

Bayesian Speckle Tracking. Part I: An Implementable Perturbation to the Likelihood Function for Ultrasound Displacement Estimation

Brett Byram, *Student Member, IEEE*, Gregg E. Trahey, *Member, IEEE*,
and Mark Palmeri, *Member, IEEE*

Abstract—Accurate and precise displacement estimation has been a hallmark of clinical ultrasound. Displacement estimation accuracy has largely been considered to be limited by the Cramer–Rao lower bound (CRLB). However, the CRLB only describes the minimum variance obtainable from unbiased estimators. Unbiased estimators are generally implemented using Bayes’ theorem, which requires a likelihood function. The classic likelihood function for the displacement estimation problem is not discriminative and is difficult to implement for clinically relevant ultrasound with diffuse scattering. Because the classic likelihood function is not effective, a perturbation is proposed.

The proposed likelihood function was evaluated and compared against the classic likelihood function by converting both to posterior probability density functions (PDFs) using a non-informative prior. Example results are reported for bulk motion simulations using a 6λ tracking kernel and 30 dB SNR for 1000 data realizations. The canonical likelihood function assigned the true displacement a mean probability of only 0.070 ± 0.020 , whereas the new likelihood function assigned the true displacement a much higher probability of 0.22 ± 0.16 . The new likelihood function shows improvements at least for bulk motion, acoustic radiation force induced motion, and compressive motion, and at least for SNRs greater than 10 dB and kernel lengths between 1.5 and 12λ .

I. INTRODUCTION

ONE of the hallmarks of clinical ultrasound has been accurate and precise displacement estimation. The effectiveness of ultrasonic displacement estimation has allowed it to become the dominant modality for imaging blood flow and has largely made new ultrasound-based imaging modalities such as acoustic radiation force impulse (ARFI) imaging and strain imaging possible.

Generally, further development of clinical ultrasonic displacement estimators focuses on some combination of improving computational efficiency [1]–[4] and expanding the parameter space over which estimators approach the Cramer–Rao lower bound (CRLB) [5]–[7]. The notion of a CRLB for the ultrasound displacement estimation problem has existed at least since Embree [8]. Although the

CRLB is a useful benchmark for displacement estimation, it only describes the limit of a minimum variance unbiased estimator (MVUE) rather than a fundamental limit for all displacement estimators. Estimators with a small amount of algorithmic bias can be designed to produce estimates that have a significantly lower mean-square error (MSE) than that described by the CRLB [9], [10].

Estimators surpassing performance defined by the CRLB have not been described in the ultrasound displacement estimation literature; however, the notion of a biased estimator does exist. Several groups have produced estimation schemes in which the displacement search region is shifted and reduced based on displacements measured at adjacent positions [1], [2], [11]. (These approaches can be seen as specific realizations of the methods developed here.) Another method that has recently gained attention specifically for strain imaging is regularized displacement estimation [12]. Regularized displacement estimation attempts to maximize (or minimize) the combination of a traditional signal similarity factor (e.g., normalized cross-correlation, sum absolute difference, etc.) and a continuity constraint. The continuity constraint is usually constructed to penalize large displacement gradients and acts to bias the final displacement estimate. Although regularized displacement estimation is appropriately conceptualized as a biased algorithm, the publications featuring these algorithms do not provide sufficiently detailed results to determine whether these algorithms achieve an improved MSE over an MVUE [12]–[15].¹

Ultrasound displacement estimation (and specifically biased estimation) needs an implementable likelihood function. Likelihood functions are extremely useful because they can be employed as part of a Bayesian framework to express knowledge of a parameter (such as tissue displacement) as a probability density function (PDF) [9]. These PDFs can then, in turn, be used to create biased parameter estimates. Biased parameter estimates have some attractive properties and will be the subject of the companion paper. This paper will focus on demonstrating an effective and implementable likelihood function for ultrasonic displacement estimation problems.

Manuscript received May 11, 2012; accepted August 30, 2012. This research was funded by National Institutes of Health grants R37HL096023 and T32EB001040.

The authors are with the Department of Biomedical Engineering, Duke University, Durham, NC (e-mail: bcb16@duke.edu).

DOI <http://dx.doi.org/10.1109/TUFFC.2013.2545>

¹It is suspected that the specific implementations of regularized displacement estimation are too coarsely sampled to achieve any improvements of the mean-square error.

II. METHODS

A. Overview

Bayes' theorem is a simple equation describing the appropriate mechanism for combining previously and newly acquired information [9]. Bayes' theorem will be used, in the case of the ultrasound displacement estimation problem, to appropriately combine information from a local similarity function and prior information about the displacement to provide a better estimate for the current displacement estimate. For the purposes of this paper, Bayes' theorem is expressed as

$$p_m(\tau_0 | x) = \frac{p_m(x | \tau_0)p_m(\tau_0)}{\int p_m(x | \tau_0)p_m(\tau_0)d\tau_0}, \quad (1)$$

where x is the data, τ_0 is the displacement, and m indexes kernel location axially. The term $p_m(x | \tau_0)$ denotes the likelihood function for the m th kernel axially, $p_m(\tau_0)$ is the prior distribution, and $p_m(\tau_0 | x)$ is the posterior distribution. The likelihood function is the means by which data are incorporated into the final PDF describing τ_0 . The prior PDF expresses previous knowledge of the displacement independent of the data. The posterior is the final PDF and represents the final state of knowledge after the current data and previous knowledge have been combined. The likelihood function, $p_m(x | \tau_0)$, will be the focus of the rest of this paper.

B. Likelihood Function for Ultrasonic Time-Delay Estimation

The likelihood function for the generic displacement estimation problem can be found in several sources [10], [16]–[18], and is

$$\begin{aligned} p_m(x | \tau_0) &= \frac{1}{(4\pi\sigma_{\text{noise}}^2)^{N/2}} \exp\left[-\frac{1}{4\sigma_{\text{noise}}^2} \sum_{n=0}^{N-1} s_1(n\Delta)^2\right] \\ &\times \exp\left[-\frac{1}{4\sigma_{\text{noise}}^2} \sum_{n=\tau_0/\Delta}^{\tau_0/\Delta+M-1} -2s_1(n\Delta)s_2(n\Delta + \tau_0)\right] \\ &\times \exp\left[-\frac{1}{4\sigma_{\text{noise}}^2} \sum_{n=\tau_0/\Delta}^{\tau_0/\Delta+M-1} s_2(n\Delta + \tau_0)^2\right], \end{aligned} \quad (2)$$

where $s_1()$ is the original signal, $s_2()$ is the displaced signal, τ_0 is the displacement between signals, Δ is the sampling period, M is the number of samples in the data record (i.e., kernel length), and σ^2 is the noise power.² The likelihood function in (2) is often used to derive the maximum likelihood estimator, which has a form that resembles the

similarity metrics commonly used for ultrasound displacement estimation.³

For use with Bayes' theorem to calculate displacement probabilities (and eventually displacement estimates), the likelihood function can be expressed in a reduced form because terms that are not a function of τ_0 will be eliminated when the posterior distribution is calculated using Bayes' theorem (1). Without losing any relevant information, the likelihood function can be represented as

$$p_m(x | \tau_0) \propto \exp\left[\frac{1}{2\sigma_{\text{noise}}^2} \sum_{n=\tau_0/\Delta}^{\tau_0/\Delta+M-1} s_{m1}(n\Delta)s_{m2}(n\Delta + \tau_0)\right]. \quad (3)$$

The argument of (3) is nearly identical to the argument of the MLE in footnote 3. The sole additional term is a scaling parameter that modulates the width of the likelihood function based on the thermal noise power. The noise power scaling is very intuitive. It allows the likelihood function to be continuously transformed between a δ -function as the noise power approaches zero and a rectangle function the size of the specified search region as the noise power approaches ∞ .

