in parentheses are percentages of findings at colonoscopy.

3.4   Discussion

Findings in the present study show that CT colonography with the unfolded cube display enables time-efficient inspection (about 15–20 minutes per patient) and comprehensive visibility (99.5%) that is superior to that with conventional 3-D display. Sensitivity and specificity of the technique are high with a cutoff value of 5 mm, a threshold that is commonly applied to distinguish possibly relevant polyps.

Various alternative display techniques have been introduced to evaluate CT colonography data. Initially, the complementary use of 2-D and 3-D images was reported to provide the best sensitivity [16][17]. Later, evidence was given that reformatted 2-D images can be just as effective as fly-through 3-D images [18][19]. We opted for a 3-D method because significantly better sensitivity was found for the 3-D modes after correction for lesion visibility [5].

In previous articles, evaluation times per patient are 20 –26 minutes with only transverse reformatted images [17][19], 15 minutes with only 3-D images [17], and 16–35 minutes with a combined approach (primarily 2-D and 3-D display for problem solving) [18][19]. Evaluation times were not specifically recorded by Hara et al [16], but no data set evaluation lasted longer than 10 minutes. Beaulieu et al [5] report that the exploration per data set was 8 minutes with 2-D display and 12 minutes 5 seconds with 3-D cine images.

In the current study, outcomes per patient are in the same range for the unfolded cube display (17–22 minutes) but are significantly slower with the conventional 3-D display (31–39 minutes). Evaluation time could be reduced with a faster fixed frame rate or with a user-controlled speed. However, this will probably not change the relative time difference between the display methods (the same frame rate was used for both techniques). To our knowledge, in no previously published study was the influence of frame rate on sensitivity explored specifically. We opted to use a "save" frame rate of one image per second at 4-mm intervals to yield optimal surface visibility and to ensure detection of lesions.

The methods that we used differ slightly from methods described in the literature. In our method, 3-D images are used primarily for inspection, while the 2-D reformatted images are used for confirmation. Evaluation time is longer with conventional 3-D display versus that with unfolded cube display; this is attributed to increased image distortion near the edge of the conventional 3-D display (as a result of the larger viewing angle: 120° vs. 90°). For closer observation of these areas, interactive adjustment of the viewing direction is required, which inherently leads to longer evaluation time.

Initialization time, such as the generation of cine images, was disregarded because it was performed off line in batch mode. Preparation for evaluation (specifically, generation of the unfolded cube and 3-D cine loops) was initiated by technicians. Exclusion from the analysis of the time for generation is justified because it is done in the background. On average, it took us 35 minutes per patient to render the unfolded cube displays. Conventional 3-D displays were generated in approximately 12 minutes (one-third of 35 minutes). Reduced resolution would speed up the generation of displays. If resolution is reduced, however, sensitivity may also be reduced. Other initialization times, such as those for image loading and path tracking, are on the order of a few minutes.

Our results for surface visibility confirm earlier findings (with nine data sets) [12]. Paik et al [12] reported approximately 75% visibility with either forward or reverse viewing only, 95% visibility with both, and 98% visibility with Mercator projection. Clearly, visibility is important for its direct influence on sensitivity. Beaulieu et al [5] show that sensitivity improves significantly after correction for lesion visibility.

Several methods have been explored for optimization of the amount of colon surface coming into view. The viewing angle may simply be increased, but this is at the expense of severe deformation toward the edges of the image. A so-called flattening method can be used to straighten the colon mathematically [9]. Thus, images of large surface areas of

the colon are generated that are similar to gross pathology specimens. Unfortunately, such an approach may yield severe distortion, which causes lesions to appear more than once in different areas [9][12]. Such repeated appearances are reported to arise whenever the central path makes a sharp turn [12].

Beaulieu et al [5] propose acquisition of a sequence of unfolded 60° views perpendicular to the central path (panoramic viewing). The latter technique results in lateral images from the colon wall. Panoramic viewing results in higher sensitivity than that with a display of transverse CT images and endoluminal 3-D cine loops [5]. The method may not yield full visibility, because the forward and backward viewing directions are not included.

Alternatively, Paik et al [12] studied the use of Mercator and stereography projections. They concluded that true-positive findings are better distinguished from the environment with Mercator projection compared with panoramic or stereographic viewing modes because of the extended field of view in Mercator projection, which enables feature tracking over a longer distance. A generally recognized drawback of Mercator projection, however, is that the image becomes distorted away from the equatorial region. Such image deformation would restrict the sensitivity in areas where the colon is strongly curved.

The unfolded cube display that we evaluated in the present study is most comparable to the panoramic technique by Beaulieu et al. [5]. Note that the unfolded cube mode includes the forward and backward viewing directions, as well as the lateral directions. Consequently, the full field of view is covered, and a feature can be tracked further as it moves through the image planes. At the same time, image distortion is limited by the rather small viewing angle.

Sensitivity per patient in the present study is in the high range compared with values from the literature (65%–94% for polyps larger than 5 mm and 25%–57% for any polyp size) [1]-[24]. Sensitivity per polyp for polyps larger than 5 mm was reported previously in the range of 46%– 90% [1]-[24]. Data from our observers are in the upper range. Sensitivities for smaller polyps (middle row in Table 3.3) are in concordance with the literature (15%–55% [1]-[24]) and confirm the poor results with CT colonography for such (perhaps not relevant) lesions.

Specificities for both observers in the present study are low compared with values reported earlier (62%–91% for all polyps and 74%–96% for polyps larger than 5mm [1]-[24]). The relatively low specificity in our study may be a result of the rather strict definitions of true-and false-positive findings that we used. We compared findings at CT colonography with those at video colonoscopy on the basis of the anatomic interrelation to haustral folds, anatomic segment, size, and morphology. In most studies, a less strict criterion concerning location was used: A finding was considered to be true-positive if the lesion was found in the same colonic segment at both colonoscopy and CT colonography. Since a colonic segment is between 15 and 40 cm long, this may cause

erroneous interpretation of CT colonography findings as true-positive while in fact they are false-positive because they match other lesions in the colonic segment or they are in fact residual stool. Thus, in our opinion, this strategy may result in overestimation of the number of true-positive results and underestimation of the number of false-positive results at CT colonography.

Another explanation for the high rate of false-positive findings is the fact that 27% of small polyps (< 5 mm) are missed at colonoscopy [25]. Therefore, some small lesions detected during CT colonography may in fact be true-positive findings. Nevertheless, the benefit of detecting lesions smaller than 5 mm in a screening setting is dubious because very small polyps are known to rarely contain malignant tissue [26].

Both the conventional 3-D and unfolded cube display yielded good interobserver agreement. Several previous studies were performed with more than one independent observer ([5][6][12][16][19][22]) with the same evaluation method. Agreement between observers with κ statistics, however, was reported only by McFarland et al .[6] ($\kappa = 0.53$– 1.00) and Pescatore et al. [22] ($\kappa = 0.56$–0.72). The values in the present study are in the same range.

The patient population in the present study included approximately equivalent numbers of patients with and without polyps. This is appropriate for our primary objective— namely, to compare two display methods regarding time efficiency and surface visibility. The sample size (power) was calculated to meet the primary aim. The number of patients is relatively small for addressing the (anticipated) small difference in sensitivity and specificity. A limitation is that the outcome may be different in a screening population. Note that the probability that a polyp resides in an invisible area is only 6.2% with a conventional 3-D display because 93.8% of the surface area is visible (the probability is 0.5% with the unfolded cube display). Consequently, a much larger population or a specifically selected population is needed to demonstrate a significant difference in sensitivity. The sample size sufficed to study the unfolded cube display on the basis of time efficiency and surface visibility.

The future role of CT colonography in cancer screening depends on improvement in issues such as imaging efficiency, patient acceptance, and effective radiation dose. Currently, one of the main drawbacks of CT colonography is the long evaluation time. Computer-aided diagnosis is an important development that could support the practical use of CT colonography. Although positive early results were reported with automatic polyp detection [27], further research is warranted. We foresee a scheme in which potential lesion sites, suggested by the computer algorithm, are checked by a human observer. Therefore, a primary 3-D display method may be superior to a primary 2-D technique (as findings in several studies indicate). The unfolded cube display method may contribute to .such an evaluation strategy. Consequently, it could facilitate the implementation of CT colonography in colorectal cancer screening.

In conclusion, the unfolded cube display is an alternative method to evaluate CT colonography data. The evaluation time was 19.5–20.0 minutes, during which 99.5% of the colon wall was inspected. The method is more time efficient and yields better surface visibility than does a conventional 3-D technique. The sensitivity for patients with medium and large polyps was eight of eight patients, and the specificity for patients without medium and large polyps was 15–19 of 22 patients. The method facilitates good agreement between observers ($\kappa = 0.692$).

The unfolded cube display successfully combines time efficiency and high accuracy; thus, it improves the 3-D display for CT colonography.

3.5   Acknowledgment

3.6   References

[1]   F.M. Vos, R.E. van Gelder, I.W.O. Serlie, J. Florie, C.Y. Nio, A.S. Glas, F.H. Post, R. Truyen, F.A. Gerritsen, J. Stoker, Three-dimensional display modes for CT colonography: conventional 3D virtual colonoscopy versus unfolded cube projection, *in: Radiology*, 228(3), 2003, pp. 878–885.

[2]   I.W.O. Serlie, F.M. Vos, R. van Gelder, J. Stoker, R. Truyen, F. Gerritsen, Y. Nio, F.H. Post, Improved Visualization in Virtual Colonoscopy Using Image-Based Rendering, *in: Data Visualization, proc. of the Joint Eurographics and IEEE TCVG Symposium on Visualization*, vol. 3, 2001, pp. 137–146.

[3]   H.M. Fenlon, D.P. Nunes, P.C. Schroy, M.A. Barish, P.D. Clarke, J.T. Ferrucci, A comparison of virtual and conventional colonoscopy for the detection of colorectal polyps, *in: N Engl J Med*, vol. 341, 1999, pp. 1496–1503.

[4]   F.G. Fletcher, C.D.J. Johnson, T.J. Welch, *et al.*, Optimization of CT colonography technique: prospective trial in 180 patients, *in: Radiology*, vol. 216, 2000, pp. 704–711.

[5]   C.F. Beaulieu, R.B. Jeffrey, C. Karadi, D.S. Paik, S. Napel, Display modes for CT colonography. II. Blinded comparison of axial CT and virtual endoscopic and panoramic endoscopic volume-rendered studies, *in: Radiology*, vol. 212, 1999, pp. 203–212.

[6]   E. McFarland, J. Brink, T. Pilgram, *et al.*, Spiral CT colonography: reader agreement and diagnostic performance with two-and three-dimensional image-display techniques, *in: Radiology*, vol. 218, 2001, pp. 375–383.

[7]   J.G. Fletcher, C.D. Johnson, Computed tomographic colonography: current and future status for colorectal cancer screening, *in: Semin Roentgenology*, vol. 35, 2000, pp. 385–393.

[8]   J. Reed, C.D. Johnson, Virtual pathology: a new paradigm for interpretation of computed tomographic colonography, *in: Proc SPIE*, vol. 3335, 1999, pp. 439–449.

[9]   S.B. Dave, G. Wang, B.P. Brown, E.G. McFarland, Z. Zhan, M.W. Vannier, Straightening the colon with curved cross sections: an approach to CT colonography, *in: Acad Radiol,* vol. 6, 1999, pp. 398–410.

[10]   G. Wang, E. McFarland, B. Brown, M. Vannier, GI tract unraveling with curved cross sections, *in: IEEE Trans Med Imaging,* vol. 17, 1998, pp. 318–322.

[11]   S. Haker, S. Angenent, A. Tannenbaum, R. Kikinis, Nondistorting flattening maps and the 3-D visualization of colon CT images, *in: IEEE Trans Med Imaging,* vol. 19, 2000, pp. 665–670.

[12]   D.S. Paik, C.F. Beaulieu, R.B. Jeffrey, C.A. Karadi, S. Napel, Visualization modes for CT colonography using cylindrical and planar map projections, *in: J Comput Assist Tomogr,* vol. 24, 2000, pp. 179–188.

[13]   R.E. van Gelder, H.W. Venema, I.W.O. Serlie, *et al.,* CT colonography at different radiation dose levels: feasibility of dose reduction, *in: Radiology,* vol. 224, 2002, pp. 25–33.

[14]   D.A. Lieberman, D.G. Weiss, J.H. Bond, D.J. Ahnen, H. Garewal, G. Chejfec, Use of colonoscopy to screen asymptomatic adults for colorectal cancer, Veterans Affairs Cooperative Study Group 380, *in: N Engl J Med,* vol. 343, 2000, pp. 162-168.

[15]   R. Truyen, P. Lefere, S. Gryspeerdt, T. Deschamps, Speed and robustness of (semi-) automatic path tracking (abstr), *in: Proceedings of the 2nd International Symposium on Virtual Colonoscopy, Boston,* 2000, vol. 102

[16]   A.K. Hara, C.D. Johnson, J.E. Reed, R.L. Ehman, D.M. Ilstrup, Colorectal polyp detection with CT colography: two versus three dimensional techniques, *in: Radiology,* vol. 200, 1996, pp. 49–54.

[17]   A.P. Royster, H.M. Fenlon, P.D. Clarke, D.P. Nunes, J.T. Ferrucci, CT colonoscopy of colorectal neoplasms: two-dimensional and three-dimensional virtual-reality techniques with colonoscopic correlation, *in: AJR Am J Roentgenol,* vol. 169, 1997, pp. 1237–1242.

[18]   M. Macari, A. Milano, M. Lavelle, P. Berman, A.J. Megibow, Comparison of time efficient colonography with two and three-dimensional colonic evaluation for detecting colorectal polyps, *in: AJR Am J Roentgenol,* vol. 174, 2000, 1543–1549.

[19]   A.H. Dachman, J.K. Kuniyoshi, C.M. Boyle, *et al.,* CT colonography with three dimensional problem solving for detection of colonic polyps, *in: AJR Am J Roentgenol,* vol. 171, 1998, pp. 989–995.

[20]   A.K. Hara, C.D. Johnson, J.E. Reed, *et al.,* Detection of colorectal polyps with CT colography: initial assessment of sensitivity and specificity, *in: Radiology,* vol. 205, 1997, pp. 59–65.

[21]   D.K. Rex, D. Vining, K.K. Kopecky, An initial experience with screening for colon polyps using spiral CT with and without CT colography (virtual colonoscopy), *in: Gastrointest Endosc,* vol. 50, 1999, pp. 309–313.

[22]   P. Pescatore, T. Glucker, J. Delarive, *et al.,* Diagnostic accuracy and interobserver agreement of CT colonography (virtual colonoscopy). *in: Gut,* vol. 47, 2000, pp. 126–130.

[23]   J. Yee, G.A. Akerkar, R.K. Hung, A.M. Steinauer-Gebauer, S.D. Wall, K.R. McQuaid, Colorectal neoplasia: performance characteristics of CT colonography for detection in 300 patients, *in: Radiology,* vol. 219, 2001, pp. 685–692.

[24] G. Spinzi, G. Belloni, A. Martegani, A. Sangiovanni, C. Del Favero, G. Minoli, Computed tomographic colonography and conventional colonoscopy for colon diseases: a prospective, blinded study, *in: Am J Gastroenterol*, vol. 96, 2001, pp. 394–400.

[25] D.K. Rex, C.S. Cutler, G.T. Lemmel, *et al.*, Colonoscopic miss rates of adenomas determined by back-to-back colonoscopies., *in: Gastroenterology*, vol. 112, 1997, pp. 24–28.

[26] J.D. Waye, B.S. Lewis, A. Frankel, S.A. Geller, Small colon polyps, *in: Am J Gastroenterol*, vol. 83, 1988, pp. 120–122.

[27] H. Yoshida, Y. Masutani, P. MacEneany, D.T. Rubin, A.H. Dachman, Computerized detection of colonic polyps at CT colonography on the basis of volumetric features: a pilot study, *in: Radiology*, vol. 222, 2002, pp. 327–336.

Chapter Four

Classifying CT Image Data into Material Fractions by a Scale and Rotation Invariant Edge Model

A fully automated method is presented to classify three-dimensional CT data into material fractions. An analytical scale-invariant description relating the data value to derivatives around Gaussian blurred step-edges – arch model - is applied to uniquely combine robustness to noise, global signal fluctuations, anisotropic scale, non-cubic voxels and ease of use via a straightforward segmentation of 3-D CT images through material fractions. Projection of noisy data value and derivatives onto the arch yields a robust alternative to the standard computed Gaussian derivatives. This results in a superior precision of the method. The arch-model parameters are derived from a small, but over-determined, set of measurements (data values and derivatives) along a path following the gradient uphill and downhill starting at an edge voxel. The model is first used to identify the expected values of the two pure materials (named L and H) and thereby classify the boundary. Second, the model is used to approximate the underlying noise-free material fractions for each noisy measurement. An iso-surface of constant material fraction accurately delineates the material boundary in the presence of noise and global signal fluctuations. This approach enables straightforward segmentation of 3-D CT images into objects of interest for computer-aided diagnosis and offers an easy tool for the design of otherwise complicated transfer functions in high quality visualizations. The method is applied to segment a tooth-volume for visualization and electronic cleansing for virtual colonoscopy.

4.1   Introduction

Segmentation isolates and delineates objects and structures of interest from their surroundings, e.g. an organ, the colon or an arterial tree in a 3-D medical CT image. It is a fundamental task in image processing and a requirement for quantification and high quality visualization. Segmentation is often complicated by the limited and anisotropic resolution of the image modality at hand. The resolution of a multi-slice spiral CT scanner is limited by its configuration (size of the detector elements) and the reconstruction algorithm [2]. The anisotropic space-variant point spread function (PSF) resembles spiral pasta [3], but is often modeled by an anisotropic 3-D Gaussian PSF. We have shown that the edge spread across tissue transitions can be accurately modeled by the erf-function and hence support the use of a 3-D Gaussian PSF (Chapter 2). Modeling the PSF by a Gaussian also permits accurate edge detection of curved surfaces [4]. The finite resolution causes that contributions of different materials are combined into the value of a single voxel. This is generally referred to as the partial volume effect (PVE) [5]. It results in blurred boundaries and hampers the detection of small or thin structures.

In this paper, we present a novel method that models the PVE to estimate material fractions in the edge region. The method deals with two-material transitions based on locally estimated derivative values. We extended previous work [6] by incorporating the invariance to anisotropic noise and anisotropic scale of the data and the generalization to an arbitrary order of derivatives. Projection of noisy data value and derivatives onto the appropriate arch model yields a robust alternative to the standard computed Gaussian derivatives. The method allows slowly varying material intensities at both sides of the transition and small structures because pure material voxels are not required to estimate model parameters. We will demonstrate how this approach may be used to segment and visualize complicated structures of interest in a reproducible and simple way. It also facilitates electronic cleansing for virtual colonoscopy.

4.1.1   Related work

The method presented here is inspired by the work of Kindlmann [7] and Kniss [8]. Kindlmann creates a histogram of the data value I and the gradient magnitude $|\partial I/\partial w| \triangleq |I_w|$ (w is along the gradient direction). This yields arch-shaped point clouds for edge regions (Figure 4.1a). Figure 4.1b shows such point clouds for a three-material phantom scanned with anisotropic resolution. The arch connects the two materials at the base line and its height depends on the scale across the edge. The second derivative in the gradient direction $\partial^2 I/\partial w^2 \triangleq I_{ww}$ is added as a third dimension. Kindlmann estimates the first and second derivatives $I_w(I)$ and $I_{ww}(I)$ as a function of data value I from the 3-D histogram by slicing it at data value I and finding the centroid of the scatterplot of I_w and I_{ww} at that value. Then they apply a mapping of the signal value onto the distance w to the nearest edge using $w(I) = -\sigma^2 I_{ww}(I)/I_w(I)$ and the estimated derivatives. The scale across edges is obtained from the histogram using $\sigma^2 = -w \max(I_w)/\max(I_{ww})$.
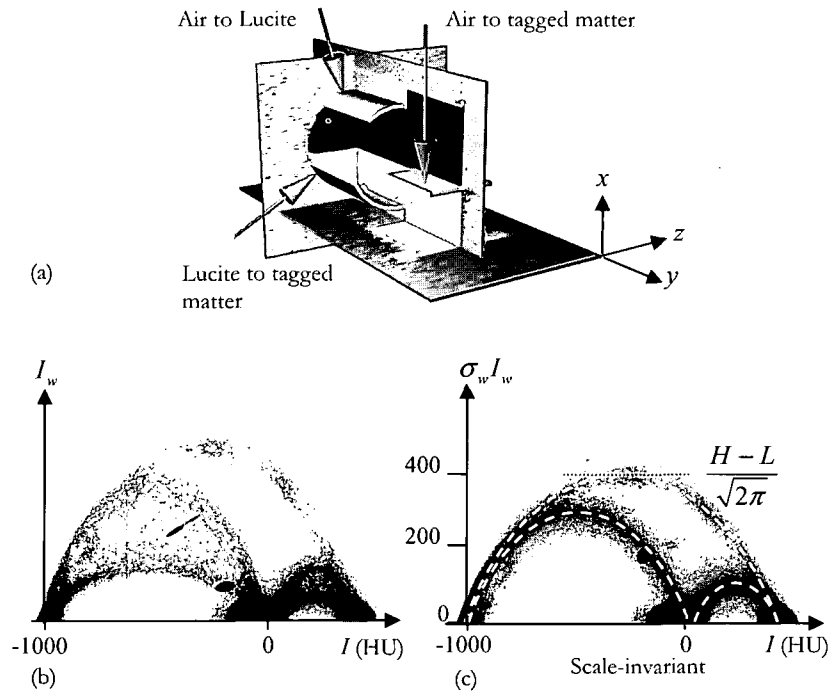
Figure 4.1   (a) Schematic overview of a three material phantom using three multi-planar reformatted images through a CT volume. The axial resolution (z-axis) is lower compared to the lateral resolution (x,y-plane). (b) Scatter-plot of intensity and gradient magnitude. (c) Scatter-plot of intensity and scale-invariant gradient magnitude. Three instantiations of the *arch* model are superimposed corresponding to the three types of material transitions.

Kindlmann and Kniss use the histogram to visualize boundary information and guide the user in designing an opacity transfer function for volume visualization. The user may select parts of the histogram to avoid rendering all edges and may specify the opacity as a function of distance to the nearest edge. The objective is to visualize regions close to selected material transitions as opaque. Rather than rendering the volume based on a transfer function that depends on only the measured signal value, Kindlmann's transfer function is derived from the triplet $\{I, I_w, I_{ww}\}$. The height of the arch-shaped point cloud is spread over a wide range due to anisotropic resolution of the scanner. Moreover, a global fluctuation in data value hampers the estimation of derivatives from the histogram. Consequently, the approximation of Kindlmann's method is limited to data of isotropic resolution without global fluctuation in data value.

Laidlaw [9] proposed a supervised Bayesian method for classification of partial volume voxels into material fractions by fitting basis functions to local histograms of voxel values. For each voxel, the relative contribution of each basis function yields the material fractions. A disadvantage of this method is that the voxel is modeled as a cubic region

not explicitly modeling the scale of the data or the blurring operator. Additionally, processing a single voxel is susceptible to noise.

De Vries [10] aims to classify image-data into materials and related interfaces. Radial basis functions are used to model the probability of an image intensity to occur given a type of material. Material-fractions are estimated by the a posteriori probability, which is expressed in the radial basis functions by Bayes' rule. The priors are estimated from a local histogram of increasing size until stabilization occurs. Finally, a partial volume measurement is classified into pure materials based on the predicted edge position. A disadvantage of this method is that it requires sufficient voxels that are not disturbed by the PVE to obtain sensible priors for the pure materials in the Bayes rule. Hence, the size of the neighborhood may be unfavorably large or the method does not stabilize at all. Small objects are likely to be misclassified because the method might stabilize on surrounding materials.

Several problems remain when adopting the methods described above:
First, problems related to the resolution of the data remain. The size of the structures of interest is often of the same scale as the resolution function, the PSF. Hence, it is favorable if the image processing does not further degrade the scale. At the same time, filtering is usually applied to cope with the noise in the data. A disadvantage of Kindlmann's method is that it implicitly assumes data of isotropic resolution. If the data is not isotropic, filtering is required to adjust the smallest scale to the largest scale.

Second, the segmentation of a single material connecting to more than one other material cannot be based on a scalar value alone, e.g. due to the PVE. An advantage of Kindlmann's method is that it discriminates between different types of boundary voxels that share the same range of data values using the higher order image structure as obtained by image derivatives. However, the method still requires several transfer functions - one for each adjacent material - to visualize one object. In addition to this, manually tuned transfer functions are dependent on the expertise of the operator and may not be used if reproducible and accurate object delineation is required.

Third, a typical problem with volume rendering methods is that the transfer function is directly applied to the data value. Since all images are hampered by noise, this noise may be amplified by the transfer-function in a non-linear manner.

4.1.2   Objective

Our objective is to automatically classify scalar-valued 3-D CT images into material-fractions. Our approach uniquely combines robustness to noise, global signal fluctuations, anisotropic resolution, non-cubic voxels and ease of use. An iso-surfaces of constant material fraction provides a straightforward way to represent the boundary of a material in a reproducible manner. This facilitates straightforward segmentation of images into objects of interest for subsequent quantification in CAD or high quality scientific visualization. The latter is used to illustrate the benefits of this approach.
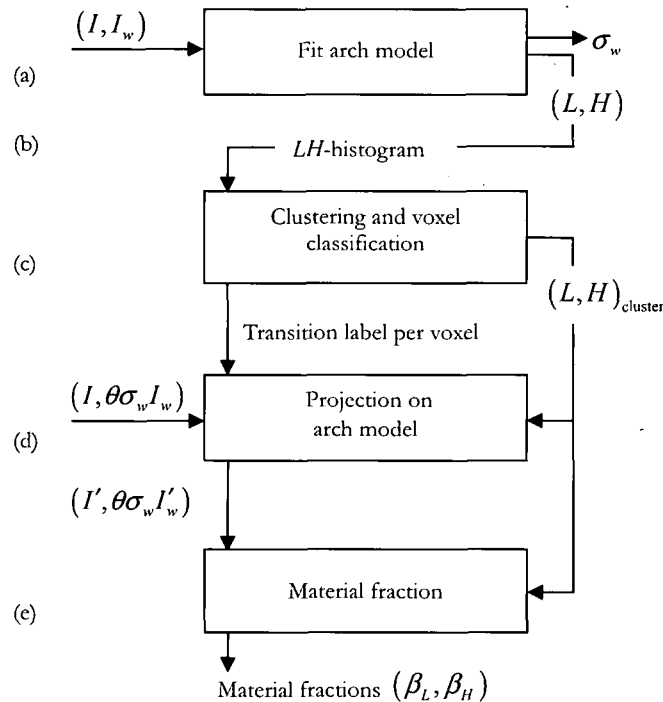
Figure 4.2   (a) Model parameters (L,H) are obtained by fitting the arch model to the measurement pair (scaled to have isotropic noise). (b) Estimated model parameters (L,H) are stored in the LH-histogram. (c) Clustering of the LH-histogram yields material transitions and is used to classify voxels into transition type. (d) Projection of the noisy data value and derivatives onto the edge-specific arch model yields a robust alternative to the standard computed Gaussian derivatives. (e) Projected measurements are mapped on the material fractions β_L β_H corresponding to material L and H.

Our approach is rotation-invariant, uses a narrow strip of voxels in the edge region that is smaller than the PSF footprint, does not rely on additional blurring for noise suppression and can therefore be applied to segment small and narrow structures.

The work of Kindlmann et al. is limited to data of isotropic resolution without global signal fluctuations, because it is based upon estimating derivatives from a histogram. Laidlaw et al. classify data into material fractions. However, a voxel is modeled as a cubic region, not explicitly modeling the scale of the data. Additionally, processing a single voxel is susceptible to noise. De Vries et al. classify data into material fractions using a neighborhood of the voxel. However, the method relies on sufficient voxels that are not hampered by the PVE.

Standard methods for edge detection use the maximum edge strength (Canny) or zero-crossing contours of the Laplacian-of-Gaussians (Marr-Hildreth). Many methods for 3-D edge detection are based on convolutions with Gaussian operators. This can be described

as local fitting of a Gaussian with Gaussian blurred edges with the maximum response at the edge position. However, with these methods the matched model is described as a function of position. Our approach differs from these methods, because it is based upon a model that describes Gaussian derivatives as a function of data value or material fraction. This enables straightforward extraction of contours at a constant material fraction. In addition, fitting the arch-model to the data does not require an accurate estimate of the gradient direction, because the distance to the edge is not a parameter. By including knowledge of the expected data value of pure materials our method generates contours that border on a specific material. Still, it uses the best of previous methods, because it is based upon the Gaussian.

4.2   Methods

4.2.1   Outline

The proposed method (Figure 4.2) employs an analytical expression called *arch* model, which relates the scale-invariant n^{th} order derivative to the data value *and* hence the material fractions along transitions.

A single arch function is parameterized by the expected pure material intensities at opposite sides of the edge (L, H) and a scale parameter, the standard deviation σ of the apparent Gaussian PSF (depends on the edge orientation for anisotropic PSF of the scanner). The parameters and L and H denote respectively the low and high material intensities. The apparent scale σ allows us to account for the space-variant, orientation dependent resolution of the scanner. The model has been constructed such that it directly describes Gaussian derivatives without having the distance to the edge as a parameter. In addition, local fitting of this model results in material intensities as function parameters (Figure 4.2a). The model is fitted to a set of measurement acquired by applying orthogonal Gaussian operators to a set of edge voxels. These edge voxels form a path along the gradient direction inside the support of the PSF. This yields local estimates for model parameters L, H and σ. The LH-parameter space is represented by an LH-histogram (Figure 4.2b). A peak in the LH-histogram constitutes one type of material transition, i.e. between the L and H values of the cluster. Cluster membership is used to classify edge voxels into transition types (Figure 4.2c).

The measured data value and gradient magnitude for a single voxel are independent, but display different noise variances. First, we make them invariant to the edge-orientation dependent apparent scale of the data and second we scale them in such a way to obtain isotropic noise. Previous work (e.g. [6]) disregards edge-orientation dependent apparent scale, which is caused by the anisotropic resolution of the scanner and does not model the noise properly. The scaled measurement pair is projected onto the model it has been assigned to in the voxel classification step (Figure 4.2d). The projection provides an estimate of the underlying noise-free data value and the true gradient magnitude that are

less sensitive to noise than an estimate obtained by Gaussian derivative filters of the same scale. The relative position of the estimated data value between the local L and H yields the material fractions (Figure 4.2e).

4.2.2 Transition model

A two-material transition (Figure 4.3a) is modeled by a unit step-function u (4.2) that is convolved with a 1-D Gaussian edge spread function (ESF) g (4.3) resulting in a cumulative Gaussian distribution G (4.1) (Figure 4.3b). The true edge-location is defined at $x = 0$. It has been shown that the cumulative Gaussian is an excellent model to describe the CT values across a two-material transition (Chapter 2). For a given direction, the ESF is approximately constant over the image and, therefore, has not to be re-estimated for every voxel (Chapter 2).

$$G(x;\sigma) = u(x) * g(x;\sigma) = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{x} \exp\left(\frac{-\tau^2}{2\sigma^2}\right) d\tau = \frac{1}{2} + \frac{1}{2}\text{erf}\left(\frac{x}{\sigma\sqrt{2}}\right) \quad (4.1)$$

with

$$u(x) = \begin{cases} 0 & x < 0 \\ 1 & x > 0 \end{cases} \quad (4.2)$$

$$g(x;\sigma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(\frac{-x^2}{2\sigma^2}\right) \quad (4.3)$$

$$\text{erf}(x) = \frac{2}{\sqrt{\pi}} \int_{0}^{x} \exp\left(-t^2\right) dt \quad (4.4)$$

A compact description of edges in 3-D images is obtained using gauge coordinates, a local Cartesian coordinate-system with axis that are aligned to the intrinsic local image coordinates. Let w represent the gradient direction, $v = w_\perp$ the basis of the isophote surface and σ_w the scale of the Gaussian function along w. Notice that a description of transitions in gauge coordinates is by definition both rotation and translation-invariant.

We assume that (a) materials are pure and only produce mixtures as a result of the convolution with the PSF, (b) the scale of the edge spread is known, for instance by calibration. Initially, we will assume that (c) the expected data values at opposite sides of the transition are 0 and 1 such that $I(w;\sigma_w) \triangleq G(w;\sigma_w)$ denotes the data value (4.1) and $I_w(w;\sigma_w) \triangleq g(w;\sigma_w)$ denotes the gradient magnitude (4.3) (in 4.2.5 we generalize the model). In the remainder of the text, we occasionally drop the (space-occupying) position information for the sake of clarity. When plotting the gradient magnitude $I_w(w)$ as a function of data value $I(w)$, arch-shaped point-clouds appear (Figure 4.1b). This representation is used by Kindlmann and Kniss for visualization purposes [7][8].

4.2.3 Scale-normalization

Scanners with a significantly anisotropic PSF cause that the apparent edge scale σ_w and therefore the observed gradient magnitude depend heavily on the edge orientation. Figure 4.1a contains a three-material phantom after scanning with an anisotropic PSF. Consequently, the scatter plot of intensity and gradient magnitude (I, I_w) as depicted in Figure 4.1b yields a wide range of arches between data values 0 and 1000. All arches share the same base along the horizontal axis, but have a height that is inversely proportional to σ_w. Let $\sigma_w I_w$ be the scale-invariant gradient magnitude [11] across a transition as depicted in Figure 4.3c. Plotting $(I, \sigma_w I_w)$ yields a single, scale and rotation invariant arch (Figure 4.1c, Figure 4.3d) of height $(2\pi)^{-1/2}$. The dashed lines in Figure 4.1c are generated by the model. The spread that remains is caused by noise.

4.2.4 Analytical expression

In this section, we will derive an analytical expression for the scale-invariant arch-function (Figure 4.3d). The arch-function describes the $(I, \sigma_w I_w)$-relation around the transition between two materials in a 3-D image, irrespective of the edge orientation, even in the case of an anisotropic PSF and non-cubic voxels. A first step is to determine the inverse cumulative Gaussian function G^{-1}. It is obtained by inserting (4.1) in $G\left(G^{-1}(x)\right) = x$ for $x \in [0,1]$. Solving for $G^{-1}(x)$ yields

$$\frac{1}{2} + \frac{1}{2}\mathrm{erf}\left(\frac{G^{-1}(x;\sigma)}{\sigma\sqrt{2}}\right) = x \Leftrightarrow$$

$$\frac{G^{-1}(x;\sigma)}{\sigma\sqrt{2}} = \mathrm{erf}^{-1}(2x-1) \Leftrightarrow G^{-1}(x;\sigma) = \sigma\sqrt{2}\,\mathrm{erf}^{-1}(2x-1) \tag{4.5}$$

The final step consists of multiplying (4.3) with σ to make the gradient-magnitude scale-invariant and substitution of x with $G^{-1}(x)$ of (4.5). This yields $\mathrm{arch}(x)$:

$$\mathrm{arch}(x) \triangleq \sigma g\left(G^{-1}(x;\sigma);\sigma\right)$$

$$= \frac{1}{\sqrt{2\pi}} \exp\left(-\left\{\mathrm{erf}^{-1}(2x-1)\right\}^2\right) \tag{4.6}$$

The arch-function is used to describe the scale-invariant gradient magnitude $\sigma_w I_w$ as a function of I. Both I and I_w are measured at scale σ_w. Note that the function $\mathrm{arch}(I)$ does *not* depend on a scale parameter. Therefore, the arch efficiently describes scale-invariant measurements: an advantage that is also inherited by the *LH*-histogram (method section 4.2.10).

In general, $\mathrm{arch}_n(x)$ is the n^{th} order scale-invariant derivative (Figure 4.3e) as a function of the Gaussian filtered data value (Figure 4.3f). It is obtained using a modified version of the Hermite polynomial of order n, H_n. The scale-invariant n^{th} order derivative of the cumulative Gaussian distribution G becomes:

$$\sigma^n \frac{d^n}{dx^n} G(x;\sigma) = \sigma^n \frac{d^{n-1}}{dx^{n-1}} g(x;\sigma) = \sigma^n H_{n-1}(x;\sigma) g(x;\sigma)$$

$$\text{with } H_n(x;\sigma) \triangleq (-1)^n \exp\left(\frac{x^2}{2\sigma^2}\right) \frac{d^n}{dx^n} \exp\left(\frac{-x^2}{2\sigma^2}\right)$$

(4.7)

Substitution of x in (4.7) with (4.5) gives $\mathrm{arch}_n(x)$.

$$\mathrm{arch}_n(x) \triangleq \sigma^{n-1} H_{n-1}\left(G^{-1}(x;\sigma);\sigma\right) \sigma g\left(G^{-1}(x;\sigma);\sigma\right)$$

$$= \sigma^{n-1} H_{n-1}\left(G^{-1}(x;\sigma);\sigma\right) \mathrm{arch}(x)_{_{\cup}}$$

(4.8)

$$= S_n(x)\, \mathrm{arch}(x)$$

with $S_n(x) = \sigma^{n-1} H_{n-1}\left(G^{-1}(x;\sigma);\sigma\right)$. The $\mathrm{arch}_n(x)$ is scale-invariant and one curve efficiently represents all measurements at a transition with the remaining spread caused by noise. The Gaussian derivatives of G up to the fourth order and the corresponding arch functions are depicted in Figure 4.3e-f. The arch-function is related to the inverse cumulative Gaussian (4.9) through its derivative (Appendix I):

$$\mathrm{arch}'(x) \triangleq \frac{d}{dx}\mathrm{arch}(x) = -G^{-1}(x,1) \Leftrightarrow$$

$$\int_{-\infty}^{x} \mathrm{erf}^{-1}(t)\,dt = \frac{-2}{\sqrt{2}}\mathrm{arch}\left(\frac{x+1}{2}\right)$$

(4.9)

It can be concluded that $\mathrm{arch}(x)$ closely resembles a parabolic function around the peak, where its derivative is approximately linear. Moving away from the peak the function deviates more and more from a parabolic function, as indicated by a rapidly changing slope of G^{-1}.

