

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

Iwo W. O. Serlie, Frans M. Vos, Roel Truyen, Frits H. Post, and Lucas J. van Vliet

Abstract—A fully automated method is presented to classify 3-D 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, noncubic 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 digital cleansing for virtual colonoscopy.

Index Terms—Anisotropic Gaussian point spread function (PSF), object segmentation, partial volume effect (PVE), transfer function for visualization, voxel classification.

I. 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

Manuscript received February 15, 2007; revised August 8, 2007. This work was supported in part by Philips Medical Systems Nederland B.V. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Scott T. Acton.

I. W. O. Serlie is with the Quantitative Imaging Group, Delft University of Technology, 2628 CJ Delft, The Netherlands, and also with the Department of Biomedical Engineering, Biomedical Imaging and Modeling, Eindhoven University of Technology, 5600 MB Eindhoven, The Netherlands (e-mail: iwo.serlie@philips.com).

F. M. Vos and L. J. van Vliet are with the Quantitative Imaging Group, Delft University of Technology, 2628 CJ Delft, The Netherlands (e-mail: f.m.vos@tudelft.nl; l.j.vanvliet@tudelft.nl).

R. Truyen is with Philips Medical Systems Nederland B.V., Healthcare Informatics CS & AD, Best, The Netherlands (e-mail: roel.truyen@philips.com).

F. H. Post is with the Computer Graphics Group, Delft University of Technology, 2628 BK Delft, The Netherlands (e-mail: frits.post@its.tudelft.nl).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIP.2007.909407

complicated by the limited and anisotropic resolution of the image modality at hand. The resolution of a multislice spiral CT scanner is limited by its configuration (size of the detector elements) and the reconstruction algorithm [1]. The anisotropic space-variant point spread function (PSF) resembles spiral pasta [2], 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 [3]. Modeling the PSF by a Gaussian also permits accurate edge detection of curved surfaces [4]. The finite resolution causes contributions of different materials 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 digital cleansing for virtual colonoscopy.

A. 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 [Fig. 1(a)]. Fig. 1(b) 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 $\sigma^2 = -w \max(I_w)/\max(I_{ww})$. Kindlmann and Kniss use the histogram to visualize boundary information

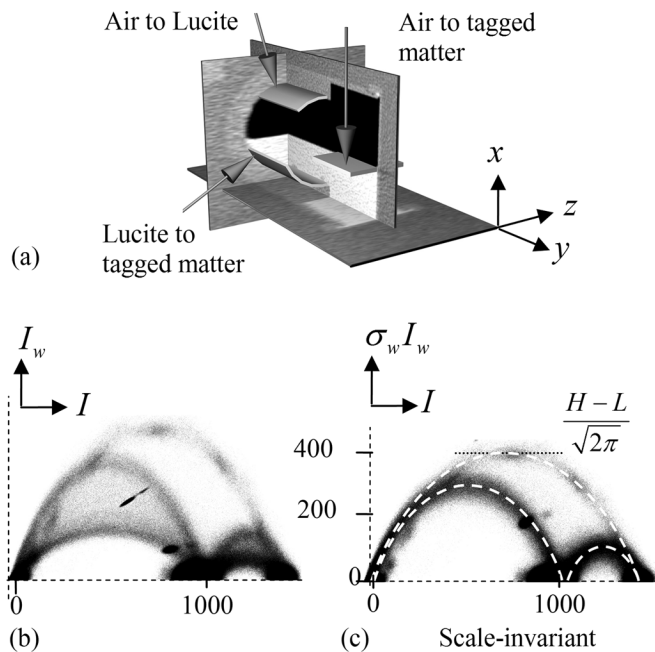


Fig. 1. (a) Schematic overview of a three material phantom using three multiplanar reformatted images through a CT-volume. The axial resolution (z axis) is lower compared to the lateral resolution (x, y plane). (b) Scatter plot of intensity and gradient magnitude. (c) Scatter plot of intensity and scale-invariant gradient magnitude. Three instantiations of the arch model are superimposed corresponding to the three types of material transitions.

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 nonlinear manner.

B. 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, noncubic voxels, and ease of use. An iso-surface 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. 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

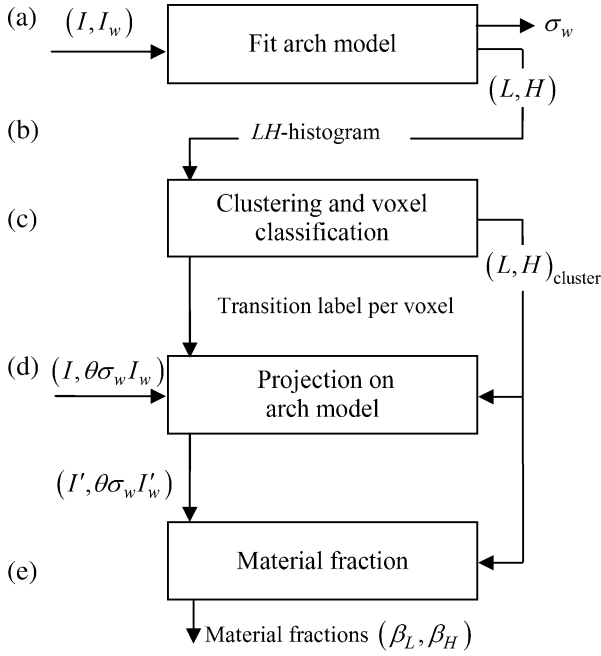


Fig. 2. (a) Model parameters (L, H) are obtained by fitting the arch model to the measurement pair (scaled to have isotropic noise). (b) Estimated model parameters (L, H) are stored in the LH histogram. (c) Clustering of the LH histogram yields material transitions and is used to classify voxels into 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 $\beta_L \beta_H$ corresponding to material L and H .

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.

II. METHODS

A. Outline

The proposed method (Fig. 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 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 [Fig. 2(a)]. The model is fitted to a set of measurements 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 [Fig. 2(b)]. 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 [Fig. 2(c)].

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 [Fig. 2(d)]. 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 [Fig. 2(e)].

B. Transition Model

A two-material transition [Fig. 3(a)] is modeled by a unit step-function u (2) that is convolved with a 1-D Gaussian edge-spread-function (ESF) g (3) resulting in a cumulative Gaussian distribution G (1) [Fig. 3(b)]. 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 [3]. For a given direction, the ESF is approximately constant over the image and, therefore, has not to be re-estimated for every voxel [3]

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

$$= \frac{1}{2} + \frac{1}{2} \operatorname{erf}\left(\frac{x}{\sigma\sqrt{2}}\right) \quad (1)$$

with

$$u(x) = \begin{cases} 0, & x < 0 \\ 1, & x > 0 \end{cases} \quad (2)$$

$$g(x; \sigma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{x^2}{2\sigma^2}\right) \quad (3)$$

$$\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x \exp(-t^2) dt. \quad (4)$$

A compact description of edges is obtained using gauge coordinates, a local Cartesian coordinate system with axes aligned