There are two practical difficulties with (3). The first problem is that estimates of electronic noise power are difficult to obtain in *in vivo* (or even phantom) scenarios and often will not even represent the dominant source of estimation noise [22]–[24]. The second problem is that unnormalized cross-correlation is known to be a poor similarity metric for ultrasound displacement estimation [25]. A possible solution to these problems is to perturb (3) to create a likelihood function more suitable for the ultrasonic displacement estimation problem. To this end, we propose the following likelihood function:

$$p_m(x | \tau_0) \propto \exp\left[\frac{\text{SNR}}{\alpha} \sum_{n=\tau_0/\Delta}^{\tau_0/\Delta+M-1} \frac{s_{m1}(n\Delta)s_{m2}(n\Delta + \tau_0)}{\sqrt{\sigma_{s_{m1}}^2 \sigma_{s_{m2}}^2}}\right], \quad (4)$$

where σ_{m1}^2 and σ_{m2}^2 are the variances of the portions of each signal used for each step of the correlation function calculation, α is a scaling term that will be empirically shown to be primarily a function of kernel length, and SNR is the signal-to-noise ratio. The proposed likelihood function is now no longer scaled by the electronic noise power but instead a generic SNR term. Additionally, cross-correlation has been replaced by normalized cross-correlation. The generic SNR term allows other noise

³The maximum likelihood estimator (MLE) is

$$\hat{\tau}_0 = \arg \max_{\tau_0} \sum_{n=\tau_0/\Delta}^{\tau_0/\Delta+M-1} s_1(n\Delta)s_2(n\Delta + \tau_0).$$

This equation describes un-normalized cross-correlation; however, usually in ultrasound, normalized cross-correlation is treated as a substitute MLE. The MLE is attractive because it asymptotically approaches the CRLB as M (the length of the data record) approaches ∞ [10]. In practice, it has been observed that performance of the MLE approaches the CRLB with kernel lengths that are 4λ or less [20], [21].

²The noise power has been doubled to reflect the noise present on both signals based on the calculations presented by Walker [19].

sources to be introduced, specifically signal correlation. The two methods of determining the SNR will be the normal method of calculating the thermal SNR,

$$\text{SNR}_{\text{thermal}} = \frac{P_{\text{signal}}}{P_{\text{noise}}}, \quad (5)$$

and the peak-correlation-coefficient-derived estimate of the SNR [26],

$$\text{SNR}_{\rho} = \frac{\rho_{\text{max}}}{1 - \rho_{\text{max}}}, \quad (6)$$

where ρ_{max} is the peak of the normalized cross-correlation function. Normalized cross-correlation is shown in (4) as everything to the right of, and including, the summation. The thermal SNR will be treated as a known value because the appropriate levels of noise added during data simulation are known. The SNR calculated from the maximum correlation will be an estimated quantity.⁴

C. Likelihood Function Evaluation

To compare the exact analytic likelihood function with the modified likelihood function, a quantitative performance metric is required. A useful quantitative metric of likelihood function performance should compare how discriminative the various likelihood functions are for a given quality of data. In terms of probabilities, this means the probability assigned to a given displacement should correspond to the quality of the data. To facilitate likelihood function evaluation, likelihood functions (which are not PDFs) can be converted to posterior probability distributions using a non-informative prior:

$$p_m(\tau_0 | x) = \begin{cases} \frac{p_m(x | \tau_0)}{\int_{\tau_0}^{\tau_\beta} p_m(x | \tau_0) d\tau_0} & \text{if } \tau_\alpha < \tau_0 < \tau_\beta \\ 0 & \text{otherwise,} \end{cases} \quad (7)$$

where τ_α and τ_β are the upper and lower limits of the traditional search region. The non-informative prior for displacement estimation is a uniform PDF corresponding to the search region that would have been used for a maximum likelihood estimate (e.g., normalized cross-correlation). Using the posterior distribution, the quality metric will be the probability of the true displacement expressed by the posterior distribution, calculated as

⁴It may seem reasonable to combine the correlation SNR and the thermal SNR into a total SNR [27] using

$$\text{SNR}_{\text{total}} = \frac{\text{SNR}_{\rho} \text{SNR}_{\text{thermal}}}{1 + \text{SNR}_{\rho} + \text{SNR}_{\text{thermal}}};$$

however, this equation is only relevant if the correlation coefficient is not corrupted by thermal noise. In all cases when both noise types are present—including simulations—the correlation coefficient reflects both noise sources.

TABLE I. BASIC SIMULATION PARAMETERS.

Parameter	Value
Center frequency	5 MHz
Bandwidth	50%
Pulse envelope	Gaussian
C	1540 m/s
depth _{max}	5 cm
Sampling frequency	10 GHz
Kernel length	3λ
SNR	20 dB
Strain	1%

$$\text{quality metric} = \prod_{m=0}^{M-1} \int_{\tau_{\text{true}} - \epsilon}^{\tau_{\text{true}} + \epsilon} p_m(\tau_0 | x) d\tau_0, \quad (8)$$

where M is the number of kernels through depth, τ_{true} is the known displacement, and ϵ describes the range of lags to integrate over. ϵ is required because the probability of a single point in a continuous distribution is always zero, and as will be seen, in some cases, the confidence in τ_{true} is several pixels.

Additionally, the quality metric has been expressed using an integral in (8) because τ_0 is appropriately modeled as continuous. In practice, the quality metric in (8) will not be realizable using the methods presented here, for which only a sampled distribution will be available. Therefore, a more practical quality metric will be used for the data in this study:

$$\text{quality metric} = \prod_{m=0}^{M-1} \sum_{n=\lfloor (\tau_{\text{true}} - \epsilon)/\Delta \rfloor}^{\lfloor (\tau_{\text{true}} + \epsilon)/\Delta \rfloor} p_m(n\Delta | x), \quad (9)$$

where $\lfloor \bullet \rfloor$ denotes the nearest integer.

To reiterate, the quality metric represents the probability of the true displacement profile through depth occurring given the data, assuming the displacements are independent (enforceable through data downsampling). This quality metric was chosen with two intentions in mind.

TABLE II. ACOUSTIC RADIATION FORCE IMPULSE (ARFI) IMAGING SIMULATION PARAMETERS.

Parameter	Value
Tracking center frequency	7 MHz
Tracking f-number	0.5
Radiation force center frequency	2.22 MHz
Radiation force f-number	2
Radiation force pulse duration	180 μs
Focal depth	2 cm
Tracking pulse transmit cycles	2
C	1540 m/s
depth _{max}	5 cm
Sampling frequency	100 MHz
Kernel length	1.5λ
SNR	20 dB
ARFI PRF	10 kHz
Young's modulus	8.5 kPa
ρ	1.0 g/cm ³
ν	0.499

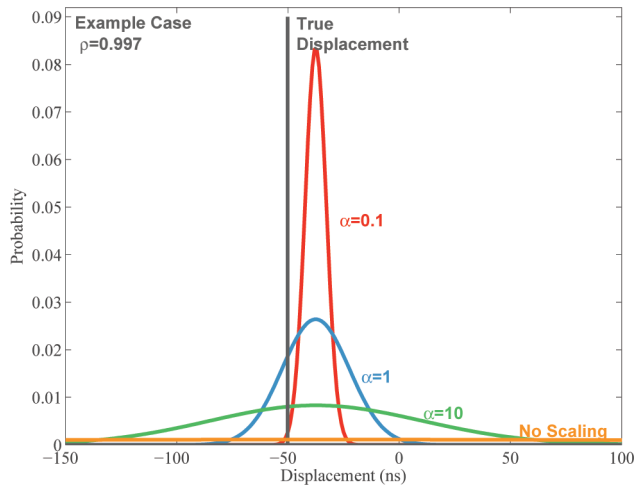


Fig. 1. This figure demonstrates how the probability density function concentrates for several values of the scaling parameter α from (4). The right concentration of the function provides the most probability to the true displacement. The probabilities are shown as posterior probability density functions, which are made using (4) and a uniform prior distribution. For the example shown, $\alpha = 1$ would give the highest probability to the true displacement. The example comes from a bulk motion displacement simulation with 20 dB of thermal noise added to the signal and a 3λ kernel.

