Statistical Region-Based Segmentation of Ultrasound Images

Greg Slabaugh [a] Gozde Unal [b] Michcal Wels [c] Tong Fang [d] Bimba Rao [e]

[a] *Research and Development Department, Medicsight, London UK*

[b] *Computer Vision and Pattern Analysis Laboratory, Sabanci University, Istanbul Turkey*

[c] *Computer Science Department, University Erlangen-Nurenberg, Erlangen Germany*

[d] *Real-Time Vision Department, Siemens Corporate Research, Princeton NJ USA*

[e] *Ultrasound Division, Siemens Medical Solutions, Mountain View CA USA*

Corresponding Author: Greg Slabaugh

Kensington Centre

66 Hammersmith Road

London, U.K.

W14 8UD

Email: greg.slabaugh@gmail.com

Telephone: +44 (0)207 605 1180

This work was completed when Greg Slabaugh, Gozde Unal, and Michael Wels were employed by Siemens Corporate Research.

Abstract

Segmentation of ultrasound images is a challenging problem due to speckle, which corrupts the image and can result in weak or missing image boundaries, poor signal to noise ratio, and diminished contrast resolution. Speckle is a random interference pattern that is characterized by an asymmetric distribution as well as significant spatial correlation. These attributes of speckle are challenging to model in a segmentation approach, so many previous ultrasound segmentation methods simplify the problem by assuming that the speckle is white and/or Gaussian distributed. Unlike these methods, in this paper we present an ultrasound-specific segmentation approach that addresses both the spatial correlation of the data as well as its intensity distribution. We first decorrelate the image and then apply a region-based active contour whose motion is derived from an appropriate parametric distribution for maximum likelihood image segmentation. We consider zero-mean complex Gaussian, Rayleigh, and Fisher-Tippett flows, which are designed to model fully formed speckle in the in-phase/quadrature (IQ), envelope detected, and display (log compressed) images, respectively. We present experimental results demonstrating the effectiveness of our method, and compare the results to other parametric and non-parametric active contours.

Key words: Ultrasound image segmentation, speckle decorrelation, zero-mean complex Gaussian flow, Fisher-Tippett distribution, Fisher-Tippett distribution, variational and level set methods

Introduction and Literature

Segmentation is a fundamental problem in ultrasound image processing and has numerous important clinical applications, including anatomic modeling, change quantification, and image-guided interventions/therapy. While ultrasound systems are continually improving by increasing their spatial and temporal resolution, a fundamental limitation to image quality is speckle, an interference pattern resulting from the coherent accumulation of random scattering in a resolution cell of the ultrasound beam. While the texture of the speckle does not correspond to any underlying structure, the brightness of the speckle pattern is related to the local echogenicity of the underlying scatterers. The speckle appears as a spatially correlated noise pattern and has a detrimental effect on the image quality and interpretability. For example, Bamber and Daft (1986) have shown that due to speckle, the detectability of lesions in ultrasound is significantly lower compared to X-ray and Magnetic Resonance (MR).

Since the speckle obfuscates the structures of interest, it also poses a difficult challenge to segmentation algorithms. Numerous ultrasound segmentation papers have appeared in the literature; a recent survey by Noble and Boukerroui (2006) reviews methods up to the year 2005. In this body of literature, many papers assume the intensities in an ultrasound image are spatially uncorrelated and/or the follow a Gaussian distribution. While these assumptions render the problem more tractable, as Michailovich and Tannenbaum (2006) argue, they are oversimplifications that are unnatural to ultrasound imaging.

We concur with Noble and Boukerroui (2006), who argue that by modeling the imaging physics of ultrasound, it is possible to derive *ultrasound-specific* segmentation techniques that are more successful than generic methods. In this class of ultrasound-specific methods, some authors have presented techniques designed for non-Gaussian image statistics, starting with the exponential distribution pioneered by Chesnaud et al (1999), as well as Rayleigh, Sarti et al (2005), Gamma, Tao and Tagare (2005), and Beta, Martin-Fernandez

and Alberola-Lopez (2005) distributions. Other related statistical image segmentation methods include Ayed et al (2005), who use Gamma distributions for SAR image segmentation, Ayed et al (2006), who consider the Weibull distribution, and nonparametric image segmentation techniques such as Kim et al (2005); Unal et al (2005). However, such methods assume that the image pixels are spatially uncorrelated, which is generally not a valid assumption for ultrasound images, and this has an impact on both the derivation of the technique as well as the effectiveness of the segmentation method.

Commercial ultrasound scanners typically employ log compression to the envelope-detected image in order produce a display image with a suitable dynamic range for presentation on a monitor. This log compression operation significantly changes the intensity distribution of the speckle; transforming fully formed speckle from a Rayleigh distribution to a Fisher-Tippett distribution, as described by Michailovich and Adam (2003); Dutt and Greenleaf (1996). Previous segmentation work based on image statistics does not address this issue, and is instead designed to work only with the envelope-detected image.

Our Contribution

In this paper we present an ultrasound-specific segmentation approach that addresses both the spatial correlation of the speckle data as well as its intensity distribution. The approach relies on two steps. First, we decorrelate the ultrasound image by applying a whitening filter. This filtering operation is designed to remove the spatial correlation of the data, while maintaining its diagnostic information. To our knowledge, this is the first paper to address the segmentation of decorrelated images. On the decorrelated image, where the assumption of spatial independence of the pixels is more appropriate, we present a unified analysis of statistical region-based segmentation algorithms for the complex Gaussian, Rayleigh, and Fisher-Tippett distributions, which correspond to fully formed speckle in the IQ, envelope detected image, and display images, respectively. Another original contribution is the derivation and use of an active contour to segment images

modeled by complex Gaussian and Fisher-Tippett distributions. We model the curve using a level set approach, which provides sub-pixel resolution and easily handles topological changes, while our flows (active contour motion equations) drive the contour to relevant structures in the image. Experimentally, we compare our ultrasound-specific segmentation approaches to other parametric and non-parametric region-based segmentation methods. Our hypothesis is that by modeling the speckle, both in terms of its spatial correlation and intensity distribution, better segmentations will be produced. This claim is justified by our experimental results, which show the improvement of our parametric flows to other region-based segmentation methods.

We note that a preliminary version of this work appeared in Slabaugh et al (2006). In this paper we expand on that work, providing more details, and presenting segmentation methods for the imaging chain after demodulation. Specifically, we consider segmentation of the complex IQ image, envelope detected image, and display image. A new feature of this paper is the complex Gaussian flow for the complex IQ image, and we show how it is similar to the Rayleigh flow for the envelope detected image. Additionally, we provide more experimental results of our methods, including comparisons with non-parametric image segmentation.

Materials and Methods

We begin with a brief review of the standard ultrasound image formation model in order to lay the foundation of our segmentation methods, which are based on the statistics of speckle. Due to space limitations this review is very brief; further details can be found in Michailovich and Adam (2003); Dutt and Greenleaf (1996); Wagner et al (1983); Goodman (2005).

As shown in Figure 1, an in-phase/quadrature (IQ) image is obtained by applying demodulation to standard radio-frequency (RF) data from the transducer. This B-mode IQ

image is complex and is the input to our system.

Intensity Distribution

Speckle is an interference pattern resulting from the coherent accumulation of random scattering in a resolution cell of the ultrasound beam. In the case of fully formed speckle, which is typically assumed when the number of scatterers per cell is greater than ten Dutt and Greenleaf (1996), it is assumed that each scatterer contributes an independent random complex component, resulting in a random walk in the complex plane. If one applies the central limit theorem to the random walk, one observes that the distribution is a zero-mean Gaussian probability density function (PDF) in the complex plane, i.e.,

$$p_Z(z) = \frac{1}{2\pi\sigma^2} e^{-|z|^2/(2\sigma^2)} \tag{1}$$

where z is complex. This PDF models the data in the IQ image. To produce a real image for display, envelope detection is performed by taking the magnitude of the IQ image. It is fairly straightforward to show that under this transformation, the distribution in the magnitude image is Rayleigh Goodman (2005), i.e.,

$$p_X(x) = \frac{x}{\sigma^2} e^{-x^2/(2\sigma^2)}, \tag{2}$$

where x is real. Typically, the magnitude image has a large dynamic range, and therefore the standard is to log-compress the image to produce an image suitable for display. Taking the natural log, i.e., Y = ln(X), one can derive the distribution in the display image,

$$p_Y(y) = p_X(x) \left| \frac{dy}{dx} \right|^{-1}, \tag{3}$$

using $dy/dx = 1/x = e^{-y}$, and normalizing, to get

$$p_Y(y) = 2\, e^{\frac{1}{2\sigma^2}} \exp\left([2y - \ln(2\sigma^2)] - \exp([2y - \ln(2\sigma^2)]) \right). \tag{4}$$

which is a doubly exponential distribution that has the form of a Fisher-Tippett distribution. This distribution therefore is the theoretical model for the image intensities for

fully formed speckle in the log-compressed magnitude IQ image. Note that in this paper, we use the terms *magnitude* image and *envelope-detected* image interchangeably, and in addition, we use the terms *log magnitude* and *display* image interchangeably.