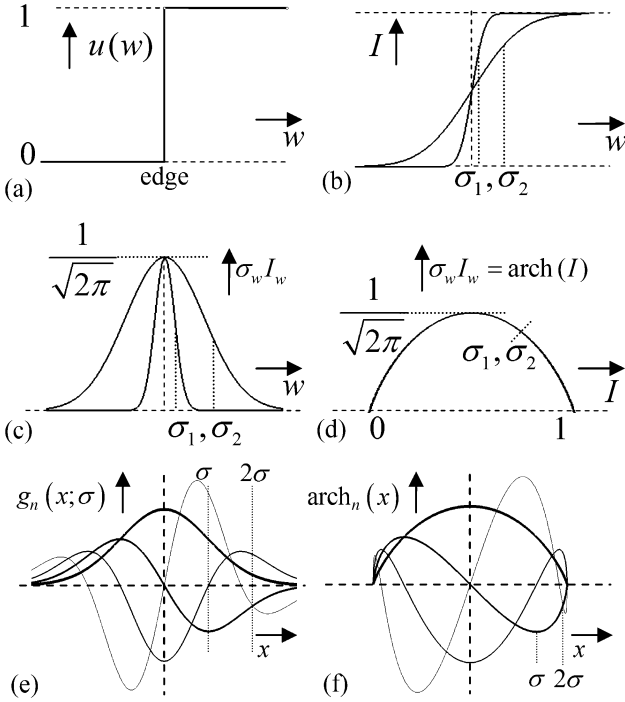


Fig. 3. (a) Material transition modeled by the unit step edge u . (b) Data values are blurred at two scales. (c) Scale-invariant gradient magnitude as a function of position. (d) Single arch is obtained upon plotting the scale-invariant gradient magnitude as a function of data value. (e) Gaussian derivatives (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.

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 assume that c) the expected data values at the transition are 0 and 1 such that $I(w; \sigma_w) \triangleq G(w; \sigma_w)$ denotes the data value (1) and $I_w(w; \sigma_w) \triangleq g(w; \sigma_w)$ denotes the gradient magnitude (3) (in E, we generalize the model). In the remainder of the text, we occasionally drop the 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 [Fig. 1(b)]. This representation is used by Kindlmann and Kniss [7], [8] for visualization purposes.

C. Scale Normalization

Scanners with a significantly anisotropic PSF cause the apparent edge scale σ_w , and, therefore, the observed gradient magnitude depend heavily on the edge orientation. Fig. 1(a) 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 Fig. 1(b) 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] along a transition as depicted in Fig. 3(c). Plotting $(I, \sigma_w I_w)$ yields a single,

scale and rotation invariant arch [Figs. 1(c) and 3(d)] of height $(2\pi)^{-1/2}$. The dashed lines in Fig. 1(c) are generated by the model. The spread that remains is caused by noise [Fig. 1(c)].

D. Analytical Expression

In this section, we will derive an analytical expression for the scale-invariant arch function [Fig. 3(d)]. 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 noncubic voxels. A first step is to determine the inverse cumulative Gaussian function G^{-1} . It is obtained by inserting (1) in $G(G^{-1}(x)) = x$ for $x \in [0, 1]$. Solving for $G^{-1}(x)$ yields

$$\begin{aligned} \frac{1}{2} + \frac{1}{2} \operatorname{erf} \left(\frac{G^{-1}(x; \sigma)}{\sigma \sqrt{2}} \right) &= x \Leftrightarrow \\ \frac{G^{-1}(x; \sigma)}{\sigma \sqrt{2}} &= \operatorname{erf}^{-1}(2x - 1) \Leftrightarrow \\ G^{-1}(x; \sigma) &= \sigma \sqrt{2} \operatorname{erf}^{-1}(2x - 1). \end{aligned} \quad (5)$$

The final step consists of multiplying (3) with σ to make the gradient-magnitude scale invariant and substitution of x with $G^{-1}(x)$ of (5). This yields $\operatorname{arch}(x)$

$$\begin{aligned} \operatorname{arch}(x) &\triangleq \sigma g(G^{-1}(x; \sigma); \sigma) \\ &= \frac{1}{\sqrt{2\pi}} \exp \left(- \left\{ \operatorname{erf}^{-1}(2x - 1) \right\}^2 \right). \end{aligned} \quad (6)$$

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 $\operatorname{arch}(I)$ 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 (Section II-J).

In general, $\operatorname{arch}_n(x)$ is the n th-order scale-invariant derivative [Fig. 3(e)] as a function of the Gaussian filtered data value [Fig. 3(f)]. 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

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

with $H_n(x; \sigma) \triangleq (-1)^n \exp \left(\frac{x^2}{2\sigma^2} \right) \frac{d^n}{dx^n} \exp \left(\frac{-x^2}{2\sigma^2} \right)$. (7)

Substitution of x in (7) with (5) gives $\operatorname{arch}_n(x)$

$$\begin{aligned} \operatorname{arch}_n(x) &\triangleq \sigma^{n-1} H_{n-1}(-G^{-1}(x; \sigma); \sigma) \sigma g(G^{-1}(x; \sigma); \sigma) \\ &= \sigma^{n-1} H_{n-1}(-G^{-1}(x; \sigma); \sigma) \operatorname{arch}(x) \\ &= S_n(x) \operatorname{arch}(x) \end{aligned} \quad (8)$$

with $S_n(x) = \sigma^{n-1} H_{n-1}(-G^{-1}(x; \sigma); \sigma)$. The $\operatorname{arch}_n(x)$ 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 Fig. 3(e) and (f).

The arch function is related to the inverse cumulative Gaussian (9) through its derivative (Appendix I)

$$\begin{aligned} \text{arch}'(x) &\triangleq \frac{d}{dx} \text{arch}(x) = -G^{-1}(x, 1) \\ \int_{-\infty}^x \text{erf}^{-1}(t) dt &= \frac{-2}{\sqrt{2}} \text{arch}\left(\frac{x+1}{2}\right). \end{aligned} \quad (9)$$

It can be concluded that $\text{arch}(x)$ 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 (6) is analytical but not in a closed form. Consequently, evaluating $\text{arch}(x)$ is cumbersome, since it requires finding the roots of $2x - 1 - \text{erf}(y) = 0$. This problem is circumvented by considering the inverse function $\text{arch}^{-1}(\text{arch}(x)) = x$ for $x \in [0, 1/2]$ that has a closed-form expression (Appendix I)

$$\text{arch}^{-1}(x) = \frac{1}{2} - \frac{1}{2} \text{erf}\left(\sqrt{-\ln(x\sqrt{2\pi})}\right) \quad \text{for } x \in \left[0, \frac{1}{\sqrt{2\pi}}\right]. \quad (10)$$

The inverse arch function describes the Gaussian filtered data value I as a function of scale-invariant gradient magnitude: $I = \text{arch}^{-1}(\sigma_w I_w)$.

E. 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

$$u(w; L, H) = L + (H - L)u(w). \quad (11)$$

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

$$I(w; \sigma_w, L, H) = L + (H - L)G(w; \sigma_w) \quad (12)$$

$$\sigma_w I_w(w; \sigma_w, L, H) = \sigma_w (H - L)g(w; \sigma_w). \quad (13)$$

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

$$\text{arch}(I; L, H) = (H - L) \text{arch}\left(\frac{I - L}{H - L}\right). \quad (14)$$

F. 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] and [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 $\text{cov}(I, \sigma_w I_w) = 0$, 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]

$$\frac{\sigma_{no}^2}{\sigma_{ni}^2} = \frac{(2D)!}{\sigma_{op}^{N+2D} \pi^{N/2} D! 2^D 2^{N+D}}. \quad (15)$$

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 $(I, \theta \sigma_w I_w)$ isotropic.

