Published in final edited form as:

IEEE Trans Ultrason Ferroelectr Freq Control. 2013 January ; 60(1): 132–143. doi:10.1109/TUFFC.2013.2545.

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]

The department of Biomedical Engineering, Durham, NC, USA

Brett Byram: bcb16@duke.edu

Abstract

A hallmark of clinical ultrasound has been accurate and precise displacement estimation. 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 hard to implement for clinically relevant ultrasound with diffuse scattering. Since 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 , while 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 like acoustic radiation force impulse (ARFI) imaging and strain imaging possible.

Generally, further development of clinical ultrasonic displacement estimators focus 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 Cramer-Rao lower bound for the ultrasound displacement estimation problem has existed at least since Embree [8]. While 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 Cramer-Rao lower bound 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 where 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 aims 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. While 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.

A missing tool for ultrasound displacement estimation (and specifically biased estimation) is 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 (like 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 to this one. This paper will focus on demonstrating an effective and implementable likelihood function for ultrasonic displacement estimation problems.

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 here 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 is 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.

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

B. Likelihood Function for Ultrasonic Time-Delay Estimation

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

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

where $s_1()$ is the original signal, and $s_2()$ 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_{noise}^2 is the noise power². Often the likelihood function in (2) is used to derive the maximum likelihood estimator, which has a form that resembles the similarity metrics commonly used for ultrasound displacement estimation³.

For use with Bayes' Theorem to calculate displacement probabilities (and eventually displacement estimates), the likelihood function can be expressed in a reduced form since 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

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

The argument of (4) is nearly identical to the argument of the maximum likelihood estimator (MLE) in (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 rectangle function the size of the specified search region as the noise power approaches ∞ .

There are two practical difficulties with (4). 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 (4) to

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

³The maximum likelihood estimator is

$$\hat{\tau}_0 = \arg \max_{\tau_0} \sum_{n=\frac{\tau_0}{\Delta}}^{\frac{\tau_0}{\Delta}+M-1} s_1(n\Delta)s_2(n\Delta+\tau_0), \quad (3)$$

This equation describes unnormalized cross-correlation, but usually in ultrasound normalized cross-correlation is treated as a substitute maximum likelihood estimator. The maximum likelihood estimator is attractive because it asymptotically approaches the Cramer-Rao lower bound 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].

create a likelihood function more suitable for the ultrasonic displacement estimation problem. To this end, the following likelihood function is proposed

$$p_m(x|\tau_0) \propto \exp \left[\frac{\text{SNR}}{\alpha} \sum_{n=\frac{\tau_0}{\Delta}}^{\frac{\tau_0}{\Delta}+M-1} \frac{s_{m1}(n\Delta)s_{m2}(n\Delta+\tau_0)}{\sqrt{\sigma_{s_{m1}}^2 \sigma_{s_{m2}}^2}} \right] \quad (5)$$

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

$$\text{SNR}_{\text{thermal}} = \frac{P_{\text{signal}}}{P_{\text{noise}}}, \quad (6)$$

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

$$\text{SNR}_{\rho} = \frac{\rho_{\max}}{1-\rho_{\max}} \quad (7)$$

where $\rho_{\max}$ is the peak of the normalized cross-correlation function. Normalized cross-correlation is shown in (5) as everything to the right of the summation including the summation. The thermal SNR will be treated as a known value since the appropriate levels of noise added during data simulation is known. The SNR calculated from the maximum correlation will be an estimated quantity ⁴.

C. Likelihood Function Evaluation

In order 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. In order to facilitate likelihood function evaluation, likelihood functions (which are not PDFs) can be converted to posterior probability distributions using a non-informative prior shown in (9).

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

$$\text{SNR}_{\text{total}} = \frac{\text{SNR}_{\rho} \text{SNR}_{\text{thermal}}}{1 + \text{SNR}_{\rho} + \text{SNR}_{\text{thermal}}}. \quad (8)$$

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.

$$p_m(\tau_0|x) = \begin{cases} \frac{p_m(x|\tau_0)}{\int_{\tau_\alpha}^{\tau_\beta} p_m(x|\tau_0) d\tau_0} & \text{if } \tau_\alpha < \tau_0 < \tau_\beta \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

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

$$\text{quality metric} = \sum_{m=0}^{M-1} \int_{\tau_{true}-\epsilon}^{\tau_{true}+\epsilon} p_m(\tau_0|x) d\tau_0 \quad (10)$$