Equation (4.6) is analytical but not in a closed-form. Consequently, evaluating $\mathrm{arch}(x)$ is cumbersome, since it requires finding the roots of $2x - 1 - \mathrm{erf}(y) = 0$. This problem is circumvented by considering the inverse function $\mathrm{arch}^{-1}(\mathrm{arch}(x)) = x$ for $x \in [0, 1/2]$ that has a closed-form expression (Appendix I)

$$\mathrm{arch}^{-1}(x) = \frac{1}{2} - \frac{1}{2}\mathrm{erf}\left(\sqrt{-\ln\left(x\sqrt{2\pi}\right)}\right) \text{ for } x \in \left[0, \frac{1}{\sqrt{2\pi}}\right]$$

(4.10)

The inverse arch-function describes the Gaussian filtered data value I as a function of scale-invariant gradient-magnitude: $I = \mathrm{arch}^{-1}(\sigma_w I_w)$.
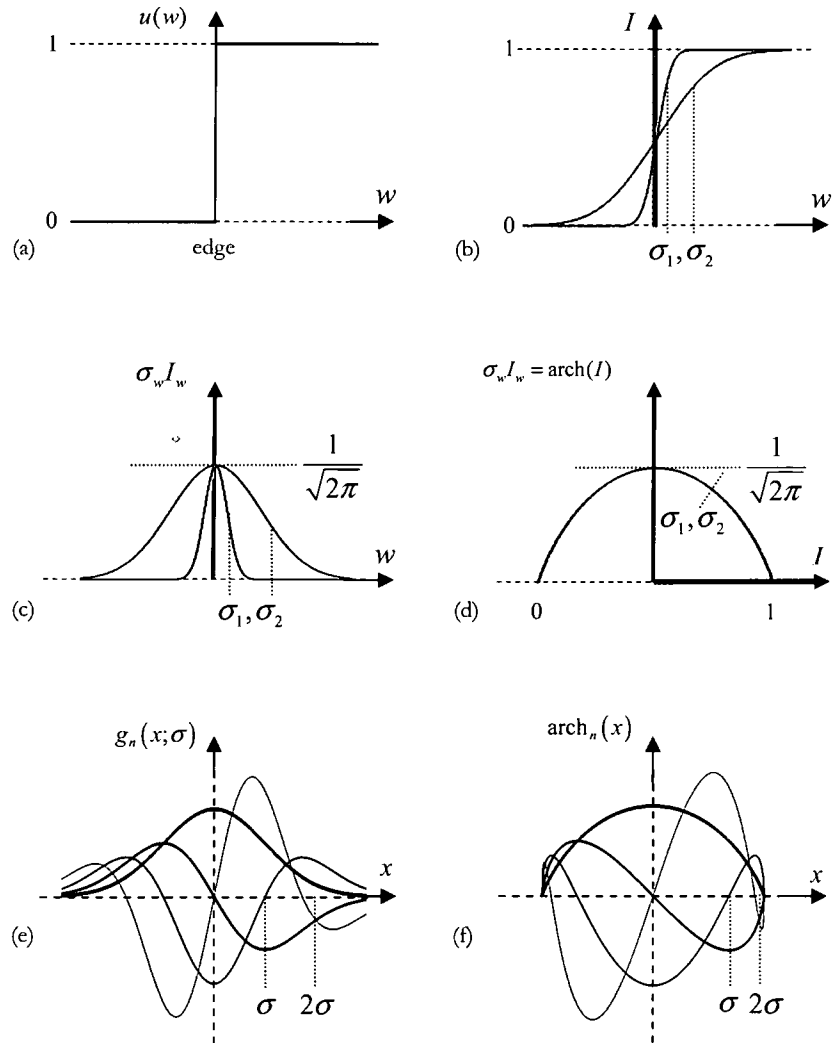
Figure 4.3   (a) A material transition modeled by the unit step edge u. (b) The data values are blurred at two scales. (c) Scale-invariant gradient magnitude as a function of position. (d) A single arch is obtained upon plotting the scale-invariant gradient magnitude as a function of data value. (e) Gaussian derivatives (order $n = 1, 2, 3, 4$) of a step-edge. (f) Arch function (order $n = 1, 2, 3, 4$) express the n^{th} Gaussian derivative as a function of the Gaussian filtered data value.

4.2.5 Generalization towards arbitrary intensity levels

Thus far, we have assumed a transition between materials with expected data values 0 and 1. The description is now generalized by adding two parameters to represent the expected data values L and H with $L < H$. A two-material edge is modeled as a scaled unit-step-function:

$$u(w; L, H) = L + (H - L)u(w) \tag{4.11}$$

Let $I(w; \sigma_w, L, H)$ represent the Gaussian filtered step-edge at scale σ_w and $\sigma_w I_w(w; \sigma_w, L, H)$ the scale-invariant gradient magnitude at the transition:

$$I(w; \sigma_w, L, H) = L + (H - L)G(w; \sigma_w) \tag{4.12}$$

$$\sigma_w I_w(w; \sigma_w, L, H) = \sigma_w(H - L)g(w; \sigma_w) \tag{4.13}$$

The generalized arch-function describes the scale-normalized gradient magnitude $\sigma_w I_w$ as a function of the intensity I and the expected data values L and H:

$$\operatorname{arch}(I; L, H) = (H - L)\operatorname{arch}\left(\frac{I - L}{H - L}\right) \tag{4.14}$$

4.2.6 Noise isotropy

Measurements, including $(I, \sigma_w I_w)$, yield the noise-free values (step-edge convolved by Gaussian PSF) contaminated by noise. The noise is assumed to be Gaussian distributed with zero mean and variance σ_{ni}^2 (like in [4][12]). An estimate of these noise-free values is obtained by mapping the measured values of $(I, \sigma_w I_w)$ onto the 'closest' point on the corresponding arch. The distance metric to be used depends on the covariance matrix of the noise. The two measurements are obtained by orthogonal operators; hence these measurements have $\operatorname{cov}(I, \sigma_w I_w) = 0$, but may display different *variances*. An isotropic (Euclidean) metric can be used if the derivative is scaled by a factor θ such that the noise in $(I, \theta\sigma_w I_w)$ is isotropic. In that case we can use the orthogonal projection from the point $(I, \theta\sigma_w I_w)$ onto the θ-weighted arch.

The relation between the variances (of the noise) before σ_{ni}^2 and after convolution σ_{no}^2 with a D-order Gaussian derivative of scale σ_{op} in N-dimensional space is [4]:

$$\frac{\sigma_{no}^2}{\sigma_{ni}^2} = \frac{(2D)!}{\sigma_{op}^{N+2D}\pi^{N/2}D!2^D 2^{N+D}} \tag{4.15}$$

Typically, for medical images, the sampling along the scanner's z-axis (axial, slice pitch, out-of-plane) is often lower when compared to the x- and y-dimensions (lateral or in-plane). We would like to use Gaussian derivative filters that are not sampled isotropically to minimize additional blurring. Let Δ denote the sampling pitch of the signal. As a rule of thumb, the Gaussian operator should obey $\sigma_{op} \geq 0.9\Delta$ to meet the Nyquist sampling criterion [13]. Using smaller scales requires interpolation of the data, which reduces Δ to

satisfy the sampling criterion. Analogous to the PSF, we do not restrict the operator to be isotropic in Δ (i.e. anisotropic sampling of the operator). In three steps we (1) compute the variance after anisotropic Gaussian filtering, (2) compute the variance of the gradient magnitude as a function of anisotropic Gaussian filtering and edge orientation, and (3) increase the gradient magnitude by a scale factor θ to make the noise in $\left(I, \theta\sigma_w I_w\right)$ isotropic.

First, consider the variance of the noise after (0^{th} order) Gaussian filtering: the first dimension of $\left(I, \theta\sigma_w I_w\right)$ and the independent variable of $\text{arch}(x)$. Let $\sigma_{op,z}$ be the axial scale and $\sigma_{op,\perp z}$ the lateral scale of the operator with respect to the z-direction. Let σ_I^2 be the variance of the noise on I after filtering. The noise-isotropy is a prerequisite to using a Euclidean metric to obtain the closest point on the model (arch). The relation between the measurements of 0^{th} and n^{th} order is modeled using arch_n. Hence, absolute noise measurements are not required and the variance σ_{ni}^2 on the measured input data is not needed. Decomposition of the Gaussian filter into an axial and a lateral component requires that we apply (4.15) with $\left(\sigma_{op,z}, D=0, N=1\right)$ and $\left(\sigma_{op,\perp z}, D=0, N=2\right)$ respectively:

$$\frac{\sigma_{no,z}^2}{\sigma_{ni}^2} = \frac{1}{\sigma_{op,z} 2\sqrt{\pi}}, \quad \frac{\sigma_{no,\perp z}^2}{\sigma_{ni}^2} = \frac{1}{\sigma_{op,\perp z}^2 4\pi} \tag{4.16}$$

Given that the two convolutions are applied in series (in arbitrary order since the convolution operator is commutative), σ_{no} of the first pass is substituted for σ_{ni} of the second pass. This gives a fixed variance after filtering in three-dimensions irrespective of the orientation of the edge w:

$$\sigma_{no}^2 = \sigma_I^2 = \frac{\sigma_{ni}^2}{2^3 \pi^{3/2} \sigma_{op,z} \sigma_{op,\perp z}^2} \tag{4.17}$$

Second, consider the variance of the noise when measuring the gradient magnitude: the second dimension of $\left(I, \theta\sigma_w I_w\right)$ and the result of $\text{arch}(x)$. This 3-D operation can be decomposed into a 1-D first Gaussian derivative filter in the gradient direction w and a 2-D Gaussian filter in the plane perpendicular to w. Let $\sigma_{op,w}$ be the effective scale of the operator in the gradient direction w as a function of the angle α_Δ between z and w.

$$\sigma_{op,w} = \sqrt{\left(\sin\left(\alpha_\Delta\right)\sigma_{op,\perp z}\right)^2 + \left(\cos\left(\alpha_\Delta\right)\sigma_{op,z}\right)^2} \tag{4.18}$$

Applying (4.15) with $\left(\sigma_{op,w}, D=1, N=1\right)$ and $\left(\sigma_{op,\perp w}, D=0, N=2\right)$ respectively gives

$$\frac{\sigma_{no,w}^2}{\sigma_{ni}^2} = \frac{1}{\sigma_{op,w}^3 4\sqrt{\pi}}, \quad \frac{\sigma_{no,\perp w}^2}{\sigma_{ni}^2} = \frac{1}{\sigma_{op,\perp w}^2 4\pi} \tag{4.19}$$

These two convolutions applied in series provide the variance of the noise after filtering in 3-D:

$$\sigma_{no}{}^2 = \sigma_{I_w}{}^2 = \frac{\sigma_{ni}{}^2}{16\pi^{3/2}\sigma^2{}_{op,\perp w}\sigma^3{}_{op,w}} \tag{4.20}$$

Note that the variance of the gradient-magnitude remains a function of the edge orientation w. Finally, using (4.17) and (4.20), the noise in $\left(I, \theta\sigma_w I_w\right)$ is made isotropic with

$$\theta = \frac{1}{\sigma_w}\left(\frac{\sigma_I}{\sigma_{I_w}}\right) = \frac{1}{\sigma_w}\left(\frac{\sqrt{2}\sigma_{op,w}^{3/2}\sigma_{op,\perp w}}{\sqrt{\sigma_{op,z}}\sigma_{op\perp z}}\right) \tag{4.21}$$

Suppose, for example, that a Gaussian operator isotropic in Δ is used to measure derivatives with $\sigma_{op,w} = \sigma_{op,\perp w} = \sigma_{op,z} = \sigma_{op,\perp z}$. Then the relation in (4.21) above is simplified considerably such that $\theta = (\sigma_{op}\sqrt{2})/\sigma_w$. Assuming the previous isotropy of the kernel, the noisy measurements $\left(I, \sigma_{op}I_w\sqrt{2}\right)$ are projected onto $\left(\sigma_{op}/\sigma_w\right)\operatorname{arch}(I)$. Remember that σ_w is the overall (apparent) edge scale $\sigma_w^2 = \sigma_{op,w}^2 + \sigma_{PSF,w}^2$.

4.2.7   Orthogonal projection on the arch

The measurements $\left(I, \theta\sigma_w I_w\right)$, obtained by orthogonal operators, are combined by projection onto the arch. To begin with, it is assumed that $L = 0$ and $H = 1$ to keep the description simple. The orientation of the projection is steered by the derivative of the arch. For this purpose we use the closed-form inverse arch function arch^{-1} (4.10) and its derivative (4.22) as depicted in Figure 4.4a:

$$\operatorname{arch}^{-1}(x) = \frac{1}{2} - \frac{1}{2}\operatorname{erf}\left(\sqrt{-\ln\left(x\sqrt{2\pi}\right)}\right)$$
$$\frac{d}{dx}\operatorname{arch}^{-1}(x) = \frac{-1}{\sqrt{\left(-\ln(2\pi) - 2\ln x\right)}} \tag{4.22}$$

Let $y = q(x, C)$ be the line orthogonal to $\operatorname{arch}^{-1}(x)$ with slope k and y-intercept m $(0 \le m \le 1/2)$ which crosses $\operatorname{arch}^{-1}(x)$ in point $C = (x_c, y_c)$. All measurements $\left(\theta\sigma_w I_w, I\right)$ on this line are projected onto point C of the $\operatorname{arch}^{-1}(x)$:

$$k = \left(\frac{d}{dx}\operatorname{arch}^{-1}(x)\right)^{-1}\Bigg|_{x=x_c} \tag{4.23}$$

$$m = -k * x_c + \operatorname{arch}^{-1}(x_c)$$

$$y = q(x, C) = kx + m = \frac{(x_c - x)}{\dfrac{d}{dx}\operatorname{arch}^{-1}(x)\Big|_{x=x_c}} + \operatorname{arch}^{-1}(x_c) \tag{4.24}$$

For a particular measured $x = \theta\sigma_w I_w, y = I$ the unique projection is achieved by numerically solving equation (4.24) for x_c. The second coordinate y_c is found by evaluating $y_c = \mathrm{arch}^{-1}(x_c)$. In general, the projection of a measurement onto the arch parameterized by L, H and scaled by θ requires proper scaling of the axis. This operation is written as $proj_{\mathrm{arch}(L,H,\theta)}$.

4.2.8   Fitting the arch function

The projection onto the arch-function requires the expected pure data values L and H at opposite sides of a transition to be known in advance. Our method decomposes partial volume samples into L and H fractions (as opposed to regarding partial-volume values as a problem).

Let A be a set of $N+1$ measurement pairs along the gradient direction in the neighborhood of an edge sample (Figure 4.4c):

$$A_n = (I, \theta\sigma_w I_w), n \in [0, N] \tag{4.25}$$

A 3-D version of Canny's edge detector is used for initial finding of edge samples. Because the arch described derivatives as a function of data value and not as a function of position (1) an edge sample needs not to be centered exactly at the edge and (2) the strip of voxels needs not to be exactly in the gradient orientation.

Since an edge is intrinsically translation invariant in the isophote plane, we select nearby voxel locations (perpendicular to the gradient direction) rather than apply interpolation. Furthermore, let $(I', \theta\sigma_w I'_w)$ be the orthogonal projection of a measurement pair onto an arch. By minimizing the summed squared residuals between the arch model and the measurements (Figure 4.4d) using the Conjugate Gradient Method [14], the best fitting arch is obtained, which then yields the L' and H' values that we are looking for:

$$\{L', H'\} = \underset{\{L,H\}}{\arg\min} \sum_n \left\| A_n - proj_{\mathrm{arch}(L,H,\theta)} A_n \right\|^2 \tag{4.26}$$

The residual error R may be used as a measure for the quality of the fit.

4.2.9   Initial values for L and H

Local minima may occur for (4.26). Consequently, accurate initial values for the L, H parameters are required. The estimates for L and H are obtained from I, I_w, I_{ww} (all measured in one voxel) and σ_w (assumed to be known after calibration). Let t be the distance from the current voxel to the nearest edge, estimated by

(a)

Part of CT data
Pixel replication to indicate voxels
Dots denote voxel-centers

(c)

(b)

(d)
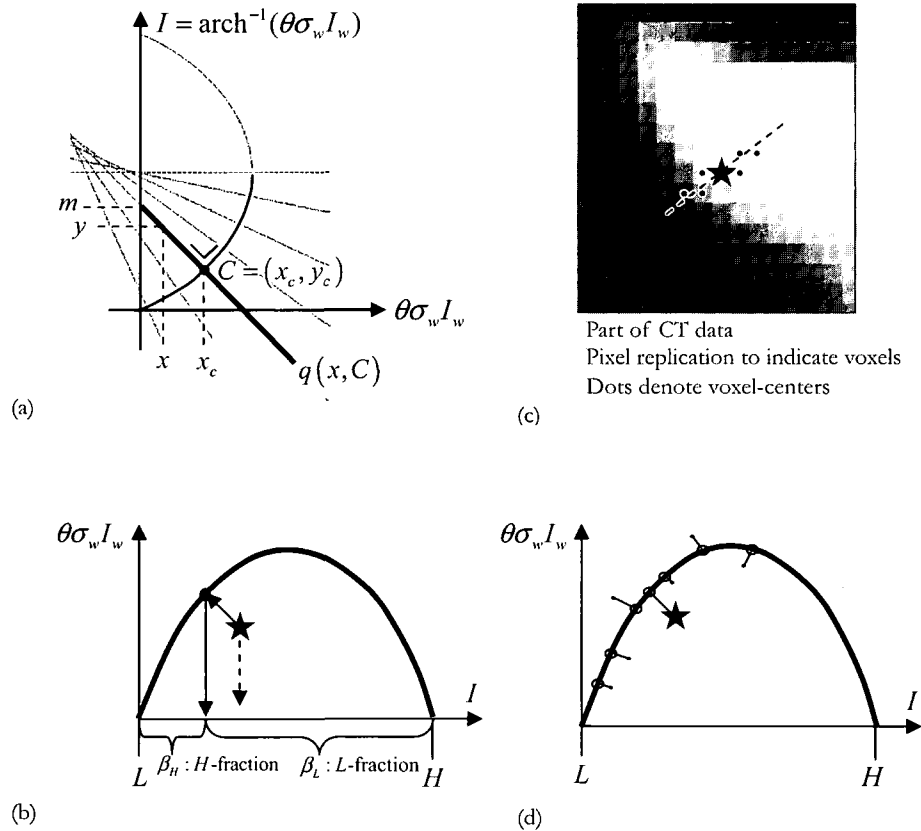
Figure 4.4   (a) Inverse arch-function showing the projection of values along a line onto a single point on the arch. (b) This method is applied to obtain material fractions. Notice that directly using intensity yields a different outcome (dashed arrow). (c) This projection is used as well to obtain L and H. A set of voxels along a path in gradient direction trajectory both uphill and downhill is shown in part of a slice from a CT volume with the voxel under investigation marked by (*). (d) L and H are obtained by fitting an arch model to the set of $(I, \theta\sigma_w I_w)$ measurements.

$$t = \frac{-\sigma_w I_{ww}}{I_w} \tag{4.27}$$

Equation (4.3) is used to compute $I'_w(t, \sigma_w)$ that yields the predicted gradient magnitude for $L = 0$ and $H = 1$. The ratio between the measured I_w and model value $I'_w(t, \sigma_w)$ defines $Q = H - L$:

$$Q = I_w / I'_w = H - L \tag{4.28}$$

Now the predicted intensity at position t for $L = 0$ and $H = Q$ is computed using (4.12)

$$I'(t, \sigma_w, L = 0, H = Q) = QG(t, \sigma_w) \tag{4.29}$$

Finally, L and H are determined from

$$L = I - I', \quad H = L + Q \tag{4.30}$$

The previous derivation leads to analytical expressions for initial guesses of L and H:

$$L(I, I_w, I_{ww}, \sigma_w) =$$
$$I + \frac{\sigma_w I_w \sqrt{\pi}}{\sqrt{2}} \left(-1 + \mathrm{erf}\left(\frac{\sigma_w I_{ww} \sqrt{2}}{2 I_w} \right) \right) \exp\left(\frac{I_{ww}^2 \sqrt{2}}{2 I_w^2} \right) \tag{4.31}$$

$$H(I, I_w, I_{ww}, \sigma_w) =$$
$$I + \frac{\sigma_w I_w \sqrt{\pi}}{\sqrt{2}} \left(1 + \mathrm{erf}\left(\frac{\sigma_w I_{ww} \sqrt{2}}{2 I_w} \right) \right) \exp\left(\frac{I_{ww}^2 \sqrt{2}}{2 I_w^2} \right) \tag{4.32}$$

4.2.10   *LH*-histogram

Fitting an arch to all sets of measurement pairs $(I, \theta \sigma_w I_w)$ yields L, H values for all voxels in the selected edge regions. All L, H values can be represented in a 2-D *LH*-histogram [6]. An *LH*-histogram provides a useful compact description of the data. The *LH*-histogram can be interpreted as the resulting parameter space of a generalized Hough transform [15]. The transform is applied to a set of measurements in $(I, \theta \sigma_w I_w)$-space. Disregarding the noise, all samples near a single transition contribute to one entry in *LH*-space.

If a material is connected to two or more other materials, their arches share one base value. Hence, these transitions cannot be separated in $(I, \theta \sigma_w I_w)$-space. Using higher order derivatives does not solve this problem because the arches still meet at a single base value. The *LH*-histogram however, does show separate peaks for each type of transition and allows identification of transition type through clustering. Figure 4.5 shows how different transitions yield separate clusters in *LH*-space. Four material samples yield crossing arches in $(I, \sigma_w I_w)$-space that can easily be separated since they map to different clusters in the *LH*-histogram.

Unfortunately, the arch-model is not valid at locations where material pockets are smaller than the size of the point-spread function as well as at T-junctions. Thin structures (smaller than the PSF) suffer heavily from the PVE (partial volume effect). None of its voxels represent pure material samples. Consequently, the thin structure's pure signal estimate (L or H) is biased towards the value of its background. A dark thin structure on a bright background will lead to horizontal lines between two clusters in the LH-histogram (a light structure on a dark background leads to a vertical line). At a three material T-junction with material intensities $M_1 < M_2 < M_3$, either L or H stays fixed depending on $M_3 - M_2$ vs. $M_2 - M_1$. In the LH-histogram, such locations are manifested as vertical (constant L) or horizontal (constant H) lines.

To discard these points from the LH-histogram, thresholds are applied to the gradient magnitude of both the L-channel and the H-channel. The thresholds are selected automatically from the histograms of the L- and H-channel gradients. First, the peak is found by searching the maximum values in both histograms. Subsequently, searching to the right, the 90% percentile is located. Because two-material transitions do occur far more frequently than three-material transitions, the left parts of both histograms include the majority of the two-material transitions. At last, those points are included in the LH-histogram, which have gradient magnitudes below the selected thresholds. The filtered LH-histogram describes the data adhering to the arch-model. Note that the LH-histogram inherited some important properties of the arch model such as translation-, rotation- and scale invariance.

4.2.11   Classification into material fractions

A simple clustering technique applied to the filtered LH-histogram allows identification of transitions [6]. This step implicitly segments the input data into transition types. Currently, we retain the locally obtained LH-values to be robust against fluctuations in signal intensity.

With $(I', \theta\sigma_w I'_w)$ the orthogonal projection of the sample onto the selected arch, β_L represent the material fraction corresponding to L and β_H the material fraction corresponding to material H (Figure 4.4b). These material fractions are obtained by:

$$\beta_L = \frac{I' - L}{H - L}, \quad \beta_H = \frac{H - I'}{H - L}, \quad \beta_L + \beta_H = 1 \qquad (4.33)$$

Material fractions remain undefined at positions where the arch model is not valid. However, the majority of edges in 3-D images are two-material edges and only a small number of applications would benefit from non two-material analysis. We may simply ignore such locations or merely select the nearest transition type and apply the previous mapping for a first order estimate of true material fractions (we adhere to the latter solution in the examples presented).
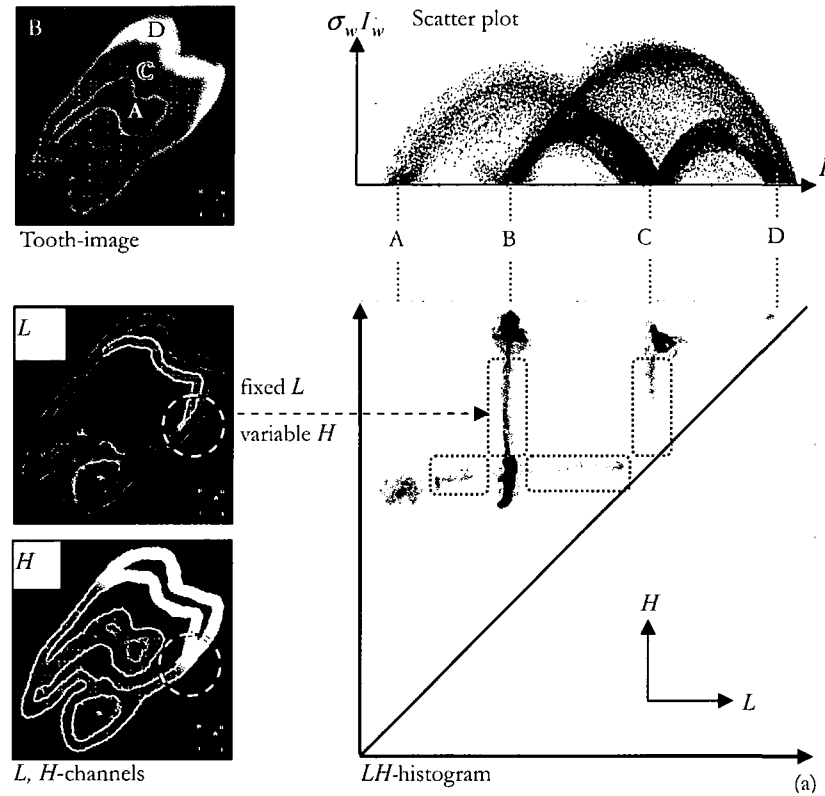
Figure 4.5   *LH*-histograms with the corresponding *L*- and *H*-channel, the scatter plot of data value
and scale-invariant gradient magnitude before (a) and after (b) filtering of the
L- and *H*-channels. (a) The dashed rectangles mark regions in the *L*- and *H*-channels
where the arch-model is not valid. (b) These invalid regions are detected automatically by
thresholding the gradient magnitude of the *L*- and *H*-channels. The clusters in the colored
circles of the *LH*-histogram have the same color in the scatter plot above.

4.2.12   Visualization

The concept of material fractions is useful to cope with a material that borders on two or
more other materials. The exact, subvoxel location of such edges depends on the type of
transition. Having a material fraction volume allows one to delineate an object at a single
isosurface at threshold 0.5.

The boundary of a single material is visualized directly by defining a surface at such a
constant material fraction. Currently available rendering engines can be used after
mapping the material fraction onto integer values. An accurate volume measurement is
obtained by summing the fractions multiplied by the volume of a voxel [16].

σ_w, I_w Scatter plot after labeling

LH-histogram
after filtering the
L-, H-channels

A   B   C   D

L

H

H

L

(b)   L-, H-channels   LH-histogram
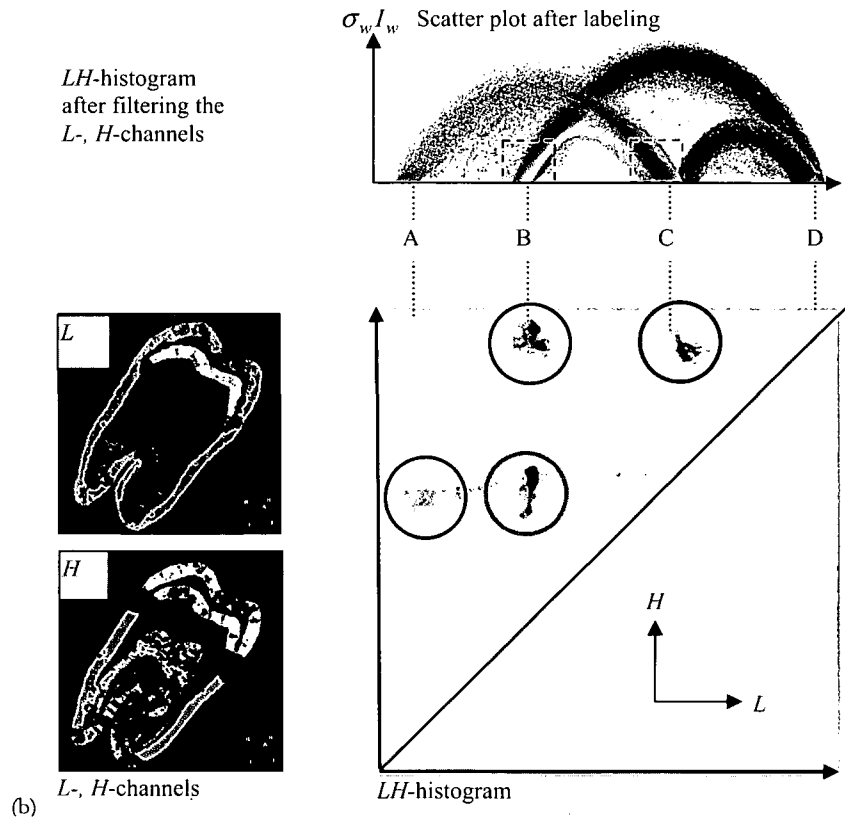
Likewise, the union of two objects may be represented by the boundary where the sum of the two material volume fractions is equal to 50%. At last, the 'intersection' (or touching) surface between two materials may be identified by the iso-surface where the difference of the fractions crosses zero, provided that the first derivative of the difference is non zero to reject positions where both fractions are zero.

4.3   Results

We will demonstrate the usefulness of the presented methods using phantom-data, a publicly available CT volume of a tooth and abdominal CT data for virtual colonoscopy. Typically, the sizes of the data were 512^3 voxels. The processing took approximately five minutes per volume on an AMD64 2.4 GHz.

4.3.1   Example I: edge localization

An important problem is the robustness of edge localization in the presence of noise and small deviations in image intensity. For instance, in CT images contrast media will never be distributed homogeneously. The accuracy and precision of the edge-localization were tested using the tube phantom (Figure 4.1). Reference data were an ultra high dose CT image (400 mAs). The raw transmission measurements were replaced by a realization of a Poisson process given a scaled version of the data value as the expected value to simulate very low dose images (20 mAs, $\sigma_{ni} =\sim 60$ HU) [17]. Moreover, the resulting low dose CT image was modified to contain a small trend in data value to represent inhomogeneities. In this way, the higher density of the contrast matter is modeled while proceeding from cecum to rectum in the CT colon images. The minimum and maximum values of contrast matter were 400 HU and 600 HU respectively. The edge-position was estimated in high dose data not containing the trend. Thereafter, the edge-position was located in low dose data with the trend added and the smallest displacement (q) was retrieved for positions on the contrast-plastic transition. The region of interest was manually indicated.

Kindlmann's method relies on obtaining estimates of $I_w(I)$ and $I_{ww}(I)$ from the histogram. The histograms obtained from a high-dose image without a trend in signal value are shown in Figure 4.6a,b (with scale 1.4 voxel). For a fair comparison with our method (see below), scale normalized derivatives were used. The histograms obtained from the low-dose images that *did* contain a trend are shown in Figure 4.6d,e. Apparently, the trend results in errors on $I_w(I)$ and $I_{ww}(I)$ due to the distortion of the arches in the histogram by noise and the trend in signal value Figure 4.6d,e. The resulting localization accuracy and precision are presented in Figure 4.6j,k (grey lines).

Alternatively, it was tested if a Gaussian mixture model fit through expectation maximization can be used to accurately determine the location of the edge in noisy data [9]. The resulting localization accuracy and precision measured in low dose data are indicated in Figure 4.6j,k by arrows. Only a single point is obtained, since there is no kernel involved in the estimation. Figure. 4.6i represents the tissue fraction in grey value (light grey means 100% Lucite). From the result it may be concluded that a Gaussian mixture model is an excellent method to make a first estimate of a voxel's material constituency. However, the edge-spread-function is not explicitly modeled, which is visible by sharp boundaries.
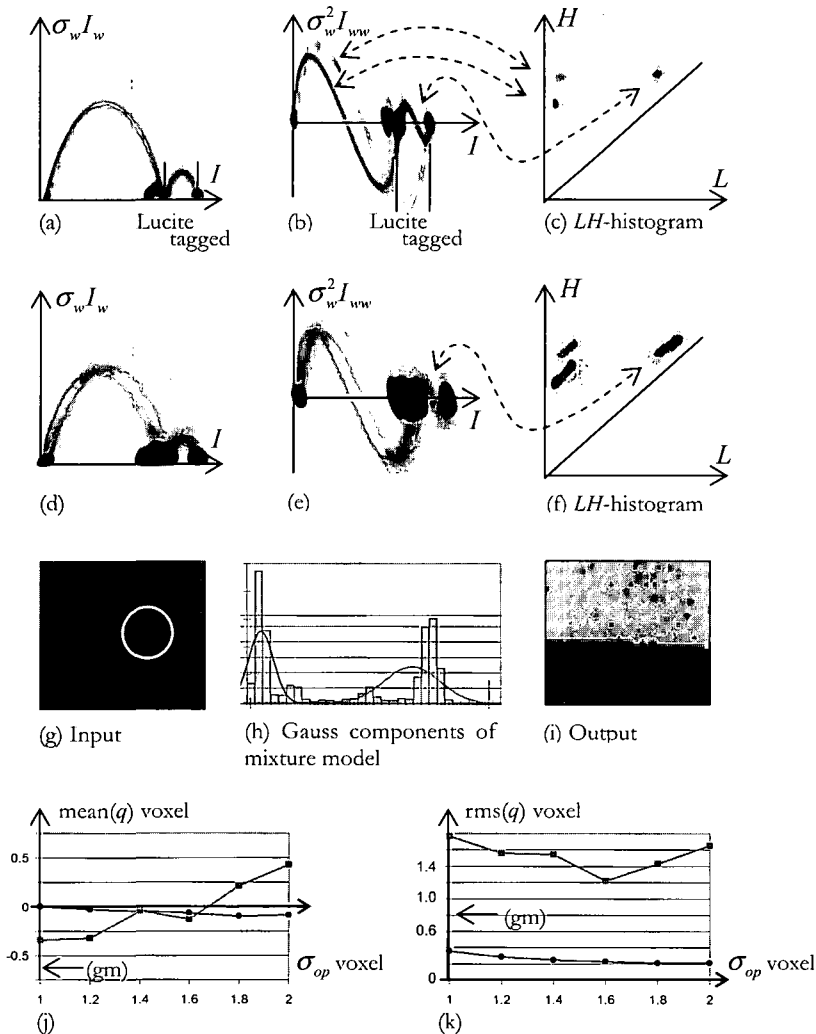
Figure 4.6 Histograms of (a) I,I_w and (b) I,I_{ww} that are used by Kindlmann to obtain estimates of $I_w(I)$ and $I_{ww}(I)$. The graphs are obtained from a high dose CT image of the phantom shown in Figure 4.1. (c) LH-histogram by fitting the arch: noise and global signal fluctuations are no problem if clusters are separated. (d-f) Similar graphs now obtained from a simulated low dose image in which the tagged matter signal values contained a trend. (g) Part of a CT slice; (h) Gaussian components of mixture model estimated from the local histogram using the expectation maximization method; (i) The estimated Lucite component is indicated in grey value (light grey corresponds to 100 % Lucite). (j-k) To compare methods, the smallest distance q between the estimated edge position in ultra high dose and low dose is measured. (j) The mean(q) and (k) root mean square(q) are plotted for (grey line) Kindlmann's method, (black line) the arch based method and (gm) the mixture model.

The arch-based method relies on being able to separate the clusters in *LH*-space (Figure 4.6c,f shows the outcome for the tube without and with the trend). Note that Kindlmann uses the scatter plots such as Figure 4.6a,b both to obtain the derivatives and to classify the data into edges. A fundamental difference with his work is that our classification is based on modeled derivatives, local *L*- and *H* values obtained from the CT data, and the *LH*-histogram. The resulting localization accuracy and precision of our method are also displayed in Figure 4.6j,k (black lines). Note that the arch-based method yields zero-bias throughout all scales and superior precision. The projection of measurements onto the 'correct' arch reduces the variance significantly without giving rise to a bias term.

