

Bayesian Speckle Tracking. Part II: Biased Ultrasound Displacement Estimation

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

Abstract—Ultrasonic displacement estimates have numerous clinical uses, including blood flow estimation, elastography, therapeutic guidance, and acoustic radiation force imaging (ARFI). These clinical tasks could be improved with better ultrasonic displacement estimates. Traditional ultrasonic displacement estimates are limited by the Cramer–Rao lower bound (CRLB). The CRLB can be surpassed using biased estimates. In this paper, a framework for biased estimation using Bayes’ theorem is described.

The Bayesian displacement estimation method is tested against simulations of several common types of motion: bulk, step, compression, and acoustic-radiation-force-induced motion. Bayesian estimation is also applied to *in vivo* ARFI of cardiac ablation lesions. The Bayesian estimators are compared with the unbiased estimator, normalized cross-correlation.

As an example, the peak displacement of the simulated acoustic radiation force response is reported because this position results in the noisiest estimates. Estimates were made with a $1.5\text{-}\lambda$ kernel and 20 dB SNR on 100 data realizations. Estimates using normalized cross-correlation and the Bayes’ estimator had mean-square errors of 17 and $7.6\ \mu\text{m}^2$, respectively, and contextualized by the true displacement magnitude, $10.9\ \mu\text{m}$. Biases for normalized cross-correlation and the Bayes’ estimator are -0.12 and $-0.28\ \mu\text{m}$, respectively. *In vivo* results show qualitative improvements. The results show that with small amounts of additional information, significantly improved performance can be realized.

I. INTRODUCTION

WITH a goal of creating displacements estimators that bypass fundamental limits on common estimators, such as normalized cross-correlation, a perturbation to the traditional likelihood function was proposed in the accompanying paper and shown to be more discriminative in the sense of appropriately concentrating probability distributions around the true displacement [1]. In addition to presenting a modified likelihood function, biased estimators were introduced in a qualitative manner. In this paper, biased estimators are implemented and it is shown that it is possible to surpass the performance limit described by the Cramer–Rao lower bound (CRLB) with relatively small amounts of additional information.

The CRLB is a general measure of information content that describes the minimum obtainable estimation error variance when using an unbiased estimator [2]. A CRLB

for displacement estimation related tasks has been derived in the ultrasound literature [3]–[7], but the derivation by Walker and Trahey [6] seems to be favored (probably because of its broad applicability):

$$\sigma(\hat{\tau}_0 - \tau_0) \geq \sqrt{\frac{3}{2f_0^3 \pi^2 T (B^3 + 12B)} \left(\frac{1}{\rho^2} \left(1 + \frac{1}{\text{SNR}^2} \right)^2 - 1 \right)} \quad (1)$$

where f_0 is the pulse center frequency, B is the pulse bandwidth, T is the kernel size, ρ is the normalized correlation between signals, and SNR is the signal-to-noise ratio.

The CRLB-limited estimator—usually referred to as the minimum-variance unbiased estimator (MVUE)—is unique because the algorithmically induced bias is zero, which can be useful. The utility of its unbiased nature is seen in situations in which many independent measurements can be acquired. In this case, the mean of the estimates will converge toward the true measurement. In most clinical scenarios, however, it is not possible to obtain multiple measurements. In *in vivo* scenarios in which only a single estimate can be acquired, the bias error is often over-emphasized as a noise mechanism because for a single estimate it is not possible to distinguish between different orthogonal error components [8]. By considering both noise mechanisms—bias and variance—a better estimator can often be realized.

The bias and variance of an estimator’s error are defined as

$$\text{bias} = \text{E}[\hat{\tau}_0 - \tau_0] \quad (2)$$

$$\text{variance} = \text{E}[(\hat{\tau}_0 - \tau_0)^2] - \text{bias}^2, \quad (3)$$

where τ_0 describes the true estimate and $\hat{\tau}_0$ describes the estimated time shift (i.e., displacement) between two signals. Both noise mechanisms can be appropriately combined in the form of the mean square error (MSE):

$$\text{MSE} = \text{E}[(\hat{\tau}_0 - \tau_0)^2] = \text{variance} + \text{bias}^2. \quad (4)$$

In many cases, allowing a small amount of bias into an estimator can lead to a drastic reduction in the estimation variance, creating an overall lower MSE. Fig. 1 shows an example of two different hypothetical biased estimators and how they might compare with an unbiased estimator that is CRLB-limited.

In the ultrasound literature, there have been no descriptions of biased displacement estimators that produce estimates with a lower MSE than would be produced by a

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.2013.2546>

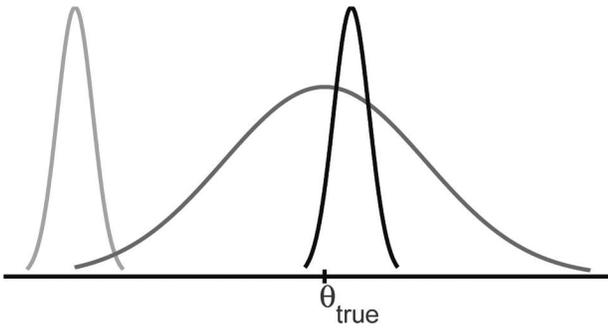


Fig. 1. Three distributions describing the error of different hypothetical estimators are shown. The broad distribution can be equated to a minimum-variance unbiased estimator (MVUE), which would be described by the Cramer–Rao lower bound. The other two estimators have a significantly smaller variance. The estimator with the small bias to the right of θ_{true} is attractive, whereas the other biased estimator is clearly not.

MVUE. Several groups have produced estimator schemes in which the displacement search region is shifted and reduced in size based on displacements measured at adjacent positions [9]–[11], but because of the way these algorithms are implemented, they only achieve improvements in computational efficiency and in reduction of peak-hopping artifacts, rather than an improvement in the actual MSE. (This type of approach will be shown to be a specific realization of the methods developed herein.)

In the rest of this paper, biased displacement estimates with MSE surpassing the limit expressed by the CRLB will be demonstrated for several different types of motion.

II. METHODS

A. Overview

It is difficult to create estimators with improved MSE characteristics relative to the MVUE using classical methods [2]. The usual alternative is to use Bayes’ theorem, which is a simple equation describing the appropriate method for combining the current data with previous knowledge about the parameter(s) to be estimated [8]. Bayes’ theorem will be used for ultrasound displacement estimation 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. To this end, Bayes’ theorem is expressed as

$$p_m(\tau_0 | x) = \frac{p_m(x | \tau_0)p_m(\tau_0)}{\int p_m(x | \tau_0)p_m(\tau_0)d\tau + 0}, \quad (5)$$

where x is the data, τ_0 is the displacement, and m indexes axial depth. The term $p_m(x | \tau_0)$ denotes the likelihood function, $p_m(\tau_0)$ is the prior probability density function (PDF), and $p_m(\tau_0 | x)$ is the posterior PDF. The likelihood function is the means by which data are incorporated into the estimate. An appropriate and implementable likeli-

hood function has been demonstrated in the companion paper [1]:

$$p_m(x | \tau_0) \propto \exp \left[\frac{\text{SNR}}{\alpha} \sum_{n=\tau_0/\Delta}^{\tau_0/\Delta+M-1} \frac{s_{m1}(n\Delta)s_{m2}(n\Delta - \tau_0)}{\sqrt{\sigma_{s_{m1}}^2 \sigma_{s_{m2}}^2}} \right], \quad (6)$$

where s_{m1} and s_{m2} are the RF A-lines from which motion will be estimated, α is an application-specific scaling term, Δ is the sampling period, and M is the length of the data record (i.e., kernel length). The quality metric for this function is normalized cross-correlation and the SNR is calculated as

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

The prior distribution expresses previous knowledge of the displacement and is a PDF. Example uses of prior PDFs would be to use the knowledge of an ARF dynamic response through depth to allow adjacent spatial locations to appropriately influence the current estimate, or similarly, in the case of blood-flow estimation, one could use fluid dynamics to influence adjacent spatial locations [12]. The posterior distribution is also a PDF and represents the final state of knowledge of the parameter(s) to be estimated after combining the new and old information. Displacement estimates at a given depth position will be made from the posterior distribution.