To verify these theoretical models, we analyzed a real ultrasound image taken of a lesion phantom (ATS Laboratories Inc, Model 539). The image, shown in Figure 2, was acquired using a Siemens Sequoia 512 system using a 6C2 transducer and frequency of 5.5 MHz. From different depths, we selected three image regions (each indicated by a white box) corresponding to the soft tissue, where the primary variation in the image intensity is due the speckle. In these regions, we formed a histogram, shown in right-most column, of the pixel intensities. Naturally, for the IQ image, this is a two-dimensional histogram formed over the real and imaginary components of the signal. Next, we estimate the standard deviation of the complex Gaussian distribution of Equation 1 using the maximum-likelihood estimator that will be described in the next section, and overlay the estimated distribution, scaled to match the histogram. The complex distribution provides an excellent fit to the histograms as demonstrated in the figure.

The envelope-detected image is shown in the left-most column of Figure 3: notice the dark appearance resulting from the large dynamic range of the intensities. We show the histogram of the selected region in the middle-left column, and fit the histogram to a Rayleigh distribution (using the maximum-likelihood estimator Sarti et al (2005)) and overlay the fit curve on top of the histogram. The Rayleigh distribution does an excellent of modeling the statistics in these examples as demonstrated in the figure.

Finally, the display image, which is a typical presentation of an ultrasound image, is shown in the right-middle column of Figure 3. For the same selected regions, we fit (using the maximum-likelihood estimator to be described in the the next section) a Fisher-Tippett distribution, which very accurately models the intensity distribution. Similar results were found for other soft tissue regions in the image. We repeat these experiments for different regions of a real ultrasound image of a carotid artery, shown in Figure 4. All the images

of carotid arteries in the paper were acquired using a Siemens Sequoia 512 system with a 8L5 transducer and a frequency of 8.0 MHz.

From this analysis we conclude that the complex Gaussian is indeed a good choice for modeling the intensities of similar regions in the IQ image, the Rayleigh model is preferred for the envelope-detected image, and the Fisher-Tippett distribution is preferred for the display image. We also note that application of the Rayleigh distribution to the display image is not ideal, as the Rayleigh and Fisher-Tippett distributions are notably different. Among other differences, the long tail of the Rayleigh distribution is located to the right of the peak (positive skew), while in the Fisher-Tippett distribution, the long tail is found on the left side of the peak (negative skew). We argue that the distribution that best matches the data should be used in the segmentation approach. Later we will derive variational flows for region-based segmentation based on these three distributions.

An issue for consideration is the sample size needed to estimate the distributions. Intuitively, one would expect that when the sample size is small, the estimated parameter of the distribution is less reliable than when the sample size is large. The phenomenon can be studied by determining confidence intervals of the estimator as a function of sample size, as shown in Figure 5 for the Rayleigh estimator using the method described in Johnson et al (1995). This plot shows the estimated parameter (solid curve) as well as 95% confidence intervals (dashed curves) as a function of samples for samples taken from a window of the carotid image in Figure 4 (g), averaged over ten realizations (sets of samples from the window) to characterize the trends that generalize independent realizations of the experiment. The figure shows that when the number of samples is very few (e.g., one to three samples), there is little confidence in the estimated parameter, as the dashed curves are far from the solid curve. However, the confidence interval becomes significantly tighter as more samples are used to estimate the distribution. Asymptotically, as the number of samples becomes infinite, the confidence interval shrinks to zero. We expect similar results for the other distributions. As an empirical rule of thumb, at least five pixels should be used to provide a reasonable estimation of the distribution.

However, more pixels provide a better statistical estimation.

Finally, we should note that non-Rayleigh scattering can occur in the magnitude image when the number of scatterers is low, their spatial locations are not independent, or the scattering is not diffuse. In these cases, numerous distributions for modeling the ultrasound image intensities have been proposed, including the Homodyned K, Rice, Nakagami, Weibull, Generalized Gaussian, and Rician-Inverse Gaussian (RiIG) distributions Eltoft (2006); Michailovich and Tannenbaum (2006). It is beyond the scope of this paper to consider all these cases, and how these distributions transform upon taking the log of the image. Indeed, many of these distributions are intractable analytically. Fortunately, in Michailovich and Tannenbaum (2006), the authors argue that the distribution for fully formed speckle is a reasonable approximation in these other cases.

Spatial Correlation

At this point we have characterized the image intensity distribution, but we have not yet addressed its spatial correlation, which renders ultrasound as arguably one of the more challenging medical imaging modalities with which to work. To understand this spatial correlation, we assume a standard image formation model where the backscattered signal and the tissue reflectivity function obey a simple relationship based on linear systems theory. Under the assumption of linear wave propagation and weak scattering, the IQ image is considered to be the result of the convolution of the point spread function (PSF) of the imaging system with the tissue reflectivity function, i.e.,

$$g(x,y) = f(x,y) * h(x,y) + u(x,y) \tag{5}$$

where $g(x,y)$, $f(x,y)$, and $h(x,y)$ denote the IQ image, the tissue reflectivity function, and the PSF, respectively. The additive term $u(x,y)$ describes measurement noise and physical phenomena that are not covered by the convolution model. In the equation above, the received IQ image $g(x,y)$ is considered to be a filtered version of the true reflectivity

function $f(x, y)$. The spatial extent of the PSF is dependent upon the size of the aperture as well as the frequency of the ultrasound imaging. Since the PSF is essentially a finite bandwidth low-pass filter, it imparts non-negligible spatial correlation to the IQ image. The correlation can be measured experimentally by calculating the half-bandwidth of the autocovariance function of the magnitude IQ image, as shown in Figure 6. This function has a notable bandwidth indicating the spatial correlation of the data; estimated sizes (computed as twice the half-bandwidth) are 4.34 and 2.45 pixels, respectively. Clearly, speckle in any real imaging situation has significant spatial correlation that should be addressed by an ultrasound segmentation method. Thus, we can improve upon previous algorithms that assume that the speckle is a white noise process. To address the spatial correlation, we first transform the IQ image using a whitening filter that decorrelates the data, resulting in another IQ image with pixels that correlate less than the original image.

Decorrelation

We perform whitening of speckled images Iraca et al (1989) by the use of a decorrelation procedure proposed in Michailovich and Tannenbaum (2006), which estimates the PSF using wavelet methods. Then, it is possible to suppress the correlation by "undoing" the effect of the PSF through deconvolution.

The speckle in the processed image has significantly less spatial correlation, as depicted in Figure 7; the half-bandwidth size has decreased to 2.36 pixels in the lateral dimension and 1.70 pixels in the range dimension. Visually, this decorrelated image appears to have a higher spatial resolution as finer details become apparent. While there may still exist some residual correlation in the image after processing, we use the term "decorrelated image" to describe the image after the decorrelation filter has been applied.

It is natural to wonder if the decorrelation affects the intensity distributions. To check,

we repeated the previous experiment of fitting Rayleigh and Fisher-Tippett distributions to histograms formed over the same soft tissue region in the phantom image. As demonstrated in Figure 8, the decorrelation does not significantly affect the distributions, so we infer that the models still hold.

Maximum Likelihood Region-Based Segmentation

In this section we introduce ultrasound-specific flows for segmentation using region-based active contours. For simplicity, we derive the variational flow for Fisher-Tippett distributions of the display image. For this, we derive the maximum likelihood estimator for this distribution as well as the maximum likelihood region-based flow for curve evolution for segmentation. Similar derivations (not presented for conciseness) for the complex Gaussian and Rayleigh distributions (corresponding to the IQ and envelope-detected image) are performed, and the resultant estimators and flows are presented. We show that the Rayleigh flow of Chesnaud et al (1999); Sarti et al (2005) is similar to the zero-mean complex Gaussian flow for the IQ image.

In a spirit similar to Chan and Vese (2001), we will evolve a contour embedded as the zero level set of a higher dimensional function based on statistical measures computed both inside and outside the contour. For this, we will need to estimate a Fisher-Tippett distribution given a set of samples from the image.

Maximum likelihood Fisher-Tippett estimator

Let $I(x, y)$ denote a pixel intensity in the display image at the location (x, y). As stated previously, the Fisher-Tippett PDF for a pixel's intensity can be written as

$$p(I(x,y)) = 2e^{\frac{1}{2\sigma^2}} e^{\left(2I(x,y) - \ln(2\sigma^2) - e^{2I(x,y) - \ln(2\sigma^2)}\right)}, \tag{6}$$

where σ^2 denotes the standard deviation parameter of the reflectivity samples. For a region Ω in the image, the log likelihood can then be expressed as

$$\ell = \int_\Omega \left(\ln 2 + \frac{1}{2\sigma^2} + 2I(x,y) - \ln(2\sigma^2) - e^{2I(x,y) - \ln(2\sigma^2)} \right) dxdy. \tag{7}$$

Next, we find an expression for σ^2 that is the maximum likelihood estimator of the FT distribution, by taking the derivative of ℓ and setting the expression equal to zero,

$$\frac{\partial \ell}{\partial \sigma} = \int_\Omega \left(-\frac{1}{\sigma^3} - \frac{4\sigma}{2\sigma^2} + \left(e^{2I(x,y) - \ln(2\sigma^2)} \right) \frac{4\sigma}{2\sigma^2} \right) dxdy = 0. \tag{8}$$

Solving for σ^2 gives

$$\sigma^2 = \frac{1}{2} \frac{\int_\Omega \left(e^{2I(x,y)} - 1 \right) dxdy}{\int_\Omega dxdy}. \tag{9}$$

Thus, given a region Ω with area given by $\int_\Omega dxdy$, we can compute the maximum likelihood value of the Fisher-Tippett distribution from the image intensities in the region. We will do this to estimate the Fisher-Tippett parameter σ^2 both inside and outside the active contour. Note that we used this equation to estimate the FT distributions shown in Figures 3, 4, and 8.