4.3.2   Example II: tooth

The input data consist of industrial CT data of a human tooth from the National Library of Medicine: http://nova.nlm.nih.gov/data/. The samples are spaced 1 mm apart within each slice and the slices are 1 mm apart. Three types of materials and background can be identified: dark root canals and pulp chamber, grey dentin and cementum and bright enamel and crown (Figure 4.7a). Our method is used to extract three material-fraction volumes corresponding to the three materials.

The *L*- and *H* parameters for each point are determined by locally fitting the arch-model. These *L*- and *H* values yield the *LH*-histogram (Figure 4.5a). Clusters in the histogram correspond to the two-material transitions between tooth-materials. In addition, the *LH*-histogram in Figure 4.5a shows horizontal and vertical lines connecting the clusters. These emanate from thin structures or junctions. An example is the junction of background (B), dentin (C) and enamel (D) marked by the dashed circles in the *L, H*-channels. This structure contains a smooth transition from one *L*-value to two *H*-values. These locations can be detected (and masked out) by a large gradient magnitude in the *L*- and/or *H*-channels in Figure 4.5a. A 'filtered' histogram is shown in Figure 4.5b. The filtered *LH*-histogram is used to segment the image into significant material transitions. The clusters within the colored circles in the *LH*-histogram of Figure 4.5b correspond to similarly colored dots in the scatter-plot, which serves to illustrate the histogram's capacity in identifying the various transition types. For example, notice how well the volume samples in overlapping areas in the labeled scatter plot (dotted squares in Figure 4.5b) are classified into separate clusters in *LH*-space. Subsequently, the local arch fit of a particular volume sample is used to map it onto material fractions as indicated in Figure 4.4b.

The three resulting material-fraction volumes of root-canal dentin and enamel are used for visualization (Figure 4.7). The delineation of the root-canal is compared with a reference method of [18] (Figure 4.8). Sereda creates an *LH*-histogram in which the estimates of *L* and *H* are obtained by the result of a local min and a local max filter respectively [19] rather than by the arch model. The data value *I* is mapped on a mixture using β_L as defined in (4.33).
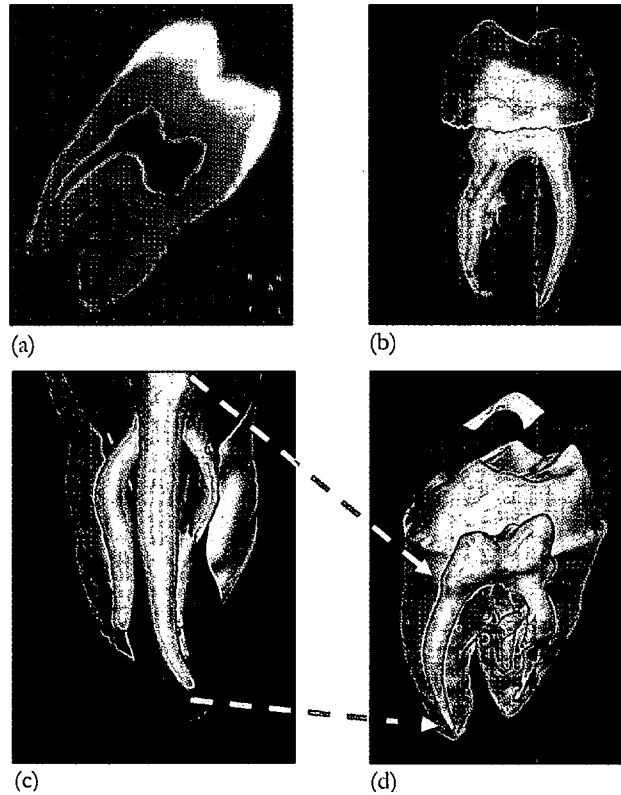
Figure 4.7   (a) Cross-section of tooth volume. (b) Delineations of the three materials (enamel, dentin
            and root canal) are combined into a single rendering. (c) and (d) show that the root-canal is
            almost completely visualized. (d) Insertion of a clipping plane on the enamel-fraction
            volume enables cutting away the crown and fully visualizing structures below. Thus, special
            pre-processing or complicated transfer function definition are not needed to create high
            quality visualizations.

Note that the estimates for L and H of small structures will be biased towards the
background value. This causes a bias in estimated material fraction, which causes serious
truncation of wedge shaped structures. A fundamental difference of our method when
compared to the reference method is that the data value does not serve as an input of
the transfer function. *First*, it is projected onto the material transition model to use this
knowledge to deal with the noise in an optimal manner. *Second*, L and H are estimated
using partial volume values without requiring 'pure material' volume samples to be
present.

Consider the root canal as a tube with decreasing diameter. A realistic surface delineation should convey this information as well. If we analyze the reference method, the distance in intensity between *H-L* becomes smaller for decreasing root-canal diameters by the PVE. When considering (4.33) the resulting mixture will become biased: not providing a good visualization (Figure 4.8a,c,e). At parts of the root canal with a decreasing diameter, a smaller mixture is found as expected (Figure 4.8b,d,f). The user may still visualize this by selecting a smaller threshold on the material mixture. The reference method however, even at a threshold of 50% root-material suffers from noise and the PVE (Figure 4.8a). It does not reliably convey the dimensions of the root canal. The user has less control of creating accurate object delineation.

4.3.3   Example III: CT colonography

CT colonography [6][20][21] (also called virtual colonoscopy) is a relatively new method to examine the colon surface for the presence of polyps. An important problem in virtual colonoscopy is to visualize the colon surface without being hampered by intraluminal remains [6]. Moreover, the partial volume effect may restrict correct polyp detection and subsequent quantification (diameter, volume) and accurate display of its morphology [22]. Tagging, via an oral contrast agent, is introduced to enhance the signal intensity of residue. We have applied our method to derive the fraction of soft tissue, air, and tagging in each edge voxel.

The data are acquired using a multi-slice CT scanner (Toshiba Aquilion). The axial scale of the resulting image is 0.95 mm and the lateral scale is 0.80 mm. Rendering opaque samples at a 50% tissue-level is assumed to yield an accurate representation of the colon surface. It reveals large parts of the colon surface that were previously obscured by fecal remains (Figure 4.9c). Analogous to the tooth example intermediate processing results are depicted in Figure 4.9a,b. The usefulness of our method is further demonstrated in Figure 4.9d, which shows the unfolded cube visualization [23] before and after processing.

4.4   Conclusion

We presented a novel approach to automatically classify scalar-valued 3-D CT images into material-fractions. Our approach uniquely combines robustness to noise, global signal fluctuations, anisotropic resolution, non-cubic voxels and ease of use. The method facilitates accurate and reproducible boundary delineation for segmentation and visualization.
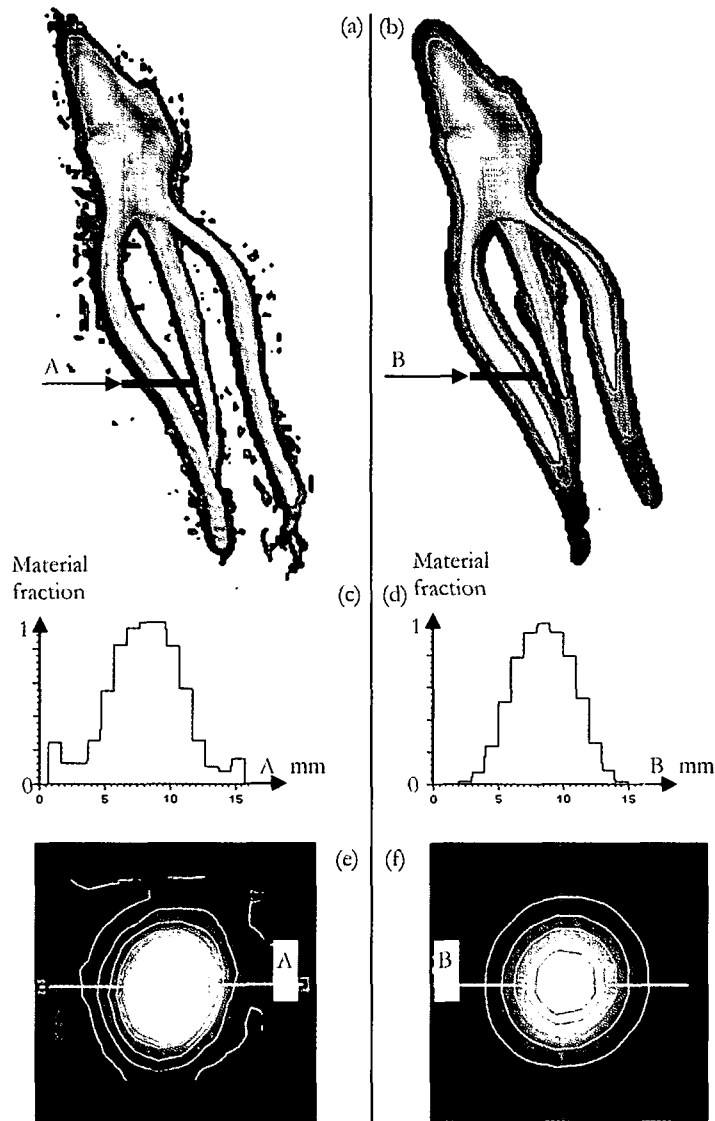
Figure 4.8   The root canal delineated using the reference method (a,c,e) and using the method described in this paper (b,d,f). The top row (a,b) shows the root canal rendering of the tooth volume using the reference method with threshold 0.2, 0.3 and 0.5 and method with threshold 0.1 (dark), 0.2 (intermediate) and 0.5 (bright). The middle row (c,d) shows the material mixture along a profile. The bottom row (e,f) shows the material fraction in grey-value with five delineations superimposed using an iso-material mixture of 0.1, 0.2, 0.3, 0.4 and 0.5 respectively.
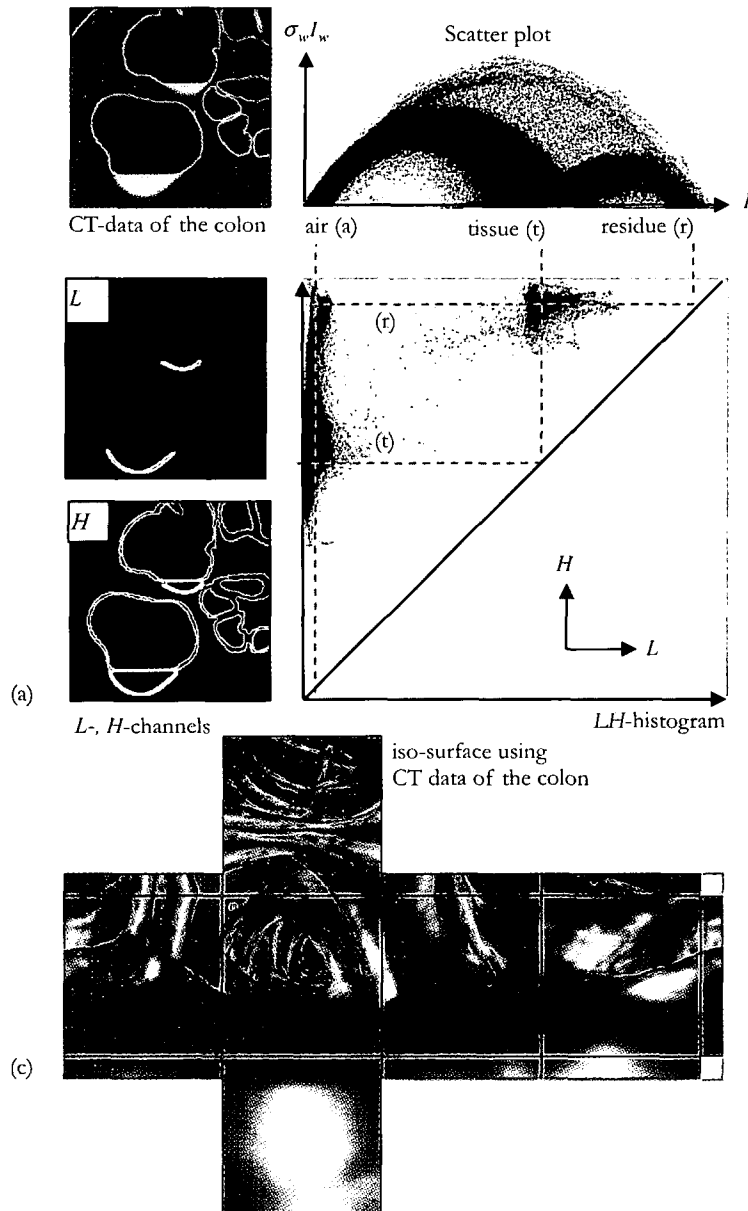
Figure 4.9   *LH*-histogram with corresponding *L*- and *H*-channel of abdominal CT data. The scatter-
            plot of the image-intensity and scale-invariant gradient magnitude before (a) and after (b)
            filtering the *L*- and *H*-channels. Regions in the *LH*-histogram where the arch-model is not
            valid are automatically suppressed by thresholding the gradient-magnitude of *L*- and *H*-
            channels. (c) Unfolded cube visualization using original data and (d) electronically cleansed
            data. Notice that the entire colon surface is presented for inspection after processing.

σ_w / I_w scatter plot after labeling

LH-histogram after filtering the L,H-channels

air (a)   tissue (t)   residue (r)

I

L

(r)

A

(t)

H

L

(b)

L-, H-channels   LH-histogram

(d)

Visualization using tissue-fraction from LH-analysis

Folds previously occluded by remains
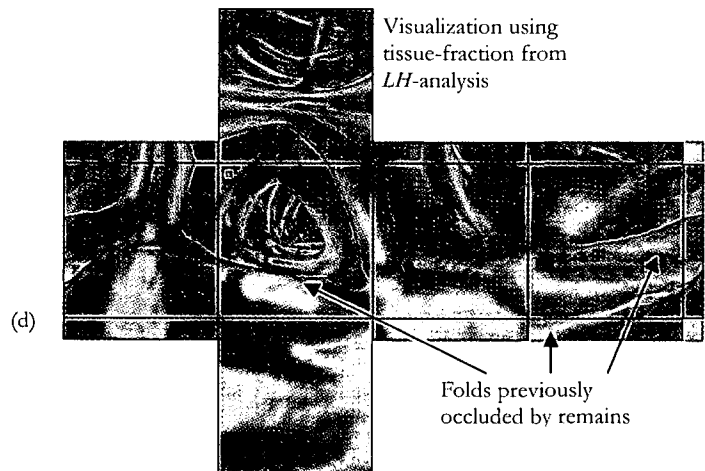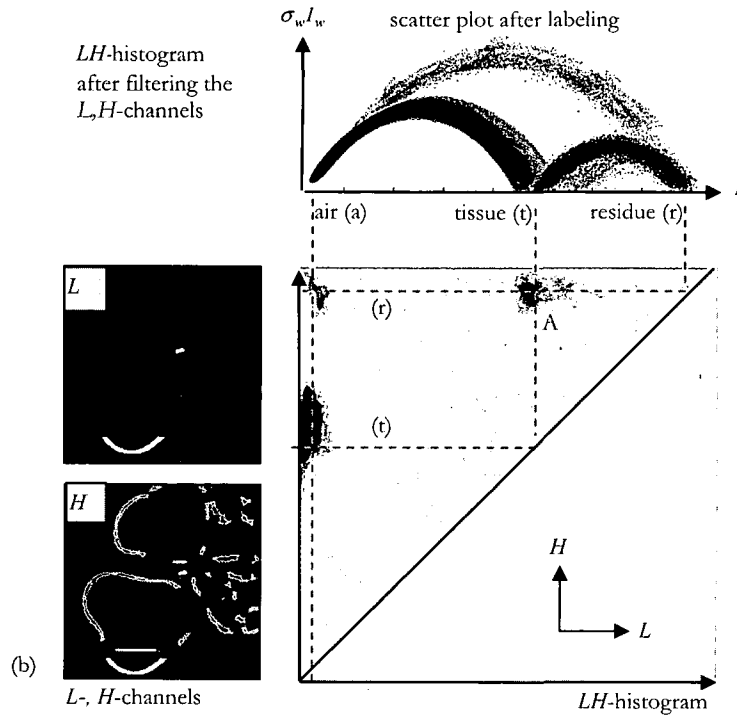
We derived an analytical expression for the relation of n^{th} order derivative as a function of data value: the arch function. It is applied to approximate the underlying noise-free material fractions or derivatives at an image position. Projecting of noisy data value and derivatives onto the arch model yields noise-free estimates of the data value and derivatives. It yields a robust alternative to the standard computed Gaussian derivatives. The arch function is rotation invariant even for anisotropic PSF. It is parameterized through the expected material data values (L and H) at opposite sides of the transition. The neighborhood of a sample is modeled by an arch trajectory, and not merely a single point. This makes the technique extremely robust against erroneous classification due to noise.

Previous work [6] did neither model the noise nor the anisotropy of the data. Including higher-order arch models one can obtain estimates of all derivatives up to this order. Both the accuracy and precision are superior to the results obtained by Gaussian derivatives. The main difference with existing methods for noise suppression such as high-order normalized convolution [22] is that an explicit edge model is used.

The LH-histogram was shown to be a useful description of the data. Overlaps occurring in the $(I, \theta\sigma_w, I_w)$ scatter plots were resolved in the LH-histogram. In addition, we demonstrated how to identify samples not adhering to the model's assumptions, e.g. three material crossings and thin layers. The 'filtered' LH-histogram constructed from masked L and H images was used to identify significant material transitions. The arch closest to a measurement was used to map it onto material fractions.

Most applications deal with two-material transitions. Voxels at multiple transition regions (more than two materials) are processed using the best two material transition. Incomplete processing is reported to leave artifacts at multiple transition regions [24]. The size of the artifact is related to the footprint of the PSF. Our current work focuses on improved image processing at these multiple transition regions.

We have demonstrated two visualization examples in which user interaction was merely required to decide which material to visualize. Objects of interest were rendered at the 50% material threshold. Thus, no complicated widgets were needed for transfer function selection. Previously described methods assume isotropic resolution. Clearly, anisotropic input data may be subject to additional blurring to meet such a requirement. It should be noted that our method does not sacrifice resolution because spurious blurring is avoided to retain the integrity of the data.

Acknowledgment

4.5   Appendix I

Evaluating $y = \text{arch}(x)$ requires calculation of $\text{erf}^{-1}(2x-1)$, which can only be computed indirectly after solving $2x - 1 - \text{erf}(y) = 0$ for y. A closed-form expression is given by the inverse-function $\text{arch}^{-1}(x)$ such that $x = \text{arch}(\text{arch}^{-1}(x))$. Substitution of (4.6)

$$x = \frac{1}{\sqrt{2\pi}} \exp\left(-\left\{\text{erf}^{-1}\left(2\,\text{arch}^{-1}(x)-1\right)\right\}^2\right) \Leftrightarrow$$

$$\ln\left(x\sqrt{2\pi}\right) = -\left\{\text{erf}^{-1}\left(2\,\text{arch}^{-1}(x)-1\right)\right\}^2 \Leftrightarrow$$

$$-\sqrt{-\ln\left(x\sqrt{2\pi}\right)} = \text{erf}^{-1}\left(2\,\text{arch}^{-1}(x)-1\right) \Leftrightarrow$$

$$-\text{erf}\left(\sqrt{-\ln\left(x\sqrt{2\pi}\right)}\right) = 2\,\text{arch}^{-1}(x)-1$$

(4.34)

yields the inverse arch-function:

$$\text{arch}^{-1}(x) = \frac{1}{2} - \frac{1}{2}\text{erf}\left(\sqrt{-\ln\left(x\sqrt{2\pi}\right)}\right) \text{ for } x \in \left[0, \frac{1}{\sqrt{2\pi}}\right]$$

(4.35)

Using the differentiation rule for inverse functions $(f^{-1})'(x) = (f'(f^{-1}(x)))^{-1}$ we find $d/dx\,\text{arch}^{-1}(x)$, the derivative of $\text{arch}(x)$

$$\frac{d}{dx}\text{arch}^{-1}(x) = \frac{1}{\dfrac{d}{dx}\text{arch}\left(\text{arch}^{-1}(x)\right)} \text{ for } x \in \left[0, \frac{1}{\sqrt{2\pi}}\right]$$

(4.36)

in which $d/dx\,\text{arch}(x)$ is proposed by a geometrical argument:

$$\frac{d}{dx}\text{arch}(x) \triangleq -G^{-1}(x;1) = -\sqrt{2}\,\text{erf}^{-1}(2x-1) \text{ for } x \in [0,1]$$

(4.37)

Inserting (4.37) in (4.36) we get

$$\frac{d}{dx}\text{arch}^{-1}(x) = \frac{1}{\sqrt{2\ln\left(x\sqrt{2\pi}\right)}} \text{ for } x \in [0,1]$$

(4.38)

4.6   Appendix II

The derivative of $\text{arch}(x)$ in (4.37) can also be derived using the differentiation rule for inverse functions with $f^{-1}(x) = \text{arch}(x)$ and $(\text{arch}^{-1}(x))'$ as in (4.38):

$$\frac{d}{dx}\text{arch}(x) = -\sqrt{-\ln(2\pi) - 2\ln(\text{arch}(x))}$$

$$= -\sqrt{-\ln(2\pi) - 2\ln\left(\frac{\exp\left(-\left\{\text{erf}^{-1}(2x-1)\right\}^2\right)}{\sqrt{2\pi}}\right)}$$

$$= -\sqrt{-\ln(2\pi) + \ln(2\pi) + 2\left(-\left\{\text{erf}^{-1}(2x-1)\right\}^2\right)} \qquad (4.39)$$

$$= -\sqrt{2}\text{erf}^{-1}(2x-1)$$

$$\triangleq -G^{-1}(x;1)$$

This may be rewritten using (4.5) and (4.6) to give the integral of the inverse error-function as an alternative to using the method of Parker [25]:

$$\int_0^x G^{-1}(t;1)dt = -\text{arch}(x) \Leftrightarrow$$

$$\int_0^x \sqrt{2}\text{erf}^{-1}(2t-1)\,dt = \frac{-1}{\sqrt{2\pi}}\exp\left(-\left\{\text{erf}^{-1}(2x-1)\right\}^2\right) \Leftrightarrow \qquad (4.40)$$

$$\sqrt{2}\int_{-1}^{2x-1}\text{erf}^{-1}(t)\,d\left(\tfrac{1}{2}t\right) = \frac{-1}{\sqrt{2\pi}}\exp\left(-\left\{\text{erf}^{-1}(2x-1)\right\}^2\right)$$

Substitution of $2x-1$ by x yields

$$\int_{-1}^x \text{erf}^{-1}(t)\,dt = \frac{-1}{\sqrt{\pi}}\exp\left(-\text{erf}^{-1}(x)^2\right)$$

$$= \frac{-2}{\sqrt{2}}\frac{1}{\sqrt{2\pi}}\exp\left(-\left\{\text{erf}^{-1}(x)\right\}^2\right) \qquad (4.41)$$

$$= \frac{-2}{\sqrt{2}}\text{arch}\left(\frac{x+1}{2}\right)$$

Hence, the derivative of the arch function is the inverse cumulative Gaussian: $\text{arch}'(x) = -\sqrt{2}\text{erf}^{-1}(2x-1)$ scaled such that the intrinsic scale is normalized to $\sigma = 1$ as opposed to the intrinsic scale of the erf that is $\sigma = \sqrt{2}$.

4.7 References

[1] I.W.O. Serlie, F.M. Vos, R. Truyen, F.H. Post, L.J. van Vliet, Classifying CT Image Data into Material Fractions by a Scale and Rotation Invariant Edge Model, *accepted for publication in IEEE Transactions on Image Processing.*

[2] G. Schwarzband and N. Kiryati, The point spread function of spiral CT, *in: Phys. Med. Biol.*, vol. 50, 2005, pp. 5307–5322.

[3] G. Wang, M.W. Vannier, M.W. Skinner, M.G.P. Cavalcanti, and G.W. Harding, Spiral CT image deblurring for cochlear implantation, *in: IEEE. Trans. Med. Imag.*, vol 17, 1998, pp. 251–262.

[4] H. Bouma, A. Vilanova, L.J. van Vliet, and F.A. Gerritsen, Correction for the dislocation of curved surfaces caused by the PSF in 2D and 3D CT images, *in: IEEE Trans. Pat. Anal. Mach. Intell.* 27, 2005, pp. 1501–1507.

[5] Y. Zou, E.Y. Sidky and X. Pan, Partial volume and aliasing artefacts in helical cone-beam CT, *Phys. Med. Biol.*, vol. 49, 2004, pp. 2365–2375.

[6] I.W.O. Serlie, R. Truyen, J. Florie, F.H. Post, L.J. van Vliet, F.M. Vos, Computed cleansing for virtual colonoscopy using a three-material transition model, *MICCAI*, LNCS-2879, 2003, pp. 175–183.

[7] G. Kindlmann, and J.W. Durkin, Semi-Automatic Generation of Transfer Functions for Direct Volume Rendering, *in: Proc. IEEE Symp. Volume Visualization*, 1998, pp. 79–86.

[8] J. Kniss, G. Kindlmann, and C. Hansen, Multi-dimensional transfer functions for interactive vol-ume rendering, *in: IEEE Trans. Visual. Comput. Graph.*, vol 8, 2002, pp. 270–285.

[9] D.H. Laidlaw, K.W. Fleischer, and A.H. Barr, Partial-volume Bayesian classification of material mixtures in MR volume data using voxel histograms, *in: IEEE Trans. Med. Im.*, vol. 17, 1998, pp. 74–86.

[10] G. de Vries, P.W. Verbeek, and U. Stelwagen, Thickness measurement of CT- imaged objects, *in: Proc. ASCI 5th Annu. Conf. Advanced School for Computing and Imaging*, 1999, pp. 179–183.

[11] L.M.J. Florack, B.M. ter Haar Romeny, J.J. Koenderink, M.A. Viergever, Scale and the differential structure of images, *in: Image Vision Comp.*, vol. 10, 1992, pp. 376–388.

[12] P. Charbonnier, L. Blanc-Feraud, G. Aubert and M. Barlaud, Deterministic Edge-Preserving Regu-larization in Computed Imaging, *in: IEEE Trans. Im. Proc.*, Vol. 6, 1997, pp. 298–311.

[13] L. J. van Vliet, Grey-Scale Measurements in Multi-Dimensional Digitized Images, *in: Ph.D. thesis, Delft Univ. Press, Stevinweg 1, Delft, The Netherlands*, 1993. www.ist.tudelft. nl/qi: group pub.: PhD.

[14] O. Axelsson, V. Barker, Finite element solution of boundary value problems, *in: AP Inc.*, 1984.

[15] D. Ballard, Generalized Hough transform to detect arbitrary patterns, *in: IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13(2), 1981, pp. 111–122.

[16] B. Rieger, F.J. Timmermans, L.J. van Vliet, and P.W. Verbeek, On curvature estimation of iso-surfaces in 3D gray-value images and the computation of shape descriptors, *in: IEEE Trans. Patt. Anal. Mach. Intell.* , vol. 26, 2004, pp. 1088–1094.

[17] J.R. Mayo, K.P. Whittall, A.N. Leung, T.E. Hartman, C.S. Park, S.L. Primack, G.K. Chambers, M.K. Limkeman, T.L. Toth, S.H. Fox. Simulated dose reduction in conventional chest CT: validation study, *in: Rad*, vol. 202, 1997, pp. 453–457.

[18] P. Sereda, A.V. Bartroli, I.W.O. Serlie and F.A. Gerritsen, Visualization of boundaries in volumetric data sets using LH histograms, *in: IEEE Trans. Vis. Comput. Graph.*, vol. 12, 2006, pp. 208–218.

[19] P.W. Verbeek, H.A. Vrooman and L.J. van Vliet, Low level image processing by max-min. filters, *in: Signal Process.*, vol. 15, 1988, pp. 249–258.

[20] J. Mandel, J. Bond, J. Church, D. Snover, Reducing Mortality from Colon Cancer Control Study, *in: N. Engl. J. Med.*, vol. 328(19), 1993, pp. 1365–1371.

[21] D.J. Vining, D.W. Gelfand, R.E. Bechtold, E.S. Scharling, E.K. Grishaw and R.Y. Shifrin, Technical Feasibility of Colon Imaging with Helical CT and Virtual Reality, *in: AJR*, 1994, vol. 162.

[22] J.J. Dijkers, C. van Wijk, F.M. Vos, J. Florie, Y.C. Nio, H.W. Venema, R. Truyen and L.J. van Vliet, Segmentation and Size Measurement of Polyps in CT Colonography, *in: MICCAI*, LNCS-3749, 2005, pp. 712-719.

[23] F.M. Vos, R.E. van Gelder, I.W.O. Serlie, J. Florie, C.Y. Nio, A.S. Glas, F.H. Post, R. Truyen, F.A. Gerritsen, and J. Stoker, Three-dimensional display modes for CT colonography: conventional 3D virtual colonoscopy versus unfolded cube projection, *in: Rad.*, v. 228, 2003, pp. 878–885.

[24] P.J. Pickhardt, J.H. Choi. "Electronic cleansing and stool tagging in CT colonography: advantages and pitfalls with primary three-dimensional evaluation", *AJR*, 181, pp. 799-805, 2003.

[25] F.D. Parker, Integrals of Inverse Functions, *in: Amer. Math. Monthly*, v. 62, 1955, pp. 439–440.

Chapter Five

Electronic Cleansing for CT Colonography by a Rotation and Scale Invariant Three Material Model

A model-based method is presented to determine material fractions around three-material transitions such as T-junctions. It enables CT colonography to be used with contrast enhanced intraluminal remains. The method is applied to partial-volume voxels and results into patient specific and reproducible segmentations. The method is scale and rotation-invariant, even for data of anisotropic resolution and non-cubic voxels. First, the mean values of the pure materials are locally identified (air, tissue and tagged material). Second, scale-invariant Gaussian derivatives are computed and scaled to produce equal variance of the noise. Third, the model is used to map the Gaussian derivatives onto material fractions for each transition voxel. Projection of noisy data value and derivatives onto the model yields a robust alternative to the standard computed Gaussian derivatives at T-junctions. Surfaces at a constant material fraction are shown to provide an accurate and precise representation of the object boundary without artifacts near junctions. It advances CT colonography to become effective with tagged data and a 3-D display mode.

5.1   Introduction

CT colonography (virtual colonoscopy) is a non-invasive alternative to colonoscopy [1] that facilitates screening for polyps. The conventional procedure is to let the patient drink a laxative on the day before the examination to remove fecal remains from the colon. Immediately before CT imaging, the patient's abdomen is inflated with CO_2-enriched air. Virtual colonoscopy images are created using perspective volume rendering to inspect the colon-surface [2]. Usually, a threshold around −650 HU is applied to all voxel values to segment the colon wall. Unfortunately, the patient preparation is never 'complete'. Remains that have a tissue-equivalent data value may occlude or mimic polyps and therefore hamper detection. Tagging the intra-luminal remains with an oral contrast agent was introduced to discriminate between tissue and remains (Figure 5.1a) [3]. The sequence of operations applied to the 3-D CT data that replace the (partial-volume) contributions of tagged remains by air is referred to as 'electronic cleansing'.

Unfortunately, positions where the 'fluid'-surface touches the colon wall complicate electronic cleansing [4]. Although only a small part of the surface maybe 'obscured' by these three-material transitions, segmentation artifacts at these locations may delay or impede a proper diagnosis (by mimicking or obscuring polyps) and not stimulate confidence by the reader (Figure 5.1a-d).

The main problem is that the point-spread function (PSF) causes contributions of several materials to be mixed into a sample. The latter effect is commonly referred to as the partial-volume effect (PVE). In general, the accurate delineation of an object having blurred edges is not a trivial task and at locations where three materials are connected, it may seem impracticable. If the blurring is not isotropic and noise is added, matters are even further complicated.

In this paper, we present a novel method that implements the inverse operation: obtaining the material fractions starting from the data value. We extend previous work [5]. by deriving an analytical expression that describes Gaussian derivatives (not as a function of space, but) as a function of the three material fractions around T-junctions. We incorporate invariance to the anisotropy of the noise (in the space spanned by Gaussian derivatives) and to the anisotropic scale of the data (in image space). Projection of noisy data value and derivatives onto the model yields a robust alternative to the standard computed Gaussian derivatives at T-junctions. The proposed method is applied to segmentation of the large bowel for CT colonography. However, it may be equally useful for other problems that involve three-material transitions.

5.1.1   Related work

Lakare et al. [6] create 23-dimensional feature vectors of local data values that are reduced to five dimensions by principal component analysis. Clustering takes place in the low
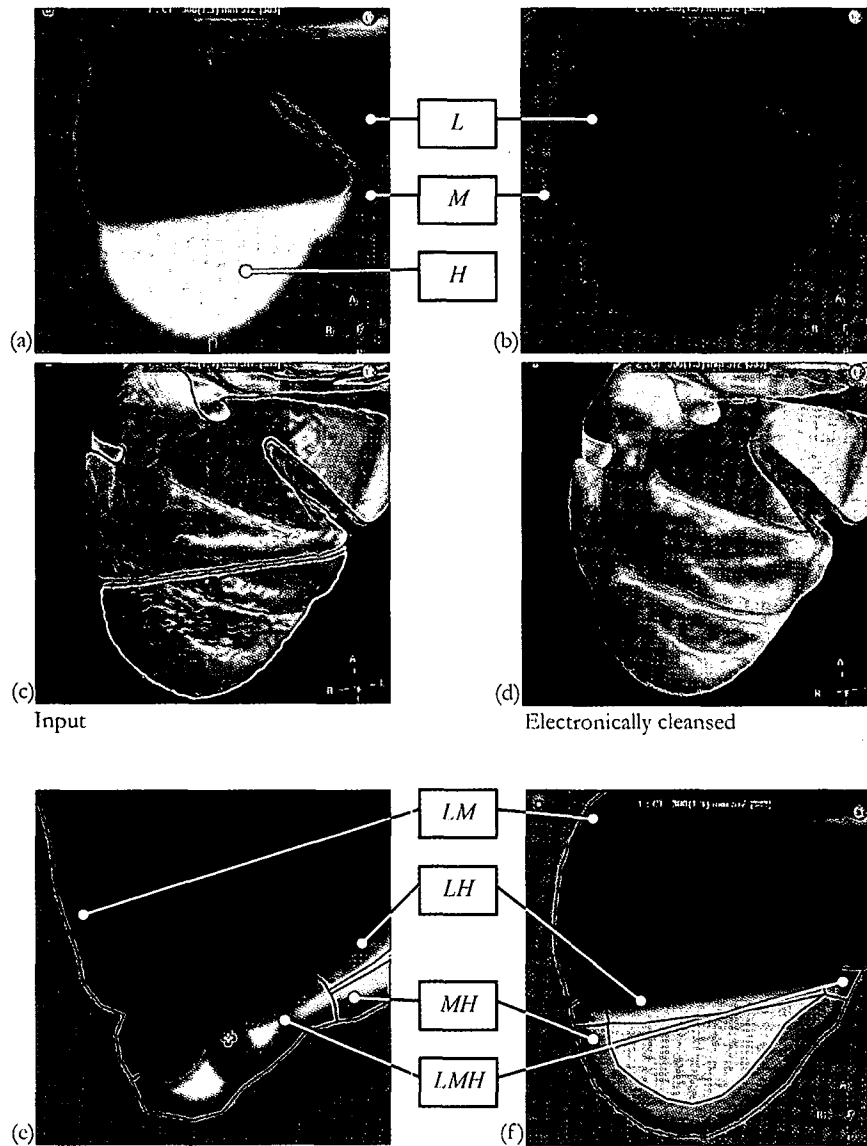
Figure 5.1 (a) Part of a slice from CT-data. *L*, *M* and *H* denote air, tissue and fluid respectively. (b) Electronic cleansing replaces *H* with *L*. (c) A transfer function based on thresholds leaves the fluid interface in (c) opaque: *the transition has a CT-value similar to tissue*. (d) If (b) were generated, a single threshold would suffice. (e,f) A significant part of the image corresponds to three-material PV-voxels (*LMH*), focused upon in this paper.

dimensional space using a vector similarity measure and a maximum number of classes. A threshold on the average intensity of each class classifies voxels to be tagged remains. The method is patient specific and deals with fluctuating contrast concentrations.

In later work, Lakare et al. [7] address a problem caused by partial-volume voxels bordering on the fluid-mask. The intensity profile is considered a unique property of each type of transition. Typical edge profiles that are present between all materials, e.g. air and tagged material, are identified by rigorously 'exploring' some 3-D CT data sets in a separate learning phase prior to the actual cleansing. During cleansing, for each edge voxel the profile is selected that best fits encountered intensity profiles. A transfer function is defined for each such profile in order to remove the partial-volume problem during rendering. The so-called segmentation ray method is fully automatic. A disadvantage may be that the segmentation rays are defined for a group of selected CT images and are as such not patient specific, which may reduce diagnostic accuracy. In addition, only two-material transitions are modeled, which may result in artifacts at positions where the fluid touches the colon-surface.

Zalis et al. [8] construct a binary subtraction mask to segment the opacified bowel content and address the partial-volume problem using a colon-surface reconstruction routine. Data values represent a distance measure derived from the subtraction mask. The surface reconstruction does not directly use the partial-volume values from the data so that the resulting transition itself is not patient specific.