Two simple methods for obtaining prior PDFs of the displacement estimate will be described. Accompanying the description of prior PDFs will be a description of appropriate estimators for resolving the final posterior distribution into estimates of displacements. A flowchart describing the entire algorithm is shown in Fig. 2.

B. Prior Probabilities

Computing prior probabilities for ultrasound displacements is an open-ended, flexible problem with many solutions. Two solutions for computing prior probabilities will be developed in this section. The intention of the proposed methods will not be to present a broadly generalizable optimal prior PDF but rather to demonstrate the promise of biased estimators and the small amount of additional information needed to surpass the CRLB. The simplest, least-informative prior for ultrasound displacement estimation is a uniform PDF with an extent equivalent to the search region. This prior is

$$p_m(\tau_0) = \begin{cases} \frac{1}{\tau_\beta - \tau_\alpha} & \text{if } \tau_\alpha < \tau_0 < \tau_\beta \\ 0 & \text{otherwise,} \end{cases} \quad (8)$$

where τ_α and τ_β denote the limits of the search region. The mean and variance of the uniform prior can be static for all displacement estimates in the field, or the uniform PDFs can be dynamic based on other information, such as

spatially adjacent displacements [9]–[11]. For the evaluation here, the uniform PDF is static with depth.

The second prior scheme that will be explored is more dynamic than the uniform PDF proposed in (8). In this scheme, the posterior distribution at the previous depth will be considered a good estimate for the prior distribution at the current depth. When the previous posterior gets too narrow based on the posterior's standard deviation, the prior will revert to a normal distribution with mean equal to the previous estimate and a standard deviation equal to a defined minimum standard deviation. That is,

$$p_m(\tau_0) = \frac{1}{\sqrt{2\pi}\sigma_m^2} e^{-(\tau_0 - \hat{\tau}_{0m-1})^2 / (2\sigma_m^2)}, \quad (9)$$

where $\hat{\tau}_{0m-1}$ is the estimate at the previous depth and the variance is a predefined σ_{\min} .

In a special case in which deterministic and stochastic implementations are compared, the prior will always be a normal distribution based on the mean and standard deviation of the posterior distribution. In this case, the standard deviation will be

$$\sigma_m^2 = \begin{cases} \sigma_{p_{m-1}(\tau_0|x)}^2 & \text{if } \sigma_{p_{m-1}(\tau_0|x)}^2 > \sigma_{\min}^2 \\ \sigma_{\min}^2 & \text{otherwise,} \end{cases} \quad (10)$$

where σ_{\min} is a defined minimum standard deviation. This is done because normal distributions are easy and efficient for taking random samples compared with arbitrary PDFs that may be encountered in more general methods. Additionally, normal distributions are useful because they represent the least informative prior when only the mean and standard deviation are known. This special case and the accompanying results are shown in the appendix.

C. Parameter Estimation From Posterior Distributions

Using the likelihood function shown in (6) and one of the approaches for computing a prior PDF just described, it becomes simple to use Bayes' theorem, shown in (5), to calculate a posterior PDF, which can be converted into a displacement estimate. Many methods exist to resolve posterior distributions into displacement estimates [2]. The two methods evaluated here are the minimum mean square estimator (MMSE) and the maximum *a posteriori* estimator (MAPE). The MMSE is

$$\hat{\tau}_0 = E(\tau_0 | x) = \int \tau_0 p(\tau_0 | x) d\tau_0. \quad (11)$$

The MMSE is the average of the posterior distribution, and as suggested by the name, it minimizes the MSE of the estimate for $\hat{\tau}_0$. This indicates a significant conceptual departure from the MVUE (limited by the CRLB), particularly because the possible improvements are available even when used with a noninformative prior.

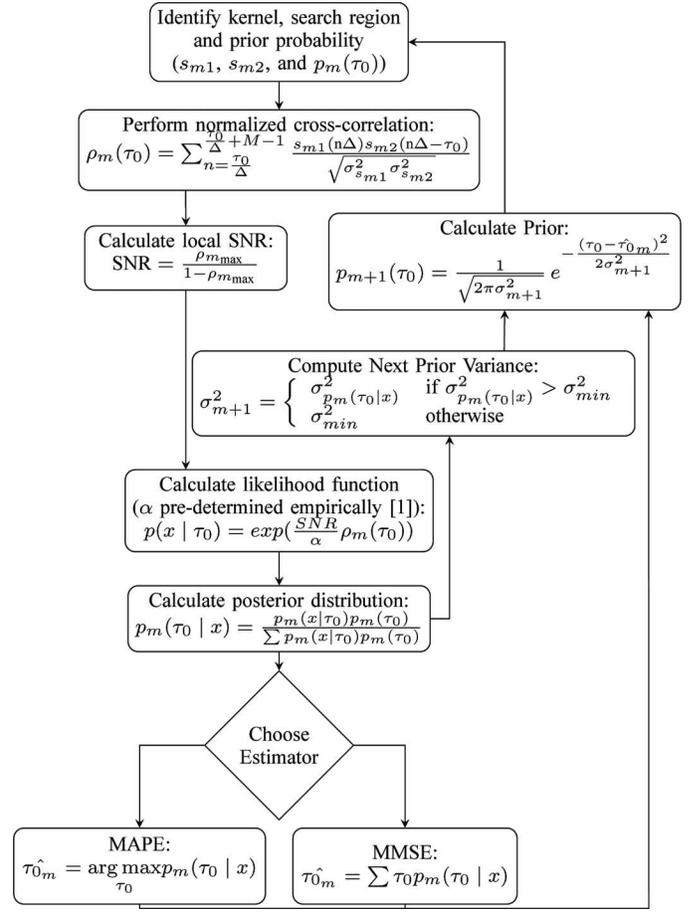


Fig. 2. A flowchart for the deterministic implementation of biased estimation described in the methods.

(It is worth stating that the MMSE is not necessarily better than the MVUE even when it yields what would be considered better estimates. This is because the optimization problem that would lead one to the MMSE is different from the problem that leads to the MVUE. A thorough discussion on the nature of the various estimators has been presented by Kay [2].) The MAPE is

$$\hat{\tau}_0 = \arg \max_{\tau_0} p(\tau_0 | x). \quad (12)$$

The MAPE resembles the methods for parameter estimation when using normalized cross-correlation, and because the conversion of the normalized cross-correlation function to a likelihood function [see (6)] is a monotonic transformation normalized cross-correlation and MAPE yield identical estimates for the case of a uniform prior [see (8)] on τ_0 [1].

In the event that the posterior distribution is normally distributed (or, more generally, in the case where the posterior distribution is symmetric about the peak) the MMSE and the MAPE will yield identical results.

Normalized cross-correlation sometimes referred to as the maximum likelihood estimator (MLE) will be the primary point of comparison for the estimators introduced.

D. Simulations

To test the new displacement estimators, four types of simulations were performed. Simulations were performed to test the estimators on bulk displacements, step displacements, compression-induced displacements (strain), and acoustic radiation force (ARF)-induced displacements.