First, consider the variance of the noise after (0th order) Gaussian filtering: the first dimension of $(I, \theta \sigma_w I_w)$ 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 0th 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 (15) with $(\sigma_{op,z}, D = 0, N = 1)$ and $(\sigma_{op,\perp z}, D = 0, N = 2)$, respectively

$$\frac{\sigma_{no,z}^2}{\sigma_{ni}^2} = \frac{1}{\sigma_{op,z} 2\sqrt{\pi}}, \quad \frac{\sigma_{no,\perp z}^2}{\sigma_{ni}^2} = \frac{1}{\sigma_{op,\perp z}^2 4\pi}. \quad (16)$$

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

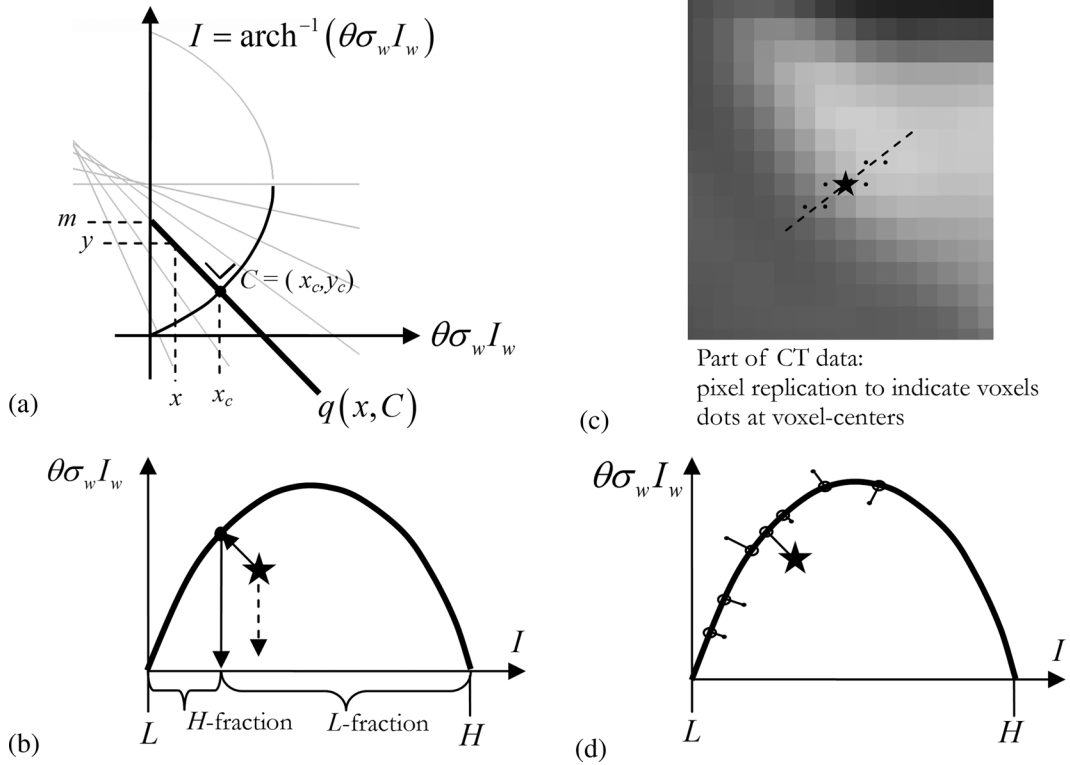


Fig. 4. (a) Inverse arch function showing the projection of values along a line onto a single point on the arch. (b) This method is applied to obtain material fractions. Notice that directly using intensity yields a different outcome (dashed arrow). (c) This projection is used as well to obtain L and 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.

gives a fixed variance after filtering in 3-D, irrespective of the orientation of the edge w

$$\sigma_{no}^2 = \sigma_I^2 = \frac{\sigma_{ni}^2}{2^3 \pi^{\frac{3}{2}} \sigma_{op,z}^2 \sigma_{op,\perp z}^2}. \quad (17)$$

Second, consider the variance of the noise when measuring the gradient magnitude: the second dimension of $(I, \theta\sigma_w I_w)$ 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

$$\sigma_{op,w} = \sqrt{(\sin(\alpha_\Delta)\sigma_{op,\perp z})^2 + (\cos(\alpha_\Delta)\sigma_{op,z})^2}. \quad (18)$$

Applying (15) with $(\sigma_{op,w}, D = 1, N = 1)$ and $(\sigma_{op,\perp w}, D = 0, N = 2)$, respectively, gives

$$\frac{\sigma_{no,w}^2}{\sigma_{ni}^2} = \frac{1}{\sigma_{op,w}^3 4\sqrt{\pi}}, \quad \frac{\sigma_{no,\perp w}^2}{\sigma_{ni}^2} = \frac{1}{\sigma_{op,\perp w}^2 4\pi}. \quad (19)$$

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

$$\sigma_{no}^2 = \sigma_{I_w}^2 = \frac{\sigma_{ni}^2}{16\pi^{\frac{3}{2}} \sigma_{op,\perp w}^2 \sigma_{op,w}^3}. \quad (20)$$

Note that the variance of the gradient-magnitude remains a function of the edge orientation w . Finally, using (17) and (20), the noise in $(I, \theta\sigma_w I_w)$ is made isotropic with

$$\theta = \frac{1}{\sigma_w} \left(\frac{\sigma_I}{\sigma_{I_w}} \right) = \frac{1}{\sigma_w} \left(\frac{\sqrt{2}\sigma_{op,w}^3 \sigma_{op,\perp w}}{\sqrt{\sigma_{op,z} \sigma_{op,\perp z}}} \right). \quad (21)$$

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 (21) is simplified considerably such that $\theta = (\sigma_{op}\sqrt{2})/\sigma_w$. Assuming the previous isotropy of the kernel, the noisy measurements $(I, \sigma_{op}I_w\sqrt{2})$ are projected onto $(\sigma_{op}/\sigma_w)\text{arch}(I)$. Remember that σ_w is the overall edge scale $\sigma_w^2 = \sigma_{op,w}^2 + \sigma_{\text{PSF},w}^2$.

G. Orthogonal Projection on the arch

The measurements $(I, \theta\sigma_w I_w)$, obtained by orthogonal operators, are combined by projection onto the arch. To begin with, we assume 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} (10) and its derivative (22) as depicted in Fig. 4(a)

$$\text{arch}^{-1}(x) = \frac{1}{2} - \frac{1}{2} \text{erf} \left(\sqrt{-\ln(x\sqrt{2\pi})} \right) \\ \frac{d}{dx} \text{arch}^{-1}(x) = (-\ln(2\pi) - 2\ln x)^{-\frac{1}{2}}. \quad (22)$$

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

$$k = \left(\frac{d}{dx} \text{arch}^{-1}(x) \right)^{-1} \Big|_{x=x_c} \quad (23)$$

$$\begin{aligned} m &= -k * x_c + \text{arch}^{-1}(x_c) \\ y &= q(x, C) = kx + m \\ &= \frac{(x_c - x)}{\frac{d}{dx} \text{arch}^{-1}(x) \Big|_{x=x_c}} + \text{arch}^{-1}(x_c). \end{aligned} \quad (24)$$

For a particular measured $(x = \theta\sigma_w I_w, y = I)$, the unique projection is achieved by numerically solving (24) for x_c . The second coordinate y_c is found by evaluating $y_c = \text{arch}^{-1}(x_c)$. In general, the projection of a measurement onto the arch parameterized by L, H and scaled by θ requires proper scaling of the axes. This operation is written as $\text{proj}_{\text{arch}(L,H,\theta)}$.