Laidlaw et al. [9] model the local image structure by intensity histograms. A voxel is treated as a volume in which the local histogram is obtained by interpolation. The region that influences the measurement is determined by the interpolation kernel. A parameterized model of histograms is created for different object geometries. Subsequently, the parameterized histogram is fitted to the local voxel histogram. The parameters define the materials and geometry in the voxel. This method is used to separate two materials. The authors argue that three material transitions might be resolved by extending the dimensionality of the histogram (by including more features). Neither the image geometry around the junction nor variations in the scale of the data are modeled explicitly.

The method presented in this paper is particularly inspired by the work of Kindlmann [10], Kniss [11] and Laidlaw [9]. Kindlmann and Kniss visualize material transitions using partial volume voxels to reduce 'jaggy' edges observed with binary techniques. A histogram of the data value I and the gradient magnitude $|\nabla I| = \partial I/\partial w \triangleq I_w$ (w is in the gradient direction) is used. This yields a unique arch-shaped point cloud for all voxels near the border between two materials. An advantage of this method is that the edge region of a particular type of transition can be discriminated from other edge regions by selecting a region close to the top of the arch in the histogram. Several problems remain: (1) arches coincide near the basis, so that the method works best at positions close to the edge region; (2) the approach is not valid near a three-material T-junction; (3) anisotropic

resolution results in broad, overlapping point clouds that complicate the correct identification of an edge type; and (4) the method is restricted by local fluctuations in data value of tagged material that spread the point clouds as well.

Inspired by Laidlaw [9] and Kindleman [10], we created an analytical model of the previously mentioned arches that are scale as well as rotation-invariant (Chapter 4) in the presence of anisotropic resolution and non-cubic voxels. However, our two-material transition model needs to be improved to describe three-material transitions. These configurations occur in applications like CT colonography, in which the colon is filled with both air and tagged fluid that border on the colon tissue.

Several problems remain with the existing methods. First, the non-isotropic scale of the data as generated by many scanners is generally not accounted for. This may be solved by low-pass filtering the data to end up with an isotropic scale. However, this is at the expense of detail, which may affect the appearance and conspicuity of polyps. Second, only two-material transitions are considered, whereas more than two materials also occur in voxels. In CT colonography, for instance, at positions where the fluid touches the surface, samples contain partial volume fractions of three materials (Figure 5.1e-f). If a polyp resides at such a position, its shape and size may not be segmented properly. Third, current methods may require considerable user-interaction or, when parameters are set beforehand, may not be patient specific. Fourth, existing methods lack a suitable geometric model, which is a necessity to exploit the local correlations. Fifth, noise is typically not taken into account, which results in unwanted effects such as amplification of the noise during rendering.

5.1.2   Objective

Our objective is to improve CT colonography to become effective with tagged data using a 3-D display mode. Therefore, we aim to classify voxels into material fractions (air, tissue, tagged material). In particular, three-material transitions must be accurately solved. The method should be invariant to anisotropic resolution (PSF) and anisotropic sampling (non-cubic voxels). In addition, the method has to be robust against local variations in image intensity of a material (e.g. the intensity of tagged material may vary) and noise. Finally, the method should require minimal user interaction to enable reproducibility.

Segmentation based upon material fractions is less complicated than based upon image intensity. It will be shown that surfaces defined at a constant material fraction provide a way to represent the object boundary for quantitative image analysis and high quality visualizations.

5.2   Methods

5.2.1   Outline

In earlier work, we proposed a method to relate material fractions to derivative features near a two-material transition (Chapter 4). The latter reference implicitly yields a method to detect voxels that may participate in three-material transitions. The two-material model produces segmentation artifacts at such locations that may obscure polyps and generate false-positives. The method presented here solves this problem.

The novelty of our work is that we derive an analytical (forward) model for the relation between the underlying material fractions and the Gaussian derivatives at three-material transitions such as T-junctions. It is made scale-invariant to account for the space-variant, orientation dependent resolution of the data. Note that an anisotropic PSF yields an apparent scale of an edge that depends on the orientation of the edge. At T-junctions we measure a Gaussian derivative jet, apply scale-normalization to compensate for the effects of the anisotropic PSF and scale the result to make the noise isotropic. After this we can employ a Euclidean metric in computing the distance between observations and our model. Hence, we find the best match between the measured Gaussian derivative jet and the jet produced by our model. The position in model space yields the material fractions we were looking for. The measurements are taken in a local environment by sampling along a path that stays parallel to the gradient direction.

5.2.2   Material fractions

During image formation, a linear combination of materials is projected onto data value (Figure 5.2a). The method presented by us enables the inverse operation (Figure 5.2b): retrieving the material fraction from the local image data.

Assume a linear imaging model in which the image intensity I can be written as a linear combination of pure material intensities μ_m with corresponding material fractions t_m.

$$I\left(t_m, \mu_m\right) = \sum t_m \mu_m \ , \ \sum t_m = 1 \tag{5.1}$$

We restrict ourselves to three material fractions that correspond to homogeneous barycentric coordinates. Thus, each instance $\left(t_1, t_2, t_3\right)$ corresponds to a point P in a triangular domain $\left(A_1, A_2, A_3\right)$ that is projected onto a data value (Figure 5.2a).

Considering the CT colonography, let $\mu_1 = L$ represent the expected (low) intensity of air, $\mu_2 = M$ the expected (medium) intensity of tissue, and $\mu_3 = H$ the expected (high) intensity of tagged residual material.

Typically, most voxels are dominated by pure material (air, tissue and residue): the corners of the triangle. Fewer voxels participate in transitions between two materials: the edges of the triangle. The smallest number of voxels participates in three-material transitions:
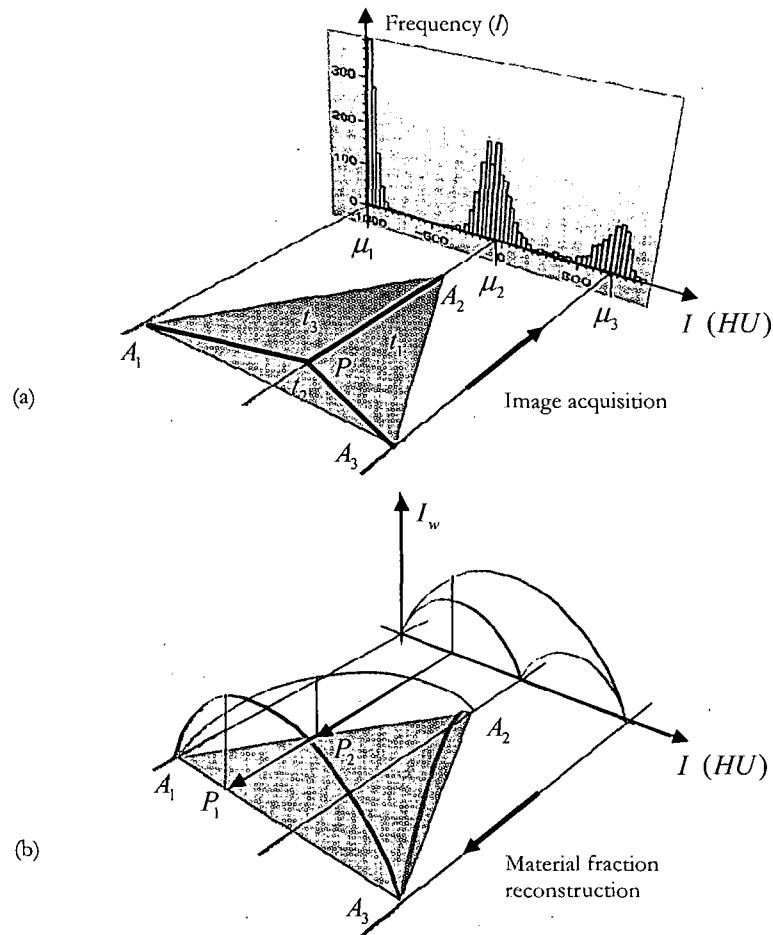
Figure 5.2   (a) Image acquisition combines contributions of different materials into the *partial volume*
CT value. We represent these contributions as barycentric positions in a triangle. (b) Our
method enables the inverse operation. The derivatives I_v and I_{vw} (not shown here) facilitate
this step.

the interior of the triangle. Still, the three-material transitions represent a significant
fraction of the voxels to be analyzed [12]. We will denote a transition type by a
combination of the letters L, M and H. Since voxels processed by the three-transition
model will border on two-material transitions, the model is required to contain a smooth
extension of a two-material transition model. Finally, trivial as it may seem, a model of
the transition between two materials should be an extension of a single-material model.
This is important in order not to produce artifacts or discontinuities.

5.2.3   Two-material transitions

Let us first consider two-material transitions. This section first summarizes methods presented in Chapter 4. The data value across a Gaussian blurred step edge and its first derivative are related by an arch-shaped curve. These arches yield the gradient magnitude as a function of the data value, c.q. a weighted sum of two materials. Therefore, the arches coincide with the edges of the triangle (Figure 5.2b). Deriving an analytical expression of the relation starts by modeling a two material transition to be a unit-step-function u (5.2) that is convolved with a Gaussian g (5.3) resulting in a cumulative Gaussian distribution G (5.4). The anisotropic space-variant point-spread function (PSF) resembles a spiral pasta [13], but it is often modeled by an anisotropic 3-D Gaussian. We found that the cumulative Gaussian is an excellent model to describe the CT values across a two-material transition (Chapter 2). Additionally, for a given direction, the edge-spread function showed to be approximately constant over the image and, therefore, has not to be re-estimated for every voxel (Chapter 2).

$$G(x;\sigma) = u(x) * g(x;\sigma) = \frac{1}{2} + \frac{1}{2}\text{erf}\left(\frac{x}{\sigma\sqrt{2}}\right) \tag{5.2}$$

with

$$u(x) = \begin{cases} 0 & x < 0 \\ 1 & x > 0 \end{cases} \tag{5.3}$$

$$g(x;\sigma) = \frac{1}{\sigma\sqrt{2\pi}}\exp\left(\frac{-x^2}{2\sigma^2}\right) \tag{5.4}$$

$$\text{erf}(x) = \frac{2}{\sqrt{\pi}}\int_0^x \exp\left(-t^2\right)dt \tag{5.5}$$

Extension of this model to edge transitions in 3-D is straightforward. Let w represent the gradient direction, $v = w\perp$ the basis of the isophote surface, and σ_w the apparent scale of the (assumed) Gaussian function along w. Notice that a description of transitions in gauge coordinates is by definition both rotation- and translation-invariant.

In general, the expected data values at opposite sides of the transition are L and H such that the intensity is defined by $I(w;\sigma_w) \triangleq (H-L)G(w;\sigma_w) + L$ and the gradient magnitude is defined by $I_w(w;\sigma_w) \triangleq (H-L)g(w;\sigma_w)$ (first derivative in gradient direction).

The arch-function describes the scale-invariant gradient magnitude $\sigma_w I_w$ as a function of I at scale σ_w:

$$\sigma_w I_w = (H - L)\,\text{arch}\left(\frac{I - L}{H - L}\right) \tag{5.6}$$

with

$$\text{arch}(x) \triangleq \sigma_w g\left(G^{-1}(x; \sigma_w); \sigma_w\right)$$
$$= \frac{1}{\sqrt{2\pi}}\exp\left(-\left\{\text{erf}^{-1}(2x-1)\right\}^2\right) \quad \text{for } x \in [0,1] \tag{5.7}$$

Notice that the function $\text{arch}(x)$ does *not* depend on a scale parameter. Therefore, it describes scale-invariant measurements. Analogously, a generalized arch function was derived that relates higher order derivatives as a function of data value (Chapter 4).

Consider three measurements $(I, \sigma_w I_w, \sigma_w^2 I_{ww})$ acquired by Gaussian derivative operators (Figure 5.3a) in the edge region around materials L and H. The measurements are projected onto the (noise-free) generalized arch functions of order 1 and 2 for given L and H (Figure 5.3b). Orthogonal projection is permitted after scaling the measurement space $\{I, \sigma_w I_w, \sigma_w^2 I_{ww}\}$ by respectively $\{\theta_0, \theta_1, \theta_2\}$ in such a way that the standard deviation of the noise on $\theta_0 I$, $\theta_1 \sigma_w I_w$ and $\theta_2 \sigma_w^2 I_{ww}$ is identical. Without loss of generality we may set $\theta_0 = 1$. The projection yields filtered values for I, I_w and I_{ww} (Figure 5.3c). The filter result for I (given L and H) yields the unique two-material fractions.

We will extend the arch-functions to three-material T-junctions to be a function of the barycentric coordinates (Figure 5.3d). We find the best match between the scale- and noise-normalized measurements and the three-material model. The best match will yield the underlying material fractions.
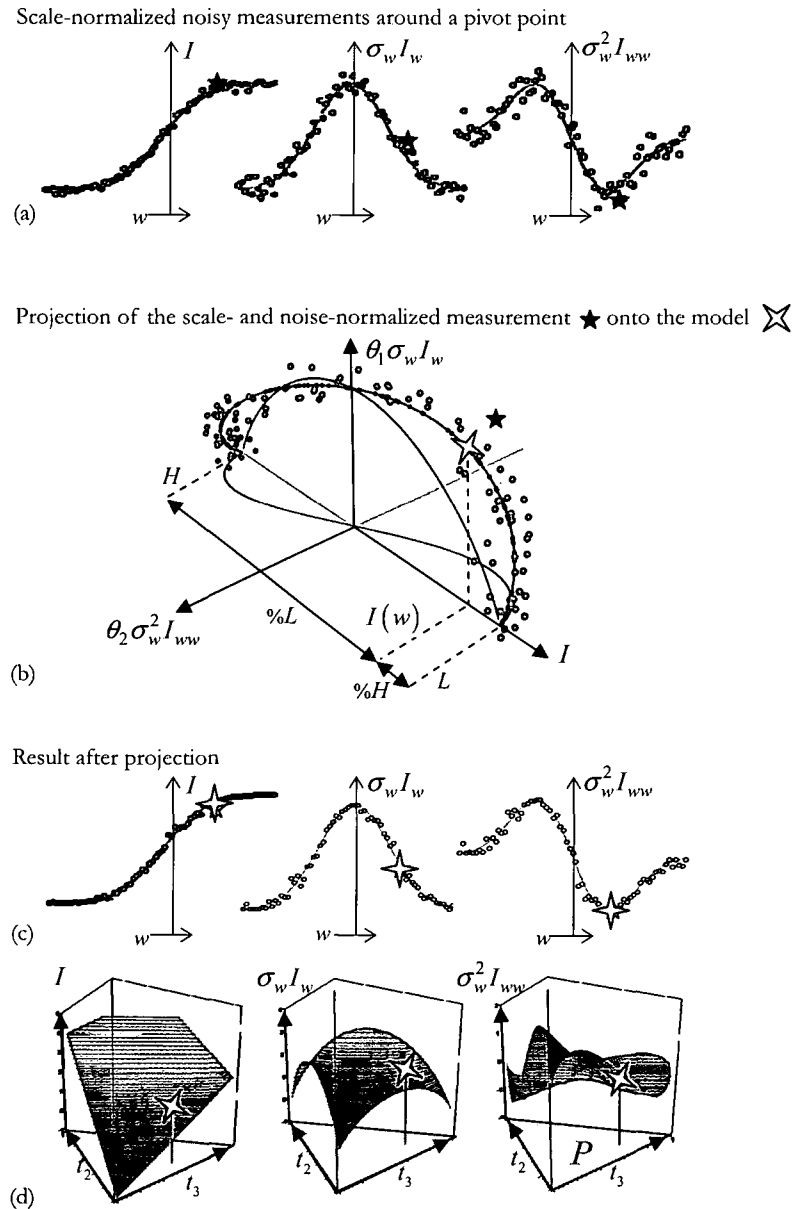
Scale-normalized noisy measurements around a pivot point

(a)

Projection of the scale- and noise-normalized measurement ★ onto the model ✕

(b)

Result after projection

(c)

(d)

Figure 5.3   (a) A Set of measurements along a path in gradient direction. The measurement under investigation is marked by '*' (b) Projection onto the arch model. (c) The small deviations from the solid lines line correspond to the error in estimated arch position. (d) Derivatives as a function of barycentric coordinates near a three-material transition. This data was generated using $L = -0.5$, $M = 0$, $H = 0.5$, noise STD $= 0.05$, input $\sigma = 1.7$ and operator $\sigma = 1.4$.
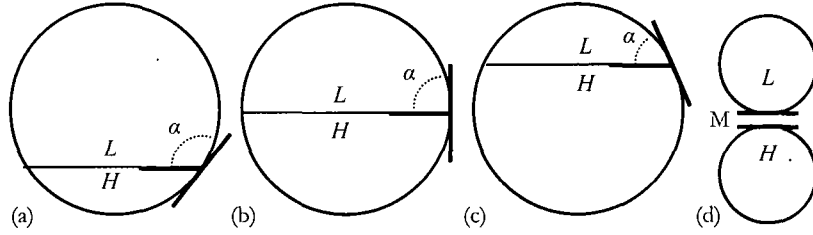
Figure 5.4   (a-c) The angle α at which the fluid touches the colon varies. (d) The thin layer between
two neighboring colon segments forms a three-material transition as well. One typical T-
junction (b) is selected to explain our method. Other angles are analyzed using the same
paradigm.

5.2.4   Three-material transitions

The LH-surface meets the colon surface at a variety of angles (Figure 5.4a-c) to form a
three-material transition. We will initially use a T-junction with $\alpha = 90$ deg to illustrate
our method. Other angles are analyzed by the same paradigm. Let the junction-model v
be the intersection of three edges (Figure 5.5a):

$$v(x,y,z) = \begin{cases} L & x < 0, y \geq 0 \\ M & x < 0, y \geq 0 \\ H & x \geq 0 \end{cases}$$

$$= L + (H - L)u(-x)u(y) + (M - L)u(x) \tag{5.8}$$

Furthermore, let V represent the junction-model after convolution with a 3-D Gaussian
(Figure 5.5b). We assume that σ_x, σ_y, and σ_z denote respectively the effective scale of
the Gaussian operators in x, y, and z-direction:

$$V(x,y,z) =$$
$$L + (H - L)G(-x,\sigma_x)G(y,\sigma_y) + (M - L)G(x,\sigma_x) \tag{5.9}$$

V_w describes the gradient-magnitude at the junction (Figure 5c):

$$V_w(x,y,z) \triangleq \|\nabla V(x,y,z)\| \tag{5.10}$$

with

$$\nabla V(x,y,z) =$$
$$\begin{pmatrix} -(H-L)g(-x,\sigma_x)G(y,\sigma_y) + (M-L)g(x,\sigma_x) \\ (H-L)G(-x,\sigma_x)g(y,\sigma_y) \\ 0 \end{pmatrix} \tag{5.11}$$
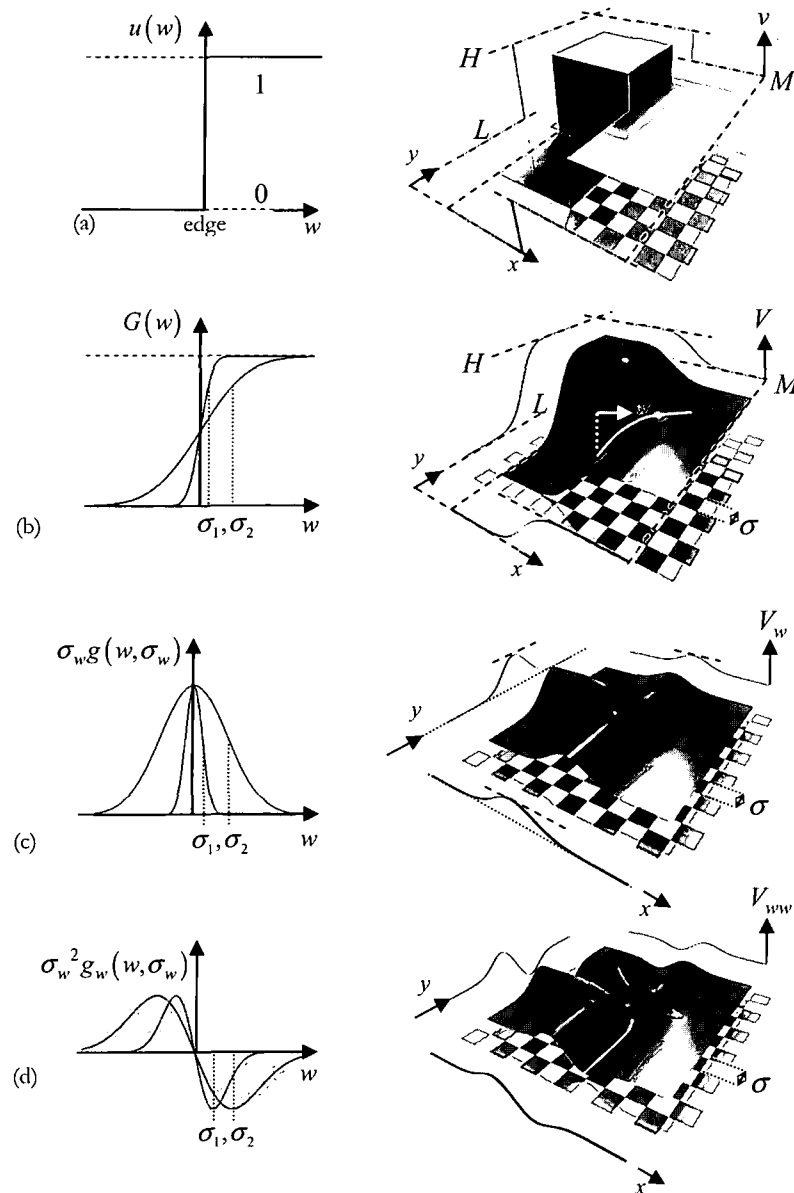
Figure 5.5   (a) A three-material transition is modeled using unit step-edges scaled with L, M and H. (b) The step edge values are low-pass filtered. (c) Gradient magnitude as a function of position. Only CT-measurements are scale-normalized; hence, the model-values are defined at $\sigma = 1$ for convenience. (d) Second directional derivative (SDGD) as a function of position.

Consider an anisotropic PSF. The anisotropy causes the gradient magnitude to vary as a function of edge orientation: an effect that complicates the procedure. It is preferred that the gradient magnitude exactly *at* the edge only varies as a function of the type of transition. One might impose additional blurring to obtain isotropic resolution. However, this goes at the expense of resolution. Scaling the gradient magnitude by σ_w yields the scale-invariant gradient magnitude. This results in identical amplitudes of the gradient magnitude at the edge for measurements of different scales. However, in the following paragraphs we will show that in addition to the scale-normalization, the derivatives have to be scaled to have isotropic noise in the space spanned by the Gaussian derivatives.

Let V_{ww} correspond to the second derivative in gradient direction of the junction-model defined in equation (5.8) (Figure 5.5d). Due to the translation-invariance along z we get:

$$V_{ww}(x,y,z) = \frac{V_x^2 V_{xx} + V_y^2 V_{yy} + 2V_x V_y V_{xy}}{V_x^2 + V_y^2} \tag{5.12}$$

Similarly to scaling of the first order derivative, multiplication with σ_w^2 yields the scale-invariant second derivative measured at the junction (Figure 5.5d). Notice that although the location of maximum gradient magnitude of blurred curved edges is displaced inwards [14], the corner point remains at its exact location [15].

5.2.5 Parachute function

The *parachute* function yields the Gaussian derivatives as a function of material fractions (in barycentric coordinates and not as a function of space). The function is named '*parachute*', because the surface created by the gradient magnitude as a function of material fractions reminds us of a parachute (Figure 5.6a). We will first derive an analytical expression of the function and then show that the second derivative is required to uniquely identify the material fraction in a three-material sample around a T-junction. An analytical expression is used to map material fractions onto a position in image space (x,y), after which the derivatives can be computed. Only two barycentic coordinates suffice as can be observed from (5.1).

First, notice that t_2 (the fraction of material M) does not depend on y (Figure 5.5a). Consequently, the inverse cumulative Gaussian relates x to t_2:

$$x = G^{-1}(t_2, \sigma_x) \tag{5.13}$$

The inverse cumulative Gaussian function G^{-1} is obtained by inserting (5.2) in $G(G^{-1}(x)) = x$ as explained in Chapter 4. Solving for $G^{-1}(x)$ yields:

$$G^{-1}(x;\sigma) = \sigma\sqrt{2}\,\mathrm{erf}^{-1}(2x-1) \text{ for } x \in [0,1] \tag{5.14}$$
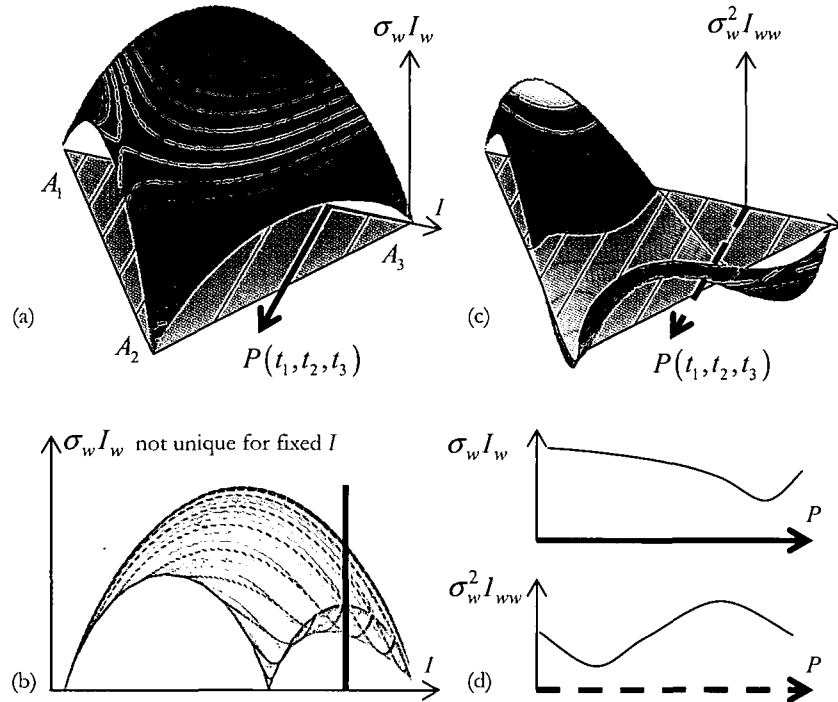
Figure 5.6   (a) Displaying $\sigma_w I_w$ as function of barycentric position reveals a parachute-shaped surface. A problem is that probing this surface along fixed intensity (dark line) results in equal altitude for two positions; *the surface folds back.* (b) Projecting the parachute upon $\left(I, \sigma_w I_w\right)$ reveals the overlap-area. (c) Including the second derivative (b) in gradient direction (dashed line) solves this problem. (d) This can be seen as equal magnitudes have a different second derivative in gradient direction.

Secondly, t_3 is regarded. Considering that

$$t_3 = G\left(-x, \sigma_x\right) G\left(y, \sigma_y\right) \tag{5.15}$$

and

$$t_2 = G\left(x, \sigma_x\right) \quad \Leftrightarrow \quad 1 - t_2 = G\left(-x, \sigma_x\right) \tag{5.16}$$

equation (5.15) is rewritten into

$$t_3 = \left(1 - t_2\right) G\left(y, \sigma_y\right) \Leftrightarrow G\left(y, \sigma_y\right) = \frac{t_3}{\left(1 - t_2\right)} \tag{5.17}$$

from which we derive

$$y = G^{-1}\left(\frac{t_3}{1-t_2}, \sigma_y\right)$$
(5.18)

Third, the generalized parachute-function describes the scale-invariant n^{th} derivative along the gradient direction as a function of material fraction by inserting equations (5.13) and (5.18) in either equation (5.9), (5.10) or (5.12) respectively for $n = 0$, 1 or 2:

$$parachute(0, t_1, t_2, t_3) = V(x, y, z)$$

$$parachute(1, t_1, t_2, t_3) = \sigma_w V_w(x, y, z)$$
(5.19)

$$parachute(2, t_1, t_2, t_3) = \sigma_w^2 V_{ww}(x, y, z)$$

with x and y functions of the material fractions as presented by equations (5.13) and (5.18). Note that the orientation of the T-junction was chosen in such a way that all z-dependencies have vanished. Analogous to (5.19), higher-order derivatives may be computed as well. Figure 5.6 a illustrates the first-order *parachute*-function, showing V_w as a function of barycentric coordinates $(1 - t_2 - t_3, t_2, t_3)$.

Properties of the first-order *parachute*-function are:

- the corners of the triangle (nodes) denote pure materials with all partial derivatives equal to zero;

- the edges of the triangle denote mixtures of two materials, the scale-invariant gradient magnitude is given by the arch-function (Chapter 4);

- the interior of the triangle denote three-material mixtures;

- the model of scale-invariant Gaussian derivatives as a function of material fractions around T-junctions;

- the measurement pair $(I, \sigma_w I_w)$ leads to ambiguity in material fractions for certain combinations. The second derivative is unique here and solves the confusion (Figure. 5.6);

- the parachute-model is scale-invariant. Additional blurring of the data only lets measurements travel over the parachute surface (towards the center).

Let us bear in mind that our objective is to classify samples into material fractions, based upon measurements from the Gaussian derivative jet.

5.2.6   Noise isotropy

Typically, a measurement triplet such as $\left(I, \sigma_w I_w, \sigma_w{}^2 I_{ww}\right)$ is contaminated with (filtered) noise in all three channels. Previous work disregards this aspect of the noise [5]. An approximation of the noise-free signal values can be obtained by finding the 'closest' point on the model. The distance metric to be used depends on the covariance matrix of the noise. The measurements are obtained by Gaussian derivative operators and may display different *variances*. The Euclidean metric can be used if the noise on the features is isotropic. In that case one can use an orthogonal projection of the observed measurements onto the model to find the 'best' guess for the signal features. Since the measurements are produced by filtering the acquired anisotropic volumetric data, we need to find the relation between the noise before and after filtering. This relation can be obtained by comparing the variances of the signal before σ_{ni}^2 and after σ_{no}^2 filtering, i.e. convolution with a D-order Gaussian derivative with scale σ_{op} in N dimensions [16]:

$$\frac{\sigma_{no}^2}{\sigma_{ni}^2} = \frac{(2D)!}{\sigma_{op}^{N+2D} \pi^{N/2} D! 2^D 2^{N+D}} \tag{5.20}$$

The result is used to predict the variances of noise in $\left(I, \sigma_w I_w, \sigma_w{}^2 I_{ww}\right)$ and to compensate for the differences. The scaling of the first axis remains unaltered, i.e. $\theta_0 = 1$. The recipe to compute the scaling parameter θ_1 for the first derivative is described in Chapter 4. The method takes the anisotropic resolution of the data into account. Analogously, a second scaling parameter θ_2 for the second derivative is calculated. Therefore, the scaled triplet $\left(I, \theta_1 \sigma_w I_w, \theta_2 \sigma_w{}^2 I_{ww}\right)$ displays isotropic noise. Projecting the noisy measurements $\left(I, \sigma_w I_w, \sigma_w{}^2 I_{ww}\right)$ onto the parachute model yields a robust alternative to the standard computed Gaussian derivatives across T-junctions.

5.2.7   Mapping image measurements onto material fractions

Our algorithm for mapping the image-based measurements onto material fractions involves the following steps.

1. Detect three-material transitions by thresholding the gradient magnitude of the L,H-channels, which are produced by our two-material edge model (Chapter 4). This step can be skipped if one wishes to apply the three-material model to all voxels.

2. Estimate the local values L, M and H of pure materials (Chapter 4).

3. Compute the Gaussian derivative jet up to the second order in the gradient direction. Optionally we compute the same triplet of measurements at points along a trace that follows the gradient both uphill as well as downhill.

4. Apply scale- and noise-normalization to all measurements.

5. Find the 'closest' point (or trace) on the corresponding parachute models for the given values of *L*, *M*, and *H*. This yields the noise-filtered values of the measurements.

6. Assign the (barycentric) coordinates of the 'closest' point found to the material fractions.

A single measurement triplet may be misclassified due to noise. This happens in particularly at the basis of the parachute, i.e. near the corners of the triangle at positions in which one material dominates (Figure 5.7 a). To solve this problem, a neighborhood of samples along w is used. Let $\{I, \theta_1 \sigma_w I_w, \theta_2 \sigma_w^2 I_{ww}\}(p)$ with $p = 1, ..., P$ be a set of P measurements along the orientation of the gradient, with q indicating the measurement at the starting voxel. Note that the partial derivatives are normalized by multiplication with $\theta_1 \sigma_w$ and $\theta_2 \sigma_w^2$. First, we need to find the trajectory at the parachute that yields the shortest Euclidean distance to the set of measurements. Second, the measurement is projected onto the selected 1-D parachute trajectory and not onto the 2-D surface. The trajectory that best matches the measurements is selected by minimizing the summed square residuals:

$$\mathbf{t} = \begin{pmatrix} t_{q,1} & t_{q,2} & t_{q,3} \end{pmatrix}^T =$$

$$\underset{t_{q,1}, t_{q,2}, t_{q,3}}{\arg \min} \sum_{p=1}^{P} \sum_{n=0}^{2} \theta_n^2 \left\| \sigma_w^n I^{(n)}(p) - \text{parachute}(n, t_{p,1}, t_{p,2}, t_{p,3}) \right\|^2 \qquad (5.21)$$

with $I^{(n)}(p)$ the n^{th} Gaussian derivative in gradient direction at point p along the trace, $t_{p,m}$ the material fraction of material m at point p at the trajectory at the parachute, and x and y the functions of t_2 and t_3 as given by equations (5.13) and (5.18) with $t_1 = 1 - t_2 - t_3$.

We implemented this minimization problem in a numerical manner as follows. Figure 5.7b shows the streamlines generated by uniform sampling of the T-junction (Figure 5.4b) followed by tracing the gradient uphill and downhill for a number of steps. Each such line corresponds to a curve in barycentric coordinate space via (5.15) and (5.16): Figure 5.7c. Notice that the lines cluster along the edges of the triangle, reflecting that the samples are rapidly dominated by two-material transitions while moving from the center of the T-junction. Alternatively, in Figure 5.7d and Figure 5.7e, similar lines are created by maintaining a minimum distance in barycentric coordinate space. The latter procedure yields a much more homogeneous distribution of the streamlines. Likewise, it allows us to control the resolution in material fraction space (Figure 5.7).

The latter approach was taken to generate a 'code book' of traces. Each entry consists of model-computed triplets along a trace of length P, with $P = 1, 3,$ or 5. Subsequently, we find the best match between image features (our measurements) and the entries in the

code book (representation of the model). The material fractions that are associated with that entry yield the desired material fractions.

5.2.8   Generalization and visualization

The concept introduced above is generalized by creating separate 'code books' for a variety of angles at which the fluid touches the colon (as in Figure 5.5). In the subsequent section, we will go into aspects such as how many angles are required and how to cope with more than one angle.

The 'code books' for other angles can be generated in a similar fashion. We have chosen to generate these code books numerically. A representation of the selected geometry is generated and the Gaussian derivative jet is computed up to order two for single positions and traces as mentioned earlier. For each sample (or trace), we know the true underlying material fractions by computing the overlap between the Gaussian PSF and the model. We store the set of traces that fill the barycentric space in the code book.

The exact, subvoxel localization of an edge usually depends on the type of transition. The concept of material fractions allows a straightforward segmentation of a material that borders on two or more other materials. Having a material fraction volume allows one to delineate the object via a single iso-surface (e.g. a material fraction threshold 0.5). Currently available rendering engines can be used after mapping the material fraction onto image values.

5.3   Results

We will demonstrate the efficiency of electronic cleansing by application of the presented three-material model to CT-data from the Walter Reed Army Medical Center. The National Institute of health offers a CT colonography image database complete with associated colonoscopy findings (www.vcscreen.com).