Fisher-Tippett flow

We would like to deform a curve $\mathbf{C}$ in order to achieve a maximum likelihood segmentation of the data. Since the log function is monotonic, we can equivalently maximize the log likelihood Chesnaud et al (1999); Sarti et al (2005), using the probability inside and outside the curve, P_i and P_o, where P_i and P_o. In the display image, we model P_i and P_o with Fisher-Tippett distributions inside and outside the contour, respectively. Specifically, the data-driven part of the curve evolution is derived as

$$\frac{\partial \mathbf{C}}{\partial t} = (\ln P_o - \ln P_i)\,\mathbf{N}$$
$$= \left(\ln\left[2e^{\frac{1}{2\sigma_o^2}} e^{\left(2I(x,y) - \ln(2\sigma_o^2) - e^{2I(x,y)\ \ln(2\sigma_o^2)}\right)}\right] - \ln\left[2e^{\frac{1}{2\sigma_i^2}} e^{\left(2I(x,y) - \ln(2\sigma_i^2) - e^{2I(x,y) - \ln(2\sigma_i^2)}\right)}\right]\right)\mathbf{N}$$
$$= \left(\ln\frac{\sigma_i^2}{\sigma_o^2} + \frac{e^{2I(x,y)} - 1}{2\sigma_i^2} - \frac{e^{2I(x,y)} - 1}{2\sigma_o^2}\right)\mathbf{N}, \tag{10}$$

where $\mathbf{N}$ is the outward-bound normal to the curve. As is typical with curve evolution methods, we add a regularization term designed to keep the evolving contour smooth, yielding

$$\frac{\partial \mathbf{C}}{\partial t} - \left(\ln\frac{\sigma_i^2}{\sigma_o^2} + \frac{e^{2I(x,y)} - 1}{2\sigma_i^2} - \frac{e^{2I(x,y)} - 1}{2\sigma_o^2} + \alpha\kappa\right)\mathbf{N}. \tag{11}$$

where κ is the curvature, α is a constant. Performing similar derivations for the complex Gaussian case (IQ image) and Rayleigh case (Magnitude image). we present flows for these distributions in Table 1. Note that the complex Gaussian and Rayleigh flows and estimators have the same form.

Implementation

The curve evolutions presented above are totally general in that they apply to any closed contour representation. be it a spline. polygon. Fourier descriptor curve, or other such representation. The method requires an initial closed contour $\mathbf{C}$, which in this paper is a small square positioned by the user. For the region inside $\mathbf{C}$, we compute the maximum likelihood parameter σ_i^2 to characterize the distribution of intensities. Similarly, for the region outside the contour, we compute σ_o^2. With these parameters estimated. we can then move the points on the contour in the direction specified by $\frac{\partial \mathbf{C}}{\partial t}$ along its normal direction using the corresponding flow.

We choose to implement the technique using level set methods, which provide subpixel resolution and easily accommodate topological changes of the contour. The level set method models the evolving curve as the zero-level set of a higher-dimensional signed distance function, $\phi(x, y)$, which is negative inside the contour and positive outside. The values of $\phi(x, y)$ are updated as a function of time as $\frac{\partial \phi(x,y)}{\partial t} = -F\,|\nabla\phi|$, using a forward

Euler numerical scheme, where $F = \frac{\partial \mathbf{C}}{\partial t} \cdot \mathbf{N}$. For further details, please refer to Sethian (1999).

In Slabaugh et al (2006), we demonstrated that for log-compressed images, the Fisher-Tippett flow produces better results, both qualitatively and quantitatively, than the complex Gaussian/Rayleigh flow. Intuitively, this result is expected as one should choose the flow that best matches the data. For log-compressed images, this suggests that the Fisher-Tippett flow would be optimal. In our results section, we will compare the Fisher-Tippett flow to two other popular flows that have appeared in the literature: the Chan-Vese flow Chan and Vese (2001) and the non-parametric flow Cremers et al (2007). To review briefly, the Chan-Vese flow is based on the means inside and outside the contour, i.e.,

$$\frac{\partial \mathbf{C}}{\partial t} = \left((I - \mu_i)^2 - (I - \mu_o)^2 + \alpha\kappa \right) \mathbf{N}, \tag{12}$$

where μ_i and μ_o are the means inside and outside the contour, respectively. The non-parametric flow is given by

$$\frac{\partial \mathbf{C}}{\partial l} = \left(\ln \left(\frac{p_i(I)}{p_o(I)} \right) + \alpha\kappa \right) \mathbf{N}, \tag{13}$$

where $p_i(I)$ and $p_o(I)$ are the probabilities of a pixel with intensity I being inside and outside the contour, respectively, determined from Parzen windowing the histogram of pixels inside and outside the contour, respectively.

Results

In Figure 9, we show comparisons between the Chan-Vese, non-parametric, and Fisher-Tippett flows for synthetic images with decreasing contrast. The images were generated using a speckle noise generator, which multiplies zero-mean complex Gaussian noise to an image and then simulates the PSF by convolving with a Gaussian-weighted lowpass filter. This produced a synthetic IQ image for subsequent processing. The image in this example consists of four light gray targets on a dark gray background. By changing

the background color, we produced three images with decreasing contrast. For all three images. the standard deviation of intensities in a homogeneous region was approximately 23.2 units, and the contrast ratio (CR), defined as the ratio between the means of the bright and dark regions, was 1.77, 1.55, and 1.4, respectively. Each image was decorrelated, producing a total of six images. Going from left to right in Figure 9, the images were the original image with contrast ratio 1.77, its decorrelated version, the image with contrast ratio 1.55, its decorrelated version, the image with contrast ratio 1.4, and its decorrelated version. For each image, we applied the Chan-Vese (top row of the figure), non-parametric (middle row), and Fisher-Tippett (bottom row) flows, using an initialization of a small square (11x11 pixels) in the center of the target.

Qualitatively, the results in Figure 9 show some interesting trends. First, for higher contrast images (left two columns), all methods perform reasonably well on both the original image and decorrelated image. This is because there is adequate separation between the estimated distributions (as measured by the mean, Parzen-windowed density, or FT parameter) inside and outside of the evolving contour for the curve to propagate to the target boundaries. There are a few places in the original image (leftmost column) where the methods break apart around large speckles; however, performing a connected component analysis on the signed distance function, one could easily remove these. For the lowered contrast images (middle two columns), the combination of the speckle and reduced contrast causes the segmentations on the original images to get stuck in local minima that do not match the target boundaries; the evolving contours break apart and follow large speckles, producing a complex irregular shape. However, on the decorrelated images, the flows achieve a much better result, and the Chan-Vese and Fisher-Tippett flows achieve the best segmentations. For the lowest contrast images (right two columns), none of the segmentations are successful on the original image, and on the decorrelated image, only the Fisher-Tippett flow is fully successful in segmenting the target. The means used in the Chan-Vese flow are not sufficient for this example, as the contrast is too low. Also, using the histograms in the non-parametric method is not ideal, as the

histograms overlap significantly due to the poor contrast. Quantitative results showing the areas inside the converged contours are presented in Table 2. The ground truth area is 12868 units. We see that both qualitatively and quantitatively, the Fisher-Tippett flow on the decorrelated image produced the best results compared to the other methods.

In Figure 10, we apply the Chan-Vese (top row), non-parametric (middle row), and Fisher-Tippett flows (bottom row) to a tumor phantom image and its decorrelated version. This phantom has targets designed to mimic tumors, which appear with different levels of contrast. Each segmentation was initialized by placing a small contour inside the target at exactly the same location and size. The segmentation of the brightest target is shown in the left two columns, for the original and decorrelated images, respectively. As observed with the experiments with synthetic data, when the contrast is very strong, all methods are successful in delineating the borders of the target. However, when the contrast decreases, the statistical modeling of the data becomes important. In the second two columns of the figure, we show the results for the lighter target with less contrast. Here, decorrelation of the data helps achieve better results for all the methods, but only the Fisher-Tippett flow applied to the decorrelated image achieves a successful result.