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 (10) because τ_0 is appropriately modeled as continuous. In practice, the quality metric above will not be realizable using the methods presented here where only a sampled distribution will be available. Therefore, a more practical quality metric for the data used in this study will be

$$\text{quality metric} = \prod_{m=0}^{M-1} \sum_{n=\lfloor \frac{\tau_{true}-\epsilon}{\Delta} \rfloor}^{\lfloor \frac{\tau_{true}+\epsilon}{\Delta} \rfloor} p_m(n\Delta|x) \quad (11)$$

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 decimation). This quality metric was chosen with two intentions in mind. First, when only relative probabilities of displacement are desired then it seems to make sense to have a probability distribution that is expected to give the most probability possible 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 then there is a strong danger that the data will be inappropriately overwhelmed by prior information.

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 1D scatterer configurations for bulk displacements and compression induced (i.e. strain) displacements. Simulations were also performed for 3D scattering geometries for acoustic radiation force (ARF) induced displacements.

The 1D scattering geometry simulations were all performed using convolution with a static point spread function and complex pulse. In order 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—12–15 scatterers are adequate but more scatterers were included to avoid gaps of scatterers for step displacements or large strains [29]. The initial set of scatterer positions for a given realization were 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 (5) calculated using $\text{SNR}_{\text{thermal}}$ and using SNR_ρ for a range of α values is compared against the classic likelihood function (4). The proposed form of the likelihood function is also compared against a likelihood function constructed from a normalized cross-correlation function but without any scaling term in order demonstrate the scaling term's importance. The likelihood functions are compared using the quality metric shown in (10). 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 the results section.

The bulk displacement simulations to evaluate likelihood function efficacy were performed first. The bulk motion simulation data was comprised of 1000 1D 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 $\frac{\lambda}{20}$ in order 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 ($\sigma = \frac{\lambda}{30}$) 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). Since 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 .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 metric for the likelihood functions for strained data are calculated for the full profile through depth. The depth locations of the posterior distributions were decimated to allow independence among the posterior distributions. Since the strain and displacement profile is known exactly through depth an ϵ of one sampling period (Fig. 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 measureable by relevant transducer configurations underestimate the true displacements by as much as 50% [29]. Because of this systematic bias due to 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 the realizations with displacements estimated using normalized cross-correlation⁵, and ϵ is set to be four sampling periods reflecting the standard deviation of all 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 the effect of modulating the scaling constant α in (5). 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 scalings change the amount of probability assigned to the true displacement. This example is from a bulk displacement example with a small search region. A larger search region example would show multiple peaks for intermediate values of α but still resolve itself into a uniform PDF or a δ -function at the scaling limits. Additionally, since 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 Sec. II-D. The figures shows the posterior distribution formed with and without any scaling in order 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 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 under and over scaling of the exponential argument in (5). 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 (5). 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, while 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, 6, 8 and 10 the dependent axis is the median probability, and it is

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

beneficial for the curves to have high median probability, which indicates that the posterior distribution is appropriately concentrated. Additionally, in Figs. 4, 6, 8, and 10 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. The implication of this not being so 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 Figs. 5, 7, 9, and 11 overlap exactly.

The bulk motion results comparing the methods for computing the likelihood function are shown in Fig. 4. 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 since 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 in order to show a degeneracy and 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 $\tau_{true} \pm \epsilon$ 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.

While 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 since 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 $\sigma_\rho^2 \approx \frac{(1-\rho^2)^2}{N}$, 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 (6) and (7) as the kernel length increases.

The scaling that produces the peak median value is plotted as a function of kernel length is shown in Fig. 5. 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 are 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 to the results of bulk motion. Despite this change the best α values are still generally tightly clustered compared to the domain of the graphs in Fig. 6. The plots in Fig. 6 also show broader peaks. This suggests that the probability for strain cases are less sensitive to the specific value of α , and the PDFs for strain are probably more diffuse when compared to 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 lengths of 6λ and 20 dB SNR 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 12λ kernel for a strain of 10%. For this case α is significantly higher than the “baseline” values. This is not particularly concerning since 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 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 (α) that produces 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 significant 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 probabilities⁶.

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” of 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 towards an optimal value using Gibbs sampling as an example.)

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.