Results are presented along the steps involved in electronic cleansing. First, the colon-volume is segmented to provide a region of interest (ROI) (5.3.1). Second, the parameters used by the three-material model are automatically determined from the data (5.3.2). Third, the method is applied to synthetic image data generated using a PSF and noise level comparable to real data. This is done to assess the influence of trace length and model angle (paragraph 5.3.3). Fourth, electronic cleansing based on the three-material model is applied to the CT data and results are shown using the unfolded cube display mode [2][17] (paragraph 5.3.4).
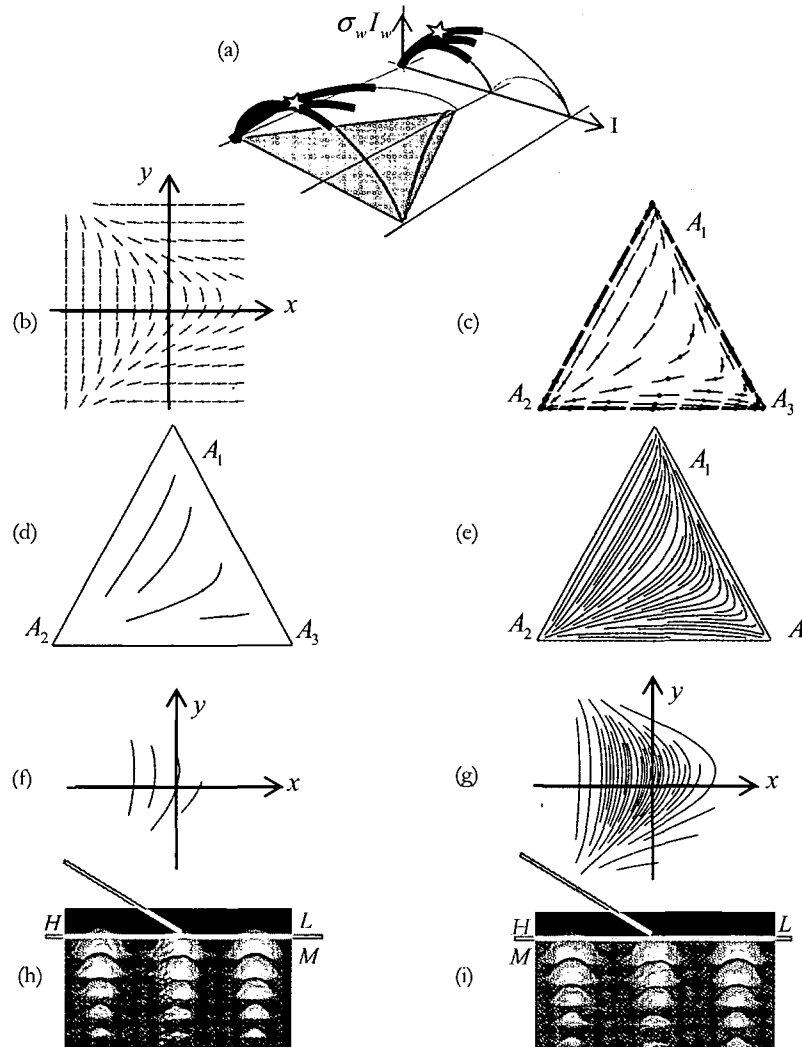
Figure 5.7   (a) Specifically near the basis of the parachute a measurement may be easily misinterpreted. Streamlines in *w*-direction are used to efficiently sample the model and properly project noisy data obtained along *w* onto the model. (b) Streamlines in *w*-direction, starting from a grid of points on the T-junction *in image space* and (c) *in barycentric space*. Over-sampling occurs in barycentric space, which is evident by the high density of lines near the edges of the triangle. (d)(e) Streamlines along *w*-direction projected into *barycentric space* until the distance between the lines is below a certain minimum: 0.2 respectively 0.05 (distances in material fraction). (f)(g) The corresponding streamlines in image space. (h)(i) Typical artifacts from undersampling disappear when changing sampling from 0.2 to 0.05 (distances in material fraction). Artificially generated polyps are used with sizes 3,5,7,9,11 mm and $\sigma = 1$ mm. The left part of middle polyps connects to fluid (*H*) and right part connects to air (*L*).

5.3.1   ROI identification

The colon volume is automatically segmented [18]. Clearly, an incomplete segmentation causes surface-parts to be missed. Figure 5.8 a,b demonstrate the amount of colon volume added by the electronic cleansing for a single case. In previous work, we demonstrated that in the set of 20 WRAMC patients on average 20% of the colon volume consisted of tagged material (Chapter 6). Moreover, the mean surface area obscured by tagged material was 23%. A more detailed analysis revealed that 7% of the surface area was identified to participate in a three-material junction.

5.3.2   Electronic cleansing parameters

Parameters of the parachute model are the mean intensities of the three materials near the junction (L, M and H), the (anisotropic) scale of the input data and the scale of the operator to calculate the Gaussian derivative jet.

The intensity of tagged material specifically fluctuates, e.g. due to inhomogeneous mixing of the contrast material. To compensate for such variations, the expected material intensities are estimated locally. Figure 8c shows for three patients how the intensity varies along the path through the colon. The graphs were generated by determining the average intensity of tagged material at the colon surface perpendicular to the central path.

The 3-D PSF of the CT scanner may be known in advance or, otherwise, measured from the data. We have found that for a given direction, the edge-spread function (ESF) is approximately constant over the image and, that the cumulative Gaussian is an excellent model to describe the CT values across a two-material transition (Chapter 2). We measured the edge spread by fitting this function to edge profiles at the colon surface (thereby assuming an anisotropic, Gaussian PSF). It was found that the in-plane resolution of the WRAMC CT-data was $\sigma_{xy} \approx 0.56$ mm and the slice-resolution was $\sigma_z \approx 0.82$ mm. The STD on the estimated scale was $STD(\sigma_{xyz}) \approx 0.03$ mm.

The local values L, M and H of pure materials are estimated using the image-data to define the parachute model that describes the local image structure [5]. However, the angle at which the fluid surface borders on the colon-surface may vary. The influence of variation in model angle is studied in 5.3.3.
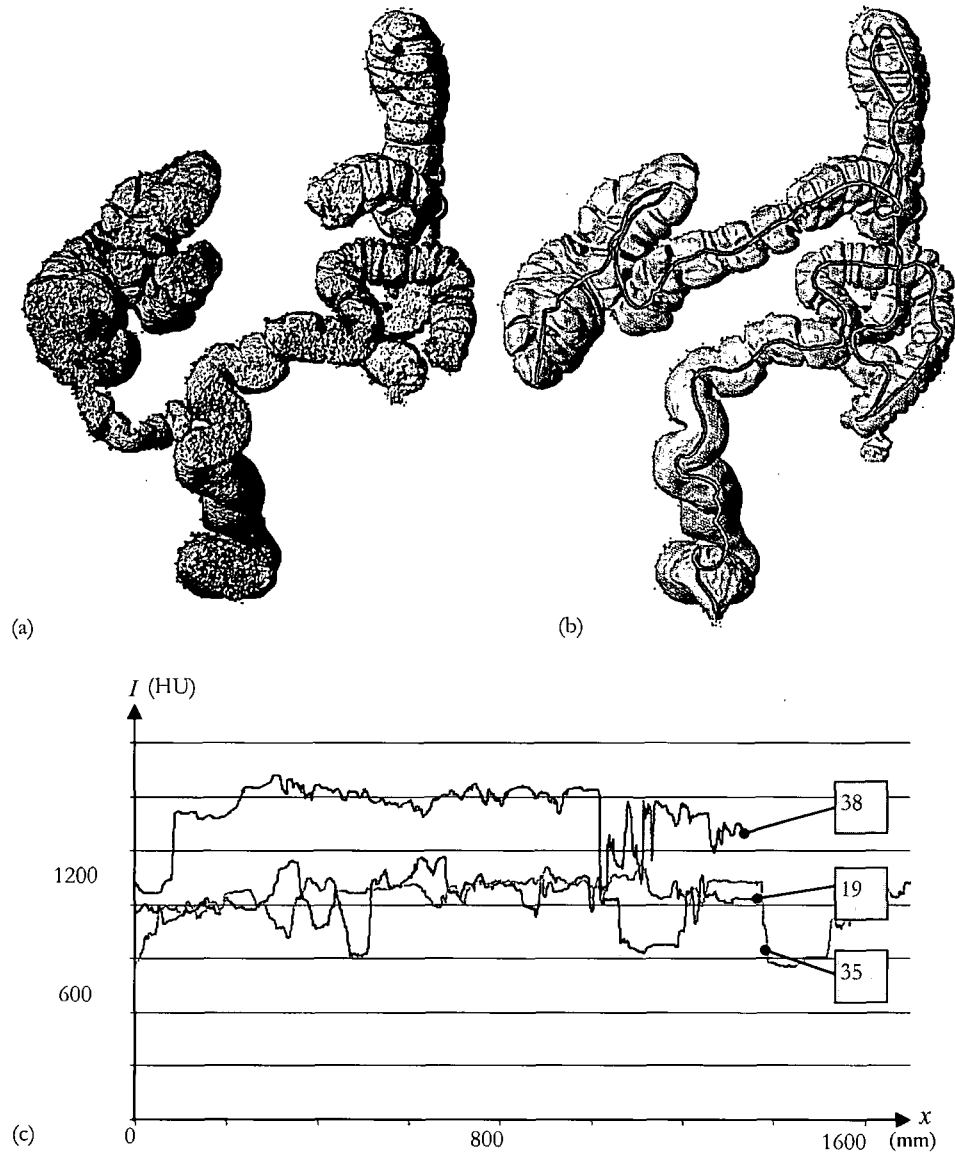
(a) (b)

(c)

Figure 5.8   (a),(b) A typical colon volume is given (a) before and (b) after electronic cleansing. (c) The
expected value of residue along a path through the colon as a function of path-position
for three different WRAMC data-volumes referred to by exam-number.

5.3.3   Model evaluation: trace length and number of angles

The proposed procedure was applied to synthetic low-pass filtered T-junctions in which the angle was varied from 0 to π and to which Gaussian noise was added (Figure 5.9a). The algorithm used models of three different angles (45, 90, and 135 degrees) and three sizes of the trace length (1, 3, and 5 points separated by σ_w). After electronic cleansing, the largest artifact was determined. This artifact was identified by the largest deviation from the true surface (ε_1, indicated in black). Additionally, its maximum diameter (ε_2, indicated in grey) was determined by the full width halfway the bulge in the edge profile, which is the standard way in colonography [19]. Figure 5.9b-d gives the size of the largest artifact (in pixels) after cleansing using each model (45, 90, and 135 degrees) as a function of the angle in the synthetic T-junction separately and combined by using the best fit to one of the models. The standard deviations on the graphs were obtained by repeating the procedure with different noise realizations. Applying the three models simultaneously and letting the best match win substantially reduces the size of the artifacts found (Figure 5.9e). Figure 5.9f demonstrates that indeed the model is applied of which the angle is closest to the angle present in the synthetic data.

5.3.4   Visualization

Figure 5.10a illustrates that simple thresholding is not sufficient to segment of the colon-surface. Intra-luminal remains obscure significant parts of the colon-surface since at the fluid-surface, tissue-equivalent CT-values are found due to the partial-volume effect. A two-material based model as proposed earlier leaves artifacts along the line where the liquid-air interface touches the colon surface (Figure 5.10b). The three-material method renders the colon surface free of artifacts (Figure 5.10c) with Figure 5.10d as a reference.

At thin layers of residue as depicted in Figure 5.11, cleansing also gives a satisfying result. The additional thin layer model as illustrated in Figure 5.4d prevents that a hole is formed through an extremely thin tissue layer that separates two loops of the colon.

The virtual colonoscopy interface is demonstrated in Figure 5.11. In addition, two polyps at T-junctions are depicted in Figure 5.12a-c without cleansing and Figure 5.12b-d with cleansing. We demonstrated that using electronically cleansed data rather than the original data yields improved observer confidence, time-efficient inspection and no difference in the conspicuity of cleansed and uncleansed polyps (Chapter 6).
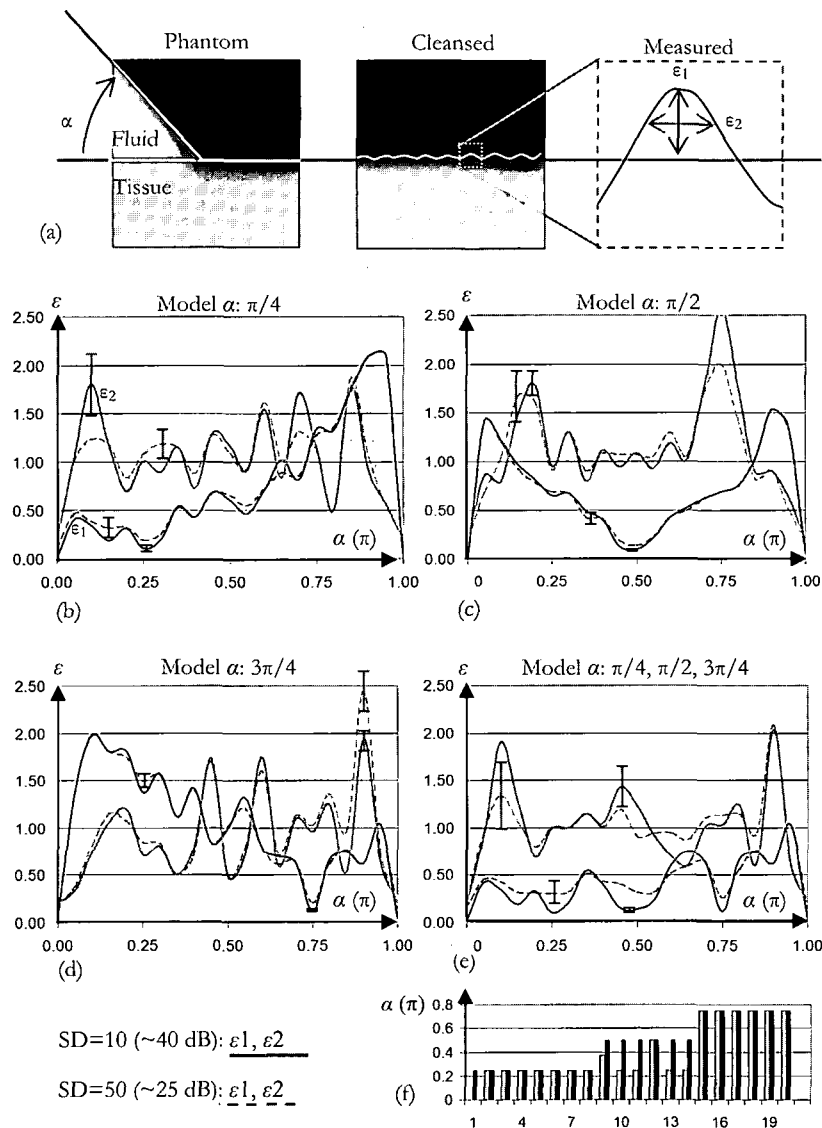
Figure 5.9   (a) An electronic phantom was created in which the angle of the tagged material varied.
              Gaussian distributed noise was added with SD=10 (best case) and SD=50 (worst case).
              The size of the largest artifact after cleansing was measured by way of $\varepsilon 1$, $\varepsilon 2$. (b) Result
              using the proposed method as a model of a single angle at 45 degrees is used, (c) similarly
              using a single angle at 90 degrees, and (d) 135 degrees. Vertically is the angle of the
              synthetic junction, horizontally the error expressed in pixels. (e) Result using the best
              match in code books of all three models. (f) Winning code book identified by its angle
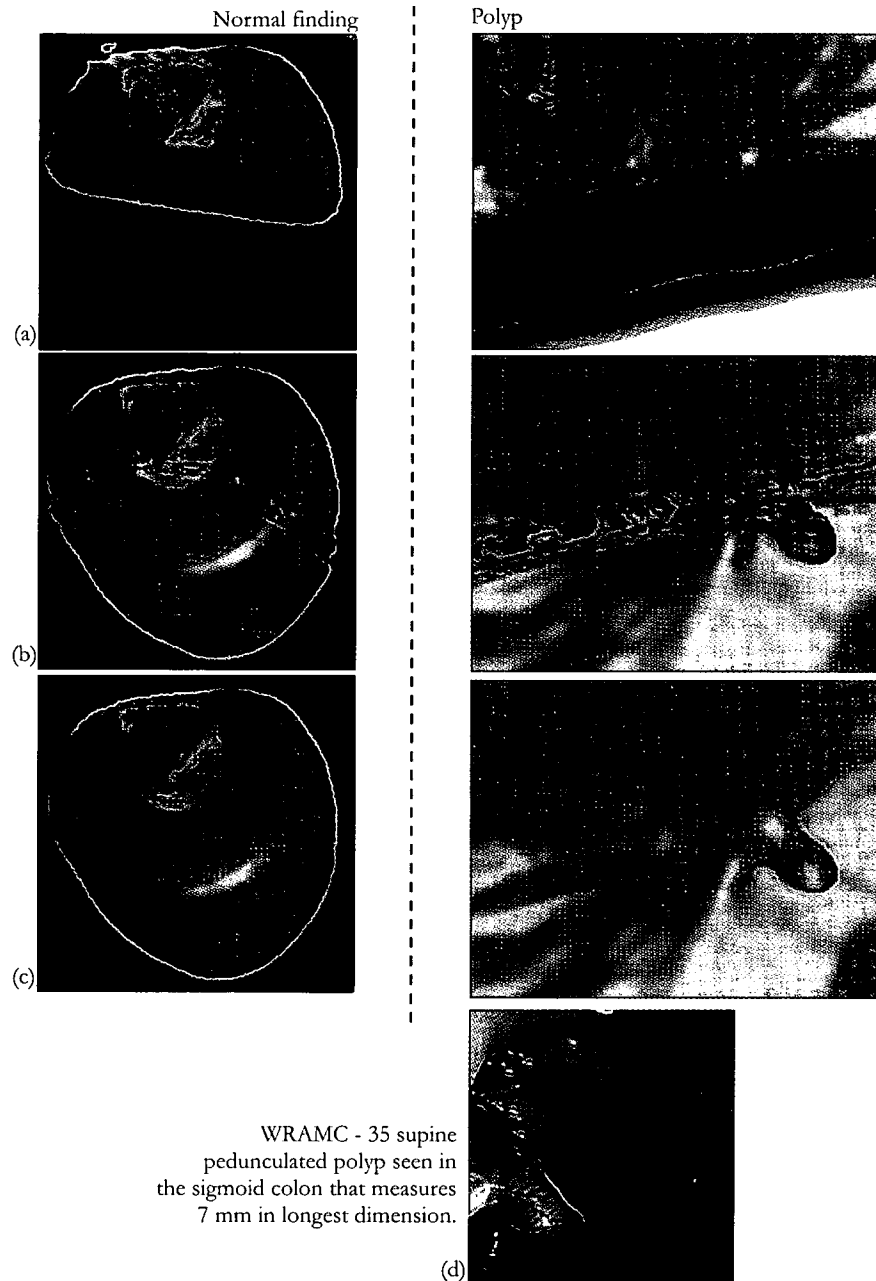   (vertically).

Figure 5.10   (a) before cleansing, (b) with two-material cleansing and (c) with three-material cleansing
(left) normal colon and (right) a 7 mm polyp together with its (d) optical colonoscopy
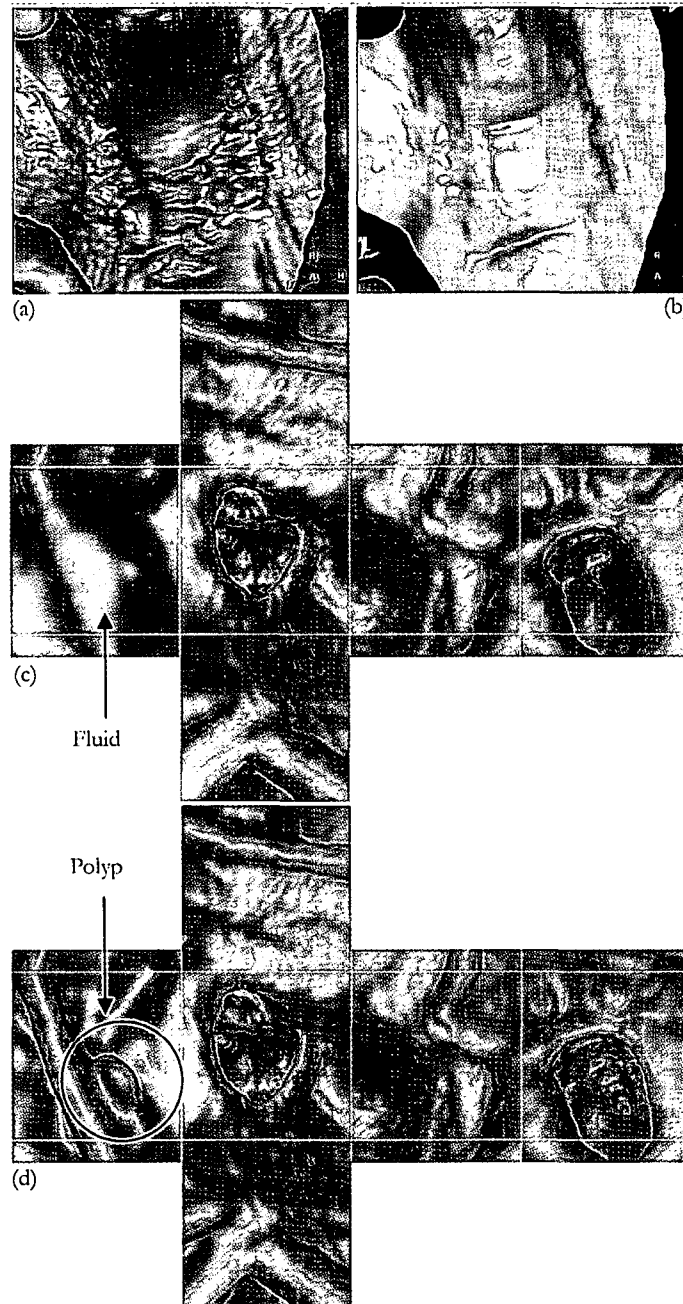image just at the moment of removing it.

Figure 5.11 (a) Virtual endoscopic view of part of the colon with a thin layer of fluid. (b) Fluid segmentation using the three-material transition-model reveals detailed colon-anatomy. (c) Virtual endoscopic enhanced 'unfolded' visualization before and (d) after electronic cleansing.
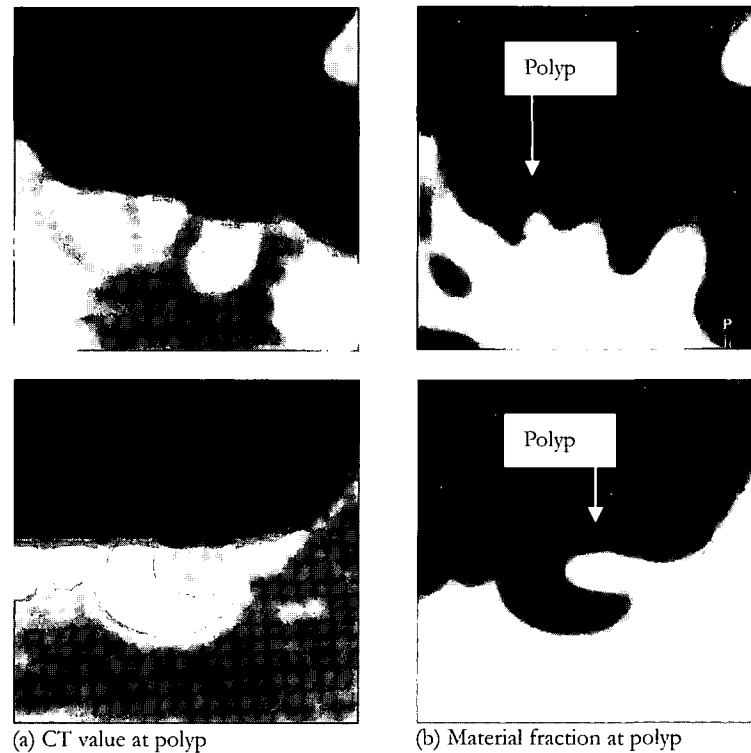
(a) CT value at polyp (b) Material fraction at polyp

Figure 5.12 Two polyps from the Walter Reed AMC data set: nr 333, and 357. (a) Part of the data showing the CT value at the polyp and (b) material fraction at the polyp. Visualization of (c) original data and (d) using electronic cleansing.

5.4   Conclusion

We proposed a model-based approach to determine material fractions in partial-volume voxels at three-material transitions. It solves the partial-volume effect that hampers accurate delineation of materials. The method is translation-, rotation- and scale-invariant – even for data with an anisotropic resolution and non-cubic voxels. All model parameters are estimated from the data such that the result is patient-specific and specifically adapts to the variability regarding the intensity of the intra-luminal remains.

We derived an analytical expression relating Gaussian derivatives (the features) in gradient direction to material fractions at three-material transitions: the parachute model. Triplets of Gaussian derivatives were computed from the noisy input image around a voxel in the proximity of a three-material transition. The measurements were gathered along a path that follows the gradient both uphill and downhill from a voxel under consideration. All measurements were properly scaled to compensate for orientation-dependent resolution,
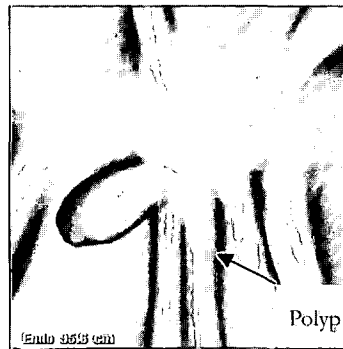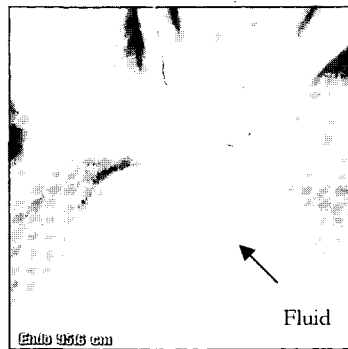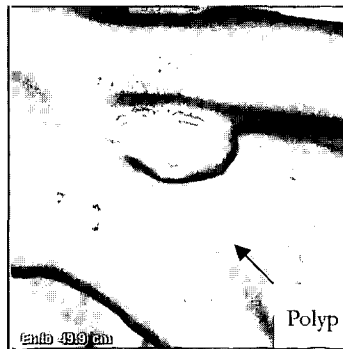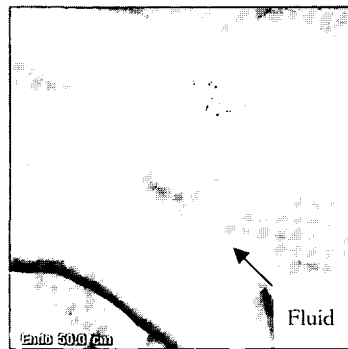
(c)   Iso-surface   visualization   using   (d)   Iso-surface using constant material
      constant CT value in original data   fraction in cleansed data

sampling and anisotropy of the noise between the derivatives of different order. The scale/noise normalized measurements are compared against 'code books' generated from modeled junctions in which the angle is varied. The entry that produces the best match to our observations is selected. This yields noise-suppressed Gaussian derivatives as a robust alternative to the standard computed Gaussian derivatives. The relative position emanating from the matching yields an approximation to the underlying noise-free material fractions at three-material transitions in a barycentric coordinate system

Previously described methods [5] only provided an analytical expression of the two-material transition and did not model the anisotropy of the data. Clearly, anisotropic input data may be blurred to meet such a requirement. It should be noticed that our method does not sacrifice resolution like that.

In addition, previous work disregarded noise. This work shows how to approximate the underlying noise-free material fraction at a three-material transition. Noise is often disregarded by visualization-methods. An advantage of projecting measured data onto a model is that noise is dealt with before proceeding into the visualization process. Finally, we show how to properly sample the three-material transition model for practical use. Typically, the data are 512^3. However, the method is edge based and needs only to be applied to positions not dealt with by a two-material-transition model. On average, our implementation adds two minutes to the processing to deal with the junctions.

The proposed method does not require complicated widgets for transfer function selection. The visualization examples involved user interaction merely to decide which material to visualize. Objects of interest were rendered at a constant, 0.5 material-fraction threshold.

The proposed method facilitates CT colonography to become effective with tagged data in combination with a 3-D display mode. The method is currently being evaluated in a clinic for diagnostic purposes. Our current work focuses on the application of the three-material model outside the scope of CT colonoscopy.

5.5 Acknowledgement

5.6 References

[1]   D.J. Vining, D.W. Gelfand, R.E. Bechtold, E.S. Scharling, E.K. Grishaw and R.Y. Shifrin, Technical Feasibility of Colon Imaging with Helical CT and Virtual Reality, *in: AJR*, vol. 162, 1994.

[2]   I.W.O. Serlie, F.M. Vos, R. van Gelder, J. Stoker, R. Truyen, F. Gerritsen, Y. Nio, F.H. Post, Improved Visualization in Virtual Colonoscopy Using Image-Based Rendering, *in: IEEE TCVG Symp. on Vis.*, Conf. Proc., 2001, pp. 137–146.

[3]   M.R. Wax, Virtual colonoscopy - CT contrast agents, *in: 2nd Int. Symp. VC. Boston*, 1998, pp. 66–68.

[4]   P.J. Pickhardt, J.H. Choi, Electronic cleansing and stool tagging in CT colonography: advantages and pitfalls with primary three-dimensional evaluation, *in: AJR*, vol. 181(3), 2003, pp. 799–805.

[5]   I.W.O. Serlie, R. Truyen, J. Florie, F.H. Post, L.J. van Vliet, F.M. Vos, Computed cleansing for virtual colonoscopy using a three-material transition model, *in: MICCAI, LNCS-2879*, 2003, pp. 175–83.

[6]   S. Lakare, D. Chen, L. Li, A. E. Kaufman, M. Wax and Z. Liang, Robust Colon Residue Detection Using Vector Quantization Based Classification for Virtual Colonoscopy, *in: SPIE Med. Im., Physiology and Function from Multidimensional Images*, San Diego, CA, USA, 2003.

[7]   S. Lakare, D. Chen, L. Li, A. E. Kaufman, M. Wax, Z. Liang, Electronic Colon Cleansing Using Segmentation Rays for Virtual Colonoscopy, *in: SPIE Med. Im., Physiology and Function from Multidimensional Images*, Vol. 4683, pp. 412–418.

[8]   M.E. Zalis, J. Perumpillichira, C. Del Frate, P.F. Hahn, CT colonography: digital subtraction bowel cleansing with mucosal reconstruction initial observations, *in: Radiology*, vol. 226: 2003, pp. 911–917.

[9]   D.H. Laidlaw, K.W. Fleischer, and A.H. Barr, Partial-volume Bayesian classification of material mixtures in MR volume data using voxel histograms, *in: IEEE Trans. Med. Im.*, vol. 17, 1998, pp. 74–86.

[10]   G. Kindlmann, and J.W. Durkin, Semi-Automatic Generation of Transfer Functions for Direct Volume Rendering, *in: Proc. IEEE Symp. Vol. Vis.*, 1998, pp. 79–86.

[11]   J. Kniss, G. Kindlmann, and C. Hansen, Multi-dimensional transfer functions for interactive vol-ume rendering, *in: IEEE Trans. Visual. Comput. Graph.*, vol 8, 2002, pp. 270–285.

[12]   P.J. Pickhardt, J.R. Choi, I. Hwang, J.A. Butler, M.L. Puckett, H.A. Hildebrandt, R.K. Wong, P.A. Nugent, P.A. Mysliwiec, W.R. Schindler, Computed tomographic virtual colonoscopy to screen for colorectal neoplasia in asymptomatic adults, *in: N. Engl. J. Med.*, vol. 349, 2003, pp. 2191–2200.

[13]   G. Wang, M.W. Vannier, M.W. Skinner, M.G.P. Cavalcanti, G.W. Harding, Spiral CT image deblurring for cochlear implantation, *in: IEEE. Trans. Med. Imag.*, vol 17, 1998, pp. 251–262.

[14]   P.W. Verbeek and L.J. van Vliet, On the location error of curved edges in low-pass filtered 2-D and 3-D images, *in: IEEE Trans. Pat. Anal. & Mach. Intell.*, vol. 16, no. 7, 1994, pp. 726–733.

[15]   V. Berzin, Accuracy of Laplacian edge detectors, *in: Comp. Vis. Graph. & Im. Proc.*, vol. 27, 1984, pp. 195–210.

[16]   H. Bouma, A. Vilanova, L.J. van Vliet, F.A. Gerritsen, Correction for the dislocation of curved surfaces caused by the PSF in 2D and 3D CT images, *in: IEEE Trans. Pat. Anal. & Mach. Intell.*, vol. 27, 2005, pp. 1501–1507.

[17]   F.M. Vos, R.E. van Gelder, I.W.O. Serlie, J. Florie, C.Y. Nio, A.S. Glas, F.H. Post, R. Truyen, F.A. Gerritsen, J. Stoker, Three-dimensional display modes for CT colonography: conventional 3-D virtual colonoscopy versus unfolded cube projection, *in: Radiology*, vol. 228, 2003, pp. 878–885.

[18]   R. Truyen, T. Deschamps, L.D. Cohen. Clinical evaluation of an automatic path tracker for virtual colonoscopy, *in: Proc. Med. Im. Comp. and Comp.-Assist. Int.*, MICCAI 2001.

[19]   J.J. Dijkers, C. van Wijk, F.M. Vos, J. Florie, Y.C. Nio, H.W. Venema, R. Truyen, L.J. van Vliet, Segmentation and Size Measurement of Polyps in CT Colonography, *in: LNCS-3749, MICCAI*, 2005, pp. 712–719.

h

Chapter Six

Lesion Conspicuity and Efficiency of CT Colonography with Electronic Cleansing Based on a Three-material Transition Model

This chapter is based upon a publication that is submitted [1].

This chapter reports on the effect on conspicuity of polyps as well as the practical efficiency of an electronic cleansing algorithm for CT colonography.

All patients were included from a public database; they all adhered to an extensive bowel preparation scheme. Patient group I consisted of all patients harboring polyps larger than 6 mm. This group served to assess the effect of the algorithm on the conspicuity of polyps. There were 129 images of polyps in total; 59 partly or completely resided in tagged material. Patient group II consisted of 19 randomly chosen patients from the same database to test the algorithm's efficiency: 10 with polyps larger than 5 mm and 9 without such polyps. All polyps from patient group I were rated regarding their conspicuity by one experienced observer. We used an enhanced 3-D visualization (the "unfolded cube display"). Two experienced observers evaluated all patients from patient group II using the enhanced 3-D display both before and after cleansing.

On average 20% of the colon volume contained fecal material or fluid. The observer rated the conspicuity of polyps residing in air not differently from 'cleansed' polyps that were partly or fully covered by tagged intraluminal remains ($p = 0.6$). Median evaluation time per patient for the cleansed data was significantly shorter for both observers than for the original data (12 minutes versus 17-20 minutes per patient; $p < 0.004$). Both observers rated the assessment effort for the cleansed data significantly smaller ($p < 0.001$) and observer confidence was rated significantly larger for the cleansed data ($p < 0.007$).

The conspicuity of polyps is identical for partially-covered polyps that were uncovered by the proposed electronic cleansing and polyps bordering on air that did not need electronic cleansing. The proposed cleansing method sustains a lower evaluation time, lower assessment effort and larger observer confidence than a conventional evaluation method without cleansing.

6.1   Introduction

Computed Tomography (CT) colonography is a minimally invasive procedure that is advocated for polyp screening [2]. Residual fecal material and fluid are well known to hinder CT colonography evaluation, especially when using a 3-D display mode. It may cover lesions (preventing detection), or reversely, mimic polyps (necessitating superfluous 2-D verification and possibly reduce specificity). Fecal tagging was introduced to discriminate between tissue and fecal material [3][4]. It facilitates a limited bowel preparation, which may contribute to a better patient compliance.

Electronic cleansing aims at replacing tagged materials (i.e. fecal remains and fluid) by air. The operation is a prerequisite to evaluate limited bowel preparation data in a primary 3-D display and it may aid 3-D problem solving with a primary 2-D reading.

Several electronic cleansing algorithms based on the increased attenuation of tagged material were described previously [5][6]. In our opinion, the variety of presently proposed algorithms reflect that a perfect solution has not been found yet. Accordingly, incomplete processing is still reported to leave artifacts [7]. A specifically noticeable problem is posed by the distracting 'bumps' emanating from locations where air, soft tissue and tagged material meet [7]. An electronic cleansing algorithm was devised aiming to improve the accuracy at such three-material junctions [8]. The method adapts to patient specific conditions, such as the local variation of the tagged material density, and is automated.

This chapter studies the effect on lesion shape as well as the practical efficiency of the electronic cleansing algorithm. The method is primarily tested regarding lesion conspicuity, cleansed surface and cleansed volume, evaluation time, assessment effort and observer confidence by employing an enhanced, primary 3-D display method without and with cleansing. We included patients from a public database that all had undergone an extensive preparation [9]. We consider an excellent electronic cleansing performance under such circumstances a prerequisite before proceeding to limited preparation CT colonography.

We hypothesize that in a 3-D display the conspicuity of polyps is identical for partially-covered polyps that were uncovered by electronic cleansing and polyps surrounded by air (obviating electronic cleansing). Additionally, we expect that electronic cleansing facilitates comprehensive visibility and time efficient inspection in comparison to using uncleansed data. Due to this, we anticipate that the assessment effort of observers is lower and confidence of observers in their diagnosis is significantly higher.
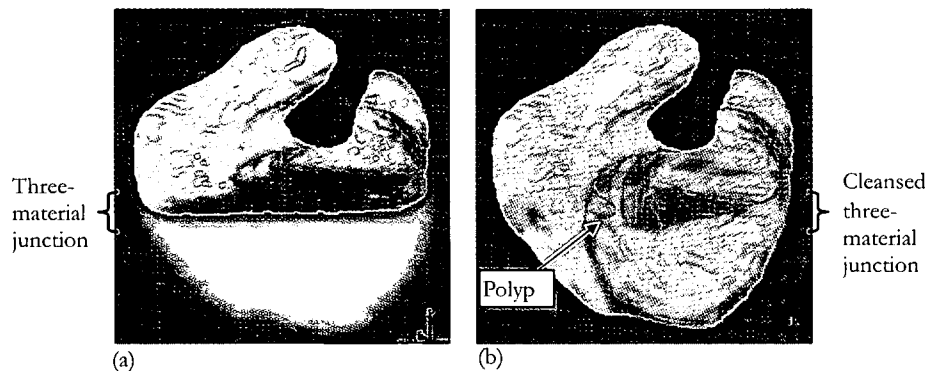
Figure 6.1   Images from WRAMC study case 35, prone position. (a) Detail of the CT colonography data with air, tissue and tagged material. (b) Electronic cleansing replaces tagged material by air. The colors are merely added for illustration purposes to separately indicate the air soft-tissue interface and the tagged material-soft tissue interface. The electronic cleansing method particularly focuses on three-material junctions. The iso-value is -650 HU. The window/level setting of the images is (w = -1250, l = -50).