First, when only relative probabilities of displacement are desired, it makes sense to have a probability distribution that is expected to give the highest possible probability to the true displacement. Second, when biased estimates are implemented in the companion paper [28], the distribution that maximizes the probability of the true displacement should provide the most appropriate representation of the data relative to the prior distribution. Without appropriate representation of the data, there is a strong danger that the data will be inappropriately overwhelmed by prior information.

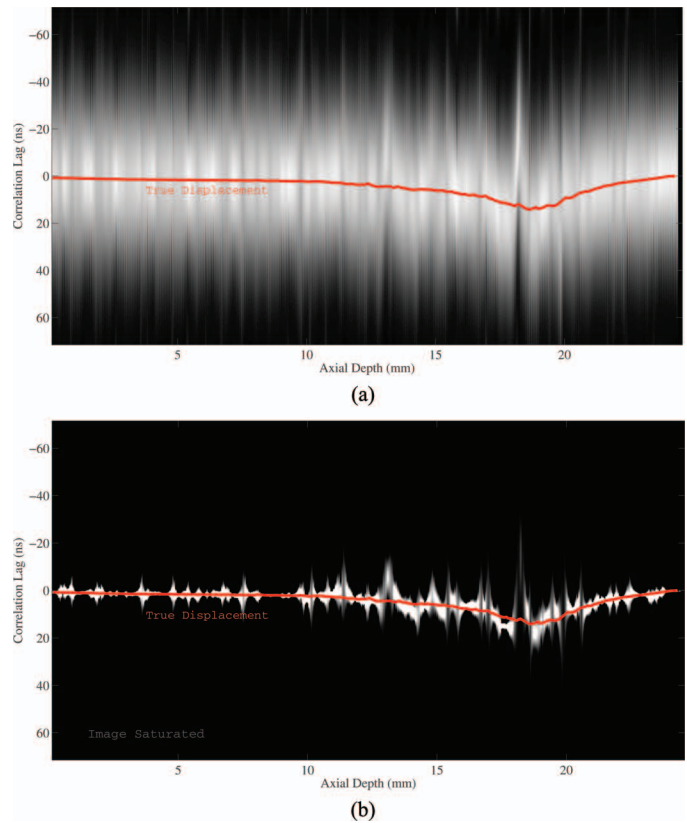


Fig. 2. These images show the final acoustic radiation force impulse (ARFI) posterior distributions for each depth (a) without and (b) with similarity metric scaling. The true displacement is also shown in red. The example demonstrates that scaling the similarity metric significantly concentrates the probability about the true displacement. The image with appropriate scaling is saturated because the dynamic range of the probability density functions (PDFs) through depth is too high to visualize otherwise.




Fig. 3. This figure compares the mean versus the median as the statistic describing the effectiveness of various likelihood functions. The mean of all of the probabilities for each value of the scaling parameter α is shown in green. The median value is the red line on the box plots. The box plots show the 25% and 75% percentiles and the box plot whiskers show 1.5 times the interquartile distance beyond the closest quartile. The red points show the probability that lay outside the 1.5 interquartile distance. The figure shows that the median is discriminative, but the mean is misleading.