The first three sets of simulations (bulk, step, and compressional displacements) are entirely 1-D. The 1-D data were simulated using an axial convolution method. The convolution was performed between a complex Gaussian pulse and a randomly distributed scatterer field. For each scatterer in the field, the pulse was given a bulk offset in pixels and a subsample phase rotation based on the difference between the nearest pixel location and its actual location. This approach to the typical convolution data simulation method allowed continuous displacements to be simulated. The final simulation result for each scatterer field was a single A-line of RF data obtained by taking the real part of the complex convolution.

For the purposes of the simulation, 35 scatterers were used per resolution cell. Thirty-five scatterers per resolution cell is excessive—12 to 15 scatterers are adequate to achieve second-order speckle statistics [13]; the large number of scatterers was chosen to ensure sufficient scatterer density regardless of the simulated displacement, including displacement fields induced by large step displacements or strains.

For all of the simulations—including the ARF simulations not yet described—thermal noise was modeled as an additive random process that was band-limited based on the simulated pulse’s characteristics. All of the simulations used the parameters shown in Table I unless specified otherwise. Specifically notable is the sampling frequency. A sampling frequency of 10 GHz was deemed sufficient to avoid using subsample estimation for the MAPE or MLE and to obtain an adequate error distribution for almost all cases. (Using a subsample estimator is undesirable in this case because it adds an additional source of bias.) The sampling frequency was also sufficient to allow numerical integration for the MMSE—with and without a dynamic prior—to be implemented as a simple summation.

All four estimators (the MLE, the MMSE with a non-informative prior, the MAPE with a dynamic prior, and the MMSE with a dynamic prior) will be evaluated on each type of simulated motion.

For the bulk displacement simulations, the effect of frequency, bandwidth, kernel length, and SNR were evaluated. These factors all influence the CRLB shown in (1). The only other factor influencing the CRLB is signal correlation. Signal correlation effects can be mechanism-dependent, so several mechanisms of inducing decorrelation will be evaluated in the simulations for step, compressional, and ARF displacements. For the bulk displacement simulations, 1000 realizations of scatterer and noise realizations were evaluated. Additionally, the displacement for each realization was drawn from a normal distribution with zero-mean and a standard deviation of $\lambda/20$. This

TABLE I. SIMULATION PARAMETERS.

Parameter	Value
Center frequency	5 MHz
Bandwidth	50%
Sampling frequency	10 GHz
Kernel length	3λ
Kernel overlap	80%
Strain-specific parameter	
Strain	1%
ARF-specific parameters	
Tracking center frequency	7 MHz
Tracking f-number	0.5
Radiation force center frequency	2.22 MHz
Radiation force duration	180 μ s
Radiation force f-number	2
Sampling frequency	100 MHz
Focus depth	2 cm
Tracking pulse repetition frequency	10 kHz
Young’s modulus	8.5 kPa
Kernel length	1.5λ
ρ	1.0 g/cm ³
ν	0.499

avoided any biases that could be caused by a regular pattern of subsample displacements and was computationally efficient because the necessary search region was small.

For the step displacement simulations, a step displacement of 20% of the center frequency’s wavelength was created at 2 cm of axial depth. The effect of a minimum standard deviation for the normal prior distribution was evaluated. The step displacement simulation results for biased estimators were also used to calculate frequency responses. Frequency responses can be obtained by differentiating the step displacement and calculating the Fourier transform of the result. The frequency response will also include the frequency response of a 47th-order low-pass filter. This particular filter was the shortest passive filter designed using the windowed linear-phase finite impulse response (FIR) method [14] [as implemented in Matlab (The MathWorks Inc., Natick, MA)] that has a lower MSE for an ARF-induced displacement compared with any of the biased estimators. Minimum prior standard deviations for the dynamic prior between 0.1 ns and 100 ns were evaluated. For the step displacement simulations, 100 realizations of data were evaluated.

The ARF results from displacement estimation using biased estimators are also compared with filtering the ARF displacements estimated using normalized cross-correlation combined with low-pass filters. The filters were designed using the windowed linear-phase FIR method [14] (as implemented in Matlab).

Strain displacements were simulated mathematically by compressing the location of the scatterers in the simulated field. Simulations were performed for compressions between 0.01% and 10%. For the strain simulations, SNRs between 10 and ∞ dB and adjacent kernel overlap between 0% and 99% were evaluated. Simulations were also performed for minimum prior standard deviations between 0.1 ns and 1 μ s. Strain was estimated from the

compression-induced displacements using only the derivative without any averaging or median filtering.

For the strain comparisons, some of the plots will be displayed as strain filters conceptualized by Varghese and Ophir [15]. The strain filter plots the SNR of the strain estimates as a function of strain. The results here will calculate the strain SNR as

$$\text{SNR}_{\text{strain}} = \frac{s_{\text{true}}}{\sqrt{\text{MSE}}}. \quad (13)$$

One-hundred realizations of scatterer distribution and noise for each level of strain were evaluated.

ARF simulations were also performed to evaluate estimator performance within the context of a more sophisticated modeling framework. ARF simulations were performed using the methods laid out by Palmeri *et al.* [16], [17], which couple finite-element simulation of the motion induced by ARF and clinically relevant beamforming implemented using Field II [18], [19]. For the ARF simulations, comparisons were made against SNRs between 10 and ∞ dB for minimum standard deviations between 0.1 ns and 1 μ s for the normal prior distribution, and for changes in tracking kernel overlap between 0 and 99%. Because ARF-induced displacement estimates systematically underestimate the tissue displacement by as much as 50% (and even in the best case underestimate by about 10%) [17], the displacement values used to calculate error metrics are the mean displacement values of all realizations tracked with normalized cross-correlation. When analyzing a specific realization, the mean profile comprised all realizations except the one in question.

E. In Vivo Example

To demonstrate basic *in vivo* feasibility, normalized cross-correlation and MAPE are compared for a cryoablation lesion visualized with ARFI. The results were acquired using an open-chested canine preparation. The cryoablation was formed on the epicardial surface of the heart using a Brymill Cry-Ac Tracker with a 3-mm Mini Probe (Brymill Cryogenic Systems, Ellington, CT). The data were acquired using a Sonoline Antares ultrasound system and VF10-5 linear array transducer (Siemens Healthcare, Ultrasound Business Unit, Mountain View, CA). The data were acquired at baseband at 8.9 MHz. The baseband data were interpolated to 142 MHz and reconstructed to radio-frequency data. For each ARF-induced displacement, 4 adjacent A-lines were acquired in parallel. The center frequency for the data used to estimate the dynamic response is 8 MHz. The pulse repetition frequency for measuring each dynamic response is 8.9 kHz. The α for (6) is 4, the minimum standard deviation for the prior was 5 ns, and the kernel length was 1.5 λ . For both implementations of the displacement estimation—normalized cross-correlation and MAPE—parabolic subsample estimation was used.

III. RESULTS

A. Bulk Motion Displacement Simulations

Results of the bulk displacement simulations are shown for parameters (excluding signal correlation) impacting the Cramer–Rao lower bound. The bias and variance are plotted as $\sqrt{\text{bias}^2}$ and $\sqrt{\sigma^2}$, respectively, to demonstrate the result in micrometers, which is often more concise, whereas the MSE is always displayed in units of squared micrometers. Results are shown for all 4 estimators for frequency and bandwidth as a function of the parameter in Fig. 3. Results for varying SNR and kernel length are shown as a function of depth for the MAPE and normalized cross-correlation (i.e., the MLE) in Figs. 4 and 5.