6.2   Material and methods[1]

6.2.1   Patient population

We included patients from publicly accessible study data from the Walter Reed Army Medical Center (WRAMC) made available by the National Cancer Institute (www.vcscreen.com) [9][10].

Patient group I contained all patients from the database harboring polyps larger than 6 mm in diameter (measured during colonoscopy), irrespective of shape or location. There were 36 polyps between 6 and 10 mm in diameter in 32/51 patients and 29 polyps larger than or equal to 10 mm in 24/51 patients amounting to 65 polyps in total and 129 images of polyps in total. The median age of the patients was 57 years (range: 48 - 76 years), 36 were male and 15 female. Patient group I served to explore the effect of electronic cleansing on lesion conspicuity.

Patient group II consisted of 10 randomly selected patients from the aforementioned database that contained polyps larger than 5 mm diameter, as well as 9 randomly selected patients without such polyps, see Table 6.1. There were 9 polyps between 6 and 10 mm in diameter in 7/19 patients and 3 polyps larger than or equal to 10 mm in 3/19 patients

amounting to 12 polyps in total. The median age of the patients was 55 years (range: 50 - 78 years), 13 were male and 6 female. Patient group II was taken to investigate the efficiency of the electronic cleansing algorithm.

6.2.2   CT colonography

The patients had undergone a standard, extensive colonic preparation. Patients consumed 500 ml of barium (2.1 percent by weight; Scan C, Lafayette Pharmaceuticals, Tyco Healthcare/Mallinckrodt, St Louis, MO) and 120 ml of diatrizoate meglumine and diatrizoate sodium (Gastrografin, Bracco Diagnostics, Princeton, New Jersey, USA). Distension was achieved through patient-controlled insufflation of room air.

6.2.3   Colonoscopy

CT colonography findings performed at Walter Reed Army Medical Center were disclosed during colonoscopy using segmental unblinding. The colonoscopy findings were available via a report that included photographs of each polyp, size measurement, morphology (pedunculated, sessile, flat) and location: cecum, ascending colon, transverse colon, descending colon, sigmoid and rectum (Figure 1.1a).

6.2.4   Electronic cleansing

The electronic cleansing method assumes that a measured CT value arises due to a combination of three materials: soft tissue, air and tagged material. Initially, the percentage of materials is determined at each voxel. Subsequently, the partial volume of tagged material is replaced by air and a new cleansed CT value is calculated. At last, an enhanced 3-D method visualizes the colon from an endoluminal perspective as if there are no fecal remains (Figure 6.1).

Consider the profile drawn in (Figure 6.2a). Suppose that one proceeds along this line and at each position the data value is recorded as well as the local 'steepness' of the transition (mathematically: the gradient or first derivative of the grey value). Plotting this in a graph results in arch shaped curves (Figure 6.2b). Thus, each transition type (1,2,3) yields its own arch. Consider for instance the arch corresponding to the transition from air to tagged material (transition 2). It starts at the Hounsfield unit of air, proceeds over the transition, and ends at the Hounsfield unit of tagged material. Moreover, the arch from air to tagged material is higher than the arch from air to soft tissue, reflecting the steeper data value transition between these materials.

The material 'constitution' of a voxel is estimated in the way sketched in Figure 6.2c, [8]. Initially, the *type* of transition is identified (i.e. air-tagged material). Subsequently, the nearest point on this arch is taken in order to obtain the material percentages. The latter procedure takes into account that the image values of the materials may vary (this specifically goes for tagged material).
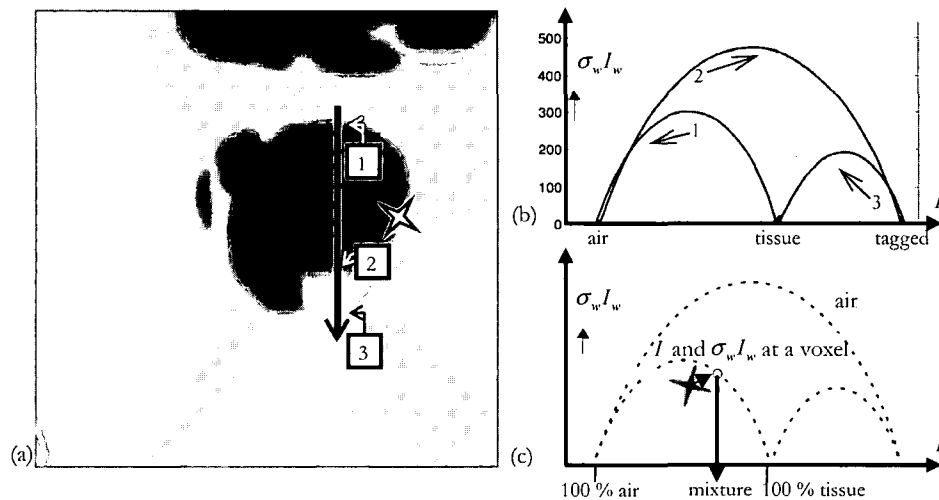
Figure 6.2   (a) An arbitrary line in a CT colonography image. (b) Going along this line in the indicated
order, the image intensity is drawn against the local 'steepness' of the data value profile
(mathematically: the derivative of the data value profile). (c) A voxel's constitution is
estimated by first identifying the proper arch. Then, the closest point on the arch is
selected, which, in turn, enables us to estimate the material percentages. A more detailed
explanation of the method is given in the electronic cleansing section.

Simply applying this algorithm to a sample dataset (Figure 6.3a) leads to the 3-D
endoluminal view of Figure 6.3b. The well visible artifact emanates from a band-shaped
area where three materials meet, also described in [7]. A separate model was developed
for such three-material junctions that, again, relates image value, local 'steepness' of the
transition, and other local intensity properties to a voxel's constitution [8]. Effectively, the
arch model for two materials is extended to three materials. The junction artifacts are
solved by first detecting these locations and then applying the separate model to estimate
the material percentages [8]. Figure 6.3c shows the same rendering after also processing
the three-material transitions.

The cleansing method was implemented by us on a proprietary, experimental version of
the ViewForum workstation (Philips Medical Systems, Best, The Netherlands). For a
thorough description of the method the observer may consult [8].

6.2.5   Observers

An abdominal radiologist (observer 1) and a research fellow (observer 2) were involved in
reading the data. Observer 1 had a previous experience of approximately 1200
colonoscopy verified CT colonography examinations at the start of the study. Observer 2
had a previous experience of 350 such examinations.

6.2.6   Image review to study polyp conspicuity

The conspicuity study focused on all true polyps larger than or equal 6 mm in as indicated in the WRAMC colonography reports from patient group I. Prone and supine positions were considered as separate cases. Each case was assessed by a researcher who had a background in designing electronic cleansing algorithms. This research fellow determined whether the polyp resided in air, or was partly/fully covered with tagged intraluminal remains.

Observer 2 reviewed all the polyps residing in air (that did not require cleansing) and the ones that were uncovered by the electronic cleansing. The cases were presented in a random order and the observer was not informed if the polyp had been electronically cleansed.

Each polyp was marked and directly presented to the observer by means of the 3-D display (unfolded cube display: ViewForum, Philips Medical Systems) [11]. The position along the centerline was under control by the observer. 2-D reformatted views of the original CT data were initialized by clicking on a position in the 3-D endoluminal images. The window/level setting of the reformats was by default ($w = 1250$, $l = -50$), but this could be freely adapted by the observer.

6.2.7   Evaluating polyp conspicuity

Initially, the observer indicated for each case on a five-point scale whether he suspected that the polyp did not require cleansing or was uncovered by electronic cleansing (e.g. by noticing a smoother polyp-surface due to image processing). On this scale 1 corresponded with a polyp definitely considered not to be 'cleansed' and 5 with a polyp strongly suspected to be 'cleansed'.

Subsequently, each polyp was scored with respect to 'conspicuity' on a 5 point Likert scale using both the enhanced 3-D display and 2-D reformats:

1) Inadequate: the lesion is not visible. This lesion will not prompt a 2-D inspection.

2) Moderate to questionable: the lesion is hardly visible. This location might lead to 2-D inspection.

3) Average: it is expected that the lesion is detected and that the location will prompt a 2-D inspection.

4) Good: the lesion is well visible and this location should provoke a 2-D inspection.

5) Excellent: the lesion is very well visible and this location will certainly lead to a 2-D inspection.

Finally, the observer had to indicate for the cases scored as 'inadequate' or 'average' whether in his opinion the cleansing algorithm caused the difficulty or if there was another cause: 1 inappropriate cleansing; 2 other (e.g. CT artifact, flat lesion). In the latter case he was asked to indicate what in his opinion was the reason of the difficulty.
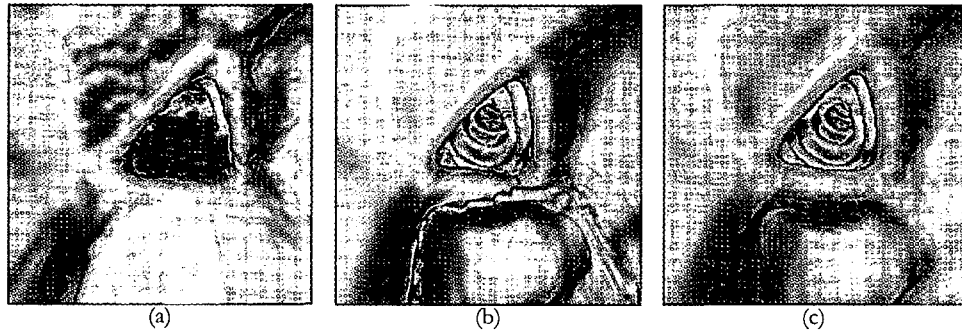
(a)   (b)   (c)

Figure 6.3   (a) 3-D endoluminal rendering without applying electronic cleansing (WRAMC study case 416). (b) Result after applying cleansing by means of the method of Figure 6.2. Notice the artifact at the three-material junction in the purple shaded band, since the method is only suitable for two-material edges. (c) Result after applying the described extended cleansing method that is applied to the three-material voxels indicated by the shaded band in image (b).

6.2.8   Outcome parameters and statistical analysis on lesion conspicuity

The polyps that were scored to reside in air were compared with those partly or completely submerged in fecal material. The difference in the consideration whether a polyp was 'cleansed' was tested by means of the Mann-Whitney test (Wilcoxon Rank-sum test). Differences in conspicuity between cleansed and uncleansed polyp-data were also statistically tested by means of the Mann-Whitney test. Two-sided tests were used to assess both these differences. A p-value less than or equal to 0.05 was considered to correspond to a statistically significant difference

6.2.9   Image review to study efficiency

The original, uncleansed data from patient group II was reviewed as follows:

Part 1: Enhanced three-dimensional cine loops (unfolded cube display) were examined in prone and supine position, starting with the former (ViewForum, Philips Medical Systems) [11]. The two positions were electronically linked: after clicking in one unfolded cube image, the corresponding unfolded cube image of the other position was also displayed [12]. The frame rate was under control by the observers. 2-D reformatted views of the original CT data became accessible by clicking on a position in the 3-D endoluminal images. The window/level setting of the reformats was by default ($w = 1250$, $l = -50$), but this could be freely adapted by the observers. Lesion size was measured by means of electronic calipers on the 3-D view (we used polyp size in matching the encountered lesions to the reference standard).

Part 2: After reviewing the unfolded cube images, the original, axial CT slices were inspected to verify surface parts obscured by fecal material and fluid. Again, the applied window/level setting was by default ($w = 1250$, $l = -50$), but this could be freely adapted.

The cleansed data of a patient was reviewed in a similar manner:

Part 1: 3-D unfolded cube images were generated from the *cleansed* volume. These images were inspected in the same way as the unfolded cube images from the *uncleansed* volume. However, the reformatted views were based on the original (uncleansed) data as proposed in [7].

Part 2: After reviewing the unfolded cube images, the original, *uncleansed* axial CT slices were inspected in the same way as described previously, again inspired by [7].

Figure 6.4 illustrates the interface to read the data.

6.2.10   Evaluating efficiency

Observer 1 and observer 2 independently evaluated the uncleansed data and the cleansed data from patient group II. Both were blinded to findings during colonoscopy, findings by themselves for the same patient (original versus cleansed) as well as each other's findings. Additionally, they were unaware of the prevalence of polyps. The interval period between two evaluations of the same patient was at least 4 weeks. The evaluations using the uncleansed data preceded those with cleansed data. The cases were presented in random order. Each surmised polyp was scored with respect to size, morphology, location (colon segment), and lesion confidence. The lesion confidence was qualified on a five-point Likert scale: 0 not a lesion and 4 absolute confidence of a lesion. The observers were aware that in a clinical setting a confidence of 2 or more would indicate a relevant lesion.

The reviewers were aware that electronic cleansing algorithms may leave artifacts, specifically clouds of debris and lines of ridges at junctions. They were instructed to report such artifacts and indicate whether it might preclude polyp detection.

The observers rated the assessment effort per colon segment on a four-point Likert scale: 1, extremely easy; 2, good; 3, difficult; 4, extremely difficult. Additionally, they rated their confidence in the reading per segment on a three-point Likert scale: 1, confident; 2, in doubt; 3, extremely doubtful. All these measures were recorded after complete evaluation of a patient.

6.2.11   Reference standard to study efficiency

Lesions detected during CT colonography evaluation were matched with the findings in the colonoscopy report (see above). This was done by a research nurse (experience > 500 such matchings) under the supervision of a research fellow (experience > 350 CT colonoscopy examinations). Both were not involved in reading the CT colonography studies.
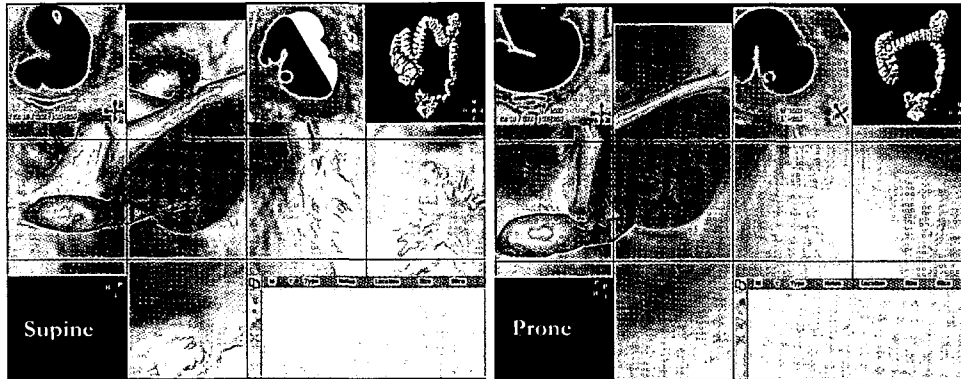
(a) 3-D primary unfolded endoluminal imaging. The insufflation catheter is visible.
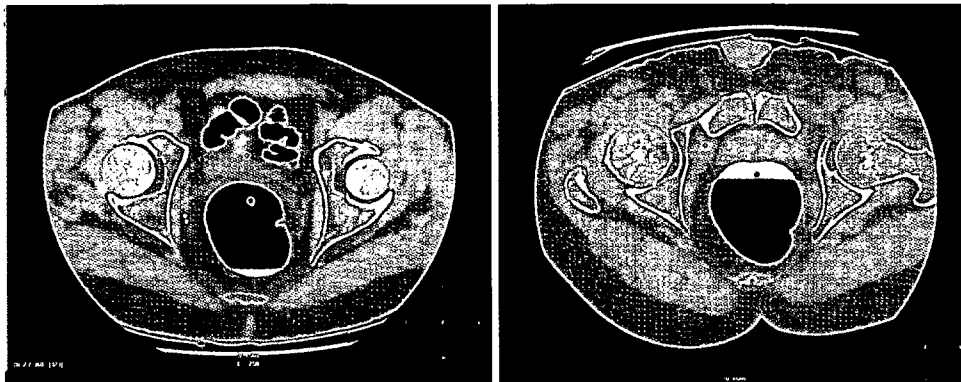
(b) 2-D secondary axial imaging.

Figure 6.4   Interface for the CT colonography evaluation (WRAMC study case 393). (a) Initially the enhanced 3-D unfolded cube display mode was used. Prone (left) and supine (right) views were linked using the colon trajectory. (b) Subsequently, axial CT images were inspected while tracking the colon's centerline. The same interface was employed to read both the original and the cleansed data.

A polyp detected via our CT colonography evaluation was considered true positive if it matched a colonoscopy finding with regard to size, morphology and location (i.e. same segment). A deviation in size by at most 5 mm was accepted to accommodate with the inherent inaccuracy of colonoscopic size measurement. A false negative finding was defined as a polyp detected during colonoscopy that did not match any CT colonography finding. A CT colonography finding with a lesion confidence of at least 2, larger than or equal to 6 mm *and* not matching a polyp detected during colonoscopy was considered to be false positive. All detections smaller than 6 mm or scored with an observer confidence lower that 2 were not considered relevant and discarded.

6.2.12   Cleansed surface fraction and cleansed volume

We approximated the 'true' colon's interior volume by thresholding the cleansed volume at −650 HU. The 'true' surface voxels border on this volume. Likewise, thresholding the original volume at −650 HU identifies the 'uncleansed' colon's interior volume; the 'uncleansed' voxels border on the latter interior volume. 'Cleansed' surface voxels are those that *are* part of 'true' surface voxels and *are not* part of the 'uncleansed' surface voxels. These conditions aim to select voxels at the tagged-material/soft-tissue transition. Similarly, the cleansed volume voxels are those that are part of the 'true' interior volume and not part of the 'uncleansed' volume.

6.2.13   Outcome parameters and statistical analysis on efficiency

Evaluation time

The evaluation time per patient was measured for both observers, for the cleansed and the original data. The time for evaluating the enhanced 3-D cine loop (Part 1) and the axial slices (Part 2) was recorded separately. The measured times included annotation of the polyps, i.e. size measurement and indications of location, morphology and lesion confidence.

All evaluation times for the original data were compared with the corresponding times for the cleansed data using Wilcoxon's matched pairs signed ranks test. Differences were considered significant at $P < 0.05$. The initialization time (e.g. for loading the data) was disregarded in the analysis.

Assessment effort and observer confidence

Differences in assessment effort and observer confidence between the original data and the cleansed data were statistically tested per segment by means of Wilcoxon's matched pairs signed ranks test. Two-tailed P values were used and a P value less than or equal to 0.05 was considered statistically significant.

Sensitivity/specificity

We consider a statistical comparison of sensitivity and specificity between the two methods based on the set of 19 patients not possible. Consequently, the measures of sensitivity and specificity by itself merely serve as descriptive statistics for polyps > 5 mm.

The sensitivity of both display methods was determined on a per polyp basis. The specificity of both display methods was determined by the false positive rate.

Cleansed surface fraction and cleansed volume fraction

The mean cleansed surface fraction is defined as the average number of 'cleansed' surface voxels divided by the average number of 'true' surface voxels. It is considered to reflect the fraction of 'true' surface added by the cleansing algorithm.
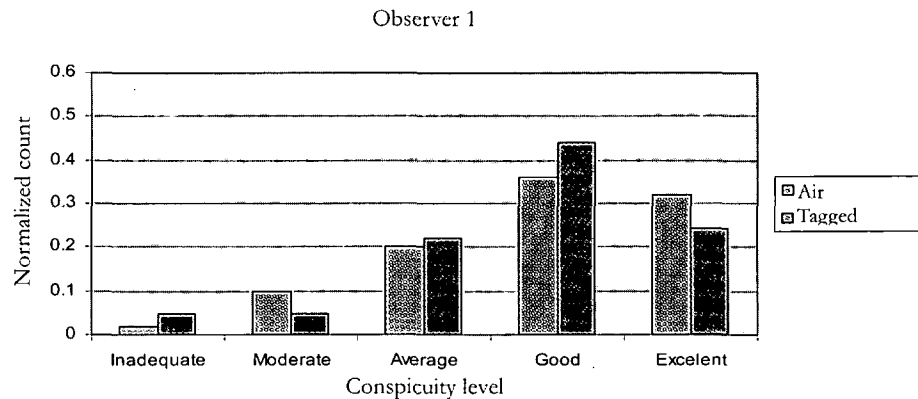
Observer 1

Figure 6.5   Histograms of polyp conspicuity. The difference between polyps residing in air and
those partly or fully residing in tagged fecal material was not significant ($p = 0.6$).

The mean cleansed volume fraction is defined by the average 'cleansed' volume divided
by the average 'true' volume. It is taken to represent the fraction of additionally exposed
'true' colon volume.

6.3   Results

70 (images of) polyps were completely surrounded by air, 18 polyps completely resided in
tagged intraluminal remains and 41 polyps partly resided in remains amounting to 59 in
total that border on tagged material. Two patients (WRAMC: 32, 32) were excluded from
the study due to inadequate diagnostic quality. Four patients (WRAMC: 26, 98, 220, 256)
were excluded due to incomplete data.

6.3.1   Conspicuity

The observer rated polyps residing in air significantly lower on scale whether he
suspected that the polyp partly or fully resided in fecal matter or fluid ($p < 0.05$). The
reader reported a slight 'noise' suppression of electronic cleansing, which was apparent
by a less 'bumpy' surface (see also Figure 6.8I and II).

The conspicuity was rated identically for partially-covered polyps that were uncovered by
electronic cleansing and polyps bordering on air that did not need electronic cleansing
($p = 0.6$). Figure 6.5 shows the normalized histogram of the conspicuity. Figure 6.8I and
II show examples.

The polyps scored with 'inadequate' or 'moderate to questionable' conspicuity were:
WRAMC 35 supine, 311 supine and prone and 185 supine. The reader attributed the
difficulty of all to the complex shape of the polyps (i.e. 'other' and not 'inappropriate

cleansing'). One polyp (i.e. WRAMC 35S) was at a fold. The other polyps were at the ileocecal valve.

Images of all polyps are accessible via the Internet [13]. The images are sorted according to the conspicuity rated by the observer.

6.3.2   Evaluation time

The median evaluation time per patient for inspecting the 3-D unfolded cube images (i.e. Part 1) significantly differed between the original (uncleansed) data and the cleansed data for observer 1 ($p = 0.037$) (Table 2). The median evaluation time for Part I was not significantly different for observer 2 ($p = 0.396$). The median time per patient for reviewing the axial 2-D slices (i.e. Part 2) was significantly larger for the original (uncleansed) data than for the cleansed data for both observers (for both: $p < 0.001$). The median total evaluation time per patient using the original data (observer 1: 20'45" min; observer 2: 17'14") was also significantly larger than using the cleansed data (observer 1: 12'25" min; observer 2: 11'8") for both observers (observer 1: $p < 0.001$; observer 2: $p < 0.004$).

6.3.3   Assessment effort and observer confidence

Both reviewers reported a "cloud of debris" artifact on one location emanating from acquisition artifacts. It was not considered to preclude polyp detection. No distracting lines or ridges at junctions were encountered.

Both observers rated the assessment effort for inspecting the original data significantly larger than for the cleansed data in each segment. Accordingly, the assessment effort over all segments was rated significantly larger for the original data than for the cleansed data by both observers (for both $p < 0.0001$).

Figure 8.6 shows histograms of the assessment effort for both observers.

The observer confidence over all patients was rated significantly smaller for the original data than for the cleansed data (observer 1: $p < 0.007$; observer 2: $p < 0.0002$). Figure 6.7 contains histograms of the observer confidence for both observers.

6.3.4   Sensitivity and Specificity

The sensitivity per polyp was identical for the original data and the cleansed data and for both observer: $10/12 = 0.83$. Observer 1 had, in total, five false positive detections in the original data and four false positive detections in the cleansed data. Observer 2 had two false positive detections in the original data and four in the cleansed data.
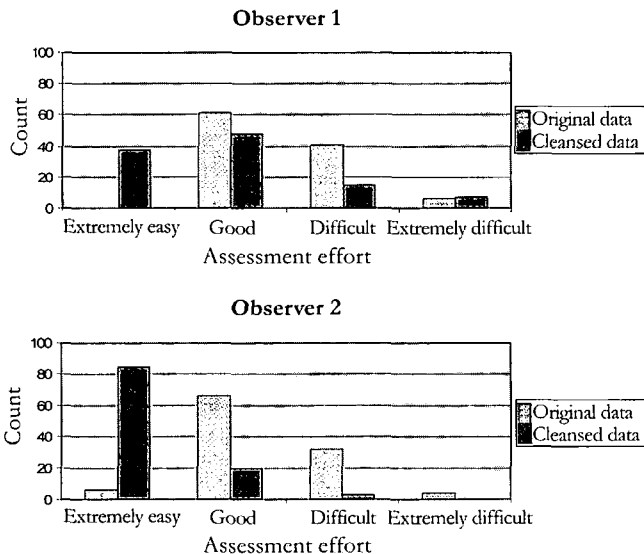
Observer 1

Observer 2

Figure 6.6   Histograms of the assessment effort of the CT colonography inspection for both observers over all segments and patients. Both observers had less effort to inspect electronically cleansed data.
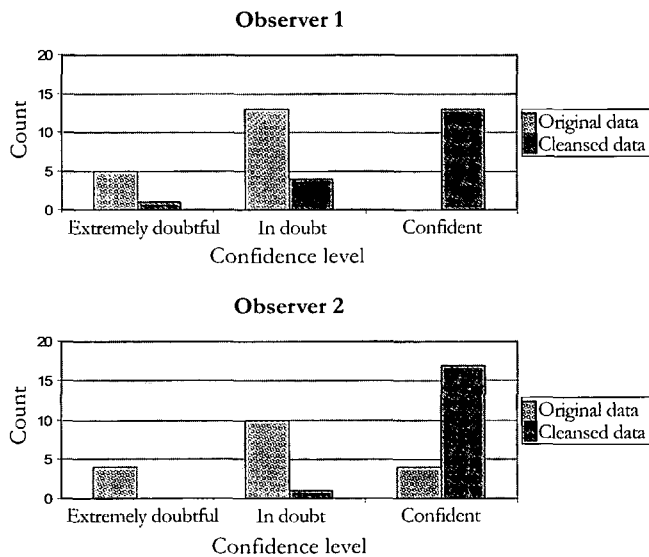
Observer 1

Observer 2

Figure 6.7   Histograms of observer confidence in the primary 3-D display using original and cleansed data for both observers over all patients. The confidence was significantly increased using electronically cleansed data.

excellent

good

average

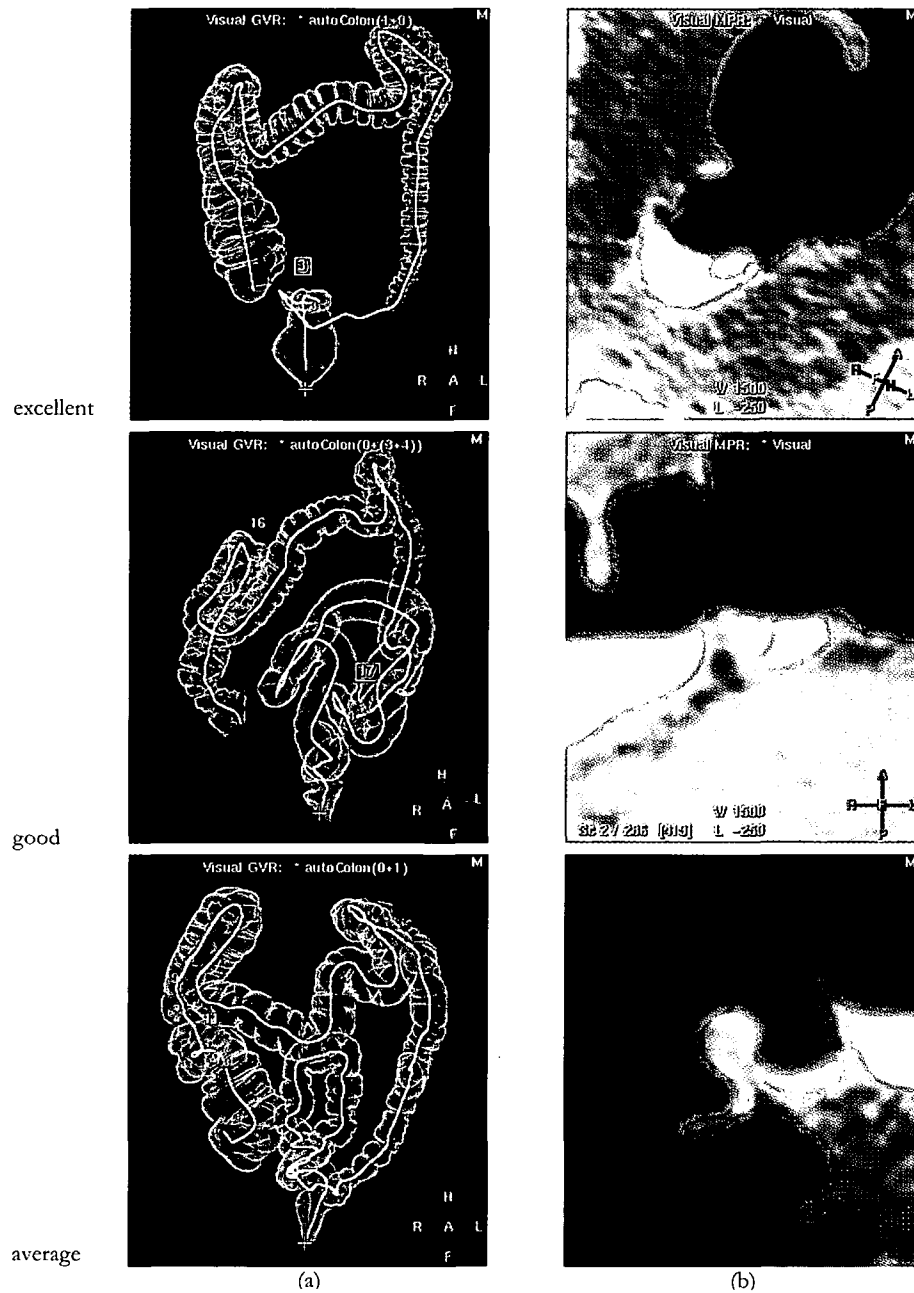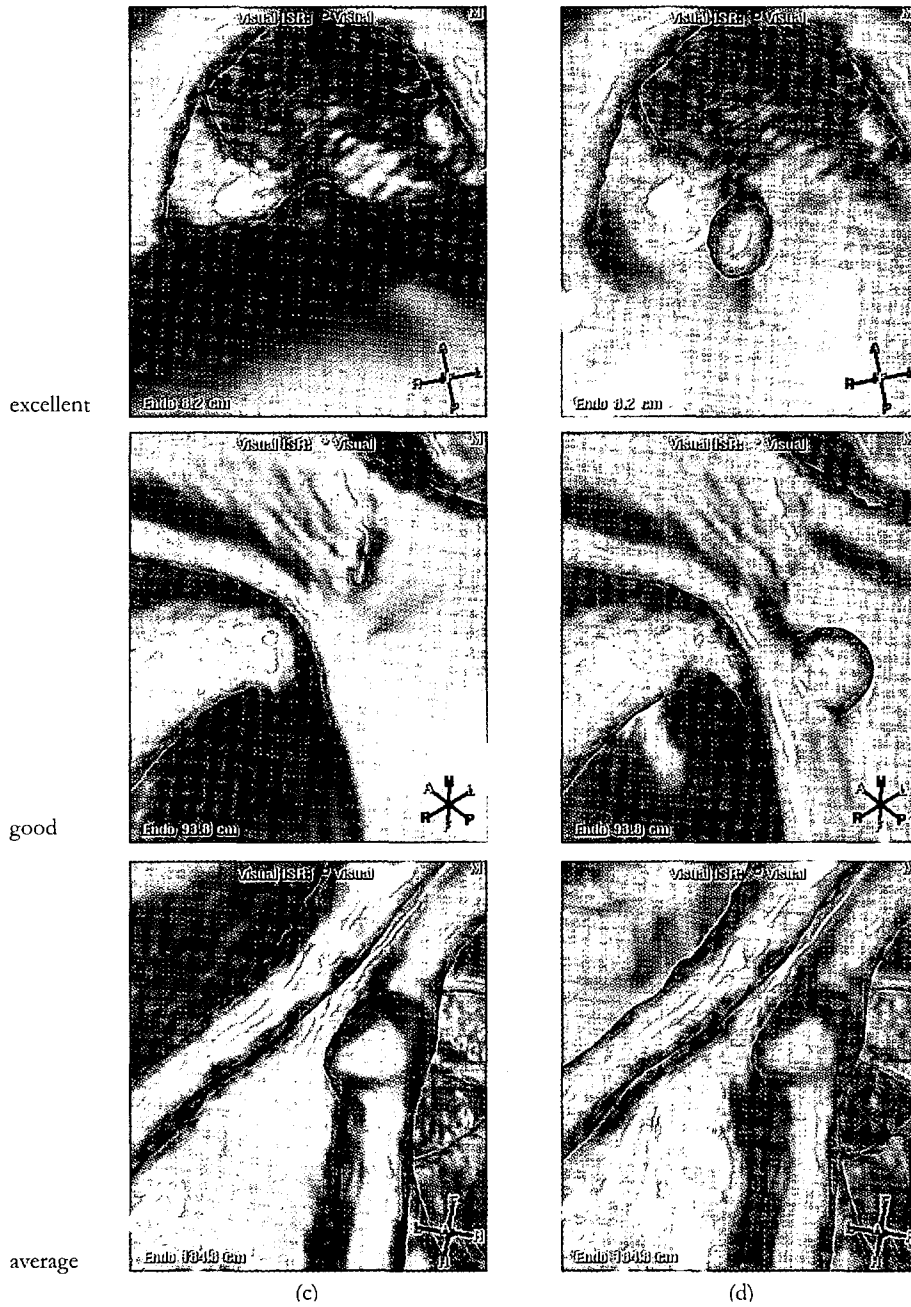(a)   (b)

Figure 6.8I   Examples of polyps before (a, b) and after (c, d) electronic cleansing from WRAMC: study
cases 38S (excellent), 192S (good), 226S (average). The rate of conspicuity refers to the
cleansed data. Figure 6.8II contains 35S (questionable) and 185S (inadequate conspicuity).

excellent

good

average

(c)   (d)

questionable

inadequate

(a)   (b)
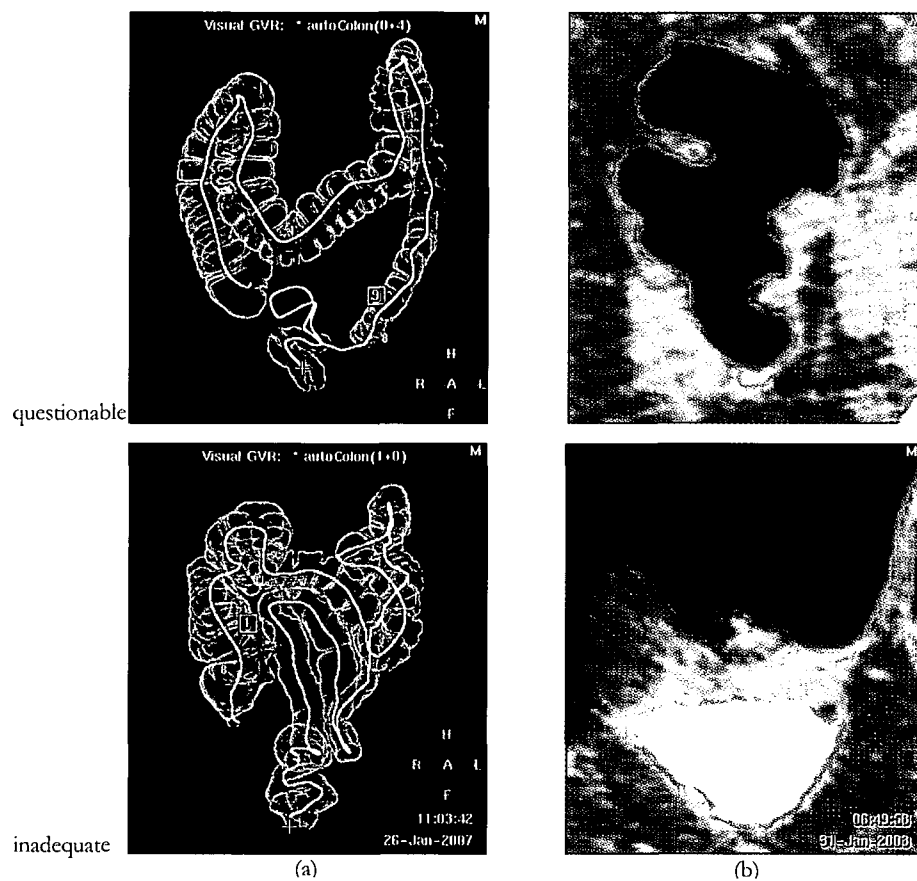
Figure 6.8II   Examples of polyps before (a, b) and after (c, d) electronic cleansing from WRAMC study
cases 35S (questionable) and 185S (inadequate conspicuity). The rate of conspicuity refers
to the cleansed data.