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

Acknowledgments

The authors would like to thank Ned Danieleley and Marko Jakovljevic.

This work was supported by NIH grants R37HL096023 and T32EB001040.

References

1. Chen L, Housden RJ, Treece GM, Gee AH, Prager RW. A hybrid displacement estimation method for ultrasonic elasticity imaging. *IEEE Trans Ultrason Ferroelectr Freq Control*. Apr; 2010 57(4): 866–882. [PubMed: 20378449]
2. Zahiri-Azar R, Salcudean S. Motion estimation in ultrasound images using time domain cross correlation with prior estimates. *IEEE Trans Biomedical Engineering*. Oct; 2006 53(10):1990–2000.
3. Loupas T, Powers JT, Gill RW. 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 Ferroelect Freq Contr*. 1995; 42(4):672–688.
4. Kasai C, Namekawa K, Koyano A, Omoto R. Real-time two-dimensional blood flow imaging using an autocorrelation technique. *IEEE Trans Ultrason Ferroelect Freq Contr*. May; 1985 SU-32(3): 458–464.
5. Pesavento A, Perrey C, Krueger M, Ermert H. A time-efficient and accurate strain estimation concept for ultrasonic elastography using iterative phase zero estimation. *IEEE Trans Ultrason Ferroelectr Freq Control*. 1999; 46(5):1057–67. [PubMed: 18244299]
6. Chaturvedi P, Insana MF, Hall TJ. 2-d companding for noise reduction in strain imaging. *IEEE Trans Ultrason Ferroelect Freq Contr*. Jan; 1998 45(1):179–191.
7. Konofagou EE, Ophir J. 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*. 1998;1183–1199. [PubMed: 9833588]
8. Embree, PM. PhD dissertation. University of Illinois; Urbana, Ill: 1985. The accurate ultrasonic measurement of volume flow of blood by time-domain correlation.
9. Jaynes, ET. *Probability Theory: The Logic of Science*. Vol. 1. Cambridge University Press; Apr. 2003
10. Kay, SM. *Fundamentals of statistical signal processing: estimation theory*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc; 1993.