In Fig. 3, where the frequency and bandwidth affects on bulk motion estimation are shown, it is clear that the MSE for biased estimation is less than that for unbiased estimation. In addition, there are several other key trends. The first trend is that normalized cross-correlation and the MMSE with a non-informative prior perform almost identically for bulk motion. The same can be said for the MAPE and MMSE when implemented with the dynamic prior for bulk motion. The second trend is that the functional behavior of the frequency and the bandwidth for the biased estimators is very similar to the behavior of the unbiased estimator. The third trend, which is very significant, is that for bulk motion, the biased estimator has lower bias than the algorithmically unbiased normalized cross-correlation. Additionally, the bias always represents a significantly smaller portion of the MSE than the variance for all estimators.

The effect of SNR on bulk motion estimates is shown in Fig. 4 as a function of depth to show how improvements evolve. The results are only shown for normalized cross-correlation and the MAPE because the other estimators were shown to perform almost identically. The results generally follow the trends seen for frequency and bandwidth. These results add additional perspective and show that the region of most rapid improvement—compared with the MLE—is within the first 0.5 cm of axial depth with modest improvements past about 1 cm axially. This demonstrates that only small amounts of additional information are required to achieve results that surpass the CRLB. Finally, the SNR results show an artifact for the MAPE when the SNR is ∞ dB.¹ The no-noise scenario for the bulk motion simulations should be considered a degenerate case because all of the probability is concentrated to a single sample. The sample that the probability concentrates to is not the exact solution, but because the probability is a δ -function, the estimator assumes this answer is exact, resulting in an inconsistent framework. This artifact is not seen in any of the other results and is not expected in *in vivo* scenarios because of the ubiquity of thermal noise and signal decorrelation.

¹The MMSE with either prior scheme shows the same artifact.

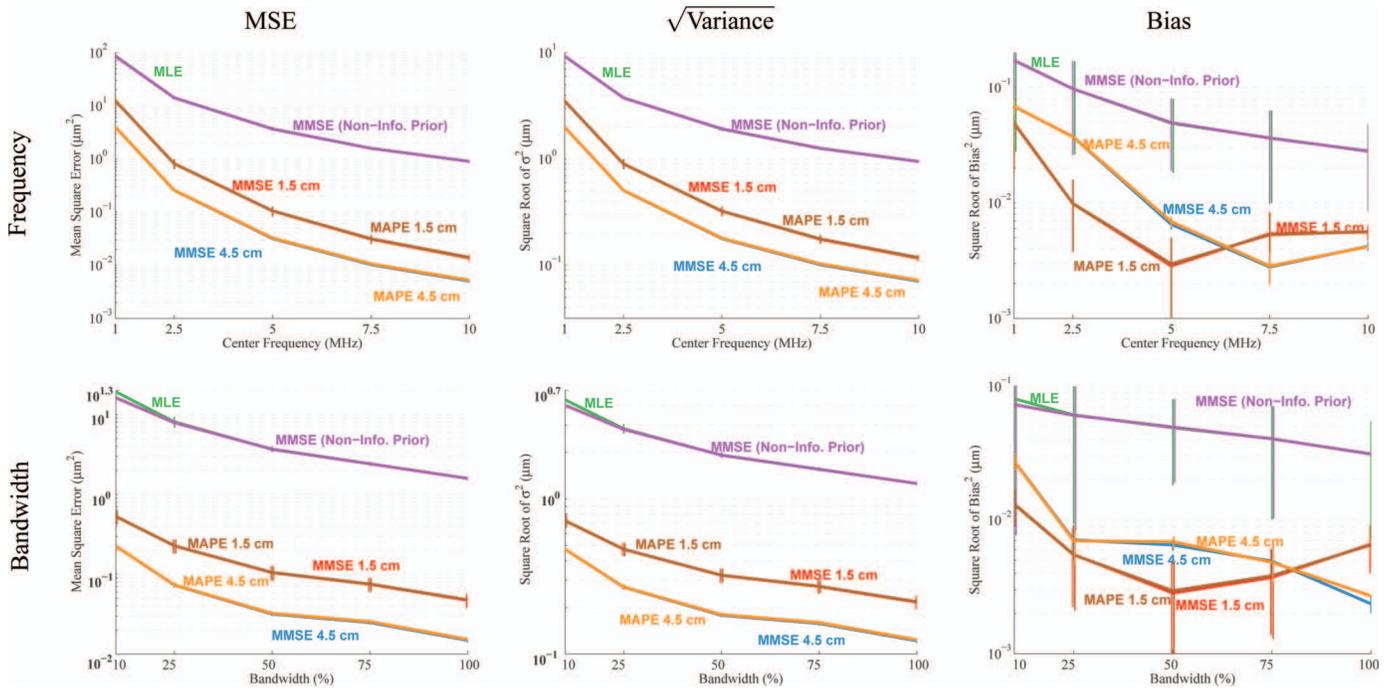


Fig. 3. Frequency and bandwidth functional responses for bulk motion. The effect of bandwidth and frequency are shown for all 4 estimators. The estimators with a depth-dependent, dynamic, Gaussian prior are shown for two depths, 1.5 cm and 4.5 cm. The mean and standard deviation displayed are calculated over a 0.5-cm window in both directions. The maximum likelihood estimator (MLE) and minimum mean square estimator (MMSE) with a non-informative prior are shown based on their statistics centered around 1.5 cm. The functional forms for frequency and bandwidth are very similar to the functional form seen for the MLE (i.e., normalized cross-correlation). Consistent with other results, the bias is lower for the biased estimator.