6.3.5   Cleansed surface fraction and cleansed volume fraction

The mean cleansed surface fraction was 269640 voxels/975378 voxels = 0.27 (range: 0.16
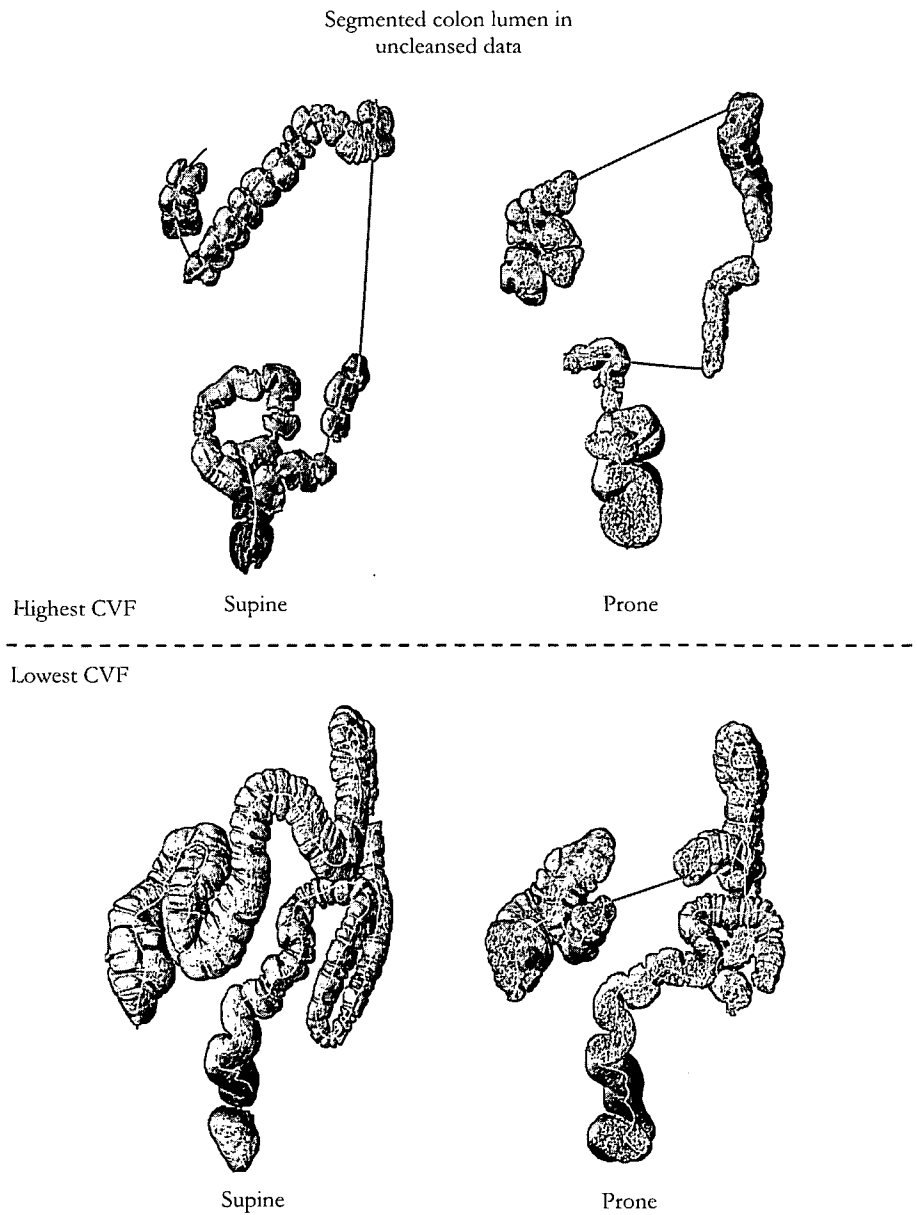to 0.46) and the mean cleansed volume fraction was 0.36 liter /1.84 liter = 0.20 (range:
0.12 to 0.45).

questionable

inadequate

(c) (d)

Figure 6.9 visualizes two examples of the air volume before and after cleansing. The top images are from the patient with the highest cleansed volume fraction in both scans (patient number: WRAMC 9; supine: 0.44 l /0.98 l = 0.45; prone: 0.44 l /1.33 l = 0.33); the bottom images are from the patient with the lowest cleansed volume fraction in both scans (patient number: WRAMC 408; supine: 0.31 l /2.33 l = 0.13; prone: 0.30 l /2.33 l = 0.13).
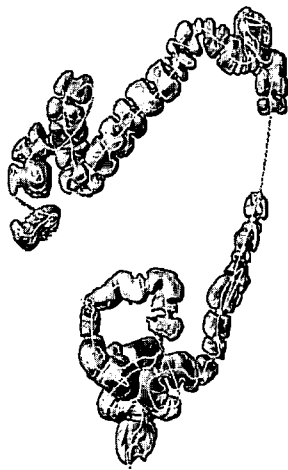
Figure 6.9   Visualization of the segmented colonic volume without (left) and with electronic cleansing (right) from the colon with the highest cleansed volume fraction (CVF) and the colon with the lowest cleansed volume fraction.
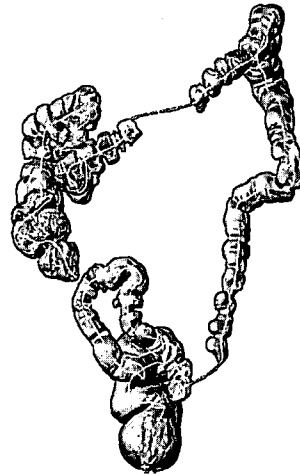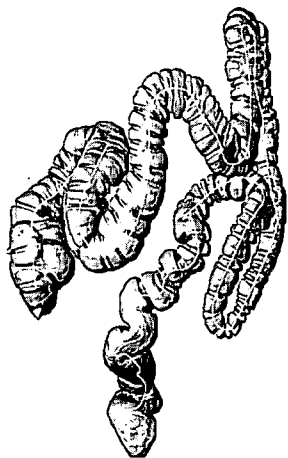
Segmented colon lumen in
electronically cleansed data
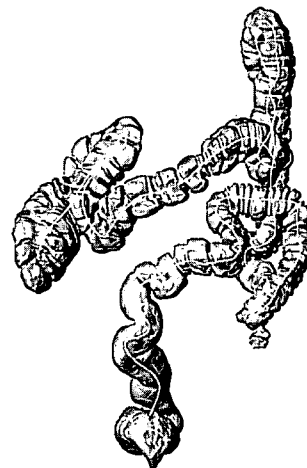

Highest CVF   Supine   Prone

- -

Lowest CVF

Supine   Prone

6.4   Discussion

This study shows that the conspicuity of polyps is identical for partially-covered polyps that were uncovered by the proposed electronic cleansing and polyps bordering on air that did not need electronic cleansing.

It is demonstrated that primary 3-D CT colonography using this cleansing method enables comprehensive visibility (mean cleansed surface fraction 27%) and time efficient inspection (about 12 minutes per patient) superior to using original data (uncleansed). Moreover, the assessment effort of the observers was low and the confidence of the observers was high. The sensitivity for polyps larger than 5 mm was high (10/12) and the false positive rate low (totally two to four false positive findings in 18 patients).

It is a complex problem to measure the performance of an electronic cleansing algorithm. Effectively, one should like to determine the extent to which the cleansing algorithm modifies polyp shape so that it is not detected anymore. To do so, one might consider determining the amount of distortion that the algorithm induces to a polyp by comparison to a manually or semi-automatically defined reference shape. We did not opt to do so since there may be a complex relation between any such deviation and the detection accuracy.

Alternatively, the accuracy of polyp detection after electronic cleansing might be compared to the detection accuracy using a primary 2-D method. However, such a comparison is confounded by the reading strategy and may not reflect the accuracy of the cleansing method.

It is for this reason that we compared polyp conspicuity of cleansed polyps and those residing in air (which are not affected by the algorithm).

Several previous papers describe technical innovations regarding electronic cleansing [5][6][14]. Zalis et al. [5] indicated that it will be essential to ascertain how artifacts still present in successfully tagged, electronically cleansed data will affect observer performance. The complex nature of electronic cleansing is also observed by Pickhardt [6] who identified pitfalls in reading electronically cleansed data. A common artifact described in the latter reference occurs at locations where air-fluid levels interface with the bowel wall. Proper action at these locations is important which is signified by 33/91 images of polyps found residing at such a location in the present study. The algorithm tested by us aims to improve electronic cleansing by finding a solution for this. Our reviewers encountered only one artifact in the study part on the efficiency which did not result from the cleansing algorithm. We hypothesize that the short average 3-D evaluation time, the high confidence, the low assessment effort, and the high specificity also signify that there were hardly any artifacts that complicated the reading.

Table 6.1 The table summarizes WRAMC study cases that were used to measure the evaluation efficiency with and without electronic cleansing. A research fellow indicated whether a polyp resided in air (A), partly resided in fecal matter (P), or fully resided in fecal matter (F).

WRAMC-patient number	Colonoscopically proven polyps	Location	
		Supine	Prone
8	6 mm, sessile, rectum	A	A
9	none	-	-
11	7 mm, pedunculated, hepatic fl.	A	A
19	10 mm, pedunculated, distal asc.;	P	A
	7 mm, pedunculated, sigmoid	A	F
20	6 mm, pedunculated, sigmoid	P	A
32	none		
35	7mm, pedunculated, sigmoid	P	P
38	10, round, sigmoid	P	A
56	8 mm, pedunculated, prox. asc.	F	A
60	none		
64	none		
81	none		
138	none		
193	8 mm, pedunculated, hepatic fl.;	A	P
	7 mm, sessile, sigmoid	A	P
180	none		
393	none		
408	12 mm, pedunculated, prox. asc.	A	A
416	none		
437	6 mm, sessile, rectum	A	A

Table 6.2 Median evaluation time (and interquartile range) per patient using original and cleansed data stratified by the primary 3-D part (unfolded cubes) and the 2-D part (axial CT slices) of the evaluation.

Evaluation part:	Original data		Cleansed data		Paired difference	
	Observer 1	Observer 2	Observer 1	Observer 2	Observer 1	Observer 2
Unfolded cubes (Part 1):	14'46" (6'44")	11'47" (4'38")	11'29" (5'12")	10'41" (5'30")	2'54" (6'29")	1'16" (8'24")
Axial CT slices (Part 2):	5'6" (2'22")	4'17" (1'46")	0'43" (0'34")	0'0" (0'00")	4'9" (2'4")	3'53" (1'34)
Total:	20'45" (6'48")	17'14" (5'49")	12'25" (5'42")	11'8" (5'34")	7'51" (7'52")	5'58" (7'49")

We did not find previous results on the effectiveness of electronic cleansing regarding cleansed surface fraction or cleansed volume fraction. A related finding on surface visibility is that on average 99.5% of the true colon surface becomes visible using the unfolded cube image sequences (95% using forward and backward looking image sequences) [11]. The average cleansed surface fraction (27%) by far exceeds the missed parts in those mentioned display modes.

We consider the applied, extensive preparation comparable to other such regimes (e.g. [15]). The reported cleansed surface fraction may therefore indicate how large an area of the colon surface remains covered by fecal matter. Still, a large cleansed surface fraction may not be clinically relevant, as it may concern a very thin layer. For that reason, we opted to also report the cleansed volume fraction, as a complimentary measure.

To the best of our knowledge, no previous paper reports reading times of *cleansed* CT colonography data using a primary 3-D display method. A reading time of approximately 14 minutes was found with an enhanced 3-D display technique similar to the approach employed currently [15]. The previous work did not apply fecal tagging, and (therefore) it did not include additional reading of axial 2-D slices for verification of obscured parts. The reading time per patient currently found (12 minutes) slightly improves compared to the previous work. The difference may be caused by the frame rate being under free control by the observers in our work (instead of fixed, as was done previously).

Earlier work using primary 2-D reading report evaluation times of 12 [16] and 16 minutes [17]. The initialization time was disregarded in the present work, because it was performed off line and initiated by technicians. On average the additional time required for the cleansing was only 5 minutes per dataset. Other initialization times such as for image loading and path tracking, are on the order of a few minutes.

A previous result on the 'readability' of cleansed CT colonography is that there is an improved rating by observers as the homogeneity of the tagged material increases [18]. The current findings on assessment effort and observer confidence confirm the accuracy of the present method perceived visually by the observers. However, there may well be interplay between the preparation regime and the cleansing regime. We observed that tagging was locally homogeneous in our population. However, along the colon differences in density of tagged material up to 800 Hounsfield Units in one patient are common (e.g. WRAMC study cases 19 or 35). One might attribute this to the extensive preparation scheme that the patients from WRAMC all had undergone. The algorithm employed by us allows for fluctuations in tagging intensity. The effect of a less homogeneous tagging warrants further research.

A limitation is the small study population. As indicated, the findings on sensitivity and specificity should merely be regarded as descriptive statistics. We have decided to include them, though, to have at least some indication in this respect and avoid the pitfall of

being out of the ordinary without notice. The reported sensitivity is in the same range as reported in [15], in which paper a similar display method is used.

Patient population II concerned an equivalent number of patients with and without polyps. This is appropriate for our primary objective, namely to compare reading cleansed versus original data in a primary 3-D fashion regarding time efficiency, surface visibility, assessment effort and observer confidence. The sample size (power) was calculated to meet the primary aim.

At last, a limitation is that our results may not be extrapolated to patients taking a limited bowel preparation. It might be expected that a rigorous preparation simplifies electronic cleansing somewhat since there may be less tagged material that typically has a homogeneous appearance. Still, the complex nature of the problem prompted us to initially study the algorithm's performance regarding extensively prepared patients. By the way, it should be noticed that a considerable cleansed volume fraction (0.20) was found, despite the extensive preparation. The volume and cleansed surface fractions might depend on the nature of the preparation.

Practical application.

The future role of CT colonography in cancer screening, in our opinion, depends on improvement on issues regarding patient compliance, efficiency, and radiation dose. Oral contrast material may enhance the patient compliance by facilitating a less rigorous purgation regime. We foresee a scheme in which a low dose CT scan of such a patient is automatically cleansed, after which a human observer checks potential lesion sites, suggested by a computer algorithm. A 3-D representation, either in a primary 3-D display or for 3-D problem solving in a primary 2-D display, could aid to do so. Although positive results were reported on automatic cleansing in several previous papers, further research is still needed. The proposed cleansing method allows for time efficiency, extensive surface visibility and a low assessment effort. Therefore, it may contribute to a practical evaluation strategy.

6.5   References

[1]   I.W.O. Serlie, A. de Vries, F.M. Vos, Y. Nio, A. Heutinck, R. Truyen, F. Post, J. Stoker, L. van Vliet, F. Vos, Lesion Conspicuity and Efficiency of CT Colonography with Electronic Cleansing Based on a Three-material Transition Model, submitted 2007.

[2]   D.J. Vining, D.W. Gelfand, R.E. Bechtold, E.S. Scharling, E.K. Grishaw and R.Y. Shifrin, "Technical Feasibility of Colon Imaging with Helical CT and Virtual Reality," *in: AJR.*, 1994, 162.

[3]   M.R. Callstrom, C.D. Johnson, J.G. Fletcher, et al., "CT colonography without cathartic preparation: feasibility study", *Rad.* 2001; 219: pp. 693–698.

[4]   P.A. Lefere, S.S. Gryspeerdt, J. Dewyspelaere, M. Baekelandt, B.G. Van Holsbeeck, "Dietary fecal tagging as a cleansing method before CT colonography: initial results– polyp detection and patient acceptance.", *Rad.* 2002; 224: pp. 393–403.

[5]   M.E. Zalis, J.J. Perumpillichira, P.F. Hahn. Digital subtraction bowel cleansing for CT
        colonography using morphological and linear filtration methods, *in: IEEE Trans Med
        Imaging,* 23, 2004, pp. 1335–1343.

[6]   D. Chen, Z. Liang, M.R. Wax, L. Li, B. Li, A.E. Kaufman, A novel approach to
        extract colon lumen from CT images for virtual colonoscopy, *in: IEEE Trans Med
        Imaging.* 2000; 19(12): pp. 1220–1226.

[7]   P.J. Pickhardt, J.H. Choi, Electronic cleansing and stool tagging in CT colonography:
        advantages and pitfalls with primary three-dimensional evaluation, *in: AJR. Am. J.
        Roentgenol.* 2003; 181: pp. 799–805.

[8]   I.W.O. Serlie, R. Truyen, J. Florie, F.H. Post, L.J. van Vliet, F.M. Vos, Computed
        cleansing for virtual colonoscopy using a three-material transition model, MICCAI
        in: Lecture Notes in Computer Science, vol. 2879, Springer Verlag, Berlin, 2003, pp. 175–
        183.

[9]   Walter Reed Army Medical Center, Virtual Colonoscopy Center – Training,
        http://wramc.vcscreen.com/training/index.html. Accessed March 1, 2006.

[10]   P.J. Pickhardt, J.R. Choi, I. Hwang, J.A. Butler, M.L. Puckett, H.A. Hildebrandt, R.K.
        Wong, P.A. Nugent, P.A. Mysliwiec, W.R. Schindler, Computed tomographic virtual
        colonoscopy to screen for colorectal neoplasia in asymptomatic adults, *in: N. Engl. J.
        Med.* Vol. 349, 2003, pp. 2191–2200.

[11]   F.M. Vos, R.E. van Gelder, I.W.O. Serlie, et al., Three-dimensional display modes for
        CT colonography: conventional 3D virtual colonoscopy versus unfolded cube
        projection, *in: Radiology,* vol. 228, 2003, pp. 878–885.

[12]   A.H. de Vries, R. Truyen, J. van der Peijl, et al., Feasibility of automated matching of
        supine and prone CT-colonography examinations, *in: Br. J. Rad.,* vol. 79:945, 2006,
        pp. 740-744.

[13]   I.W.O. Serlie, "WRAMC polyps", http:// www.ist.tudelft.nl/ qi: group publications:
        technical reports.

[14]   S. Lakare, M. Wan, M. Sata, A. Kaufman, "3D digital cleansing using segmentation
        rays", in: T. Ertl, B. Hamann, and A. Varshney, (eds.), *in: IEEE Visualization 2000*
        (Proc. Int. Conf. Utah, USA, Oct.8-13), *ACM Press,* 2002, pp. 37–44.

[15]   R.E. van Gelder, C.Y. Nio, J. Florie, et al., Computed tomographic colonography
        compared with colonoscopy in patients at increased risk for colorectal cancer,
        in: Gastroenterology, vol. 127, 2004, pp. 41–48.

[16]   M. Macari, A. Milano, M. Lavelle, P. Berman, A.J. Megibow, "Comparison of time-
        efficient CT colonography with two- and three-dimensional colonic evaluation for
        detecting colorectal polyps", *in: Am J Roentgen,* vol. 174, 2000, pp. 1543–1549.

[17]   E. Neri, F. Vannozzi, P. Vagli, A. Bardine, C. Bartolozzi, "Time efficiency of CT
        colonography: 2D vs 3D visualization", *in: Comp Med Imag Graph,* vol. 30, 2006
        pp. 175–180.

[18]   M.E. Zalis, J.J. Perumpillichira, C. Magee, G. Kohlberg, P.F. Hahn, "Tagging-based,
        electronically cleansed CT colonography: evaluation of patient comfort and image
        readability", *in: Radiology,* vol. 239, 2006, pp. 149–159.

Chapter Seven

Visualization of Noisy and Biased Volume Data Using First and Second Order Derivative Techniques

This chapter was based upon a paper presented at IEEE Visualization 2003 in Seattle, Washington, USA [1]. Other papers that were created in cooperation with the AMC hospital in Amsterdam that investigate the problem of noise due to very low doses are [2][3].

The quality of volume visualization depends strongly on the quality of the underlying data. In CT colonography, data acquired at a low radiation dose results in a low signal-to-noise ratio. Alternatively, MRI data is acquired without ionizing radiation, but suffers from noise and a bias field (global signal fluctuations). Current volume visualization techniques often do not produce good results with noisy or biased data.

This chapter describes methods for volume visualization that deal with these imperfections. The techniques are based on specially adapted edge detectors using first and second derivative filters. The filtering is integrated into the visualization process.

The first derivative method results in good quality images but suffers from a significant localization bias. The second derivative method yields excellent surface localization, especially in highly curved areas. The filter permits fast implementations and the rendering shows minimal distortion of detail resulting in a better visualization of polyps.

7.1   Introduction

In recent years, volume visualization has become increasingly accepted and used for many applications including CT colonography. Advances in volume visualization have resulted in real-time rendering with high image quality. Obviously, the quality of visualization strongly depends on the quality of the underlying data (Figure 7.1). Unfortunately, high quality input data is not always available.

Virtual colonoscopy is a non-invasive method to inspect the colon using a CT volume of the abdominal region, instead of physically inserting an optical colonoscope. As the data is acquired by CT data acquisition, this conventionally involves a significant radiation dose. The health risk incurred by ionizing radiation prompted researchers to try low dose CT data acquisition. In CT the image quality depends on the product of X-ray tube current (mA) and exposure time (s), usually expressed in mAs. A conventional dose for CT virtual colonoscopy is about 40 mAs, a low dose is in the order of 1.5 mAs. Alternatively, MRI could be used, which is free of ionizing radiation.

Unfortunately, low dose CT data acquisition yields noisy data. In addition to an increased noise level, MRI suffers from global signal fluctuations, generally referred to as the bias field. Volume rendering of an iso-surface through a constant datavalue fails to visualize true object surfaces in such data (Figure 7.2). Noise or bias can significantly disturb the results if not properly accounted for and prohibit polyp detection in images generated by direct volume rendering methods [4], as well as by isosurface volume rendering.

A straightforward solution to reduce the noise is by low-pass filtering the input data (e.g. via a Gaussian filter). However, this may results in significant loss of image detail. A modification to non-linear Gaussian filtering yields proper visualization of fine details, but still requires a fairly high dose at 20 mAs [5]. Another method is to render the intensity values semi-transparently using ray casting [6][7] that amounts to filtering along the ray. For semi-transparent volume rendering, complex multi-dimensional transfer functions can be generated, which are based on higher-order information to differentiate between tissues and emphasize material boundaries.

Bias field removal has received considerable attention in recent years [8][9]. The build-in regularization of Gaussian derivative filters avoids amplification of the noise in exchange for a lower resolution. Alternatively, the use of modern data-driven filtering techniques based upon partial differential equations and steered filters is not considered because of their high demand on computational resources. The methods should permit a fast implementation.
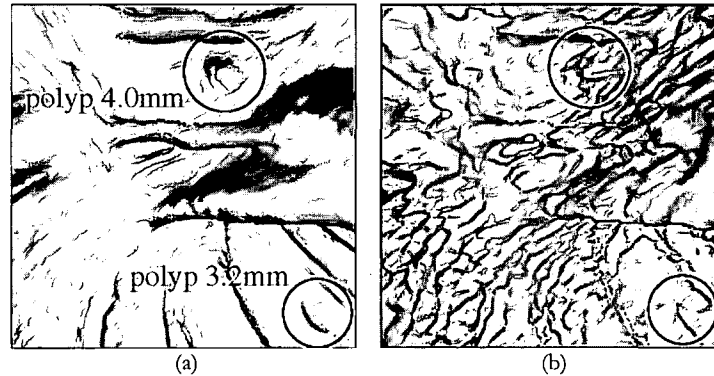
Figure 7.1   Two virtual colonoscopic images: (a) rendered using a high-dose CT volume and (b) rendered using a simulated low-dose CT volume.
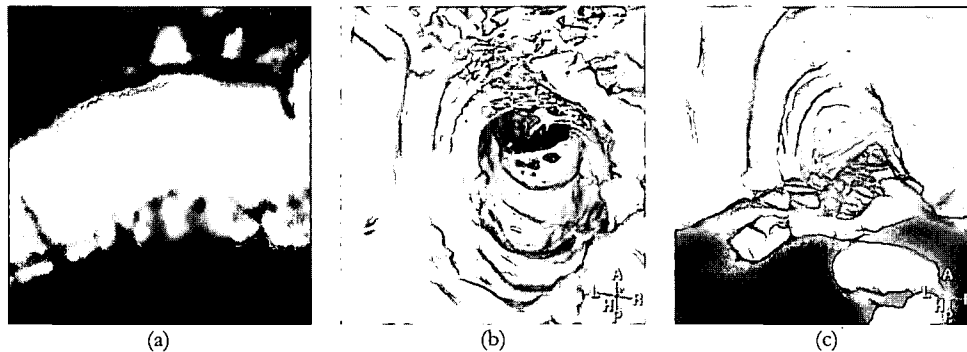

Figure 7.2   Virtual colonoscopic images based upon MRI-data. (a) Image values in a slice demonstrate a bias field. Two isosurface visualizations show that one isosurface cannot segment the colon in such data. Either the (b) bottom or the (c) top of the colon is segmented.

In our approach, the data are preprocessed using first- and second-derivative operators to eliminate the effect of bias. Regularization of image derivatives reduces the noise. As the main object of interest is the interior surface of the colon, a simple iso-surface rendering suffices to visualize the surface. Since no semi-transparent rendering is used, complex transfer functions are not needed. The target of the methods is good visualization. Hence, the visual impression is leading in the evaluation of the methods.

The methods that are presented are based on existing first- and second derivative filters. The novelty of this work is in the modification of these methods to visualize noisy and biased volumetric colonography data described in section 7.2.

7.2   Methods

Many volume-rendering techniques act by directly applying a transfer function to the measurement values (i.e. values within a certain range are rendered translucently, the others opaquely). We propose two methods that are based on first and second derivatives respectively. Henceforth, we will refer to these methods as first-order and second-order methods.

7.2.1   First-order method based on the Canny edge detector

Canny [10] created a method to detect the position of 2-D straight step edges in the presence of additive Gaussian noise. He proposed to compute the gradient magnitude using Gaussian derivatives to approximate the filter that optimizes the conflicting constraints of a high signal-to-noise ratio (requiring a large filter) and a good localization (requiring a small filter). Two problems occur when applying Canny to visualize noisy images for colonography. The first problem is that applying Canny results in a pixel-based segmentation showing jaggy edges. The second problem is that visualization of an isosurface through the top of ridges requires a robust surface-normal estimator.

The solution to the first problem is an extension of the Canny edge detector and transforms the CT volume into a normalized gradient magnitude volume in three steps.

1.   *Gradient magnitude calculation.* Gaussian first-derivative filters with a scale $\sigma_{x,y,z}$ are used to obtain the gradient at each voxel suppressing the effect of spatially uncorrelated noise.

2.   *Non-maxima suppression.* The edge is located at the maximum of the gradient magnitude $\|\nabla I\|$ To permit a fixed threshold t to visualize the local maximum of the gradient magnitude across the edge, the gradient magnitude $\|\nabla I\|$ is divided by its local maximum along the gradient (across the edge):

$$E = \left(\frac{\|\nabla I\|}{\max \|\nabla I\|} \right)^{p} \tag{7.1}$$

The search window around a voxel extends 2.5 mm in both directions along a path in the gradient direction w. The size of the search window corresponds to the size of the smallest relevant detail. Due to noise, the normalized gradient magnitude E might not exactly be one at the top of the ridges. In addition, sampling along a ray for visualization may lead to missing the top of ridges if a threshold $t = 1$ is used (Figure 7.7a). Using $t < 1$ results in a small localization error. The responses E are raised to a power p for sharpening to decrease the distance between the true edge and approximation by t (equation 7.1). The maximum sharpening-value of p is steered by the sampling that is required to properly represent the ridges in the data (Nyquist).
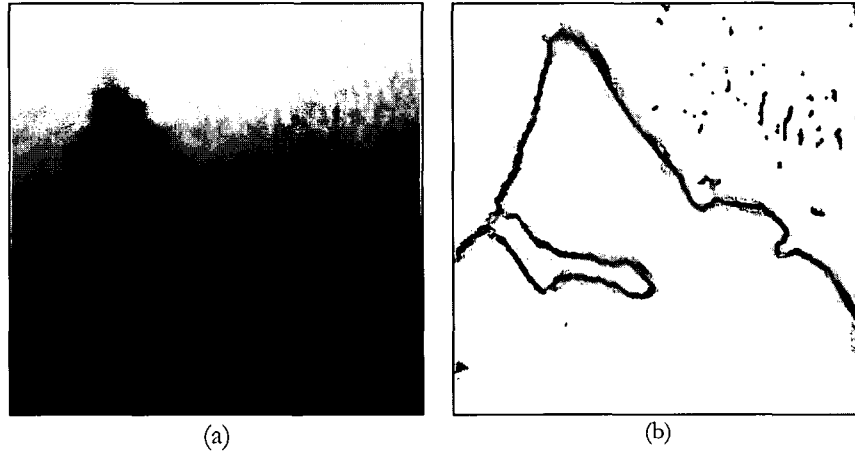
Figure 7.3   (a) Part of a slice in an MRI volume showing the bias from top to bottom and (b) the normalized gradient magnitude volume.

Assuming Gaussian blurred step-edges, the normalized gradient magnitude is modeled by the Gaussian function. It is normalized such that $g_n(0;\sigma)=1$. When expressing the scale of the ridges after filtering the data by σ in unit voxel, then the maximum allowed sharpening becomes $p=\sigma^2$; this results in a ridge with a scale of exactly one voxel.

$$g_n(x;\sigma) = \exp\left(-\left(\frac{x}{\sigma\sqrt{2}}\right)^2\right) \tag{7.2}$$

3. *Hysteresis Thresholding* is applied to all voxels after non-maxima suppression. The maximum threshold is chosen to include edges not generated by noise. The lower threshold is set higher than the noise level and is used to close the surface.

The second problem with using Canny for colonography is that visualization of the ridge requires a robust surface-normal estimator. After Gaussian first derivative filtering, the surface to be visualized is a thin sheet in 3-D (one can visualize it as a ridge in 2-D, Figure 7.3). The Gaussian derivative is not a good estimator of the surface normal. A solution is obtained using the Structure Tensor [11]. The Structure Tensor (F) is defined by the dyadic product of the gradient vector with its transpose

$$\overline{F} = \overline{(\nabla I)(\nabla I)^t} = \begin{bmatrix} \overline{I_x I_x} & \overline{I_x I_y} & \overline{I_x I_z} \\ \overline{I_x I_y} & \overline{I_y I_y} & \overline{I_y I_z} \\ \overline{I_x I_z} & \overline{I_y I_z} & \overline{I_z I_z} \end{bmatrix} \tag{7.3}$$

where I_x, I_y and I_z denote the partial derivatives of the data $I(x,y,z)$ that are obtained by a convolution with a Gaussian derivative filter. Another Gaussian is used to regularize the Structure Tensor (depicted by the overhead bar). The structure tensor $\bar{F}$ is rewritten in

$$R^T \bar{F} R = \begin{bmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_1 & 0 \\ 0 & 0 & \lambda_1 \end{bmatrix} \tag{7.4}$$

with $R = \begin{bmatrix} \hat{v}_1 & \hat{v}_2 & \hat{v}_3 \end{bmatrix}$ where λ_i are the eigenvalues and $\hat{v}_i$ the corresponding eigenvectors of unit length. The corresponding eigenvalues denote the average gradient energy along $\hat{v}_i$. Note that at a surface point, the first eigenvector points along the surface normal, the second and third eigenvectors span the surface. Starting from the camera position, positions on a ray before the first local maximum are rendered with an opacity set to zero and after the first local maximum with an opacity set to one. The orientation of $\hat{v}_1$ and the ray are used as parameters of Phong's shading model.

The first-order method has four parameters:

1. The threshold on the normalized gradient t: chosen as high as possible, but slightly smaller than 1 to reduce the localization error.

2. The (sharpening) factor p: chosen as large as the sampling criteria permits: $p = \sigma^2$.

3. The scale of the Gaussian first-derivative filters $\sigma_{x,y,z}$: it is chosen small, to correspond as much as possible to the scale of the smallest relevant detail (i.e. the smallest features of interest).

4. The degree of tensor smoothing σ_T: chosen small to visualize the surface as accurately as possible as well as to suppress influence of normals from neighboring surfaces. Remember that the normals are only used in Phong's shading model.

A clear advantage of the method is that it does not depend on the datavalue of a specific iso-surface. The parameter influence is shown in section 7.3.

7.2.2   The second-order method based on the PLUS operator

Another method for boundary detection that does not depend on a specific isosurface is to use the zero-crossing after applying a second derivative operator. Common second derivative edge detectors are the Marr-Hildreth detector, defined as the Laplacian of Gaussians (LoG), and the second derivative in the gradient direction ($SDGD$) [12]:

$$\begin{aligned} LoG(I) &= I_{xx} + I_{yy} + I_{zz} \\ SDGD(I) &= I_{ww} \end{aligned} \tag{7.5}$$
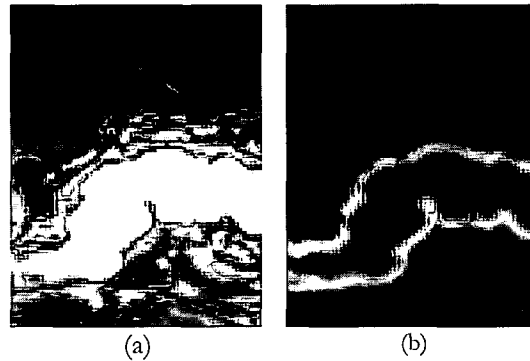
(a)   (b)

Figure 7.4   (a) Part of a slice in an MRI volume and (b) the result of the PLUS operator.

in which I_{xx}, I_{yy}, I_{zz} and I_{ww} are second derivatives of the image f calculated via Gaussian derivatives; x, y and z are the main coordinates and w is the direction of the gradient vector.

With curved object boundaries there is always a trade-off between noise suppression and surface displacement. A larger filter size reduces noise, but at the expense of increased surface displacement. First order derivative filters, such as the Gaussian derivative filter of the previous section, cause a surface displacement that is always in the direction of positive curvature. Unfortunately, it is generally impossible to compensate for the first-derivative surface displacement.

The *LoG* and *SDGD* filters have opposite surface displacements. The operators may be combined [13][14] to reduce the effect as the opposite-signed displacements cancel each other to a certain extent. The combined edge detector is generally referred to as the PLUS operator:

$$PLUS(I) = I_{xx} + I_{yy} + I_{zz} + I_{ww} \qquad (7.6)$$

Figure 7.4 presents MRI data with contrast fluid injected into the colon. The application of the PLUS operator results in a new volume (Figure 7.4) that is referred to as the PLUS volume. A transition visible is in Figure 7.4b from negative values (black) to positive values (white); in between is the zero crossing.

The PLUS *volume* is visualized by a modified ray casting algorithm. Initially, those voxels are selected that have a gradient magnitude higher than a certain threshold. Hereby, a coarse region of interest is identified. This threshold is applied to reduce noise. We have found that the value of the threshold is not very critical to the result. Subsequently, the algorithm renders the samples on a ray with an opacity set to zero before the first zero crossing and with an opacity set to one after it. The surface normals are calculated by Gaussian first derivative filters in the PLUS volume. This is possible because the gradient vectors around the transition do not have opposite directions.

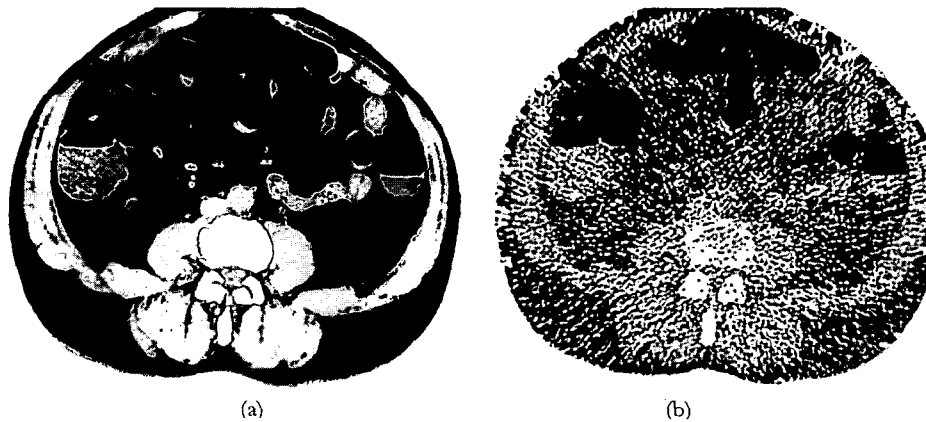(a)   (b)

Figure 7.5   (a) Slice of a high dose 40 mAs CT volume (b) and a simulated low dose 1.5 mAs CT
volume.

The second order method has two parameters:

1. the scale $\sigma_{x,y,z}$ of the Gaussian second derivative filters (Equation 7.2);

2. the width of the Gaussian first derivative filters σ_n (used to calculate surface
    normals in the PLUS volume).

Only the widths of the second derivatives affect the localization of the surface. The
second derivative parameter balances noise against surface detail. σ_n merely affects the
shading. Therefore, σ_n should be as small as possible to represent the surface as
accurately as possible.

7.3   Results

The first and second order methods were tested both visually and quantitatively. The
implementation we used preprocesses the CT volume into the normalized gradient
volume (Figure 7.3b) and the PLUS volume (Figure 7.4b). This results in a fully
automatic preprocessing step and a real-time ray casting.

7.3.1   Data acquisition

A CT volume was randomly selected from an ongoing virtual colonoscopy study at the
Academic Medical Center in Amsterdam (the Netherlands) in 2003. The original CT
scans were acquired at a level of 40 mAs to serve as a reference. The low dose CT scans
were generated from these normal dose scans.