In Figure 11, we demonstrate our method on an image of the carotid artery. As before, we show the Chan-Vese (top row), non-parametric (middle row), and Fisher-Tippett (bottom row) flows, applied to the original image (left column) and decorrelated image (right column). We initialized each segmentation with a small square located in the left-most side of the artery. The objective is for the segmentation to expand, to propagate the length of the artery and achieve a good segmentation. When applied to the original image, both the Chan-Vese and non-parametric flows get stuck in local minima and fail to evolve through the entire artery. However, the Fisher-Tippett flow in this case, with its better matching statistical model of the intensity distribution, is able to propagate the length of the artery. On the decorrelated image, all segmentations improve to various degrees. Both the Chan-Vese and non-parametric flows propagate farther; however, they still get stuck in a local minima. The best results for this experiment occur for the Fisher-Tippett

flow on the decorrelated image.

Another example of a carotid artery segmentation is shown in Figure 12. As before, we initialize the segmentation with a small square in the leftmost side of the artery. On the original image (left column), none of the flows is able to propagate the entire length of the artery; however, the Fisher-Tippett flow (bottom row), with its better modeling of the data, goes much farther along the artery than the Chan-Vese (top row) or non-parametric (middle row) flows. On the decorrelated image (right column), only the Fisher-Tippett flow propagates the entire length of the artery, and achieves a reasonable segmentation for this data.

We note that all of the flows studied in this paper are region-based flows. One could additionally add boundary-based terms Kass et al (1987); Casselles et al (1997), which would drive the active contours towards strong edges in the image. However, since the focus of this paper is region-based image segmentation, we leave this subject for future work.

Discussion and Summary

One question that has not yet been addressed is the following: which of the following produces the best results: the complex Gaussian flow, applied to the IQ image, the Rayleigh flow, applied to the envelope detected image, or the Fisher-Tippett flow, applied to the display image? The answer is that all flows produce the same result. Earlier, we showed that the complex Gaussian flow applied to the IQ image is identical to the Rayleigh flow applied to the envelope detected image. The Fisher-Tippett flow applied to the display image, also produces identical results, as we demonstrate in Figure 13. The top row shows the result of the Rayleigh flow applied to the envelope detected image (left) and display image (right), and the bottom row we show the Fisher-Tippett flow applied envelope detected image (left) and display image (right). The Rayleigh flow applied to the envelope

detected image produces the same result as the Fisher-Tippett flow on the display image. In both these cases. we chose the flow that matches the data. However, the other segmentations are unsuccessful, as the data does not match the model. More specifically, the result in the upper right part of the figure models Fisher-Tippett distributed data with a Rayleigh distribution, and the result in the lower left part of the figure models Rayleigh-distributed data with a Fisher-Tippett distribution. This experiment emphasizes the fact that one should choose the flow that best models the data.

In this paper we present ultrasound-specific methods for image segmentation. Speckle can be difficult to handle since it exhibits significant spatial correlation and does not generally follow a Gaussian distribution. Our method first decorrelates the ultrasound image using a whitening filter. We then perform maximum likelihood segmentation using region-based active contours and either the complex Gaussian, Rayleigh, or Fisher-Tippett distribution model, depending on from which stage of the imaging pipeline the image comes. We have derived the complex Gaussian flow and Fisher-Tippett flows, and have shown how the complex Gaussian flow is equivalent to the Rayleigh flow.

In this work, our piecewise constant model assumes that the image is bimodal mixture of (complex Gaussian, Rayleigh, or FT) distributions. This model is effective for a wide class of image segmentation problems, including the ones presented in this paper. However, the model can be extended in several ways. First, a multi-phase segmentation approach Vese and Chan (2002), using multiple level set functions to model different tissue classes can be used to produce multi-label segmentations. For images with statistics that vary spatially in a piecewise smooth fashion, one could adapt the Mumford-Shah model Mumford and Shah (1989) for the distributions presented in this paper. However, multi-phase implementations and piecewise smooth modeling of the image statistics are left for future work.

We believe our method may impact clinical workflows that utilize automated computational tools involving segmentation. One class of such tools is computer-aided diagnosis

applications that make measurements of anatomical structures. For example, one may be interested in the overall tumor load in a patient's anatomy, the regression or progression of lesions in follow-up scans so that accurate delineations of tumor boundaries are known, and automated tools rather than manual outlining are required for efficiency and repeatability of the analysis. A similar application would be in image-guided therapies that rely upon ultrasound guidance, as well as radiation therapies that require accurate delineation of lesion boundaries. The application of the method in such workflows is left for future work. We note that our method can be extended to 3D image volumes in a straightforward manner since our implicit framework generalizes to the 3rd dimension by adding an extra coordinate dimension in the implementation.

In all our experiments, we observed the best results when when we first decorrelate the image and then apply the flow that best matches the intensity distribution. We observe that our ultrasound-specific flows produce better results than other generic parametric or non-parametric flows. From these experimental results we conclude that the combined decorrelation and statistical region-based active contour results in improved segmentations. For future work, we are interested in more comprehensive validation studies on clinically important cases. Finally, we believe the theory underlying this paper will be useful in other applications, such as filtering, tracking, and registration; we plan on investigating these topics in the future.

References

Ayed I, Hennane N, Mitiche A. Unsupervised Variational Image Segmentation/Classification Using a Weibull Observation Model. IEEE Trans. Image Process.. 2006;15(11):3431–3439.

Ayed I, Mitiche A, Belhadj Z. Multiregion Level-Set Partitioning of Synthetic Aperture Radar Images. IEEE Trans. on Patt. Anal. and Mach. Intel.. 2005;27(5):793–800.

Bamber JC, Daft C. Adaptive Filtering for Reduction of Speckle in Ultrasonic Pulse-Echo

Images. Ultrasonics. 1986:41–44.

Casselles V, Kimmel R, Saprio G. Geodesic Active Contours. The Intl. Journal of Computer Vision. 1997;22(1):61 79.

Chan T, Vese L. Active Contours Without Edges. IEEE Trans. on Image Processing. 2001;10(2):266–277.

Chesnaud C, Réfrégier P, Boulet V. Statistical Region Snake-Based Segmentation Adapted to Different Physical Noise Models. IEEE Transactions on Pattern Analysis and Machine Intelligence. 1999;21(11):1145 1157.

Cremers D, Rousson M, Deriche R. A review of Statistical Approaches to Level Set Segmentation: Integrating Color, Texture, Motion and Shape. International Journal of Computer Vision. 2007;72:195–215.

Dutt V, Greenleaf J. Statistics of the Log-Compressed Envelope. Journal of the Acoustical Society of America. 1996;99(6):3817–3825.

Eltoft T. Modeling the amplitude statistics of ultrasonic images. IEEE Trans. on Med. Img.. 2006;25(2):229 240.

Goodman J, Speckle Phenomena: Theory and Applications, 1st Edition. Work in Progress. 2005.

Iraca D, Landini L, Verrazzani L. Power Spectrum Equalization for Ultrasonic Image Restoration. IEEE Trans. on Ultrasonics, Ferroelectrics, and Frequency Control. 1989;36(2):216 222.

Johnson N, Kotz S, Balakrishnan R, Continuous Univariate Distributions, seconds Edition. Vol. Volume 1. 1995.

Kass M, Witkin A, Terzopoulos D. Snakes: Active Contour Models. Intl. Journal of Computer Vision. 1987;1(4):321–331.

Kim J, Fisher JW, Yezzi A, Cetin M, Willsky A. A Nonparametric Statistical Method for Image Segmentation Using Information Theory and Curve Evolution. IEEE Trans. on Image Processing. 2005;14(10):1486 1502.

Martin-Fernandez M, Alberola-Lopez C. An Approach for Contour Detection of Human Kidneys from Ultrasound Images using Markov Random Fields and Active Contours.

Medical Image Analysis. 2005;9(1):1–23.

Michailovich O, Adam D. Robust Estimation of Ultrasound Pulses Using Outlier-Resistant De-Noising. IEEE Trans. on Medical Imaging. 2003;22(3):368 392.

Michailovich O, Tannenbaum A. Despeckling of Medical Ultrasound Images. IEEE Trans. on Ultrasonics, Ferroelectrics, and Frequency Control. 2006;53(1):64–78.

Mumford D, Shah J. Optimal Approximation by Piecewise Smooth Functions and Associated Variational Problems. Commun. Pure Appl. Math. 1989;42:577–685.

Noble JA, Boukerroui D. Ultrasound Image Segmentation: A Survey. IEEE Trans. on Medical Imaging. 2006;25(8):987 1010.

Sarti A, Corsi C, Mazzini E, Lamberti C. Maximum Likelihood Segmentation of Ultrasound Images with Rayleigh Distribution. IEEE Trans. on Ultrasonics, Ferroelectrics, and Freq. Control. 2005;52(6):947–960.

Sethian JA, Level Set Methods and Fast Marching Methods Evolving Interfaces in Computational Geometry. Fluid Mechanics, Computer Vision, and Materials Science. Cambridge University Press. 1999.

Slabaugh G, Unal G, Fang T, Wels M. Ultrasound-Specific Segmentation Via Decorrelation and Statistical Region-Based Active Contours. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition. 2006. pp. 45–53.