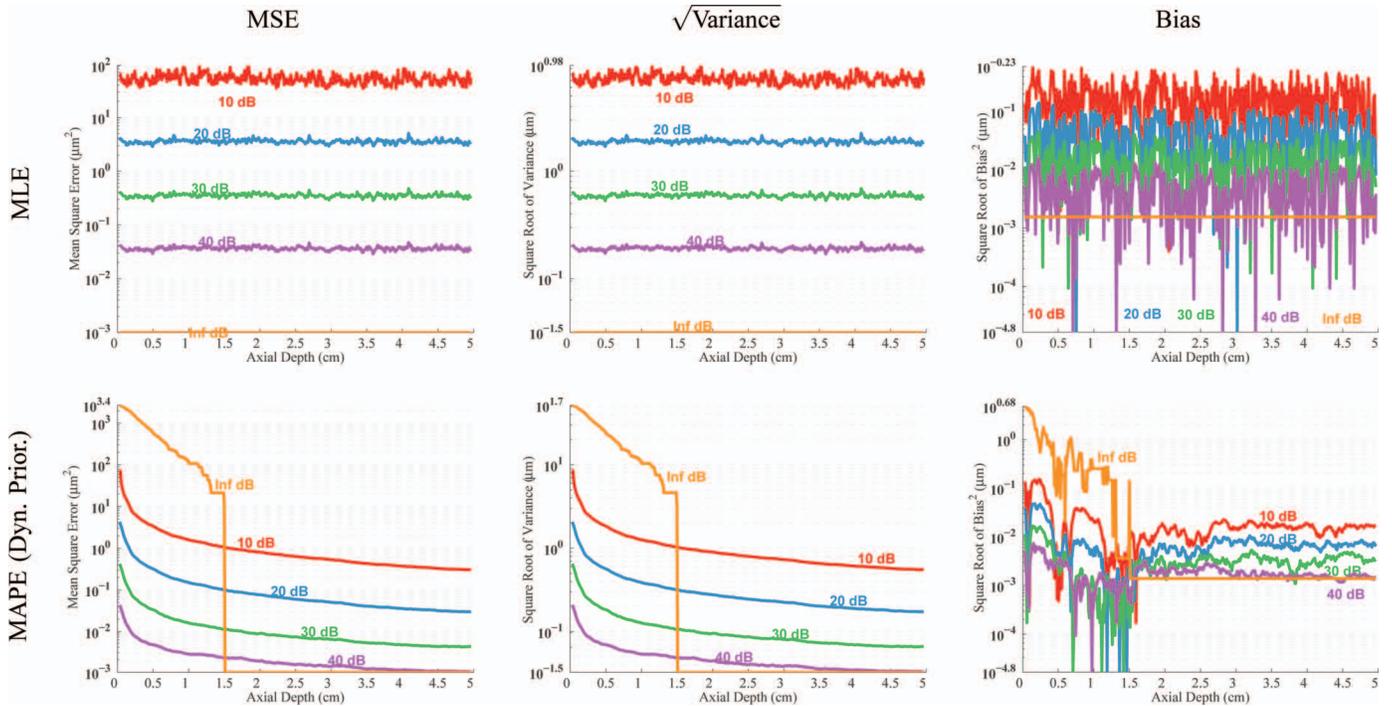


Fig. 4. SNR comparisons for bulk motion. This figure compares the maximum likelihood estimator (MLE) (normalized cross-correlation) to the maximum *a posteriori* estimator (MAPE) with the dynamic Gaussian prior distribution as a function of depth for several levels of SNR. For both cases, the mean-square error is dominated by variance. The bias is very similar between the two estimators, but the biased estimator actually has a lower bias for bulk motion when comparing a given level of SNR. The no-noise case for the MAPE (i.e., ∞ dB) should be seen as a degenerate case because the sampling frequency is not adequate for this level of probability concentration. This failure is not seen when any amount of signal decorrelation is present.

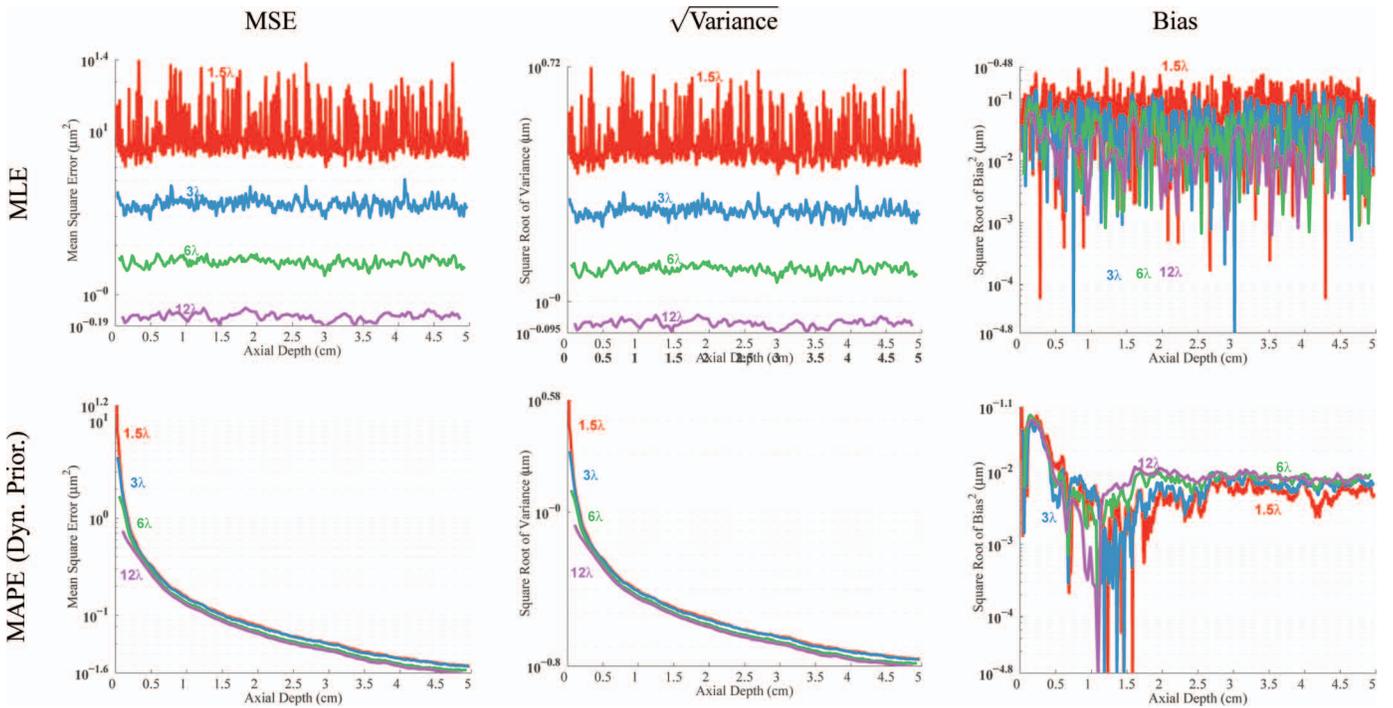


Fig. 5. Kernel size comparisons for bulk motion. This figure compares the maximum likelihood estimator (MLE) (normalized cross-correlation) to the maximum *a posteriori* estimator (MAPE) with the dynamic Gaussian prior distribution as a function of depth for several common kernel lengths. For both cases, the mean-square error is dominated by variance. The biased estimator actually has a lower bias for bulk motion when comparing a given level of SNR. For the MAPE the shortest kernel length evaluated produces the lowest bias.

The most interesting bulk motion results are for the kernel length. Kernel length results are shown as a function of depth for the MLE and MAPE in Fig. 5. The results show that for the MAPE, the kernel length has almost no effect on the MSE, particularly when compared with the behavior of normalized cross-correlation. In addition, the results for the MAPE show that the smallest kernel length evaluated (1.5λ) has the smallest bias. The results for kernel length are unexpected but intriguing. It is hypothesized that by passing information from one estimate to the next, the amount of information contained in the displacement estimate from 8 consecutive 1.5λ kernels is equivalent to the information from estimating the displacement from a single 12λ kernel.

B. Step Displacement Results

Results for the step displacement simulations are shown in Fig. 6. The figure shows the MSE of the MAPE in response to a step displacement. The results are shown for a selection of prior PDF minimum standard deviation. The broadest standard deviation shown is basically normalized cross-correlation because the prior imposes so little effect on the MSE. Additionally, one of the useful aspects of simulating a step displacement profile is the ability to construct the frequency response. Frequency responses for a selection of minimum prior standard deviations are also shown in Fig. 6. Included in the figure is the frequency response of a 47th-order low-pass filter because this was found to be the shortest filter that produced a lower MSE than any biased estimator for ARF-induced displacements.

C. Compression-Induced Displacement Results

Results for compression-induced motion estimates and the resulting strain estimates are now shown. The results show the effect of assigning an appropriate minimum standard deviation to the prior PDF, the effect of SNR, and the effect of kernel overlap.

The first results show the effect of minimum standard deviations of the prior PDF on compressional motion estimates. These results are shown in Fig. 7. The results show that there is a relatively narrow range of minimum standard deviations that produce better results than normalized cross-correlation. However, for a large range of standard deviations, the estimates are no worse than normalized cross-correlation.² Similar results are shown for the actual estimation of strain in the form of a strain filter in Fig. 8. Results are shown for a range of minimum standard deviations for the MAPE and for one minimum standard deviation case for all 4 estimators. One interesting result is that small strains in the range of 0.01 to 0.1%—normally difficult to measure—can be estimated effectively. This is probably at least in part because the prior scheme assumes bulk motion, which is also probably why the improvement in the more traditional clinical strain region is more modest.