H. 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. Let A be a set of $N + 1$ measurement pairs along the gradient direction in the neighborhood of an edge sample [Fig. 4(c)]

$$A_n = (I, \theta\sigma_w I_w), \quad n \in [0, N]. \quad (25)$$

A 3-D version of Canny's edge detector is used for initial finding of edge samples. Because the arch describes 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 and the measurements [Fig. 4(d)] using the conjugate gradient method [14], the best fitting arch is obtained. It yields the L' and H' values that we are looking for. The residual error R may be used as a measure for the quality of the fit

$$\{L', H'\} = \arg \min_{\{L, H\}} \sum_n \left\| A_n - \text{proj}_{\text{arch}(L,H,\theta)} A_n \right\|^2. \quad (26)$$

I. Initial Values for L and H

Local minima may occur for (26). Consequently, 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 voxel to the nearest edge, estimated by

$$t = -\sigma_w(\sigma^2 I_{ww})/(\sigma I_w). \quad (27)$$

Equation (3) is used to compute $I'_w(t, \sigma_w)$ that 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 $Q = H - L$

$$Q = I_w/I'_w = H - L. \quad (28)$$

Now the predicted intensity at position t for $L = 0$ and $H = Q$ is computed using (12)

$$I'(t, \sigma_w, L = 0, H = Q) = QG(t, \sigma_w). \quad (29)$$

Finally, L and H are determined from

$$L = I - I', \quad H = L + Q. \quad (30)$$

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

$$\begin{aligned} L(I, I_w, I_{ww}, \sigma_w) &= I + \frac{\sigma_w I_w \sqrt{\pi}}{\sqrt{2}} \left(-1 + \text{erf} \left(\frac{I_{ww} \sigma_w}{I_w \sqrt{2}} \right) \right) \\ &\quad \times \exp \left(\frac{I_{ww}^2 \sigma_w^2}{2I_w^2} \right) \end{aligned} \quad (31)$$

$$\begin{aligned} H(I, I_w, I_{ww}, \sigma_w) &= I + \frac{\sigma_w I_w \sqrt{\pi}}{\sqrt{2}} \left(1 + \text{erf} \left(\frac{I_{ww} \sigma_w}{I_w \sqrt{2}} \right) \right) \\ &\quad \times \exp \left(\frac{I_{ww}^2 \sigma_w^2}{2I_w^2} \right). \end{aligned} \quad (32)$$

J. LH Histogram

Fitting an arch to all sets of measurement pairs $(I, \theta\sigma_w I_w)$ yields L, H values for the voxels in the edge regions. All L, H values can be represented in a 2-D LH histogram [6]. An LH histogram provides a 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 shows separate peaks for each type of transition and allows identification of transition type through clustering. Fig. 5 shows how different transitions yield separate clusters in LH -space. Four material samples yield crossing arches in $(I, \theta\sigma_w I_w)$ -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. 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

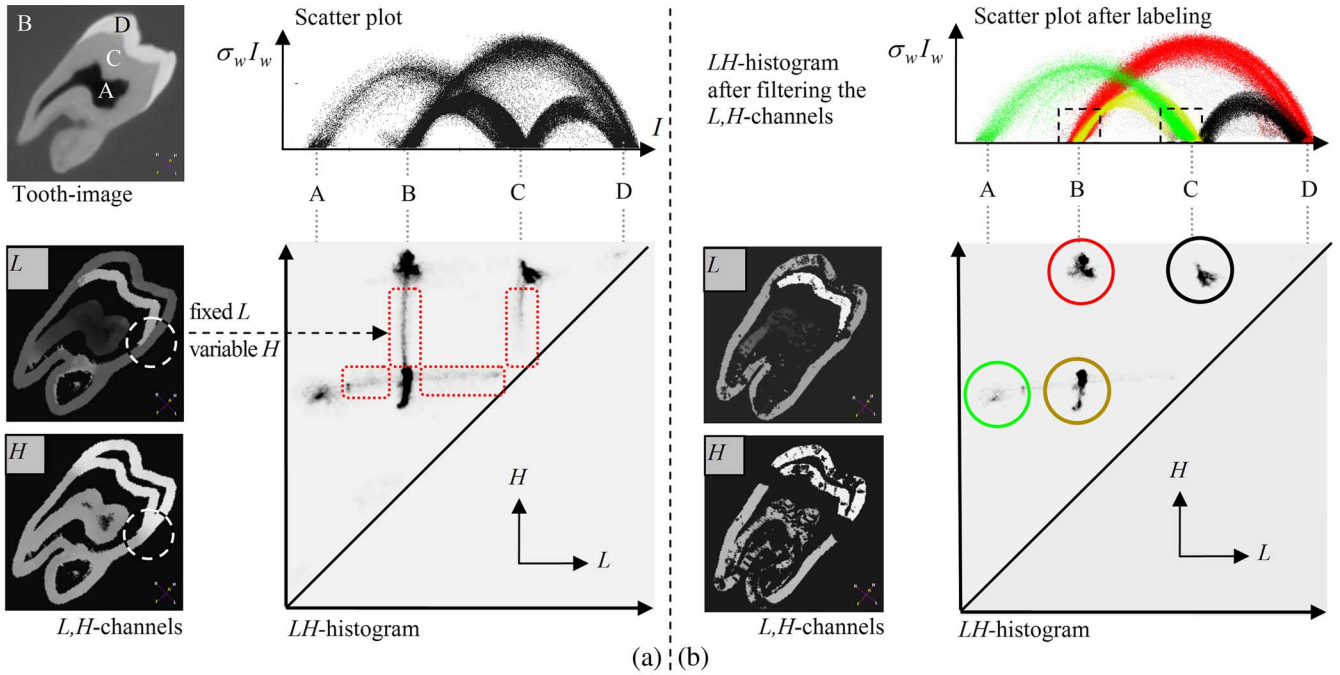


Fig. 5. LH histograms with the corresponding L - and H -channel, the scatter plot of data value and scale-invariant gradient magnitude (a) before and (b) after filtering of the L , H -channels. (a) The dashed rectangles mark regions in the L , H -channels where the arch model is not valid. (b) These 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.

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 $M_3 - M_2$ versus $M_2 - M_1$. 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.

K. 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 [Fig. 4(b)]. These material fractions are obtained by

$$\beta_L = \frac{I' - L}{H - L}, \quad \beta_H = \frac{H - I'}{H - L}, \quad \beta_L + \beta_H = 1. \quad (33)$$

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 ignore such locations or 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).