Low-dose CT scans were reconstructed at 1.5 mAs by modifying the raw transmission
data of each spiral CT scan using a simulation technique. The raw transmission
measurements were replaced by a realization of the Poisson process given a scaled
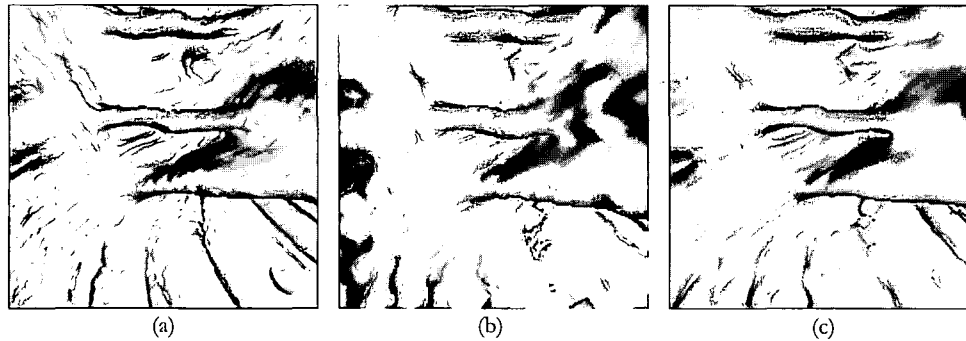
Figure 7.6   (a) Iso-surface volume rendering of high dose CT data. (b) Rendering of low dose CT data using the first order method and (c) second order method.
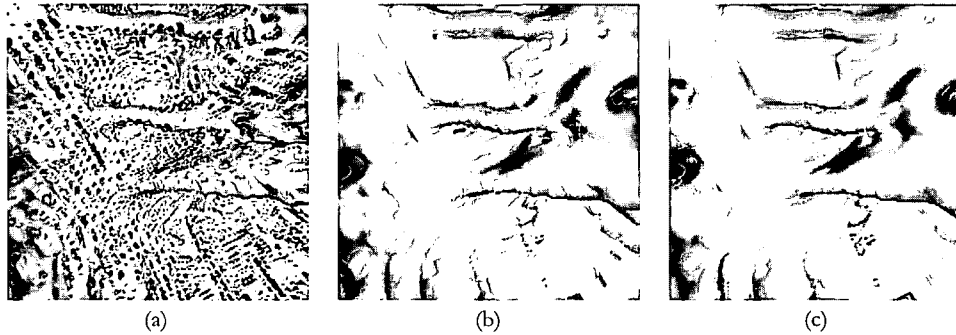

Figure 7.7   (a) Iso-surface volume renderings using the first-order method at (a) $t = 0.9$ that is too high creating holes in the surface, (b) using $t = 0.7$ and (c) $t = 0.5$ that is too low creating excessive surface displacement (all images are generated using $\sigma_{x,y,z} = 1.0$ mm and $\sigma_T = 1.5$ mm).

version of the 40 mAs datavalue as the expected value. Note that the variance and the mean values of Poisson distribution are identical. Typically, the volume consisted of 512x512x256 voxels at a resolution of 0.6x0.6x1.2 mm^3. Example images are shown in Figure 7.5 (the noise is correlated).

An MRI volume was randomly selected from an ongoing study at the Academic Medical Center in Amsterdam (the Netherlands) in 2003. The data were acquired via a Siemens 1.5 Tesla scanner. A typical volume consisted of 256x256x128 voxels sized 1.2x1.2x2.0 mm^3.

The first and second order methods were implemented on an experimentally enhanced version of the Philips EasyVision workstation.

7.3.2   Choice of parameters

The 3-D visualizations of the high-dose and low-dose CT data have been made with identical viewing parameters (Figure 7.6,7,8,9,10). Figure 7.6a shows an iso-surface volume rendering of the high dose CT volume: we consider this the reference standard.

A 'first-order' method with optimal parameter settings (t = 0.5, $\sigma_{x,y,z}$ = 1.2 mm, σ_T = 1.0 mm, p = 8) is illustrated in Figure 7.6b. A 'second-order' method with optimal parameter settings ($\sigma_{x,y,z}$ = 1.2 mm, σ_n = 1.0 mm) is shown in Figure 7.6c. The parameter settings are discussed below. Notice that the optimal first and second order images are rendered with the same filter sizes.

The effects of significant parameter adjustments in the first order rendering are depicted in Figure 7.7 and Figure 7.8. First, the results of varying the surface localization parameter t are shown. Using this parameter each row in Figure 7.8 shows the result as a single parameter is varied (from top to bottom: $t, \sigma_{x,y,z}, \sigma_T$, p respectively). The effects of significant parameter adjustments in the second order rendering are depicted in Figure 7.9. The optimum parameter setting was determined based on visual inspection of these renderings.

To study the effect of the first-order and second-order method on bias fields, two camera positions are selected in an MRI volume. Figure 7.10a shows the result for small signal fluctuations and noise, Figure 7.10b for the other camera position with a large bias. The parameter settings were again found experimentally. Both images were rendered with the same parameter settings. As expected, the iso-surface volume rendering shows holes and surface fragments, due to the signal fluctuations (Figure 7.10b, zero-order).

7.3.3   Localization accuracy

To measure the localization accuracy, we created a reference image for the CT data by iso-surface volume rendering of the (normal) 40 mAs data. The threshold selecting the iso-surface was determined by maximizing the sum of gradient magnitudes over all threshold positions. The surface displacement between the reference surface and the rendered surface was defined as the distance along viewing rays between the reference surface and the rendered surface respectively. We have opted for this approach to illustrate the localization accuracy from a typical viewpoint. A reference surface cannot be easily defined for MRI data due to lack of a proper standard. For that reason we have restricted the quantitative analysis to CT data.
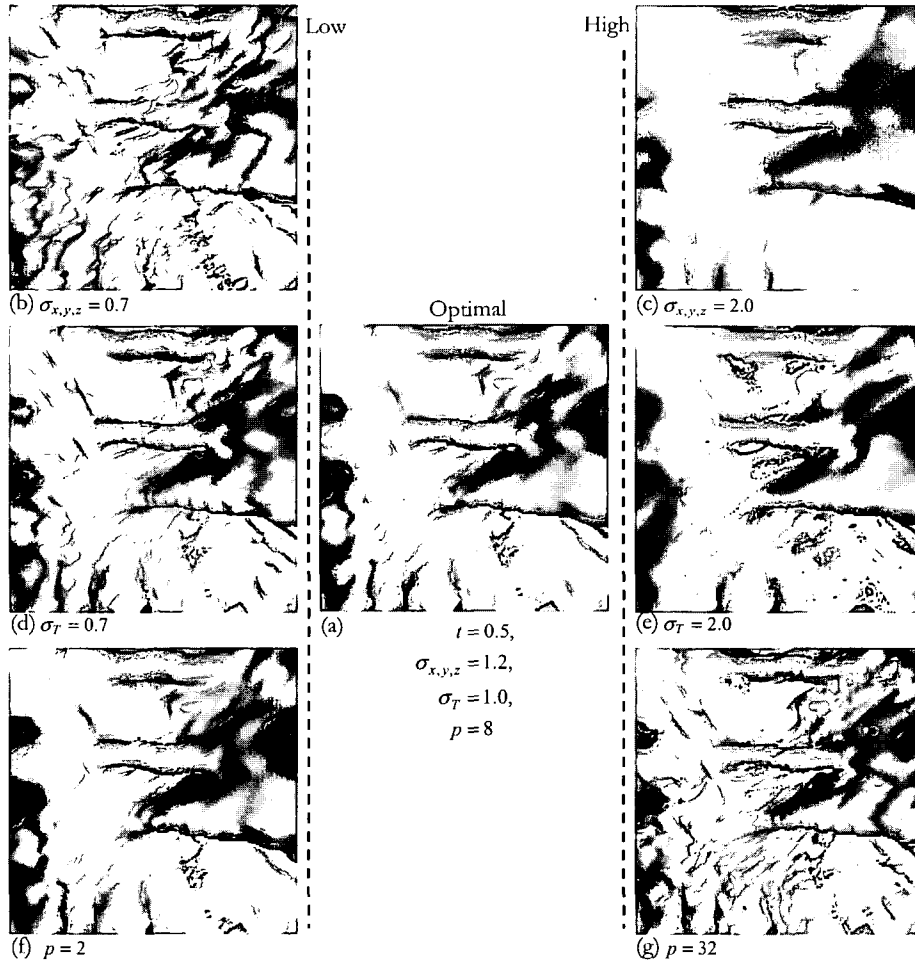
Figure 7.8   Parameter influence on first-order method (all sigma's in mm). (a) Optimal parameter settings. (b) Low $\sigma_{x,y,z}$, resulting in a noisy surface, (c) high $\sigma_{x,y,z}$ resulting in a smooth surface at the expense of detail. (d) Lower σ_T resulting in staircase effects, (e) higher σ_T resulting in false normals. (f) Low p resulting in surface displacement and (g) high p resulting in tiny holes in the surface.
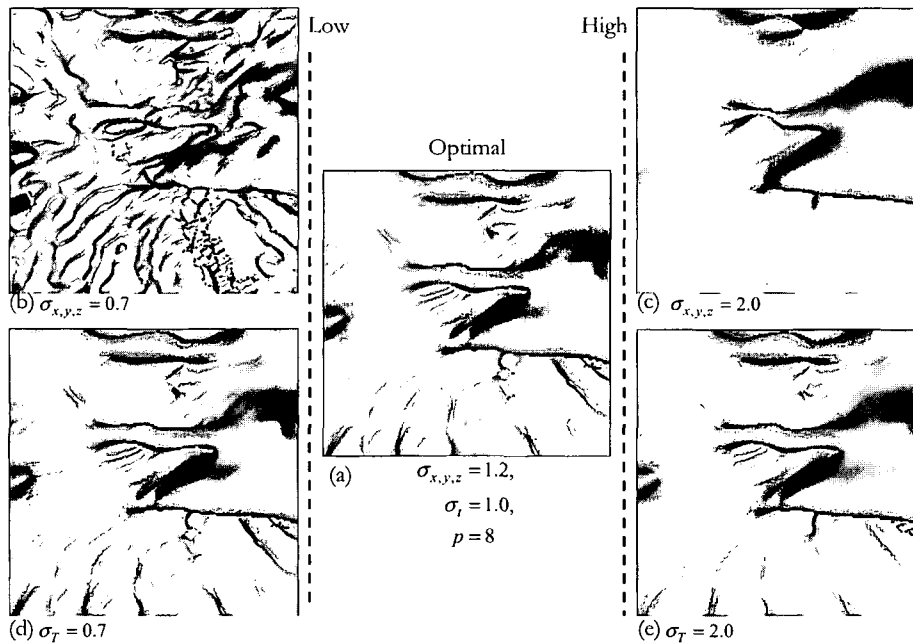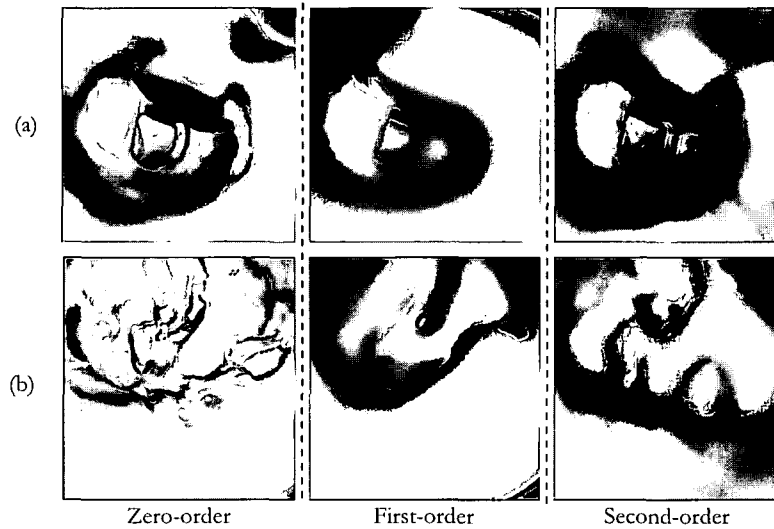
Figure 7.9   Parameter influence on second-order method (all sigma's in mm) (a) Optimal parameter settings. (b) Low $\sigma_{x,y,z}$, resulting in a noisy surface, (c) high $\sigma_{x,y,z}$ resulting in a smooth surface at the expense of detail. (d) Lower σ_T resulting in noisy normals, (e) higher σ_T resulting in loss of detail.

Figure 7.10   (a) Result on MRI in area of high contrast and low fluctuation. (b) Result on MRI in area of low contrast and high fluctuation. (left) Iso-surface volume rendering. (middle) Ray casting of normalized gradient volume results in smoothing. (right) Ray casting of PLUS volume.
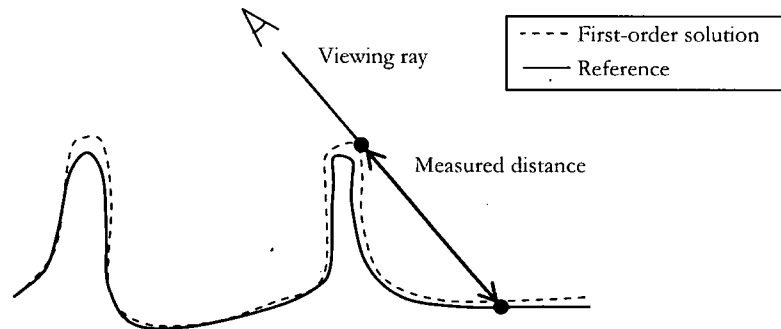
Figure 7.11   Example of a measurement artifact in the surface displacement.
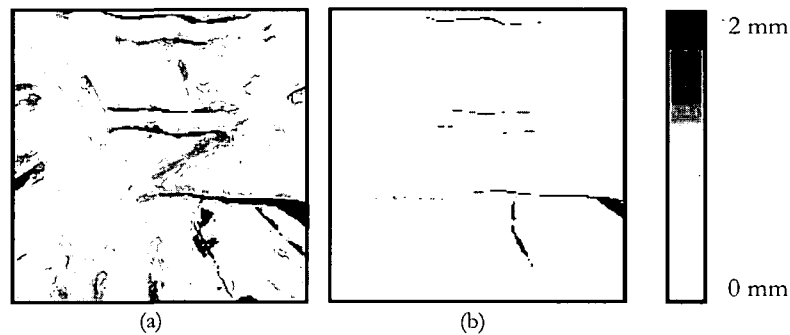
Figure 7.12   A greyscale image showing the degree of the surface displacement using the (a) first order method and (b) second order method.

A small percentage (approximately 0.5%) of very large errors was caused by small variations in fold thickness (Figure 7.11). Such artifacts were eliminated from the analysis by removing the errors larger than 10 mm.

The mean error in the first-order method was 0.6 mm towards the camera. The errors were within 3 mm in 95% of the cases. The second-order method using the PLUS-operator had a mean localization error of 0.19 mm; the errors were below 1.0 mm in 95% of the cases.

Figure 7.12 depicts the surface displacement for both methods. As expected, the first-order method has a significantly larger localization error (Figure 7.12a) than the PLUS-based second-order method (Figure 7.12b).

7.3.4   Performance

The total rendering time is composed of preprocessing time and the time needed for ray casting. The preprocessing for the first order method takes approximately five minutes

(which includes calculating the Structure Tensor for the whole volume) and the ray casting is real time. The whole data set was preprocessed for the second order method in approximately five minutes and the ray casting is real time. These numbers were obtained for the CT data using an AMD 2.6 GHz computer with 2 GB of internal memory. For other data volumes the preprocessing time is scaled in proportion with volume size.

7.4   Conclusion

This chapter reports on progress in volume visualization of data with severe noise levels and bias fields. We presented methods for volume rendering based on first and second derivatives of the input volumes.

First, a modified version of the Canny edge detector was applied, to suit application in virtual colonoscopy. The modification included a method for robust surface normal estimation. This first order technique effectively deals with noise, but introduces a surface displacement.

Next, a visualization method was presented incorporating the PLUS operator that combines the second-derivative in gradient direction and the Laplacian of Gaussians. This operator yields accurate surface localization, and it is insensitive to global signal fluctuations. The visual results show good surface quality and preservation of important shape details.

The methods were tested on simulated low dose CT data (derived from regular dose data). The new techniques were compared to iso-surface volume renderings from the regular dose data (low noise, no global fluctuations). The comparisons showed good agreement in visualization of significant details, such as polyps diagnosed by radiologists.

The results presented suggest that low dose CT and MRI may become feasible as techniques for virtual colonoscopy in the future. This may be a step towards the use of virtual colonoscopy in large-scale clinical screening for early diagnosis of colonic polyps.

7.5   Acknowledgments

7.6   References

[1]   M.P. Persoon, I.W.O. Serlie, F.H. Post, R. Truyen, F.M. Vos, "Visualization of Noisy and Biased Volume Data Using First and Second Order Derivative Techniques", *IEEE Visualization*, pp. 379-385, 2003.

Chapter Eight

Conclusion

This thesis addresses visualization and image processing methods that solve two problems that are critical to the success of CT colonography; a non-invasive method to find the precursors of colon cancer: (8.1) *first*, visualization methods for displaying the complete colon-surface after a perspective transformation; (8.2) *second*, segmentation methods for extracting the colon surface from the data, which can cope with air and tagged intraluminal remains that border on soft tissue. Especially, at intersections between three materials that are common in CT colonography, image processing methods are developed to successfully segment the colon. (8.3) In general, the model-based techniques that have been created are a robust alternative to the standard computed Gaussian derivatives across image transitions. (8.4) Being able to focus on specific material transitions has proven to be useful in CT colonography and other applications. The LH-histogram (derived from the material transition in data/ gradient magnitude space) is a valuable tool to classify voxels into types of transitions in the presence of a bias field such as present in MRI data.

8.1   Unfolded cube display

The *unfolded cube display* [1] yields an omnidirectional projection of the inside of the colon. With the viewpoint at the center, on each face of the cube an image is projected with a 90° viewing angle. The six images are unfolded onto a 2-D image to examine the six faces simultaneously. A series of unfolded cubes obtained along the colon trajectory present the bowel surface to the physician in a 'panoramic' way. This facilitates the physician to look in all directions at the same time, as well as providing an intuitive view for thorough and rapid inspection. The conventional display involves two antero- and retrograde view directions along the colon trajectory that are examined as two series. A disadvantage is that this display requires a large viewing angle of typically 120° to expose most of the colon surface, which involves excessive artifacts. A second disadvantage with the antero- and retrograde viewing directions is that the reader has to switch from one image to another and cannot stay 'immersed' in one modus of viewing.

From the clinical *evaluation* [2] it follows that the unfolded cube display substantially reduces the problem of blinded (occluded) areas. Thirty patients were included. On average, the conventional antero- and retrograde display showed 93.8% of the colon surface and the unfolded cube display showed 99.5% of the colon-surface. In addition, the mean evaluation-time when using the unfolded display (20 minutes) was significantly shorter than using the conventional display (36 minutes). Using the conventional antero- and retrograde (perspective) displays, the radiologists had to manually rotate the camera to get a view of blinded areas, never being confident of having seen everything.

A prototype has been built to examine colon-data and measure lesions using the unfolded cube as a primary display and the axial 2-D slices for secondary 2-D reading. Part of the prototype has been integrated into the commercial EasyVision software (Philips Medical Systems Nederland B.V.). Rendering of the unfolded cube display has been made real-time such that preprocessing and storing a series of images is not necessary.

8.2   Electronic cleansing

In CT colonography, three types of materials occur in the colon region: air, soft tissue and intraluminal remains that have been tagged using an oral contrast medium. Hence, a measured CT value arises due to a combination of these three materials. Electronic cleansing [3] is an image processing method that first estimates the percentage of materials that contributes to a CT value and second replaces the (partial-volume) contribution of tagged material by air. The processed CT volume is examined as if the data do not contain intraluminal remains at all. It solves the problem of (tagged) remains blocking the view at the colon surface when using a 3-D display mode such as the unfolded cube. Several previous papers did report advances regarding electronic cleansing [4][5][6][7]. Zalis et al. [4] indicated that it will be essential to ascertain how

artifacts still present in successfully tagged, electronically cleansed data will affect the observer's performance. The complex nature of electronic cleansing is observed by Pickhardt as well who identifies pitfalls in reading electronically cleansed data [7]. A common artifact, described in the latter reference, occurs at locations where the air-fluid transitions border on the colon wall, resulting in transitions between three materials. The methods that are described in this thesis solve all problems associated with the aforementioned problems in electronic cleansing. In addition, in the presence of inhomogeneities that for example occur when using tagged bowel contents, the classification is improved due to local parameter estimation.

From a *clinical evaluation* [8] of our electronic cleansing method it is concluded that it enables a more effective and efficient examination of the colon that is superior to the unfolded cube display on uncleansed data. Nineteen patients were included from publicly available study data by the Walter Reed Army Medical Center [9]. On average, the cleansed surface fraction, the surface area that is blinded in a 3-D display when not using electronic cleansing, was almost 30%. In addition, the mean evaluation-time decreased from approximately 19 to 12 minutes when using the 3-D display and the electronic cleansing method. Observers were asked to rank the assessment effort, based on artifacts visible in the rendering with and without cleansing. They gave a significantly lower effort when using the electronically cleansed data. This confirms the predicted small sizes, based on simulated data, of artifacts at locations where the air-fluid transitions border on the colon.

All images of polyps larger than 6 mm confirmed with optical colonoscopy partly or fully residing in tagged intraluminal remains in the WRAMC data [9] (48 images in total) have been electronically cleansed. The clinical evaluation shows that the conspicuity of these electronically cleansed polyps, when displayed with the unfolded cube display, is the same as the conspicuity of polyps residing in air (50 images in total)

8.3   Edge-processing

In general, the models that have been created, describe a robust alternative to the standard computed Gaussian derivatives across image transitions.

First, we have shown that the erf-function is an excellent model to describe a two-material transition in CT data. Second, we have extended the erf-model to relate the derivatives across the edge: the 'arch'-model. The function is named 'arch', because the line defined by the magnitude of the first derivative as a function of material fraction reminds of an arch or catenary. In addition, the arch-function is extended to relate the derivatives across the three-material transition: the 'parachute'-model. The function is named 'parachute', because the surface created by the first derivative as a function of material fractions reminds of a parachute. The projection onto the model of noisy image value and derivatives combines measurements to reduce the noise without having to increase the aperture of the operator.

With regard to CT colonography, the models are used to estimate the fractions of materials across transitions that contributed to a CT value. Outside the scope of CT colonography, our methods have been applied as well. The two-material transition model is fitted to the data and two parameters of the model, the expected values at opposite sides of the transition, are used to separate voxels into types of transitions. The expected values at opposite sides of the transition are summarized in the *LH*-histogram. It has been shown that the *LH*-histogram allows easier identification of boundaries than using the arches in the domain of scalar value and gradient magnitude as described by Kindlmann et al. [10]. Our *LH*-histogram has been proven a valuable tool to classify voxels into types of transitions in the presence of a bias field such as present in MRI data [11]. In addition, 2-D transfer functions for visualization have been based on the *LH*-histogram by selecting relevant areas and by assigning color and opacity to these areas. The *LH*-histogram is also successfully applied to region growing of the boundaries, based upon the assumption that it is undesirable to grow outside the current boundary, into another boundary, into areas of constant intensity, or into small noisy boundaries [11].

8.4   References

[1]   I. W. O. Serlie, F. M. Vos, R. van Gelder, J. Stoker, R. Truyen, F. Gerritsen, Y. Nio, F. H. Post, "Improved Visualization in Virtual Colonoscopy Using Image-Based Rendering", *Data Visualization 2001, proc. of the Joint Eurographics and IEEE TCVG Symposium on Visualization*, 2001, pp. 137–146.

[2]   F.M. Vos, R.E. van Gelder, I. W. O. Serlie, J. Florie, C.Y. Nio, A.S. Glas, F.H. Post, R. Truyen, F.A. Gerritsen, J. Stoker, "Three-dimensional display modes for CT colonography: conventional 3D virtual colonoscopy versus unfolded cube projection", *Radiology*, vol. 228 no. 3, 2003, pp. 878–885.

[3]   I.W.O. Serlie, R. Truyen, J. Florie, F.H. Post, L.J. van Vliet, F.M. Vos, "Computed cleansing for virtual colonoscopy using a three-material transition model", *MICCAI*, *LNCS*, vol. 2879, 2003, pp. 175–183.

[4]   M.E. Zalis, J.J Perumpillichira, P.F. Hahn, Digital subtraction bowel cleansing for CT colonography using morphological and linear filtration methods. *IEEE Trans Med Imaging*, vol. 23, 2004, pp. 1335–1343.

[5]   D. Chen, Z. Liang, M.R. Wax, L. Li, B. Li, A.E. Kaufman. „A novel approach to extract colon lumen from CT images for virtual colonoscopy", *IEEE Trans Med Imaging*., vol. 19(12), 2000, pp. 1220–1226.

[6]   S. Lakare, M. Wan, M. Sata, A. Kaufman, "3D digital cleansing using segmentation rays", T. Ertl, B. Hamann, and A. Varshney, (eds.), *IEEE Visualization 2000, ACM Press*, 2002, pp. 37–44.

[7]   P.J. Pickhardt, J.R. Choi, I. Hwang, J.A. Butler, M.L. Puckett, H.A. Hildebrandt, R.K. Wong, P.A. Nugent, P.A. Mysliwiec, W.R. Schindler, "Computed tomographic virtual colonoscopy to screen for colorectal neoplasia in asymptomatic adults", N. Engl. J. Med., vol. 349, 2003, pp. 2191–2200.

[8] I.W.O. Serlie, A.H. de Vries, C.Y. Nio, A. Heutinck, R. Truyen, F.H. Post, F.A. Gerritsen, L.J. van Vliet, J. Stoker, F.M. Vos, The efficiency of a cleansing algorithm for CT colonography based on a three material transition model, submitted.

[9] Walter Reed Army Medical Center, Virtual Colonoscopy Center – Training, http://wramc.vcscreen.com/training/index.html. Accessed March 1 2006.

[10] G. Kindlmann and J.W. Durkin, "Semi-automatic generation of transfer functions for direct volume rendering", *Proc. of IEEE Symposium on Volume Visualization*, 1998, pp. 79–86.

[11] P. Sereda, A.V. Bartroli, I.W.O. Serlie, and F.A. Gerritsen, "Visualization of boundaries in volumetric data sets using LH histograms", *IEEE Trans. Vis. Comput. Graph.*, vol. 12, 2006, pp. 208–218.

[12] M.P. Persoon, I.W.O. Serlie, F.H. Post, R. Truyen, F.M. Vos, Visualization of Noisy and Biased Volume Data Using First and Second Order Derivative Techniques, *IEEE Visualization*, 2003, pp. 379–385.

Summary

In this thesis visualization and image processing methods are proposed that solve problems that are critical to the success of CT colonography; a non-invasive method to find the precursors of colon cancer.

Chapter 1 introduces the field of CT colonography. Traditionally, the diagnosis is based upon endoscopy of the digestive system. CT colonography provides a non-invasive alternative. Fecal tagging is introduced to be able to discriminate between tissue and intraluminal remains. However, these remains may block the 3-D view of the colon wall. To solve this problem, image processing techniques were created to enable a 3-D display for CT colonography.

Chapter 2 explores basic aspects of CT imaging for colonography: the noise characteristics, the spatial resolution and the model that describes the CT value across a material transition. We show that this is a valid model for the purpose of our research.

Chapter 3 presents a method for displaying the complete colon-surface as well as its clinical evaluation. Many existing visualization methods cannot display the areas between the folds of the colon. To solve this problem we present an imaging technique that projects six overview directions onto the six sides of a cube. Our unfolded cube display significantly increases the surface in view and reduces the evaluation time with respect to the conventional method.

Chapter 4 introduces a method to approximate the surface of an object such as the colon surface. Existing methods use a surface through constant data value to approximate the object-surface of for example the colon. In 3-D colon imaging, this leads to fluid blocking the view of the colon surface. To solve this problem we propose using iso-material fractions. It allows automatic segmentation even if objects connect to more than one material. A robust alternative to the standard computed Gaussian derivatives across two-material image transitions is discussed. The method employs the arch function, which relates the nth Gaussian derivative to the data value and hence the material fractions across transitions. This alternative is translation-, rotation- and scale-invariant.

Chapter 5 deals with three-material transitions. Existing segmentation-methods of the colon surface may leave artifacts at locations where three materials meet. To solve this problem the two-material transition model that is introduced in chapter 4 is extended to three-material transitions. Surfaces through constant material fraction simplify automatic surface segmentation and artifacts that typically occur at locations where three materials meet are suppressed.

Chapter 6 reports on a clinical evaluation of electronic cleansing. First it is investigated if electronic cleansing for 3-D visualization in CT colonography affect the conspicuity of polyps that are partly or fully covered by tagged remains. In addition, two experienced observers evaluated patient data using the traditional method and the CT colonography that is improved with the new image processing methods. It is concluded that electronic cleansing yields an identical conspicuity of polyps surrounded by air as polyps partly or completely residing in tagged material. Furthermore, the proposed cleansing method sustains a lower evaluation time, lower assessment effort and larger observer confidence than a conventional evaluation method without cleansing.

Chapter 7 describes methods for volume visualization of noisy data and data with global signal fluctuations. In CT colonography, data acquired at a low radiation dose results in a low signal-to-noise ratio. Alternatively, MRI data is acquired without ionizing radiation, but suffers from noise and a bias field (global signal fluctuations). Volume visualization techniques are described that are based on specially adapted edge detectors using first and second derivative filters.

Iwo Serlie, Delft, October 2007

Curriculum vitae

Iwo Serlie was born in Woerden, on the 6th of July in 1972. In 1993 he obtained the VWO diploma at the Christelijk Lyceum Veenendaal and started a study in Computer Science at the Delft University of Technology. After graduating in the Computer Graphics and CAD/CAM Group, he received his M.Sc. degree in 2001.

In 2001 he started a Ph.D. – continuing in the field of medical visualization and image processing - at the Pattern Recognition Group (currently the Quantitative Imaging Group) of the Delft University of Technology under supervision of prof.dr.ir Lucas van Vliet, dr. Frans Vos and ir. Roel Truyen from Philips Medical Systems Nederland B.V. The research project titled "Electronic cleansing for Visualization in CT Colonography" was done in collaboration with Philips Medical Systems Nederland B.V. and the Academic Medical Center of Amsterdam providing a rapid prototyping environment and up to date medical data.

In 2006 he became a post-doc at the Technische Universiteit Eindhoven, in collaboration with Philips Medical Systems Nederland B.V. to be able to continue doing research in the medical field.
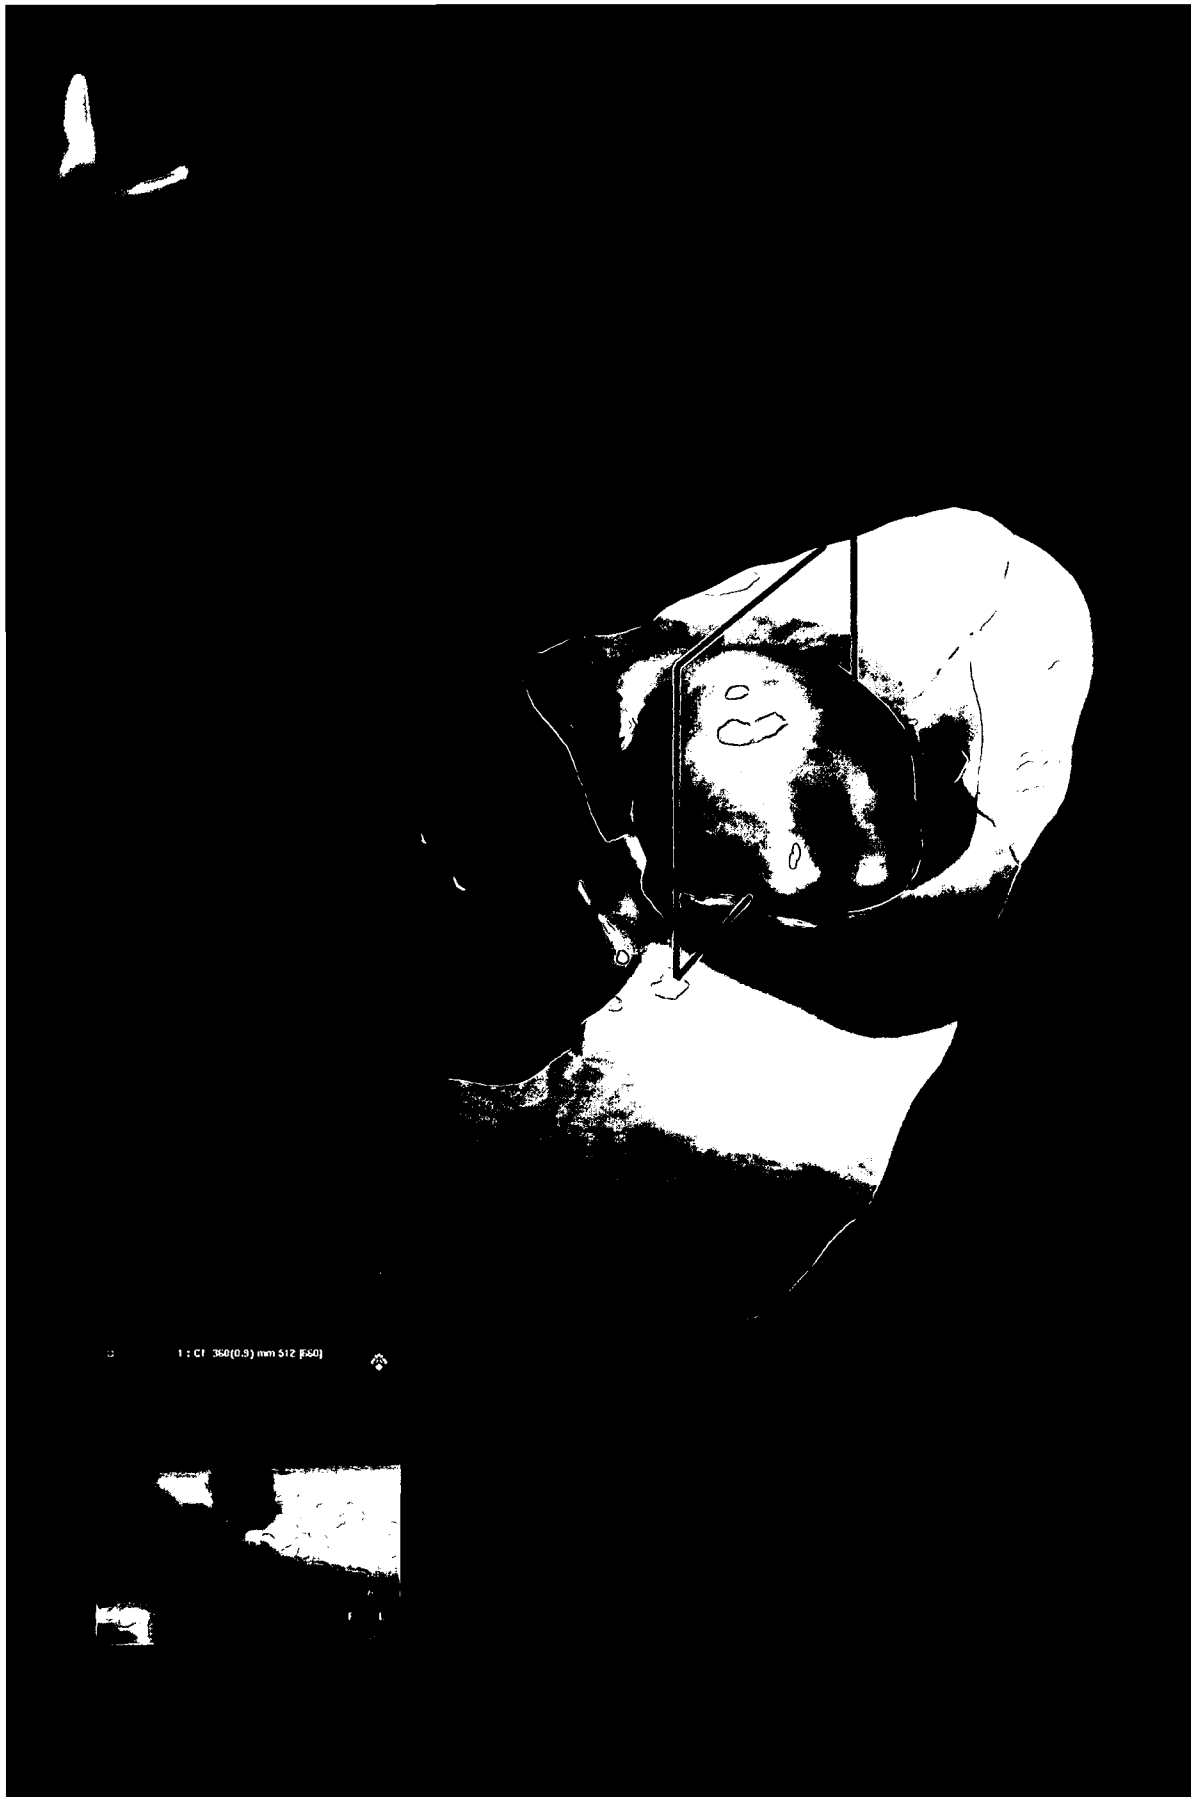
Dankwoord

1 : CT 360(0.3) mm 512 [550]