Tao Z, Tagare H. Tunneling Descent Level Set Segmentation of Ultrasound Images. In: Information Processing in Medical Imaging. 2005. pp. 750 761.

Unal G, Yezzi A, Krim H. Information-theoretic active polygons for unsupervised texture segmentation. The Intl. Journal of Computer Vision. 2005;62(3):199 220.

Vese L, Chan T. A Multiphase Level Set Framework for Image Segmentation Using the Mumford and Shah Model. International Journal of Computer Vision. 2002;50:271–293.

Wagner R, Smith S, Sandrik JM, Lopez H. Statistics of speckle in ultrasound b-scans. IEEE Transactions on Sonics and Ultrasonics. 1983;30(3):156 163.

Figure 1. Block diagram. The ultrasound imaging system produces an IQ image, which is the input provided to our approach, which appears within the dotted rectangle.

Data	Data Model	Parameter Estimation	Curve Evolution
IQ Image	Comp. Gaussian	$\sigma^2 = \frac{\int_\Omega \mathbf{I}(x,y)\|^2 dxdy}{2\int_\Omega dxdy}$	$\frac{\partial \mathbf{C}}{\partial t} = \left(\ln \frac{\sigma_i^2}{\sigma_o^2} + \frac{\|\mathbf{I}(x,y)\|^2}{2\sigma_i^2} - \frac{\mathbf{I}(x,y)\|^2}{2\sigma_o^2} + \alpha\kappa\right)\mathbf{N}$
Mag. Image	Rayleigh	$\sigma^2 = \frac{\int_\Omega I(x,y)^2 dxdy}{2\int_\Omega dxdy}$	$\frac{\partial \mathbf{C}}{\partial t} = \left(\ln \frac{\sigma_i^2}{\sigma_o^2} + \frac{I(x,y)^2}{2\sigma_i^2} - \frac{I(x,y)^2}{2\sigma_o^2} + \alpha\kappa\right)\mathbf{N}$
Display Image	FT	$\sigma^2 = \frac{\int_\Omega \left(e^{2I(x,y)}-1\right)dxdy}{2\int_\Omega dxdy}$	$\frac{\partial \mathbf{C}}{\partial t} = \left(\ln \frac{\sigma_i^2}{\sigma_o^2} + \frac{e^{2I(x,y)}-1}{2\sigma_i^2} - \frac{e^{2I(x,y)}-1}{2\sigma_o^2} + \alpha\kappa\right)\mathbf{N}$

Table 1

Flows

	Orig. CR = 1.77	Dec. CR = 1.77	Orig. CR = 1.55	Dec. CR = 1.55	Orig. CR = 1.41	Dec. CR = 1.41
Chan-Vese	13024	12898	2906	12753	3713	2380
Non-Parametric	13037	13051	5721	12481	4881	9782
Fisher-Tippett	12946	12930	6772	12831	5049	12486

Table 2

Areas of the different segmentations in Figure 9. Ground truth area is 12868 units. Contrast ratio (CR) is defined as the ratio between the means of the bright and dark regions.

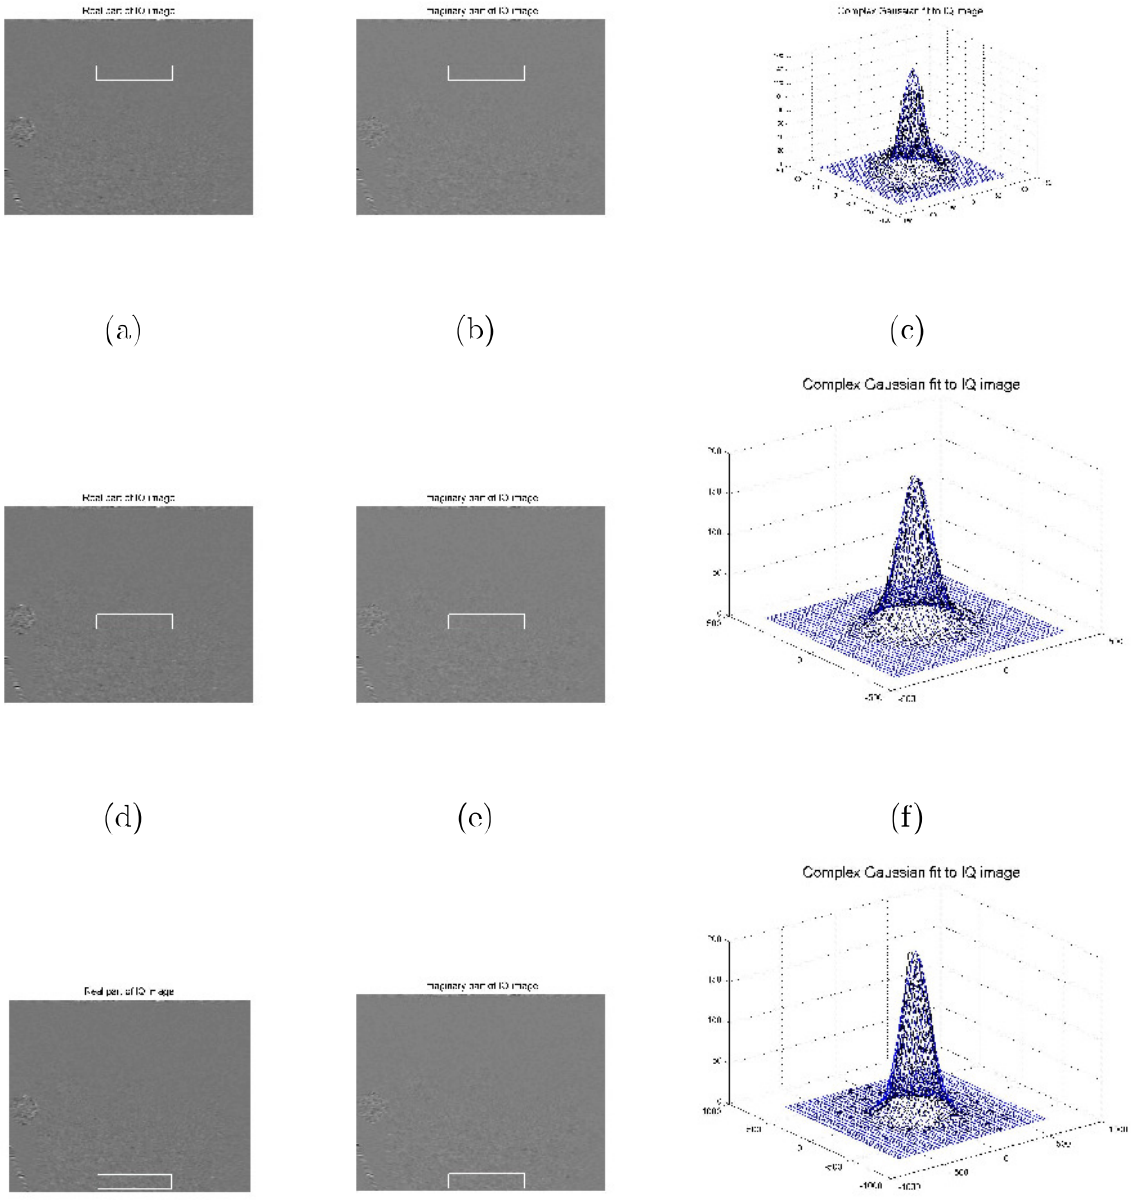

(a) (b) (c)

(d) (e) (f)

(g) (h) (i)

Figure 2. Using complex Gaussian distributions to model ultrasound image intensities in different parts of the IQ image. Left column: real part of the IQ image, middle-left column: imaginary part of the IQ image, right column: complex Gaussian distribution fit to histogram. We select regions in the near-field (top row), mid-field (middle row) and far-field (bottom row).