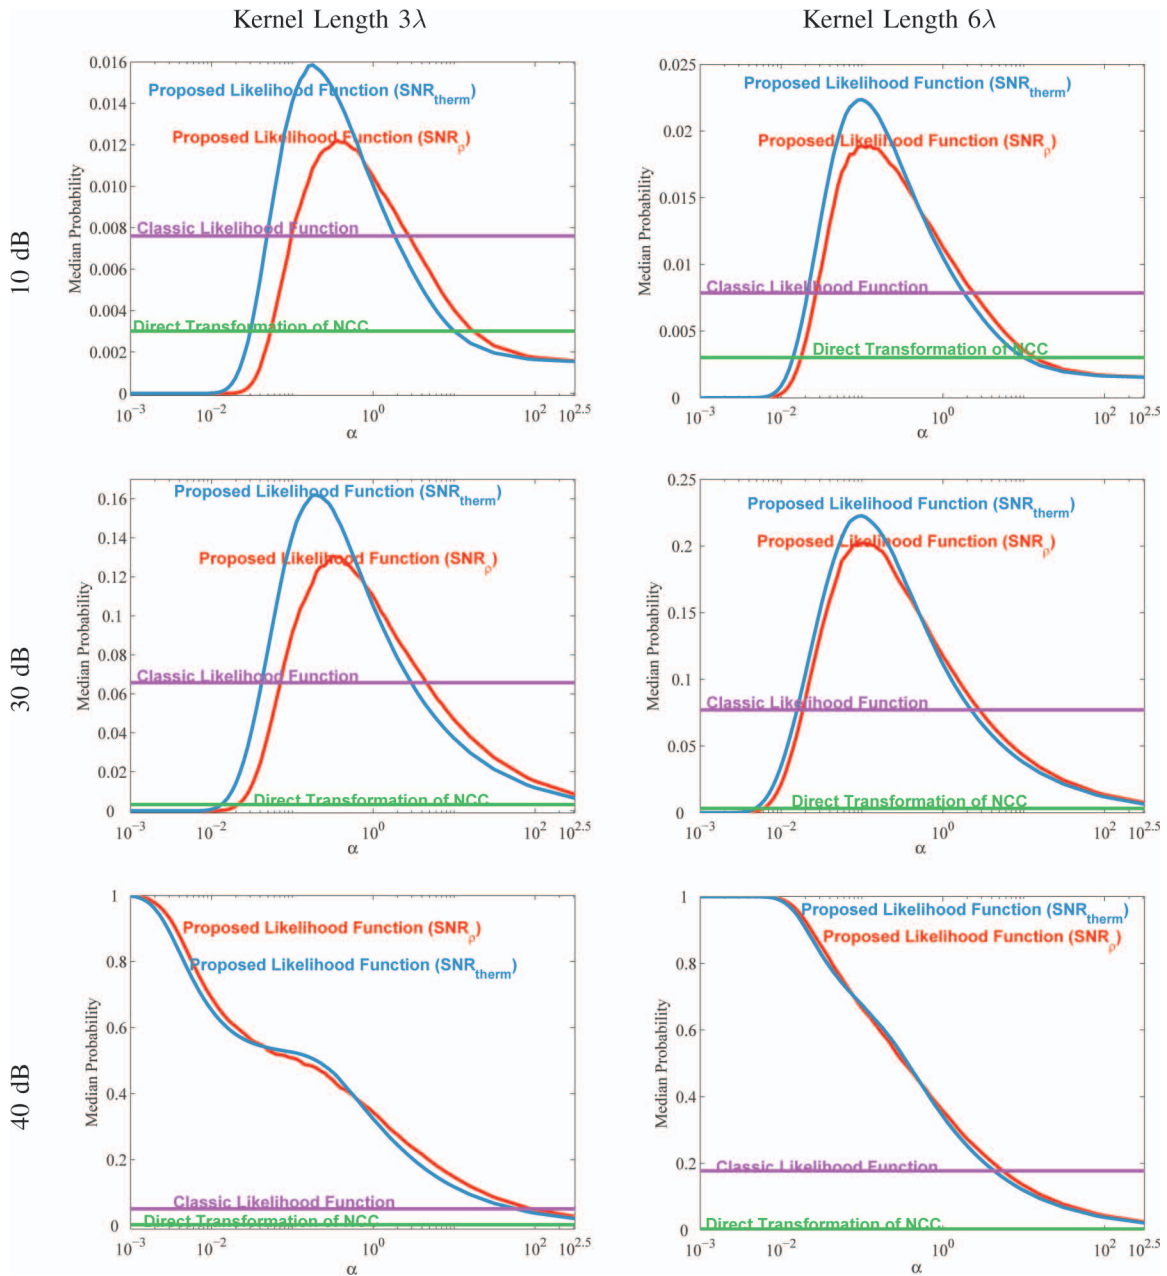


Fig. 4. This figure compares different methods for calculating the likelihood function for bulk motion-induced displacements. The various likelihood functions' performances are plotted as a function of α . The comparisons are shown for several thermal noise levels—10, 30, and 40 dB—and two different kernel sizes— 3λ and 6λ . The proposed likelihood function implemented with $\text{SNR}_{\text{therm}}$ or SNR_{ρ} is more discriminative than the classic likelihood function. The proposed likelihood function implemented with $\text{SNR}_{\text{therm}}$ always outperforms the likelihood function implemented with SNR_{ρ} . This is expected because $\text{SNR}_{\text{therm}}$ is a known quantity from the simulations as opposed to an estimated quantity like SNR_{ρ} . The difference between these two methods decreases as the kernel length increases resulting in a better estimate by SNR_{ρ} . The 40 dB case is shown for completeness but represents a degenerate case, as described in the text.

D. Likelihood Function: Simulations

The methods for computing the likelihood function are evaluated using several different sets of simulations, each with different levels of correlation noise. Simulations were performed for 1-D scatterer configurations for bulk displacements and compression-induced (i.e., strain) displacements. Simulations were also performed for 3-D scattering geometries for acoustic radiation force (ARF)-induced displacements.

The 1-D scattering geometry simulations were all performed using convolution with a static point spread function and complex pulse. To allow for arbitrary scatterer placement (rather than a fixed grid) the complex pulse was phase-shifted based on the position of each scatterer. Resulting complex signals were stripped of their imaginary components to produce RF A-lines. The simulations were performed with an average of 35 scatterers per -6 -dB resolution cell; 12 to 15 scatterers are adequate, but more scatterers were included to avoid gaps of scatterers




Fig. 5. This figure shows the value of α that maximizes the median probability. The data are displayed as a function of kernel length for several levels of thermal noise.

for step displacements or large strains [29]. The initial set of scatterer positions for a given realization was positioned randomly with a Gaussian distribution. The amplitude of each scatterer was also randomly distributed with a Gaussian distribution.

The thermal noise for all the simulations was modeled as a band-limited additive Gaussian random process. The band-limit of the noise was based on the pulse characteristics of the simulated signals.

The proposed form of the likelihood function shown in (4), calculated using $\text{SNR}_{\text{thermal}}$ and using SNR_{ρ} for a range of α values, is compared with the classic likelihood function (3). The proposed form of the likelihood function is also compared with a likelihood function constructed from a normalized cross-correlation function but without any scaling term to demonstrate the scaling term's importance. The likelihood functions are compared using the quality metric shown in (8). The actual calculation of the quality metric varies slightly for each simulation and will be addressed along with the specific description of each set of simulation data.

The statistic used to assess likelihood performance will be the median. Median statistic usage will be empirically justified as more appropriate than the mean in Section III.

The bulk displacement simulations to evaluate likelihood function efficacy were performed first. The bulk motion simulation data was comprised of 1000 1-D RF-data pairs with a constant displacement through depth between the two A-lines. The actual bulk displacement between the pairs of simulated A-lines were drawn from a normal distribution with zero-mean and a standard deviation of $\lambda/20$ to avoid any bias or artifact that could be introduced by consistent patterns of sub-sample displacements relative to the sampling frequency. The relatively small displacement range ($\sigma = \lambda/20$) was chosen to decrease the necessary search region size and the corresponding computational overhead of calculating large displacements. Unless otherwise specified, the simulation parameters are

those shown in Table I. The bulk displacement likelihood functions were evaluated using the quality metric calculated from the posterior distribution of a single displacement estimate (rather than a series of posterior distributions through depth). Because the true displacement is known exactly but will almost always lie between two samples, the value of ε for calculating the quality metric is a single sampling period.

The strain simulations analyzed compressional strains between 0.01% and 10% and at 20 dB SNR. Additionally, the likelihood functions were compared for SNRs of 10, 20, 30, 40 dB, and no noise with a strain of 1%. The quality metrics for the likelihood functions for strained data are calculated for the full profile through depth. The depth locations of the posterior distributions were downsampled to allow independence among the posterior distributions. Because the strain and displacement profile is known exactly through depth, an ε of one sampling period in (10) is used when calculating the quality metric. For the strain simulations, 250 speckle and noise realizations were simulated for each case.

ARFI simulations are used to evaluate the proposed likelihood function under more realistic displacement profiles and beamforming. The ARFI simulations utilize the method developed by Palmeri *et al.* [24]. Palmeri's method combines finite element simulation of the ARFI dynamic response in tissue with realistic ultrasound beamforming simulated using Field II [30], [31]. For the ARFI simulations, 100 independent speckle realizations in a homogeneous region were generated using the parameters shown in Table II. The likelihood functions were evaluated using the described quality metric. The quality metric summarizes the quality of the probabilities through the full depth of the simulated displacements.

For ARF-induced displacements, the displacements measurable by relevant transducer configurations underestimate the true displacements by as much as 50% [29]. Because of this systematic bias caused by scatterer shearing under the point spread function, the displacement profile used for τ_{true} to calculate the probability statistics is taken to be the mean of all of the realizations with displacements estimated using normalized cross-correlation⁵ and ε is set to be 4 sampling periods, reflecting the standard deviation of all of the ARF dynamic responses used to calculate the average displacement profile.

III. RESULTS

First, several example results are presented to add additional motivation to the problem and build intuition. The first example result is shown in Fig. 1, which shows

⁵Specifically, when calculating the error for a specific speckle and noise realization, the true displacement profile was calculated with all of the noiseless speckle realizations except the speckle realization corresponding to the speckle realization being analyzed.