²Even when MSE improvements over normalized cross-correlation are not realized, the prior can still act as a peak-hopping filter. Peak-hopping can be reduced by the prior (as compared to the non-informative prior) by slightly attenuating correlation peaks away from the prior's mean.

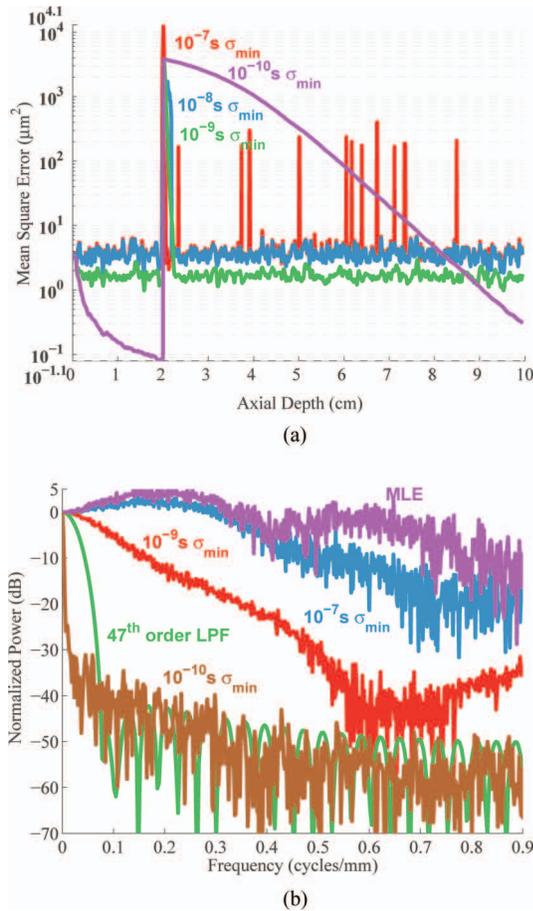


Fig. 6. Step displacement results. (a) The results of the maximum *a posteriori* estimator (MAPE) for a step displacement at 2 cm are shown for several different minimum standard deviations for the normal prior. (b) The result of using the step response to calculate the frequency response. This figure also shows the frequency response of the shortest filter that outperforms the biased estimator in a mean-square error sense.

The effect of SNR on estimating compressional motion and strain are shown in Fig. 9. The results show that the MSE is lowest when there is 20 dB of SNR. There is a small decline in performance when the SNR gets better and a larger decline in performance when the SNR gets smaller. It is not known why this is the case, but it is hypothesized that 20 dB is the right amount of SNR to keep the posterior (and by extension subsequent prior) distributions sufficiently broad to more adequately handle the gradually increasing displacement encountered in these simulations. The results for actual strain estimates do not show estimation improvement at 20 dB but do show that at SNRs 20 dB and higher, SNR has little impact on the strain MSE. This trend is not significantly different than the results derived from normalized cross-correlation based estimates.

The effect of kernel separation on displacement estimation is displayed in Fig. 10. The results show the best lag between kernels for the purpose of displacement estimation is 80% of the kernel's length. However, these results may be slightly misleading because the minimum prior standard deviation used for these estimates was optimized

for 80% kernel overlap. It may be possible to find a better combination of kernel overlap and minimum prior standard deviation. The result of estimating the strain with various levels of kernel overlap is also shown in Fig. 10. These figures show, for the biased estimators, that the amount of kernel overlap matters less when estimating strain. For these results, the MLE-derived displacements can actually lead to better estimates of strain. This seems to be true for kernel overlap that is about 50% or less. This occurs because the difference between the two estimates is higher and the noise on the estimates has less impact. The general method employed in the literature is to calculate strains with kernel overlap of 75% to 80% and then filter the results (or fit to a model), which mitigates the improvement of having no kernel overlap.

D. ARF-Induced Displacement Results

ARF-induced motion estimates are shown next. The results will show the effect of assigning an appropriate

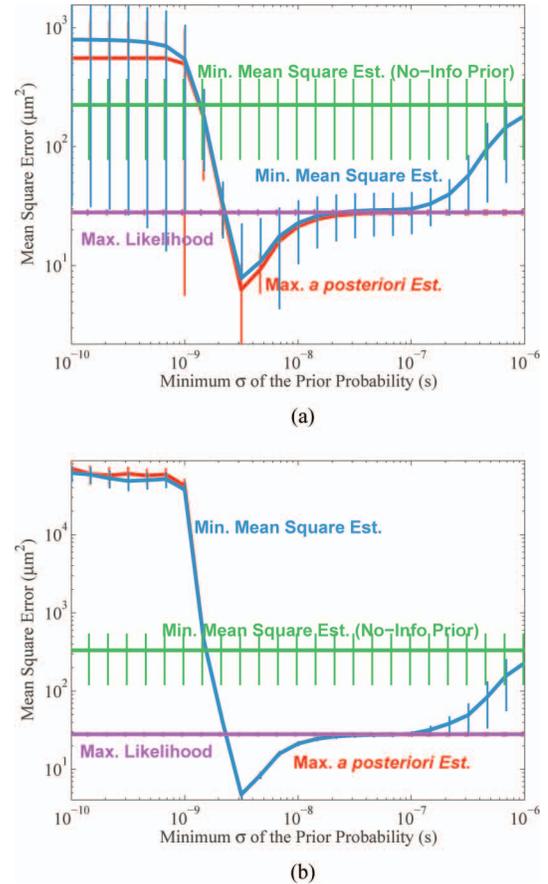


Fig. 7. Minimum prior standard deviation's effect on compressional motion estimates. The effect of minimum standard deviation on the mean-square error (MSE) is shown for all 4 estimators. Results are displayed for two ranges of depth: (a) 0 to 1 cm and (b) 2 to 3 cm axially. The standard deviation is shown (untransformed) and is displayed only on the upper side when $\mu - \sigma$ would be less than zero (because the data are displayed on a log scale). All 4 estimators are shown. There is a range of minimum standard deviations where the maximum *a posteriori* estimator (MAPE) and minimum mean square estimator (MMSE) outperform normalized cross-correlation for 1% strain.

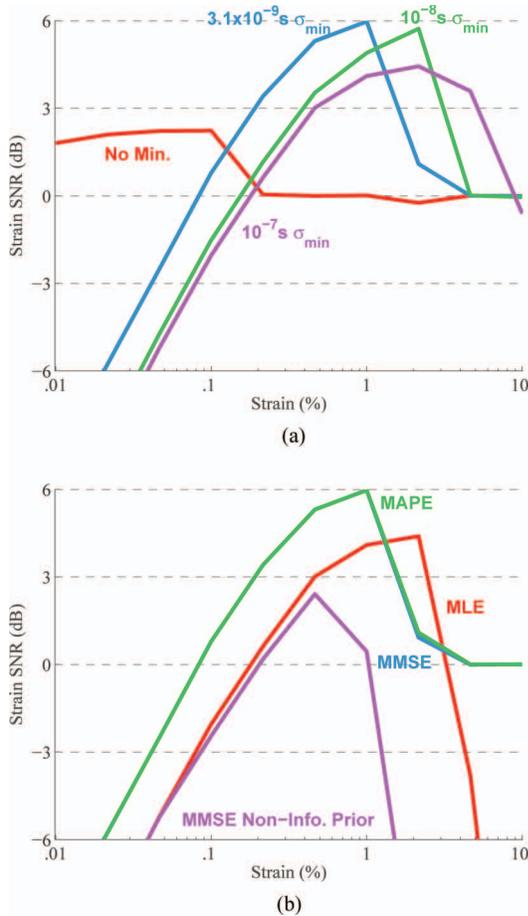


Fig. 8. (a) Several strain filters are plotted for varying levels of σ_{\min} as expressed in (9) for the maximum *a posteriori* estimator (MAPE) and the minimum mean square estimator (MMSE). (b) All of the estimators are compared (for this figure, the minimum prior standard deviation is 3.1 ns).

minimum standard deviation for the prior PDF and the effect of kernel overlap.