L. Visualization

The concept of material fractions is useful to cope with a material that borders on two or more other materials. The exact 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 by defining a surface at 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]. 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.

III. 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.

A. Example 1: 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 (Fig. 1). Reference data were an ultrahigh-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, $\sigma_{ni} \approx 60$ HU) [17]. 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 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 all points (manually indicated) on the contrast-plastic transition.

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 Fig. 6(a) and (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 Fig. 6(d) and (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 [Fig. 6(d) and (e)]. The resulting localization accuracy and precision are presented in Fig. 6(j) and (k) (gray 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 Fig. 6(j) and (k) by arrows. Only a single point is obtained, since there is no kernel involved in the estimation. Fig. 6(i) represents the tissue fraction in gray value (light gray 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.

The arch-based method relies on being able to separate the clusters in LH -space [Fig. 6(c) and (f) shows the outcome for the tube without and with the trend]. Note that Kindlmann uses the scatter plots such as Fig. 6(a) and (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

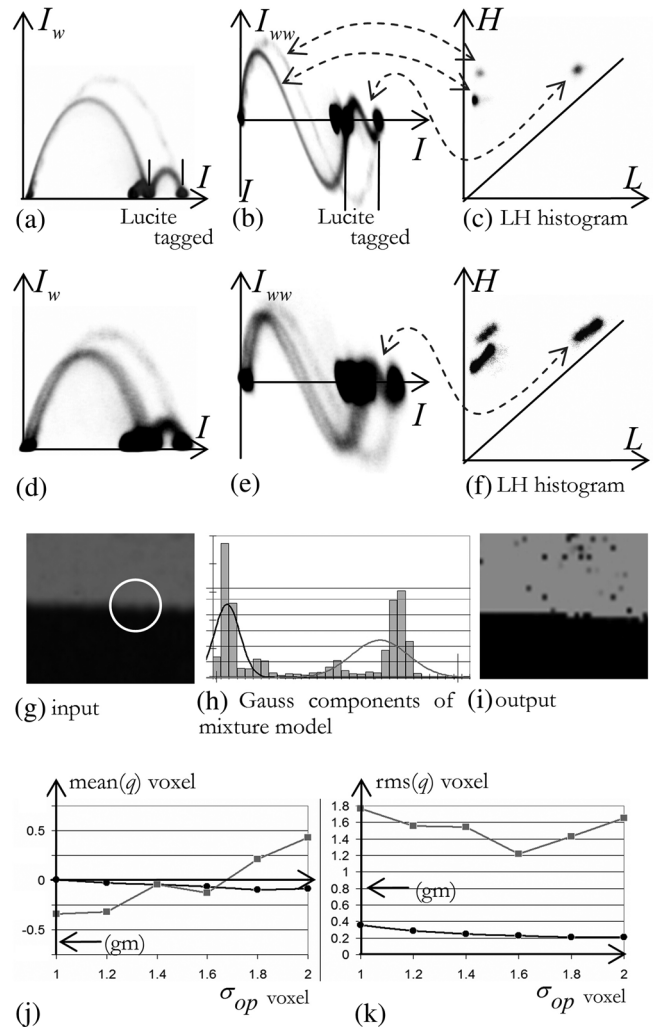


Fig. 6. Histograms of (a) I , I_w and (b) I , I_{ww} that are used by Kindlmann to obtain estimates of $I_w(I)$ and $I_{ww}(I)$. The graphs are obtained from a high-dose CT image of the phantom shown in Fig. 1. (c) LH histogram by fitting the arch: noise and global signal fluctuations are no problem 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 gray value (light gray corresponds to 100% Lucite). (j), (k) To compare methods, the smallest distance q between the estimated edge position in ultrahigh-dose and low dose is measured. (j) The mean (q) and (k) root mean square (q) are plotted for (gray line) Kindlmann's method, (black line) the arch-based method, and (gm) the mixture model.

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 Fig. 6(j) and (k) (black lines). Note that our 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.

B. Example 2: Tooth

The input data consist of industrial CT data of a human tooth from the National Library of Medicine: <http://www.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:

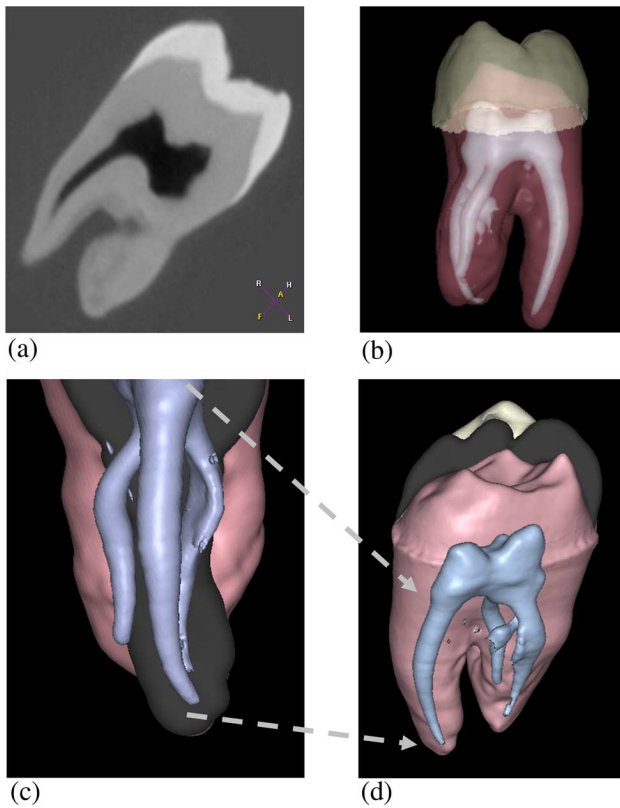


Fig. 7. (a) Cross section of tooth volume. (b) Delineations of the three materials (enamel, dentin and root canal) are combined into a single rendering; (c) and (d) show that the root canal is almost completely visualized. (d) Insertion of a clipping plane on the enamel-fraction volume enables cutting away the crown and fully visualizing structures below. Thus, special preprocessing or complicated transfer function definition are not needed to create high-quality visualizations.

dark root canals and pulp chamber, gray dentin and cementum and bright enamel and crown [Fig. 7(a)]. 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 fitting the arch model. These L and H values yield the LH histogram [Fig. 5(a)]. Clusters in the histogram correspond to the two-material transitions between tooth-materials. In addition, the LH histogram in Fig. 5(a) 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 LH 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. A “filtered” LH histogram is shown in Fig. 5(b). This histogram is used to segment the image into significant material transitions. The clusters within the colored circles in the LH histogram of Fig. 5(b) 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 Fig. 5(b)] 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 Fig. 4(b).