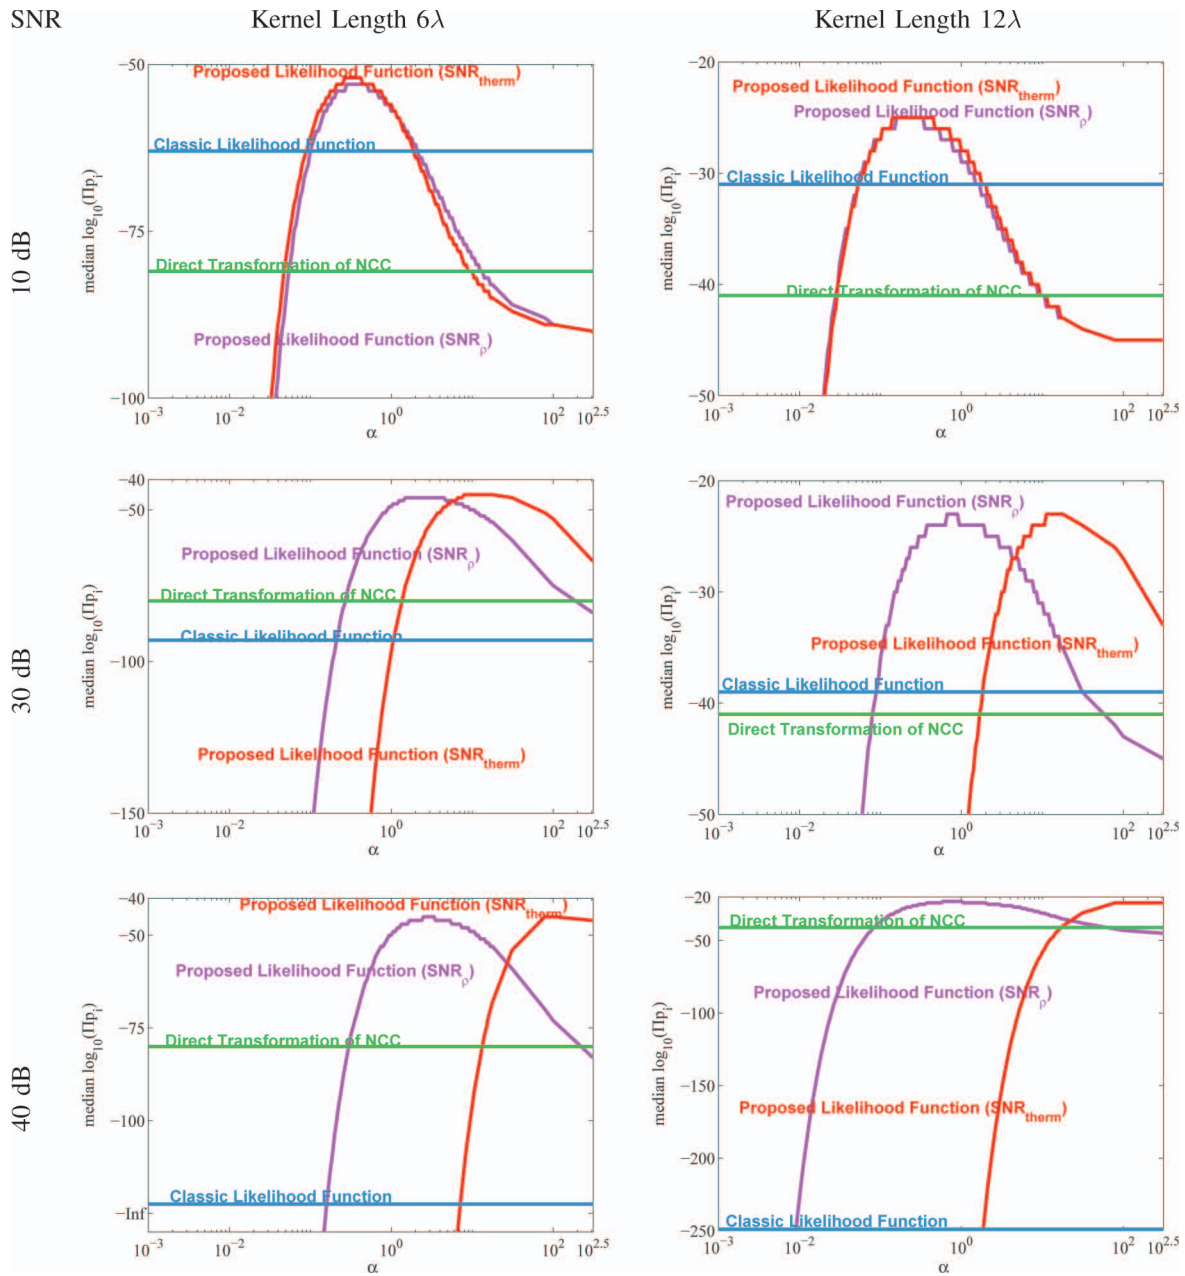


Fig. 6. This figure compares different methods of calculating likelihood functions for compression-induced displacements. The classic likelihood function, the direct transformation of normalized cross-correlation, and the proposed method with both specified and data-derived SNR are shown. The median values are displayed as the log of the respective median values. Additionally, the y -axis represents the full probability of the entire displacement profile through depth.

the effect of modulating the scaling constant α in (4). The limits of the scaling constant (0 and ∞) have already been stated to result in a rectangular window and a delta function, respectively; some intermediate forms of the posterior distribution are shown in Fig. 1. The figure demonstrates how various scaling constants change the amount of probability assigned to the true displacement. This example is from a bulk displacement example with a small search region. An example with a larger search region would show multiple peaks for intermediate values of α , but would still resolve itself into a uniform PDF or a

delta function at the scaling limits. Additionally, because the transformation from normalized cross-correlation function to likelihood function is monotonic, the location of the peak value of each curve is preserved but other characteristics of the shape change.

The second example result shows a similar effect on a full set of ARFI displacements through depth. This is shown in Fig. 2. The figure shows posterior distributions (assuming a non-informative prior) and the true ARFI displacement, as described in Section II-D. The figure shows the posterior distribution formed with and without




Fig. 7. This figure shows the value of α that maximizes the median probability for the strain-induced displacements as a function of kernel length. The data are displayed for several levels of thermal noise. The range for the best α decreases as the kernel size increases. For the kernel sizes usually used to estimate displacements from compression, the best values for α are narrowly spread.

any scaling to emphasize the importance of effective scaling for appropriate concentration of probability around the true displacement.

The statistic used to assess likelihood function discrimination will be the median. The use of the median statistic is justified by Fig. 3, which shows an example displacement scenario plotted as a function of the scaling term, α . The figure shows that the median probability from 1000 speckle realizations has a peak as a function of α , and that the median is sensitive to underscaling and overscaling of the exponential argument in (4). In contrast, the mean probability of the same speckle realizations as a function of α is nearly constant for small α values. This is because the mean is sensitive to outliers, and the mean statistic is not sensitive to overscaling the argument of (4). The mean’s insensitivity occurs because as PDFs become very narrow from the small α values most of the true displacement probabilities go to zero but a few probabilities get extremely high and approach one, pulling the mean up. In other words, considering the example shown in Fig. 1, the median is sensitive to the undesirable nature of the red and green curves, whereas the mean is only sensitive to the undesirable nature of the green curve and finds the red and blue curve equivalently palatable on average.

Before the rest of the results are given, a brief description of how to interpret the graphs is provided. In Figs. 4,