First, the results showing the impact of the minimum standard deviation of the prior are shown in Fig. 11. The figure shows three different positions along the axis of the ARF dynamic response. The figure demonstrates the absolute performance of each estimator as well as each estimator's relative performance to the others over a range of selections for σ_{\min} . The estimators that use a dynamic prior do have the lowest MSE, but they also have the highest MSE. This trend indicates the importance of proper selection of σ_{\min} . Additionally, in the region of best performance, there is no difference between the estimators using a dynamic prior. However, when non-informative priors are used, the MMSE has a slightly lower MSE for the estimate from the leading edge and peak, but slightly worse performance on the trailing edge compared with the MLE.

The results for changes in kernel separation are shown in Fig. 12. The figure shows the MSE plotted as a function of the distance between kernels in percent of kernel overlap for the same positions along depth shown in Fig. 11. One of the initial hypotheses was that more overlap would be beneficial to the biased estimators. This figure

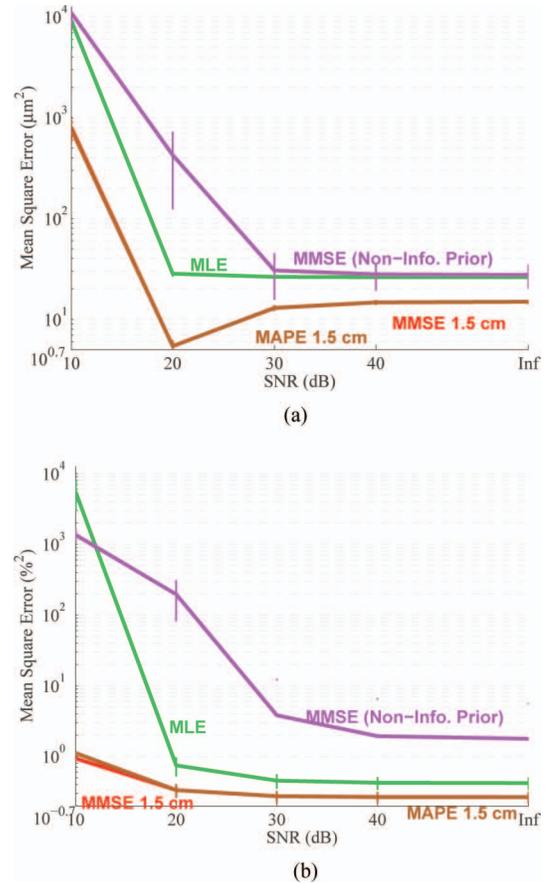


Fig. 9. SNR functional response for (a) compressional motion and (b) strain estimation. This figure shows compressional motion estimation mean-square error (MSE) as a function of SNR. The results are shown for 4 different estimators [maximum likelihood estimator (MLE), maximum *a posteriori* estimator (MAPE), minimum mean square estimator (MMSE) non-informative prior, and MMSE].

shows a trend that is distinctly in the opposite direction. The figure suggests that a kernel overlap of no more than 20% produces the best results. This is an attractive result because it decreases computational overhead.

E. Simulation Results Summary

Considering all of the results as a whole, none of the estimators' results had zero bias, including normalized cross-correlation. Because normalized cross-correlation is algorithmically unbiased, this suggests that there is bias in the data used to track the motion of diffuse scatterers. The biased estimators often had lower bias than normalized cross-correlation, and when the bias of the biased estimator was worse it was still lower than the noise from variance. These observations are consistent with the arguments and examples given by Jaynes on biased estimation [8]. (This pattern of improvement excludes minimum standard deviations that are grossly inappropriate. As an example, a prior standard deviation of 10^{-11} in the ARFI results shown in Fig. 12 would result in much worse behavior not conforming to the pattern just described.) The observation that the variance is still the dominant noise

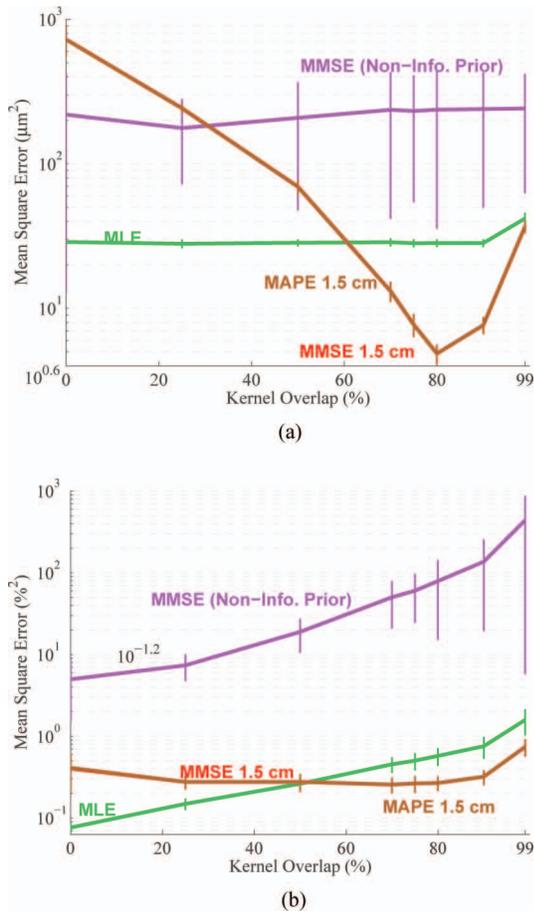


Fig. 10. Kernel overlap functional response for compressional motion and strain estimation. This figure shows compressional motion mean-square error (MSE) for displacement and strain as a function of kernel overlap. Results are shown for four different estimators [maximum likelihood estimator (MLE), maximum *a posteriori* estimator (MAPE), minimum mean square estimator (MMSE) non-informative prior, and MMSE]. (a) The lowest MSE for kernel overlaps typically found in the literature. (b) The best performance for the no-kernel-overlap case. This likely stems from the strain estimation method used here.

source is significant because it signifies that the estimators are not over-biased; or in other words, the estimators are not being dominated by prior information and are sensitive to the data over a fairly wide range of prior standard deviations.

Although these observations about the variance are generally consistent with ultrasonic displacement estimation literature, the observations regarding bias are not. Two hypotheses for the source of the displacement bias are posited. First, it is possible that bias comes from subsample displacements. This seems unlikely, based on the design of the bulk motion simulations, because the displacements were drawn from a Gaussian distribution centered about zero displacement. If the source of the bias was from the randomness of the displacements, it is expected that measured bias would be several orders of magnitude smaller than what was actually measured and would mirror the simulations without noise. Second, it is possible that the bias comes from the noise. This immediately seems to be at least plausible, because Fig. 4 shows that the bias is a

function of the SNR. In this case, the noise-induced bias is simply the orthogonal component to noise-induced variance. The higher-than-expected bias may be related to the correlation length of the noise, which in turn is related to the bandwidth of the signal.