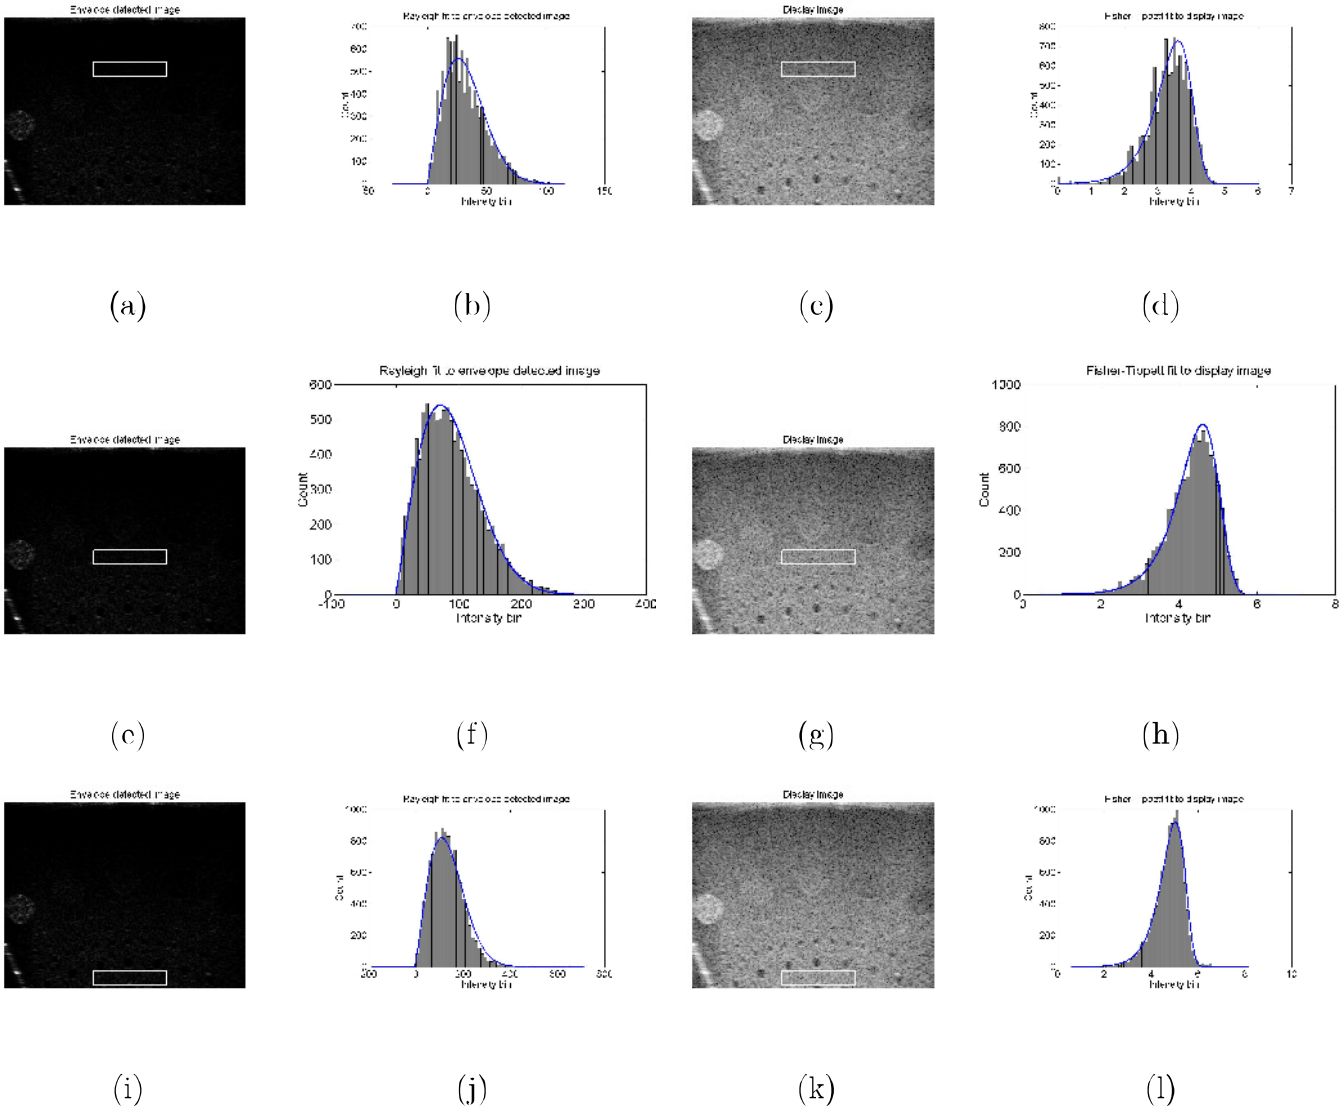

Figure 3. Using Rayleigh and Fisher-Tippett distributions to model ultrasound image intensities in different parts of the image. Left column: envelope-detected image with region selected, middle-left column: Rayleigh distribution fit to intensities in region, middle-right column: display image with region selected, right column: Fisher-Tippett distribution fit to intensities in region. We select regions in the near-field (top row), mid-field (middle row) and far-field (bottom row).

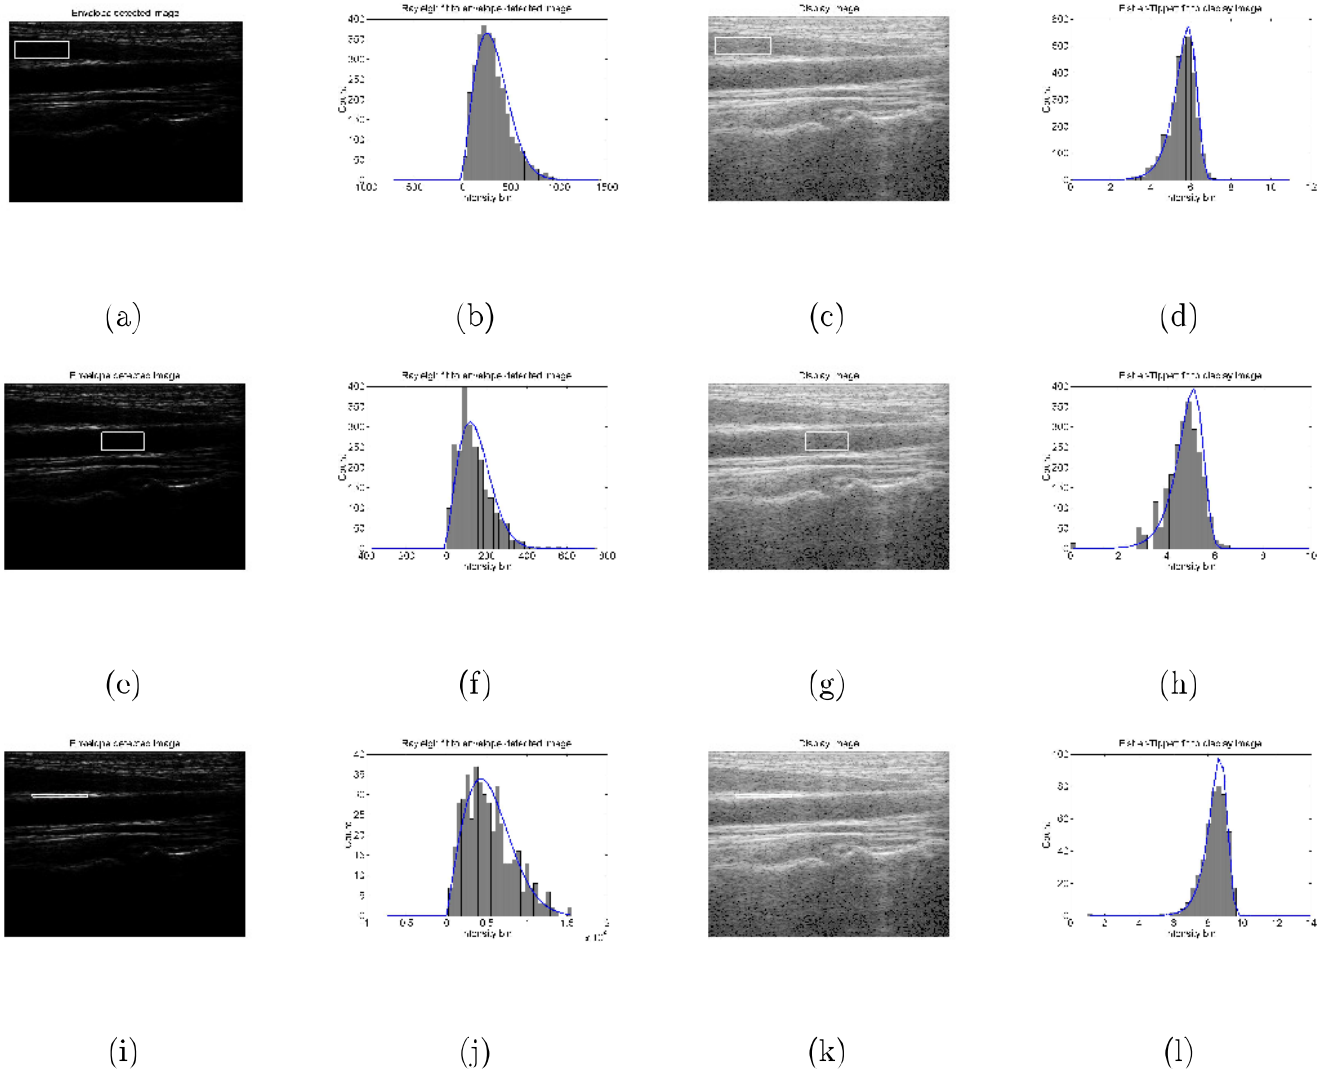

Figure 4. Fitting Rayleigh and Fisher-Tippett distributions to different regions in an ultrasound image of a carotid artery. For each row, left to right: magnitude image, Rayleigh distribution fit to the histogram, display image, FT distribution fit to the histogram.