11. Jiang J, Hall TJ. A parallelizable real-time motion tracking algorithm with applications to ultrasonic strain imaging. *Phys Med Biol*. Jul; 2007 52(13):3773–3790. [PubMed: 17664576]
12. Pellot-Barakat C, Frouin F, Insana M, Herment A. Ultrasound elastography based on multiscale estimations of regularized displacement fields. *IEEE Trans Med Imaging*. Feb; 2004 23(2):153–163. [PubMed: 14964561]
13. Jiang J, Hall TJ. A generalized speckle tracking algorithm for ultrasonic strain imaging using dynamic programming. *Ultrasound Med Biol*. 2009; 35(11):1863–79. [PubMed: 19682789]
14. Petrank Y, Huang L, O'Donnell M. Reduced peak-hopping artifacts in ultrasonic strain estimation using the viterbi algorithm. *IEEE Trans Ultrason Ferroelectr Freq Control*. 2009; 56(7):1359–67. [PubMed: 19574147]
15. McCormick M, Rubert N, Varghese T. Bayesian Regularization Applied to Ultrasound Strain Imaging. *IEEE Trans Biomed Eng*. Jun; 2011 58(6):1612–1620. [PubMed: 21245002]
16. van Trees, HL. *Detection, Estimation and Modulation Theory, Part III*. New York: Wiley; 1971.
17. Urkowitz, H. *Signal Theory and Random Processes*. Dedham: Artech House; 1983.
18. Bar-Shalom, Y.; Fortmann, TE. Multitarget tracking using joint probabilistic data association. In: Wegman, EJ.; Smith, JG., editors. *Statistical Signal Processing*. New York: Marcel Dekker; 1984. p. 353–363.
19. Walker, W. PhD dissertation. Duke University; 1995. Adaptive ultrasonic imaging performance for near-field aberrating layers.
20. Foster SG, Embree PM, O'Brien WD. Flow velocity profile via time-domain correlation error analysis and computer-simulation. *IEEE Trans Ultrason Ferroelectr Freq Control*. May; 1990 37(3):164–175. [PubMed: 18285029]
21. Walker W, Trahey G. A fundamental limit on delay estimation using partially correlated speckle signals. *IEEE Trans Ultrason Ferroelect Freq Contr*. 1995; 42(2):301–308.
22. Kallel F, Bertrand M, Meunier J. Speckle motion artifact under tissue rotation. *IEEE Trans Ultrason Ferroelect Freq Contr*. 1994; 41(1):105–122.
23. Meunier J, Bertrand M. Ultrasonic texture motion analysis: Theory and simulation. *IEEE Trans Med Imging*. 1995; 14:293–300.
24. Palmeri ML, Sharma AC, Bouchard RR, Nightingale RW, Nightingale KR. A finite-element method model of soft tissue response to impulsive acoustic radiation force. *IEEE Trans Ultrason Ferroelectr Freq Control*. Oct; 2005 52(10):1699–1712. [PubMed: 16382621]
25. Viola F, Walker WF. A comparison of the performance of time-delay estimators in medical ultrasound. *IEEE Trans Ultrason Ferroelect Freq Contr*. 2003; 50:392–401.
26. Friemel BH, Bohs LN, Nightingale KR, Trahey GE. Speckle decorrelation due to two-dimensional flow gradients. *IEEE Trans Ultrason Ferroelect Freq Contr*. Mar; 1998 45(2):317–327.
27. Varghese T, Ophir J. A theoretical framework for performance characterization of elastography: the strain filter. *IEEE Trans Ultrason Ferroelectr Freq Control*. 1997; 44(1):164–172. [Online]. Available: <http://www.hubmed.org/display.cgi?uids=18244114>. [PubMed: 18244114]
28. Byram BC, Trahey GE, Palmeri M. Bayesian speckle tracking. part ii: Biased ultrasound displacement estimation. *IEEE Trans Ultrason Ferroelect Freq Contr*. 2013; 60(1):132–143.
29. Palmeri ML, McAleavey SA, Trahey GE, Nightingale KR. Ultrasonic tracking of acoustic radiation force-induced displacements in homogeneous media. *IEEE Trans Ultrason Ferroelectr Freq Control*. Jul; 2006 53(7):1300–1313. [PubMed: 16889337]
30. Jensen, JA. Field: A program for simulating ultrasound systems. Paper presented at the 10th Nordic-Baltic Conference on Biomedical Imaging Published in *Medical & Biological Engineering & Computing*; 1996. p. 351–353.
31. Jensen JA, Svendsen NB. Calculation of pressure fields from arbitrarily shaped, apodized and excited ultrasound transducers. *IEEE Trans Ultrason Ferroelect Freq Contr*. 1992; 39:262–267.
32. Fisher RA. Frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population. *Biometrika*. 1915; 10:507–521.
33. Fisher, RA. *Statistical Methods for Research Workers*. Edinburgh: Oliver and Boyd; 1954.
34. Byram, BC. PhD dissertation. Duke University; 2011. Chronic myocardial infarct visualization using 3d ultrasound.

Fig. 1.

This figure demonstrates how the probability density function concentrates for several values of the scaling parameter α from (5). 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 (5) 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.

(a) Unscaled

(b) Scaled

Fig. 2.

These images show the final ARFI posterior distributions for each depth with and without 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 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 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 are showing the 25% and 75% percentiles and the box plot whiskers are showing 1.5 the interquartile distance beyond the closest quartile. The red points show points in the sample 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.

Watermark-text

Watermark-text

Watermark-text

Fig. 5.

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

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.

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 is 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.

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. The level of simulated strain is listed below each plot. Strains of .01%, .1%, 1%, and 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 $\text{SNR}_{\text{therm}}$ 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.

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.

Fig. 11.

This figure shows the value of α that maximizes the median probability for the ARFI induced displacements. The data is 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.

TABLE I**Basic Simulation Parameters**

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

TABLE II

ARFI Imaging Simulation Parameters

Parameter	Value
Tracking Center Frequency	7 MHz
Tracking F/#	0.5
Radiation Force Center Frequency	2.22 MHz
Radiation Force F/#	2
Radiation Force Pulse Duration	180 μ s
Focal Depth	2 cm
Tracking Pulse Transmit Cycles	2
C	1540 m/s
depth _{max}	5cm
Sampling Frequency	100 MHz
Kernel Length	1.5 λ
SNR	20dB
ARFI PRF	10 kHz
Young's Modulus	8.5 kPa
ρ	1.0 g/cm ³
ν	0.499