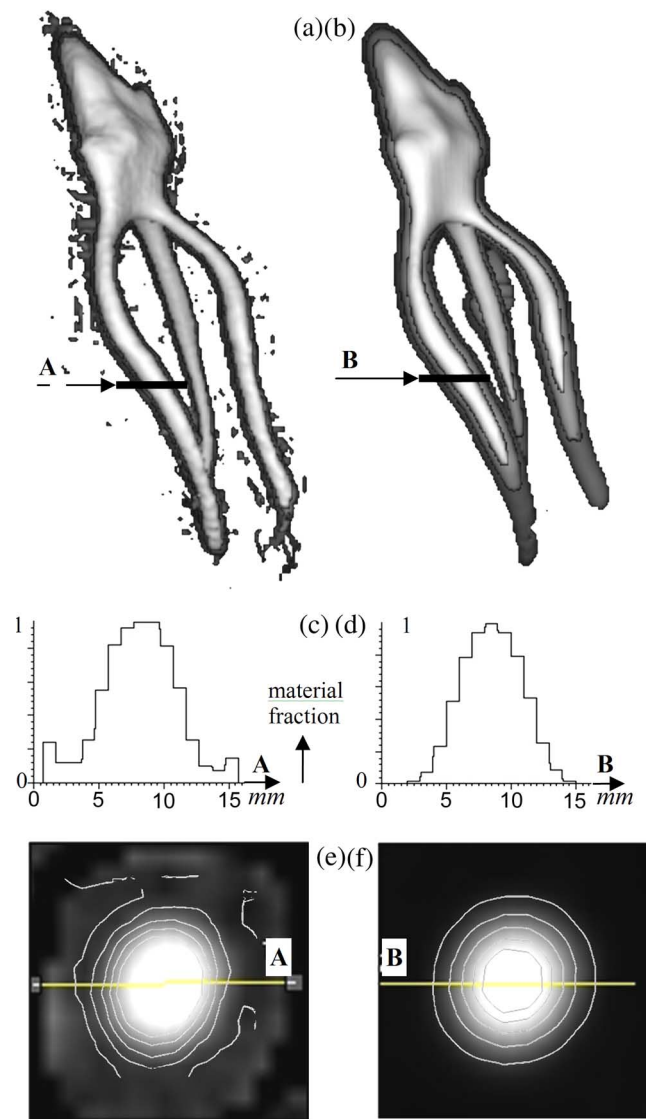


Fig. 8. Root-canal delineated using the reference method (a), (c), (e) and using the method described in this paper (b), (d), (f). The top row (a), (b) shows the root-canal rendering of the tooth volume using the reference method with threshold 0.2, 0.3, and 0.5 and method with threshold 0.1 (dark), 0.2 (intermediate), and 0.5 (bright). (c), (d) The middle row shows the material mixture along a profile. (e), (f) The bottom row shows the material fraction in gray-value with five delineations superimposed using an iso-material mixture of 0.1, 0.2, 0.3, 0.4, and 0.5, respectively.

The three resulting material-fraction volumes of root-canal dentin and enamel are used for visualization (Fig. 7). The delineation of the root-canal is compared with a reference method of [18] (Fig. 8). Sereda creates an LH histogram in which the estimates of L and H are obtained by the result of a local min and max filter, respectively, [19] rather than the arch model. The data value I is mapped on a mixture using β_L as defined in (33). 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 distortion of wedge-shaped structures. A fundamental difference of our method 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 deal with the noise in an optimal manner. Second, L and H are estimated using