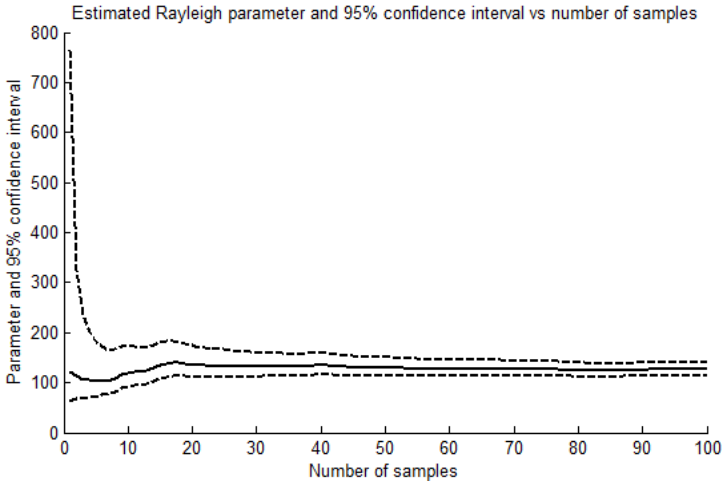

Figure 5. Confidence intervals for the Rayleigh estimator as a function of number of samples taken from the window shown in Figure 4(g).

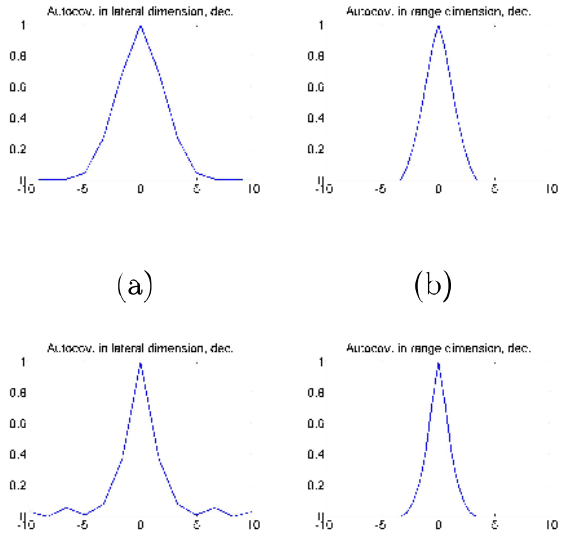

Figure 6. The autocovariance function of the selected region from Figure 3 (a) is shown in the lateral (horizontal) dimension in (a) and the range (vertical) direction (b). The corresponding result for the decorrelated image is shown in (c) and (d).

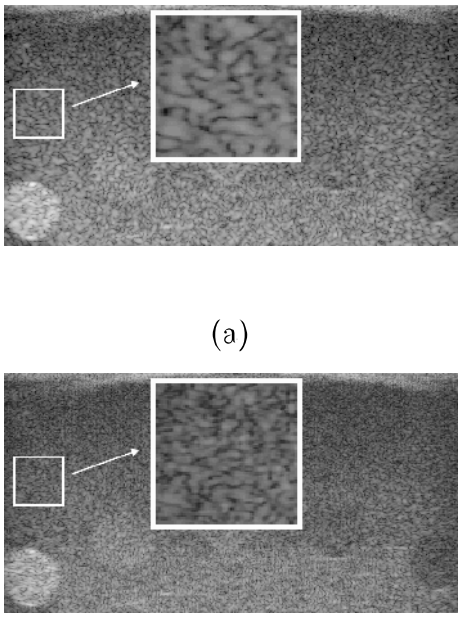

(a)

(b)

Figure 7. Decorrelation decreases speckle size. Original image (a) and decorrelated image (b).

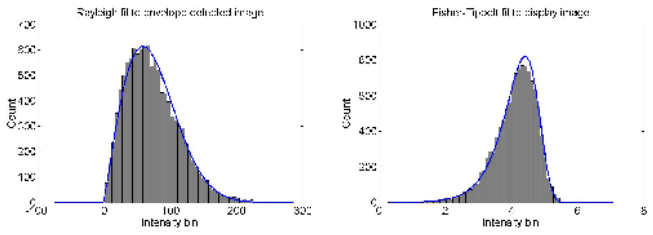

(a) (b)

Figure 8. Previous distributions apply to the decorrelated image for the region in the middle of the image (corresponding to 3 (a) and (c). The Rayleigh fit to the magnitude IQ decorrelated image is shown in (a), and the Fisher-Tippett fit to the log magnitude decorrelated IQ image is shown in (b).

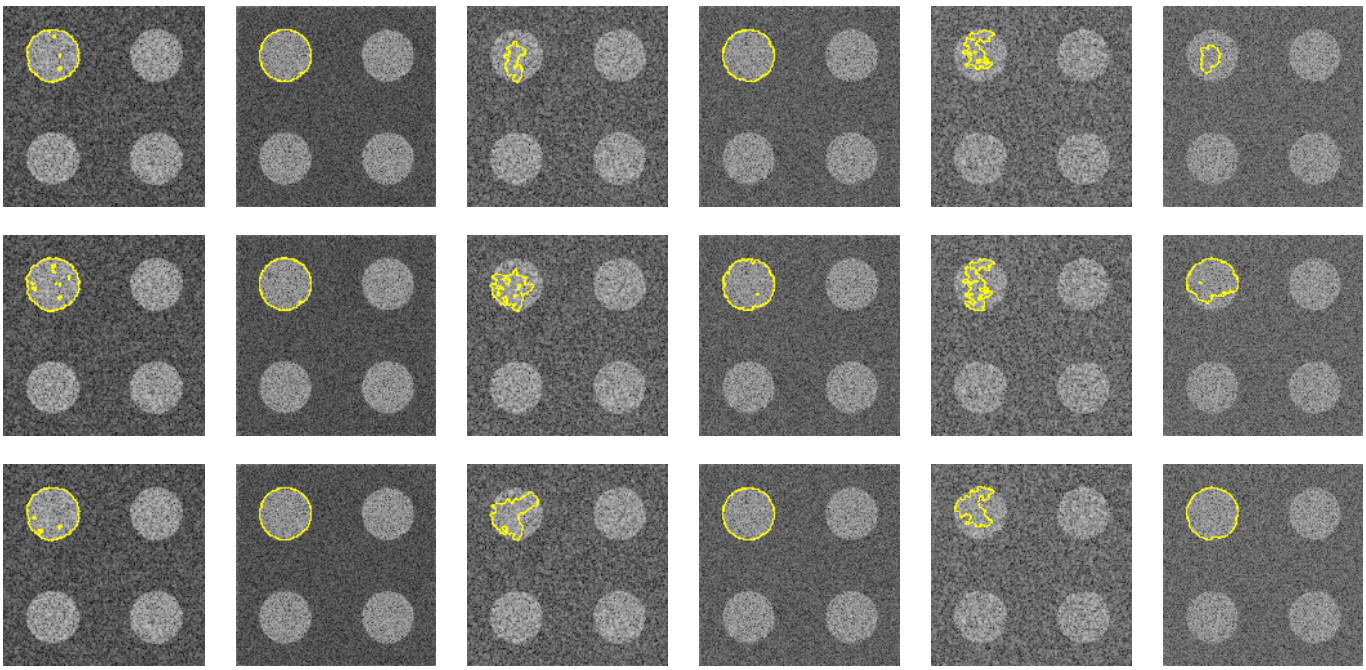

Figure 9. Comparison of segmentations of synthetically generated data. Top row: Chan-Vese flow. Middle row: Non-parametric flow. Bottom row: Fisher-Tippett flow. Left to right: We show both the original and the decorrelated image for three decreasing contrasts. For higher contrast images, all methods work reasonably well. For lowered contrast, decorrelating the image improves the results significantly. For the lowest contrast, the best result occurs for the Fisher-Tippett flow applied to the decorrelated image.