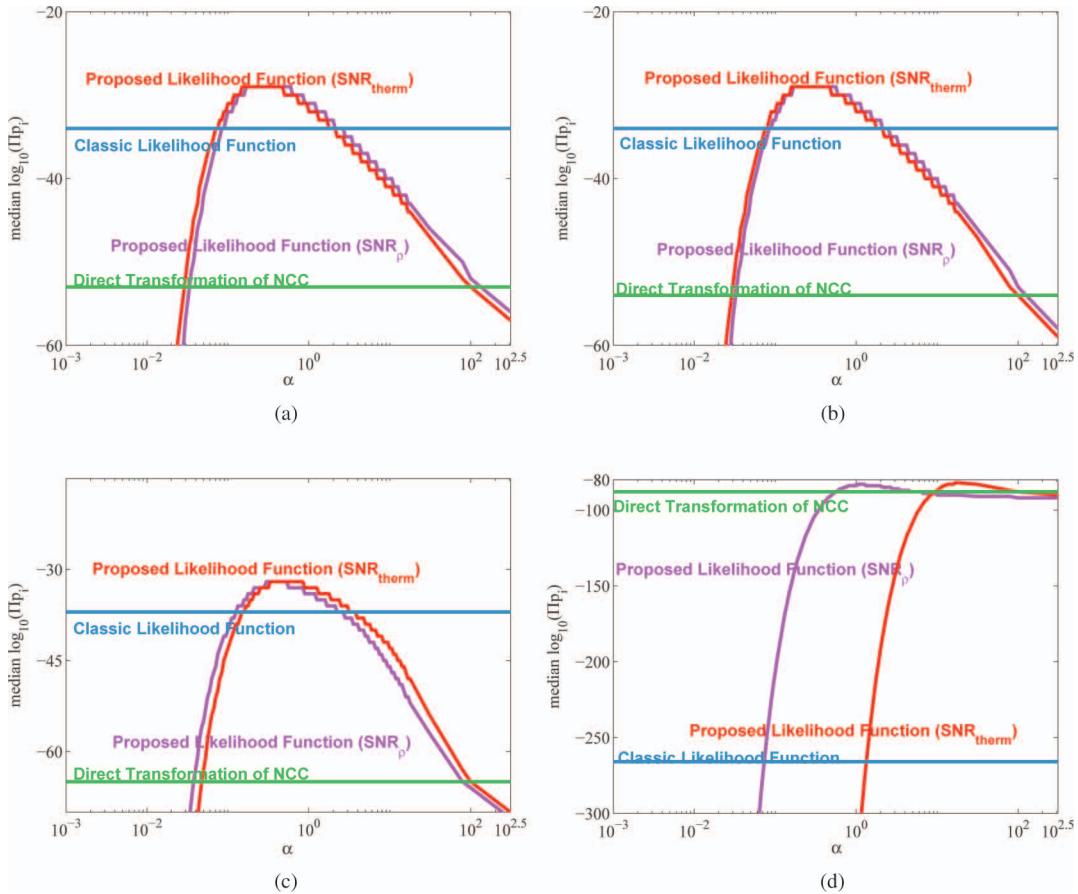


Fig. 8. The results of likelihood function discrimination for various magnitudes of compressional motion are shown for a 6λ kernel and an SNR of 20 dB. Strains of (a) 0.01%, (b) 0.1%, (c) 1%, and (d) 10% are shown. The likelihood functions computed using the classical method and an unscaled conversion of the normalized cross-correlation are shown, along with the two proposed likelihood functions. The proposed likelihood functions have similar behavior to each other when the thermal noise dominates, however, when decorrelation dominates as a noise source, the likelihood function dependent on SNR_{therm} shifts to a higher α . The likelihood function that uses SNR_{ρ} does not shift as the dominant noise sources transitions from thermal noise to decorrelation noise.





Fig. 9. This figure shows the values of α that lead to the most discriminative likelihood functions for different magnitudes of compressive displacements. Results are shown for kernel lengths of 4.5, 6, and 12 λ kernels. The dependent axis is displayed on a log scale only to accommodate the α value for the 12 λ kernel at 10% strain. 

6, 8, and 10, the dependent axis is the median probability, and it is beneficial for the curves to have high median probability, which indicates that the posterior distribution is appropriately concentrated. Additionally, in these figures, several graphs are shown; it is also beneficial if the peak of the curve for each method occurs at the same location on the independent axis. If this is not so, the implication is that α depends on one of the parameters that is being measured (or at least a parameter that is not well known). In Figs. 5, 7, 9, and 11, the dependent axis is the maximum α value. In the ideal case, all of the curves in the graphs within these figures overlap exactly.

Fig. 4 shows the bulk motion results comparing the methods for computing the likelihood function. The methods compared are the classic likelihood function, the likelihood function using normalized cross-correlation without scaling, and the proposed likelihood function for the two methods of computing the SNR. The classic likelihood function and the unscaled normalized cross-correlation are both constant because the likelihood function performance is plotted as a function of scaling. Three different levels of thermal noise are considered—10, 30, and 40 dB. The 40-dB case is included to show a degenerate case and to raise the issue of adequate sampling. The sampling for the simulations was 10 GHz, but the concentration of the posterior for 40 dB of thermal SNR was such that all the probability of the posterior distributions was concentrated within the range $\tau_{\text{true}} \pm \varepsilon$ used to evaluate the likelihood functions for bulk motion. (This will be shown to be unimportant when motion-induced signal decorrelation is present.) The most important trend of the data presented in Fig. 4 is that for significant ranges of the scaling term α , the classic likelihood function and the unscaled normalized cross-correlation likelihood function are less discriminative than the proposed likelihood function.

Although the proposed likelihood function is better, one additional important trend is that the likelihood function using the thermal SNR—which is known exactly from

the simulations—performs consistently better than the SNR derived from ρ . The performance of the SNR derived from ρ is not unexpected because it is an estimated value, and the variance of ρ is directly related to the value of ρ . The exact distribution that ρ follows has been well characterized by Fisher [32] and is not trivial, but a simple approximation is $\sigma_{\rho}^2 \approx (1 - \rho^2)^2/N$, where N is the number of independent samples used [33]. This approximation of σ_{ρ}^2 is consistent with the convergence of the performance of the likelihood functions calculated using the two different SNRs, (5) and (6), as the kernel length increases.

Fig. 5 shows the scaling that produces the peak median value, plotted as a function of kernel length. The scaling has a strong trend with kernel length, but appears to be nearly independent of SNR. (The 40-dB case does not follow the same trend, but this has already been shown to be a degenerate case.)

The first set of strain-induced displacement results is shown in Figs. 6 and 7. The figures show that there is a more significant change in α for changes in thermal noise in the presence of compressive motion compared with the results of bulk motion. Despite this change, the best α values are still generally tightly clustered compared with the domain of the graphs in Fig. 6. The plots in Fig. 6 also show broader peaks. This suggests that the probabilities for strain cases are less sensitive to the specific value of α , and the PDFs for strain are probably more diffuse when compared with the PDFs obtained for bulk motion displacements.

Next, results for constant noise and varying strain are shown. The results for computing the likelihood function for several strain magnitudes are shown for a kernel length of 6 λ and an SNR of 20 dB in Fig. 8. Results are plotted as a function of strain in Fig. 9 for 4.5, 6, and 12 λ . These results show the value of α that maximizes the probability. The results in Fig. 9 indicate that over traditionally measurable levels of strain, the value of α is nearly constant for a given kernel length. The only exception is for the 12 λ kernel for a strain of 10%. For this case, α is significantly higher than the baseline values. This is not particularly concerning because this combination of strain and kernel length is not typically encountered in elastography applications.

The ARFI results of likelihood function evaluation are shown next. The results are all reported as the probability that the true displacement profile occurred, given the data. The first figure, Fig. 10, shows the performance of the various methods of calculating the likelihood function in the context of ARFI data. The figure shows trends similar to those for the bulk motion case shown in Fig. 4. For the ARFI comparisons shown in Fig. 10, smaller kernel lengths are shown that are typical for *in vivo* ARFI motion estimation. The figure shows the classic method and the direct transformation of normalized cross-correlation as well as the two proposed methods. Both of the proposed methods are again better over a large range of α scalings, but, as with the strain results, the method that relies on