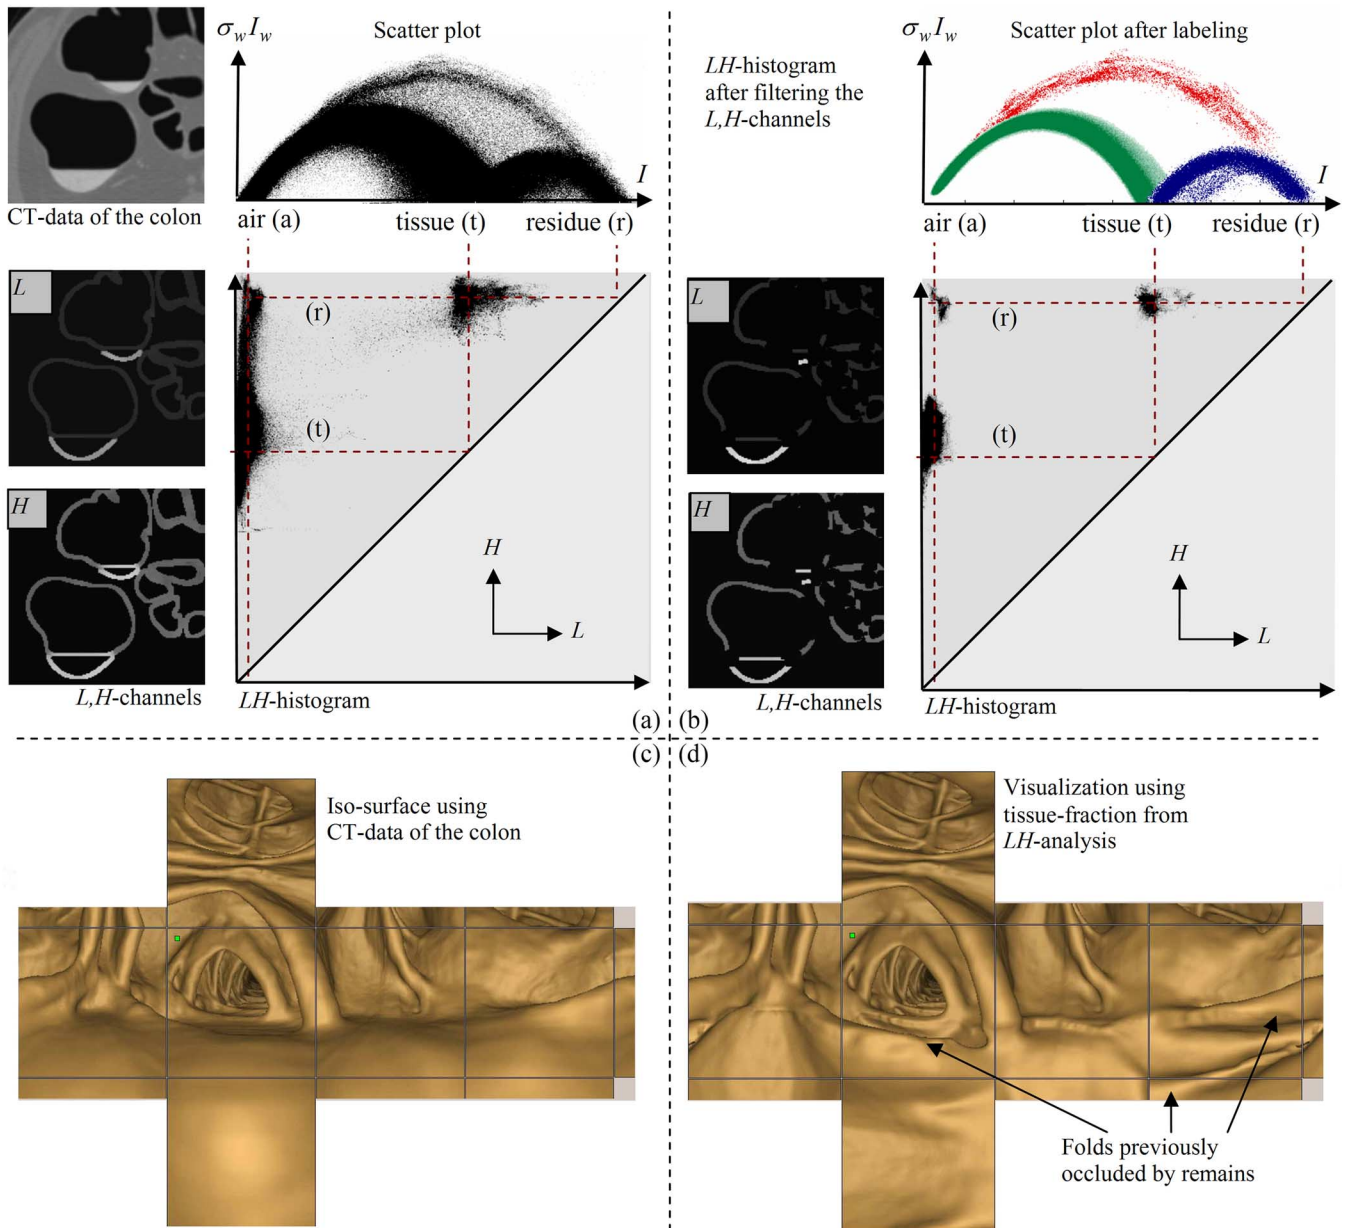


Fig. 9. LH histogram with corresponding L and H channel of abdominal CT-data. The scatter plot of the image-intensity and scale-invariant gradient magnitude (a) before and (b) after filtering the L , H -channels. Regions in the LH histogram where the arch model is not valid are automatically suppressed by thresholding the gradient-magnitude of L and H channels. (c) Unfolded cube visualization using original data and (d) digitally cleansed data. Notice that the entire colon surface is presented for inspection after processing.

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 (33) the resulting mixture will become biased: not providing a good visualization [Fig. 8(a), (c), and (e)]. At parts of the root canal with a decreasing diameter, a smaller mixture is found as expected [Fig. 8(b), (d), and (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 [Fig. 8(a)]. It does not reliably

convey the dimensions of the root canal. The user has less control of creating accurate object delineation.

C. Example 3: 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 PVE 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 data value of remains. We have applied

our method to derive the fraction of tissue, air, and tagging in each edge voxel.

The data are acquired using a multislice 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 [Fig. 9(c)]. Analogous to the tooth example intermediate processing results are depicted in Fig. 9(a) and (b). The usefulness of our method is further demonstrated in Fig. 9(d), which shows the unfolded cube visualization [23] before and after processing.

IV. 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, noncubic voxels and ease of use. The method facilitates accurate and reproducible boundary delineation for segmentation and visualization.

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 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 is [22] 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.

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 $x = \text{arch}(\text{arch}^{-1}(x))$. The substitution of (6)

$$\begin{aligned} x &= \frac{1}{\sqrt{2\pi}} \exp\left(-\{\text{erf}^{-1}(2\text{arch}^{-1}(x) - 1)\}^2\right) \Leftrightarrow \\ \ln(x\sqrt{2\pi}) &= -\{\text{erf}^{-1}(2\text{arch}^{-1}(x) - 1)\}^2 \Leftrightarrow \\ -\sqrt{-\ln(x\sqrt{2\pi})} &= \text{erf}^{-1}(2\text{arch}^{-1}(x) - 1) \Leftrightarrow \\ -\text{erf}\left(\sqrt{-\ln(x\sqrt{2\pi})}\right) &= 2\text{arch}^{-1}(x) - 1 \end{aligned} \quad (34)$$

yields the inverse arch function

$$\text{arch}^{-1}(x) = \frac{1}{2} - \frac{1}{2} \text{erf}\left(\sqrt{-\ln(x\sqrt{2\pi})}\right) \text{ for } x \in \left[0, \frac{1}{\sqrt{2\pi}}\right]. \quad (35)$$

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 $\text{arch}^{-1}(x)$

$$\frac{d}{dx} \text{arch}^{-1}(x) = \frac{1}{\frac{d}{dx} \text{arch}(\text{arch}^{-1}(x))} \text{ for } x \in \left[0, \frac{1}{\sqrt{2\pi}}\right] \quad (36)$$

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

$$\frac{d}{dx} \text{arch}(x) \triangleq -G^{-1}(x; 1) = -\sqrt{2} \text{erf}^{-1}(2x - 1) \text{ for } x \in [0, 1]. \quad (37)$$

Inserting (37) in (36), we get

$$\frac{d}{dx} \text{arch}^{-1}(x) = \frac{1}{\sqrt{-2 \ln(x\sqrt{2\pi})}}$$

for

$$x \in \left[0, \frac{1}{\sqrt{2\pi}}\right]. \quad (38)$$

APPENDIX II

The derivative of $\text{arch}(x)$ in (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 (38)

$$\begin{aligned}
 \frac{d}{dx} \text{arch}(x) &= \sqrt{-\ln(2\pi) - 2 \ln(\text{arch}(x))} \\
 &= \sqrt{-\ln(2\pi) - 2 \ln\left(\frac{\exp\left(-\{\text{erf}^{-1}(2x-1)\}^2\right)}{\sqrt{2\pi}}\right)} \\
 &= \sqrt{-\ln(2\pi) + \ln(2\pi) + 2\left(-\{\text{erf}^{-1}(2x-1)\}^2\right)} \\
 &= -\sqrt{2}\text{erf}^{-1}(2x-1) \\
 &\triangleq -G^{-1}(x; 1) \text{ for } x \in \left[0, \frac{1}{2}\right]. \tag{39}
 \end{aligned}$$

This may be rewritten using (5) and (6) to give the integral of the inverse error-function as an alternative to using the method of Parker [25]

$$\begin{aligned}
 \int_0^x G^{-1}(t; 1) dt &= -\text{arch}(x) \Leftrightarrow \\
 \int_0^x \sqrt{2}\text{erf}^{-1}(2t-1) dt &= \frac{-1}{\sqrt{2\pi}} \exp\left(-\{\text{erf}^{-1}(2x-1)\}^2\right) \Leftrightarrow \\
 \sqrt{2} \int_{-1}^{2x-1} \text{erf}^{-1}(t) d\left(\frac{1}{2}t\right) &= \frac{-1}{\sqrt{2\pi}} \exp\left(-\{\text{erf}^{-1}(2x-1)\}^2\right). \tag{40}
 \end{aligned}$$

The substitution of $2x - 1$ by x yields

$$\begin{aligned}
 \int_{-1}^x \text{erf}^{-1}(t) dt &= \frac{-1}{\sqrt{\pi}} \exp\left(-\{\text{erf}^{-1}(x)\}^2\right) \\
 &= \frac{-2}{\sqrt{2}} \frac{1}{\sqrt{2\pi}} \exp\left(-\{\text{erf}^{-1}(x)\}^2\right) \\
 &= \frac{-2}{\sqrt{2}} \text{arch}\left(\frac{x+1}{2}\right). \tag{41}
 \end{aligned}$$

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 $\sigma = \sqrt{2}$.

ACKNOWLEDGMENT

The authors would like to thank Dr. P. Rogalla, Charité Berlin, for providing them with tagged patient data, and Philips Medical Systems Nederland B.V. for providing the ViewForum prototyping software.