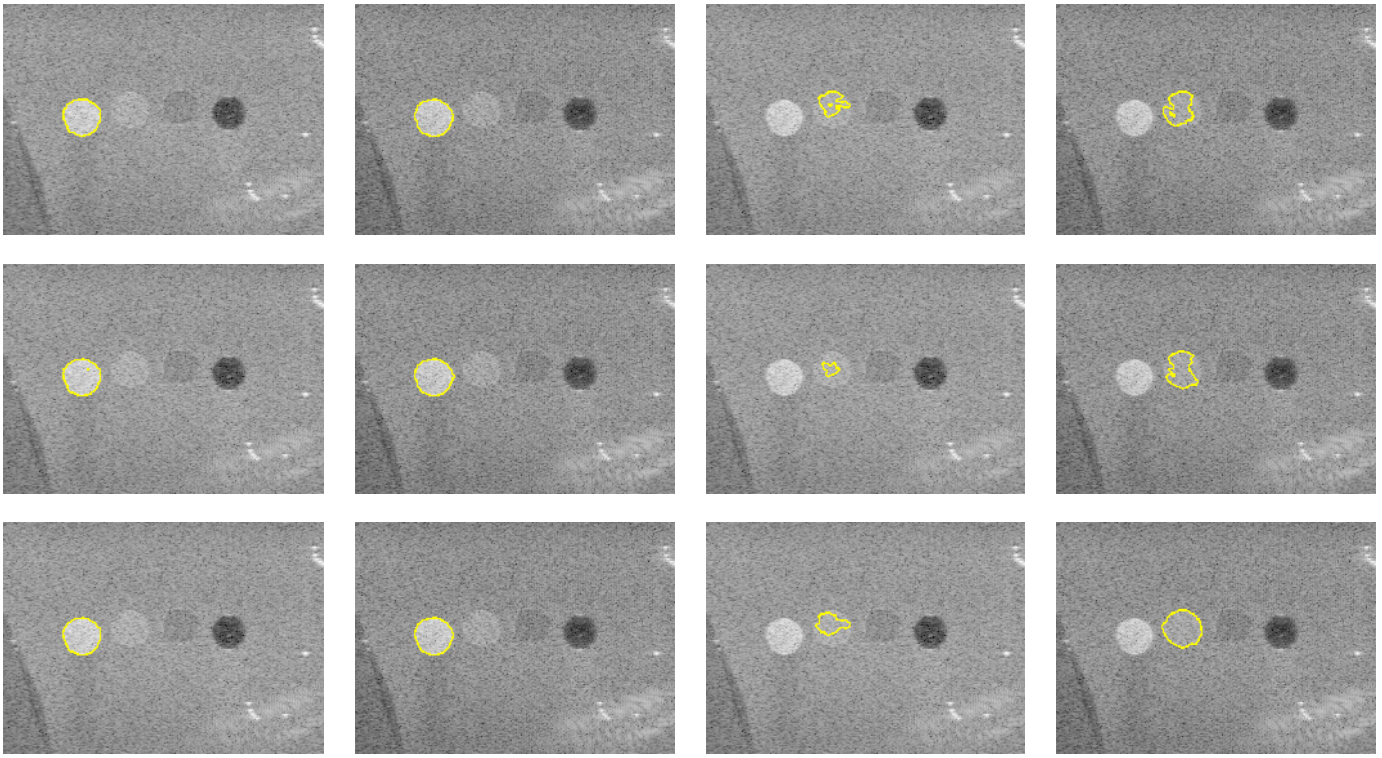

Figure 10. Chan-Vese (top row), non-parametric (middle row), and Fisher-Tippett flow (bottom row) applied to a tumor phantom image. Left two columns: segmentation of the brightest target, on the original (left) and decorrelated (right) image. Right two columns: segmentation of a target with less contrast, for the original (left) and decorrelated (right) image. When the contrast diminishes, it becomes increasingly important to decorrelate the image as well as use the most appropriate flow for the intensity distribution.

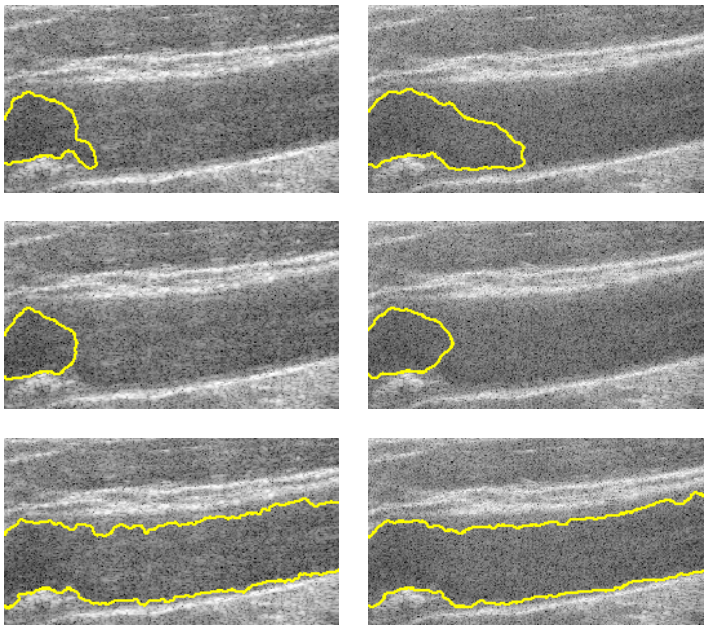

Figure 11. Segmentation of a carotid image. Chan-Vese (top row). non-parametric (middle row), and Fisher-Tippett flow (bottom row). Original (left column) and decorrelated image (right column). The best results occur for the Fisher-Tippett flow on the decorrelated image.

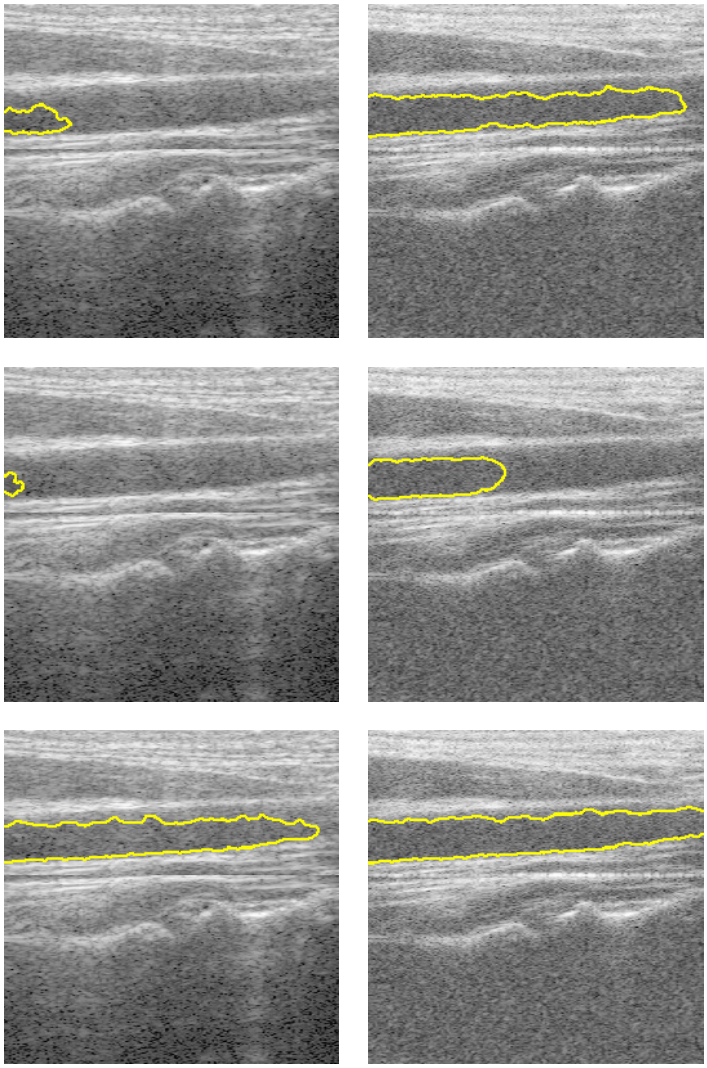

Figure 12. Another example carotid artery segmentation. Chan-Vese (top row), non-parametric (middle row), and Fisher-Tippett flow (bottom row). Original (left column) and decorrelated image (right column). The best results occur for the Fisher-Tippett flow on the decorrelated image.

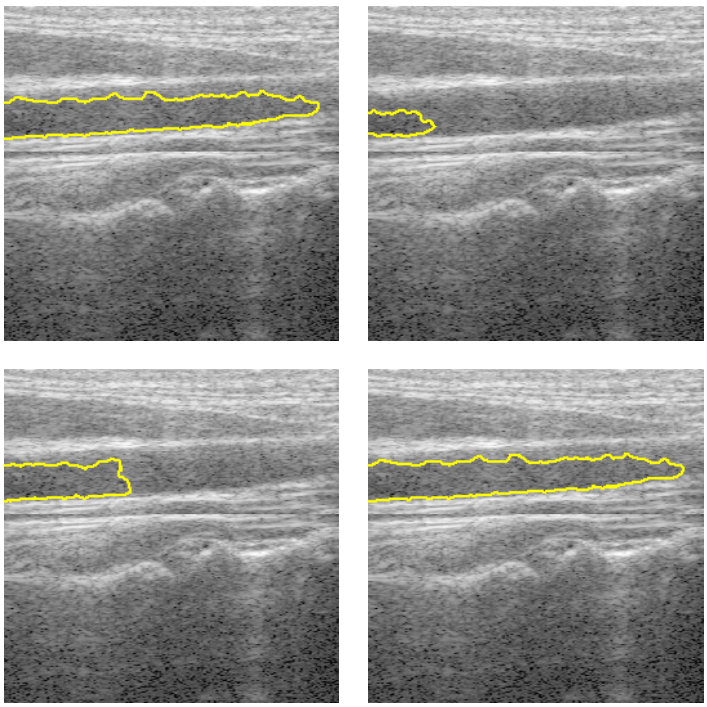

Figure 13. Choosing the model that best matches the data. In the top row, we show the complex Gaussian / Rayleigh flow applied to the envelope-detected image (left) and display image (right). Here, the model matches the data for the example on the left but not on the right. In the bottom row, we show the Fisher-Tippett flow applied to the envelope-detected image (left) and display image (right). Here, the model matches the data for the example on the right but not on the left. Furthermore, we note that the examples in the upper left and lower right of this figure produce identical results.