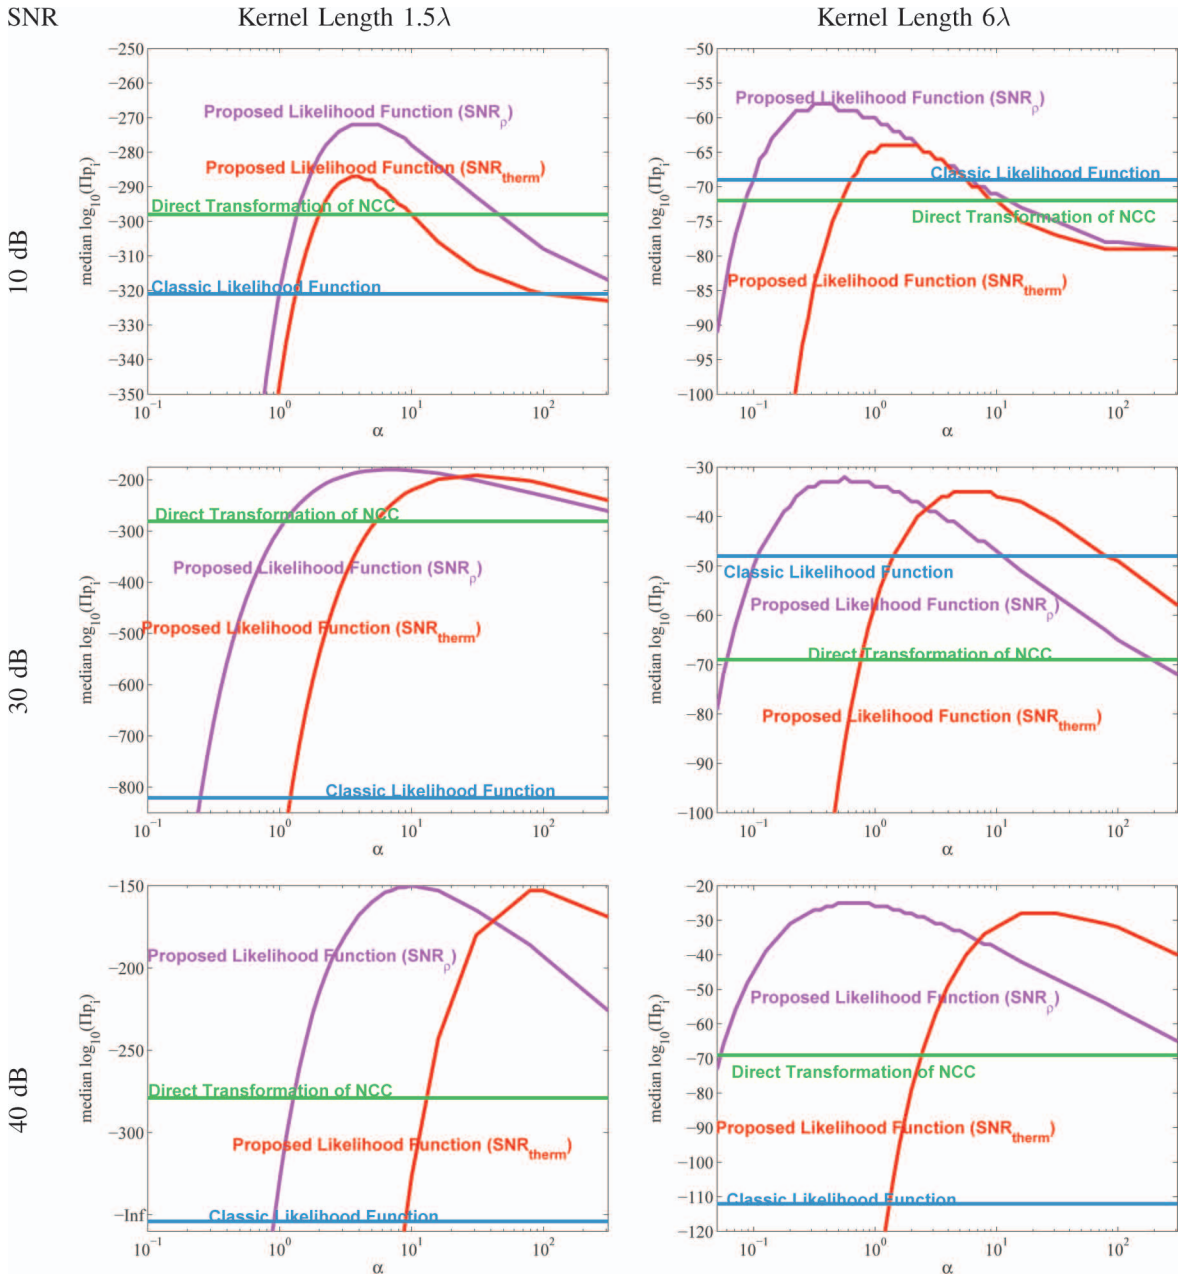


Fig. 10. This figure compares different methods of calculating likelihood functions for ARFI displacements. The classic likelihood function, the direct transformation of normalized cross-correlation, and the proposed method with specified SNR and data-derived SNR are shown. The median values are displayed as the log of the respective median values. Additionally, the y -axis represents the full probability of the entire ARFI displacement profile through depth.

the correlation coefficient to calculate the SNR is the best method. In some cases, the performance of the likelihood function only incorporating thermal noise is similar, but the position of this likelihood function is very volatile with levels of SNR. The artifact seen in the 40-dB bulk motion case is not seen in the ARFI displacements because even when the thermal noise is low, there is still a significant amount of ARFI motion-induced signal decorrelation.

For the ARFI displacements, the scaling factors (α) that produce the peak median values are plotted as a function of kernel length and shown in Fig. 11. The model line that was created from the bulk motion displacements

is also in Fig. 11. The trend of the model line exists in the ARFI results, but there is an additional scaling term. This is shown in the second graph in Fig. 11, which shows the ARFI results divided by the line modeled from the bulk motion results.

IV. DISCUSSION AND CONCLUSION

An implementable perturbation to the likelihood function has been demonstrated for several displacement scenarios. An implementable likelihood function forms a sig-





Fig. 11. This figure shows the value of α that maximizes the median probability for the ARFI-induced displacements. The data are displayed as a function of kernel length for several levels of thermal noise. The range of the best α shrinks rapidly as the kernel length increases. Additionally, SNR usually does not vary over 30 dB throughout the ARFI image region of interest. 

nificant part of the framework necessary to develop biased displacement estimates, but it is also useful to have an appropriate method for converting possible displacements into quantifiable probabilities.⁶

The derived likelihood function will be useful for any implementation of Bayesian speckle tracking. To assist in the use of the likelihood function for arbitrary displacement scenarios, the scaling factor α was introduced and was shown to be a useful way of modifying the likelihood function for different kernel lengths. This avoids the requirement for an infinitely long data record to apply the likelihood function appropriately. That being said, the likelihood function was also shown to have a reasonably shallow peak as a function of α and is not extremely sensitive to the selection of specific α values—implying that calculation of the likelihood function should be adequately robust. (If better α values are required in the future, it may be possible to treat α as a random value and iterate toward an optimal value using Gibbs sampling, for example.)

One thing that has not been provided here is an analysis of how the likelihood function can be used to get better displacement estimates. The likelihood function can be used by itself to obtain displacement estimates, or it can be used as part of a regularized estimation scheme to incorporate other knowledge of the displacement at each

⁶McCormick *et al.* [15] describe an *ad hoc* method for remapping any similarity metric so that it behaves as a probability density function. Their results—in respect to being discriminative towards the true displacement—are nearly identical to the unscaled likelihood function case displayed throughout this paper. This is consistent with their observation that their regularization method primarily works for large strains because as the strain level increases, the probability density function becomes more diffuse. We hypothesize that this systematic broadening of the likelihood function could make an algorithm inappropriately dominated by prior knowledge.

location into the final estimate. An initial investigation has been performed in the companion paper [28].

The method outlined and validated in this work should be compatible with other methods present in the literature. As an example, companding [6] should be employable in conjunction with the likelihood function to increase its discrimination.

Another result of the work performed here is that it is now possible to express ultrasound displacement estimates as PDFs. Expressing displacement estimates as PDFs provides an additional tool for exploring higher-order speckle tracking errors. For example, do correlations exist between higher-order moments (higher than the variance) of a given displacement PDF and the behavior of the estimation error?

Finally, the likelihood function has only been applied to simulated 1-D displacement fields. However, the utility of the likelihood function has been introduced and demonstrated elsewhere for 2-D and 3-D *in vivo* scenarios [34].

ACKNOWLEDGMENTS

The authors would like to thank N. Daniele and M. Jakovljevic.