REFERENCES

- [1] G. Schwarzband and N. Kiryati, "The point spread function of spiral CT," *Phys. Med. Biol.*, vol. 50, pp. 5307–5322, 2005.
- [2] G. Wang, M. W. Vannier, M. W. Skinner, M. G. P. Cavalcanti, and G. W. Harding, "Spiral CT image deblurring for cochlear implantation," *IEEE Trans. Med. Imag.*, vol. 17, no. 2, pp. 251–262, Apr. 1998.
- [3] I. W. O. Serlie, F. M. Vos, H. W. Venema, and L. J. van Vliet, CT Imaging Characteristics. [Online]. Available: <http://www.ist.tudelft.nl/qi>. Group Pub.: Tech. Rep.
- [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 2-D and 3-D CT images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 9, pp. 1501–1507, Sep. 2005.
- [5] Y. Zou, E. Y. Sidky, and X. Pan, "Partial volume and aliasing artefacts in helical cone-beam CT," *Phys. Med. Biol.*, vol. 49, pp. 2365–2375, 2004.
- [6] I. W. O. Serlie, R. Truyen, J. Florie, F. H. Post, L. J. van Vliet, and F. M. Vos, "Computed cleansing for virtual colonoscopy using a three-material transition model," in *Proc. MICCAI*, 2003, vol. 2879, 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*, Oct. 1998, pp. 79–86.
- [8] J. Kniss, G. Kindlmann, and C. Hansen, "Multi-dimensional transfer functions for interactive volume rendering," *IEEE Trans. Vis. Comput. Graph.*, vol. 8, no. 3, pp. 270–285, Jul. 2002.
- [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," *IEEE Trans. Med. Imag.*, vol. MI-17, no. 1, pp. 74–86, Feb. 1998.
- [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, and M. A. Viergever, "Scale and the differential structure of images," *Image Vis. Comput.*, vol. 10, pp. 376–388, 1992.
- [12] P. Charbonnier, L. Blanc-Feraud, G. Aubert, and M. Barlaud, "Deterministic edge-preserving regularization in computed imaging," *IEEE Trans. Image Process.*, vol. 6, no. 3, pp. 298–311, Mar. 1997.
- [13] L. J. van Vliet, "Grey-scale measurements in multi-dimensional digitized images" Ph.D. dissertation, Delft Univ. Technol., Delft, The Netherlands, 1993. [Online]. Available: <http://www.ist.tudelft.nl/qi>
- [14] O. Axelsson and V. Barker, "Finite element solution of boundary value problems," AP Inc., 1984.
- [15] D. Ballard, "Generalized Hough transform to detect arbitrary patterns," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-13, no. 2, pp. 111–122, Feb. 1981.
- [16] B. Rieger, F. J. Timmermans, L. J. van Vliet, and P. W. Verbeek, "On curvature estimation of iso-surfaces in 3-D gray-value images and the computation of shape descriptors," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, no. 8, pp. 1088–1094, Aug. 2004.
- [17] J. R. Mayo, K. P. Whittall, A. N. Leung, T. E. Hartman, C. S. Park, and S. L. Primack *et al.*, "Simulated dose reduction in conventional chest CT: Validation study," *Radiology*, vol. 202, pp. 453–457, 1997.
- [18] P. Sereida, 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, no. 2, pp. 208–218, Mar./Apr. 2006.
- [19] P. W. Verbeek, H. A. Vrooman, and L. J. van Vliet, "Low level image processing by max-min. filters," *Signal Process.*, vol. 15, pp. 249–258, 1988.
- [20] J. Mandel, J. Bond, J. Church, and D. Snover, "Reducing mortality from colon cancer control study," *N. Eng. J. Med.*, vol. 328, no. 19, pp. 1365–1371, May 1993.
- [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," *AJR*, vol. 162, 1994.
- [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," *Lecture Notes Comput. Sci.*, vol. 3749, pp. 712–719, 2005.
- [23] F. M. Vos and R. E. van Gelder *et al.*, "Three-dimensional display modes for CT colonography: Conventional 3D virtual colonoscopy versus unfolded cube projection," *Rad*, vol. 228, pp. 878–885, 2003.
- [24] P. J. Pickhardt and J. H. Choi, "Electronic cleansing and stool tagging in CT colonography: Advantages and pitfalls with primary three-dimensional evaluation," *AJR*, vol. 181, pp. 799–805, 2003.
- [25] F. D. Parker, "Integrals of inverse functions," *Amer. Math. Monthly*, vol. 62, pp. 439–440, 1955.



Iwo W. O. Serlie received the M.Sc. degree in technical informatics from the Delft University of Technology (TU Delft), Delft, The Netherlands. He is currently pursuing the Ph.D. degree in the Quantitative Imaging Group, Department of Imaging Science and Technology, TU Delft.

His research interests include medical image analysis, visualization via the unfolded cube method, and electronic cleansing for virtual colonoscopy.



Frans M. Vos received the M.Sc. degree in medical informatics and computer science from the University of Amsterdam, Amsterdam, The Netherlands, in 1993, and the Ph.D. degree from the Vrije Universiteit Amsterdam in 1998.

He was a Visiting Scientist at Yale University, New Haven, CT, in 1992. After that, he was a Research Fellow with the Pattern Recognition Group, Delft University of Technology (TU Delft), Delft, The Netherlands, from 1998 to 2003. He became an Assistant Professor in the Quantitative Imaging Group, TU Delft, in 2003. Since 2000, he has also been a staff member with the Department of Radiology, Academic Medical Center Amsterdam. His main research interests are in medical image processing and visualization, particularly focusing on virtual colonoscopy, diffusion tensor imaging, and statistical shape analysis.



Roel Truyen received the M.Sc. degree in electrical engineering in 1993 from the Catholic University, Leuven, Belgium.

He specialized in image processing at the Ecole Nationale Supérieure des Télécommunications, Paris, France. He is currently with the Clinical Science and Advanced Development Department, Healthcare Informatics, Philips Medical Systems, The Netherlands. His research interests include clinical applications and image processing in the oncology domain with a special interest in virtual

colonoscopy.



Frits H. Post is an Associate Professor of computer science (visualization) at the Delft University of Technology, Delft, The Netherlands, where he has lead a research group in data visualization since 1990. His current research interests include vector field and flow visualization, medical imaging and visualization, virtual reality, and interactive exploration of very-large time-varying data sets.

Prof. Post is the Chairman of the Eurographics Steering Committee on Data Visualisation and a Co-Founder of the Annual Joint Eurographics IEEE VGTC EuroVis Symposium. He is a Fellow of the Eurographics Association and an Associate Editor of the *ACM Transactions on Graphics*.



Lucas J. van Vliet studied applied physics and received the Ph.D. degree (*cum laude*) from the Delft University of Technology (TU Delft), Delft, The Netherlands, in 1993. His thesis entitled "Grey-scale measurements in multidimensional digitized images" presented novel methods for sampling-error-free measurements of geometric object features.

He is a Full Professor of multidimensional image processing and analysis at the TU Delft. He has worked on various sensor, image restoration, and image measurement problems in quantitative microscopy and medical imaging. He was Visiting Scientist at LLNL (1987), UCSF (1988), Amoco ATC (1989–1990), Monash University (1996), and LBNL (1996).

Dr. Vliet was awarded a fellowship from the Royal Netherlands Academy of Arts and Sciences (KNAW) in 1996.