REFERENCES

- [1] L. Chen, R. J. Housden, G. M. Treece, A. H. Gee, and R. W. Prager, “A hybrid displacement estimation method for ultrasonic elasticity imaging,” *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 57, no. 4, pp. 866–882, Apr. 2010.
- [2] R. Zahiri-Azar and S. Salcudean, “Motion estimation in ultrasound images using time domain cross correlation with prior estimates,” *IEEE Trans. Biomed. Eng.*, vol. 53, no. 10, pp. 1990–2000, Oct. 2006.
- [3] T. Loupas, J. T. Powers, and R. W. Gill, “An axial velocity estimator for ultrasound blood flow imaging, based on a full evaluation of the Doppler equation by means of a two-dimensional autocorrelation approach,” *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 42, no. 4, pp. 672–688, 1995.
- [4] C. Kasai, K. Namekawa, A. Koyano, and R. Omoto, “Real-time two-dimensional blood flow imaging using an autocorrelation technique,” *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. SU-32, no. 3, pp. 458–464, May 1985.
- [5] A. Pesavento, C. Perrey, M. Krueger, and H. Ermert, “A time-efficient and accurate strain estimation concept for ultrasonic elastography using iterative phase zero estimation,” *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 46, no. 5, pp. 1057–1067, 1999.
- [6] P. Chaturvedi, M. F. Insana, and T. J. Hall, “2-D companding for noise reduction in strain imaging,” *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 45, no. 1, pp. 179–191, Jan. 1998.
- [7] E. E. Konofagou and J. Ophir, “A new elastographic method for estimation and imaging of lateral displacements, lateral strains, corrected axial strains and Poisson’s ratios in tissues,” *Ultrasound Med. Biol.*, vol. 24, no. 8, pp. 1183–1199, 1998.
- [8] P. M. Embree, “The accurate ultrasonic measurement of volume flow of blood by time-domain correlation,” Ph.D. dissertation, Dept. of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, 1985.
- [9] E. T. Jaynes, *Probability Theory: The Logic of Science*, vol 1, New York, NY: Cambridge University Press, 2003.
- [10] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Upper Saddle River, NJ: Prentice-Hall, 1993.
- [11] J. Jiang and T. J. Hall, “A parallelizable real-time motion tracking algorithm with applications to ultrasonic strain imaging,” *Phys. Med. Biol.*, vol. 52, no. 13, pp. 3773–3790, Jul. 2007.

- [12] C. Pellot-Barakat, F. Frouin, M. Insana, and A. Herment, "Ultrasound elastography based on multiscale estimations of regularized displacement fields," *IEEE Trans. Med. Imaging*, vol. 23, no. 2, pp. 153–163, Feb. 2004.
- [13] J. Jiang and T. J. Hall, "A generalized speckle tracking algorithm for ultrasonic strain imaging using dynamic programming," *Ultrasound Med. Biol.*, vol. 35, no. 11, pp. 1863–1879, 2009.
- [14] Y. Petrank, L. Huang, and M. O'Donnell, "Reduced peak-hopping artifacts in ultrasonic strain estimation using the Viterbi algorithm," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 56, no. 7, pp. 1359–1367, 2009.
- [15] M. McCormick, N. Rubert, and T. Varghese, "Bayesian regularization applied to ultrasound strain imaging," *IEEE Trans. Biomed. Eng.*, vol. 58, no. 6, pp. 1612–1620, Jun. 2011.
- [16] H. L. van Trees, *Detection, Estimation and Modulation Theory*, pt. III, New York, NY: Wiley, 1971.
- [17] H. Urkowitz, *Signal Theory and Random Processes*. Dedham, MA: Artech House, 1983.
- [18] Y. Bar-Shalom and T. E. Fortmann, "Multitarget tracking using joint probabilistic data association," in *Statistical Signal Processing*, E. J. Wegman and J. G. Smith, Eds., New York, NY: Marcel Dekker, 1984, pp. 353–363.
- [19] W. Walker, "Adaptive ultrasonic imaging performance for near-field aberrating layers," Ph.D. dissertation, Dept. of Biomedical Engineering, Duke University, Durham, NC, 1995.
- [20] S. G. Foster, P. M. Embree, and W. D. O'Brien, "Flow velocity profile via time-domain correlation error analysis and computer-simulation," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 37, no. 3, pp. 164–175, May 1990.
- [21] W. Walker and G. Trahey, "A fundamental limit on delay estimation using partially correlated speckle signals," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 42, no. 2, pp. 301–308, 1995.
- [22] F. Kallel, M. Bertrand, and J. Meunier, "Speckle motion artifact under tissue rotation," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 41, no. 1, pp. 105–122, 1994.
- [23] J. Meunier and M. Bertrand, "Ultrasonic texture motion analysis: Theory and simulation," *IEEE Trans. Med. Imaging*, vol. 14, no. 2, pp. 293–300, 1995.
- [24] M. L. Palmeri, A. C. Sharma, R. R. Bouchard, R. W. Nightingale, and K. R. Nightingale, "A finite-element method model of soft tissue response to impulsive acoustic radiation force," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 52, no. 10, pp. 1699–1712, Oct. 2005.
- [25] F. Viola and W. F. Walker, "A comparison of the performance of time-delay estimators in medical ultrasound," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 50, no. 4, pp. 392–401, 2003.
- [26] B. H. Friemel, L. N. Bohs, K. R. Nightingale, and G. E. Trahey, "Speckle decorrelation due to two-dimensional flow gradients," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 45, no. 2, pp. 317–327, Mar. 1998.
- [27] T. Varghese and J. Ophir, "A theoretical framework for performance characterization of elastography: The strain filter," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 44, no. 1, pp. 164–172, 1997.
- [28] B. C. Byram, G. E. Trahey, and M. Palmeri, "Bayesian speckle tracking. Part II: Biased ultrasound displacement estimation," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 60, no. 1, pp. 144–157, 2013.
- [29] M. L. Palmeri, S. A. McAleavey, G. E. Trahey, and K. R. Nightingale, "Ultrasonic tracking of acoustic radiation force-induced displacements in homogeneous media," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 53, no. 7, pp. 1300–1313, Jul. 2006.
- [30] J. A. Jensen, "Field: A program for simulating ultrasound systems," *Med. Biol. Eng. Comput.*, vol. 34, suppl. 1, pt. 1, pp. 351–353, 1996.
- [31] J. A. Jensen and N. B. Svendsen, "Calculation of pressure fields from arbitrarily shaped, apodized and excited ultrasound transducers," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 39, pp. 262–267, 1992.
- [32] R. A. Fisher, "Frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population," *Biometrika*, vol. 10, pp. 507–521, 1915.
- [33] R. A. Fisher, *Statistical Methods for Research Workers*. Edinburgh, UK: Oliver and Boyd, 1954.
- [34] B. C. Byram, "Chronic myocardial infarct visualization using 3D ultrasound," Ph.D. dissertation, Dept. of Biomedical Engineering, Duke University, Durham, NC, 2011.



Brett C. Byram received the B.S. degree in biomedical engineering and math from Vanderbilt University, Nashville, TN, in 2004. He received the Ph.D. degree in biomedical engineering in 2011 from Duke University, Durham, NC. He is currently a postdoctoral associate in the biomedical engineering department at Duke University. His ultrasound research interests include beamforming, motion estimation, and other related signal processing tasks.



Gregg E. Trahey (S'83–M'85) received the B.G.S. and M.S. degrees from the University of Michigan, Ann Arbor, MI in 1975 and 1979, respectively. He received the Ph.D. degree in biomedical engineering in 1985 from Duke University. He served in the Peace Corps from 1975 to 1978 and was a project engineer at the Emergency Care Research Institute in Plymouth Meeting, PA, from 1980–1982. He currently is a Professor with the Department of Biomedical Engineering at Duke University and holds a secondary appointment with the Department of Radiology at the Duke University Medical Center. His current research interests include adaptive phase correction and beamforming, and acoustic radiation force imaging methods.



Mark L. Palmeri received his B.S. degree in biomedical and electrical engineering from Duke University, Durham, NC, in 2000. He was a James B. Duke graduate fellow and received his Ph.D. degree in biomedical engineering from Duke University in 2005 and his M.D. degree from the Duke University School of Medicine in 2007. He is currently an Assistant Research Professor of Biomedical Engineering and Anesthesiology at Duke University. His research interests include ultrasonic imaging, characterizing the mechanical properties of soft tissues, and finite element analysis of soft tissue response to acoustic radiation force excitation.