F. In Vivo Results

In vivo ARFI results of a canine right ventricle cryoablation are shown in Fig. 13. The figure shows the same data with motion estimated using normalized cross-correlation and using Bayesian speckle tracking with MAPE. The images show the stiff cryolesion as low displacements and the soft, spared tissue as high displacements. The image formed using Bayesian speckle tracking shows less variability than the image made using normalized cross-correlation. Additionally, there is detail in the MAPE image that is not present (or hard to see) in the normalized cross-correlation image. An example of the additional detail that can be visualized using MAPE is the blockiness on the right side of the MAPE image outside the lesion. This blockiness results from the position of the parallel tracking beams and the early time point after the radiation force excitation that is being shown.³ The impact of the receive beamforming appears to be present in the image formed using normalized cross-correlation, but is not as obvious because of the higher variance of the displacements.

The contrasts of the ARFI image made using normalized cross-correlation and MAPE are 0.4370 and 0.3988, respectively. The contrast-to-noise ratios (CNRs) for the ARFI image made with normalized cross-correlation and MAPE are 0.2534 and 0.6236, respectively.⁴ (The CNR calculated for the MLE-derived ARFI image may seem low, based on the visibility of the lesion, but the image has been saturated on the high-displacement end. Image saturation significantly improves the effective CNR.)

One downside side of Bayesian speckle tracking—as implemented here—is that the final result is influenced by the starting position of the prior scheme. Worse results have been seen when the prior scheme starts outside of the tissue, which could occur because it starts above the myocardium in the standoff pad, or because it starts below the myocardium inside the chamber. The dependence on start position is not seen to be a major limiting factor because the important issue seems to be starting at any tissue location.

³At later times, the blockiness goes away as the displacement propagates away from its original lateral location in the form of some transverse wave (e.g., Lamb wave, shear wave, etc.).

⁴The contrast is defined as

$$C = \frac{\mu_{\text{background}} - \mu_{\text{lesion}}}{\mu_{\text{background}}}$$

The CNR is defined as

$$\text{CNR} = \frac{|\mu_{\text{background}} - \mu_{\text{lesion}}|}{\sqrt{\sigma_{\text{background}}^2 + \sigma_{\text{lesion}}^2}}$$

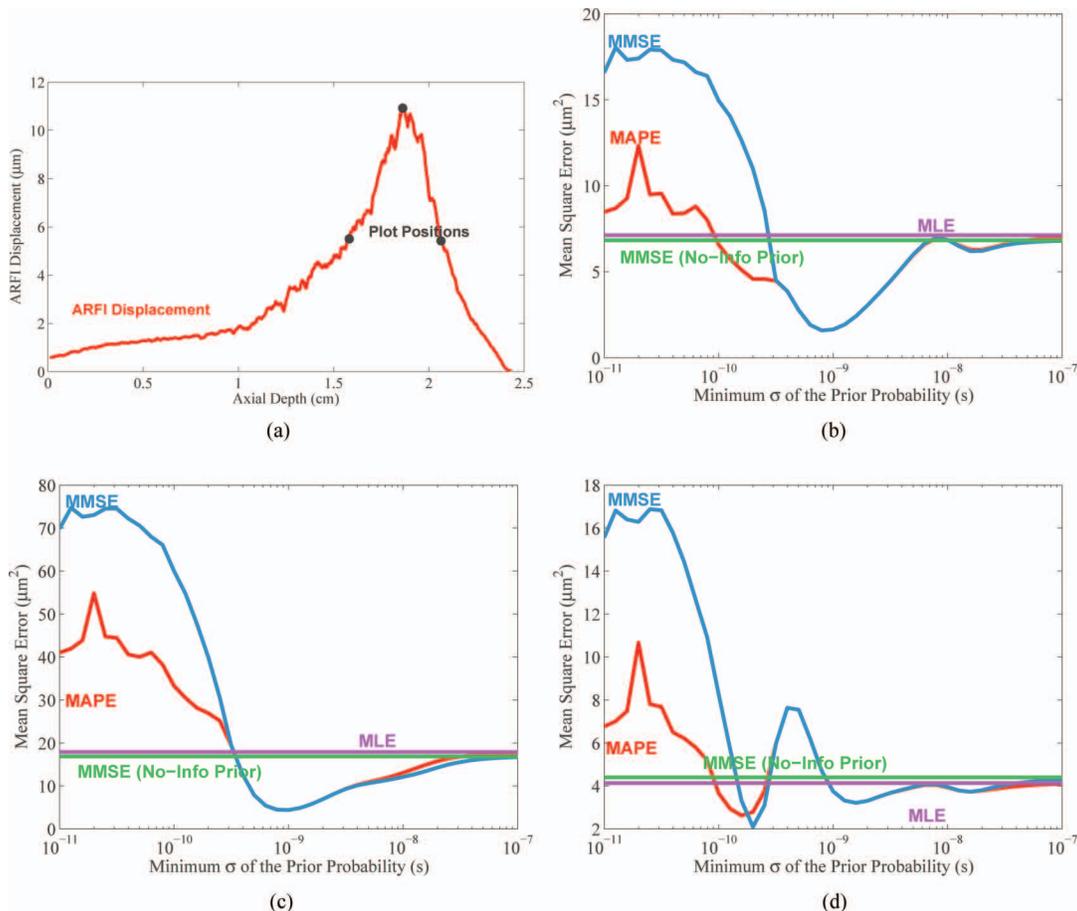


Fig. 11. Acoustic radiation force imaging (ARFI) displacement mean-square error (MSE) as a function of minimum prior width: (a) displacement profile, (b) leading edge, (c) peak, and (d) trailing edge. These data show the mean square error as a function of the minimum allowable prior distribution. The data are shown for the estimates at the positions shown in Fig. 11(a). Four different estimators are compared at each depth. The maximum *a posteriori* estimator (MAPE), the minimum mean square error (MMSE) with a dynamic prior, the MMSE with a non-informative prior, and the maximum likelihood estimate (MLE).

IV. DISCUSSION

The results presented in this paper are significant not because of the specific method used to estimate prior probabilities,⁵ but because with a simple prior scheme, the CRLB can be surpassed. It is entirely expected that within the framework laid out here and in the companion paper [1], displacement estimation performance will continue to improve. This improvement will come from increasingly sophisticated methods for computing prior information [20], [21].

As an example, one of the first ways to realize improved performance is suggested by the results seen from low-pass filtering the unbiased displacements of ARF induced

motion.⁶ The ARF induced motion estimates that were low-pass filtered by a long filter had slightly better MSE performance than any of the biased estimates. The low-pass filter appears to be advantaged (at least in part) because it was implemented non-causally (i.e., symmetrically) and, unlike the biased estimators, was not tethered by a false causality. Forcing a false causality is a common theme in advanced displacement estimation methods [22]–[24] where only estimates before the current estimate are allowed to assist in refining the current estimate. By moving away from this unnecessary constraint, estimator performance will likely improve.

In addition to developing a framework for Bayesian speckle tracking, two new estimators were introduced, the MMSE and the MAPE. The usefulness of each estimator is situational. For large search regions—such as in static elastography imaging—MAPE works well, particularly if the prior has a large variance. For smaller search regions required for estimating ARFI-induced motion, the MMSE

⁵If someone imitates the method used here for computing dynamic priors, they should know that an inherent weakness in the approach is that there is not a good mechanism for allowing the prior PDFs to become wider if the information at an adjacent location appears to be uncorrelated. This is easily observed in the results of the step displacement simulations where small values for σ_{\min} result in nearly unrecoverable performance degradation after the step displacement. This will be a function of both the step size and the standard deviation of the prior distribution.

⁶These results were primarily shown in the figure of frequency response, Fig 6. It was observed that a 47th-order low-pass filter could produce a better MSE.

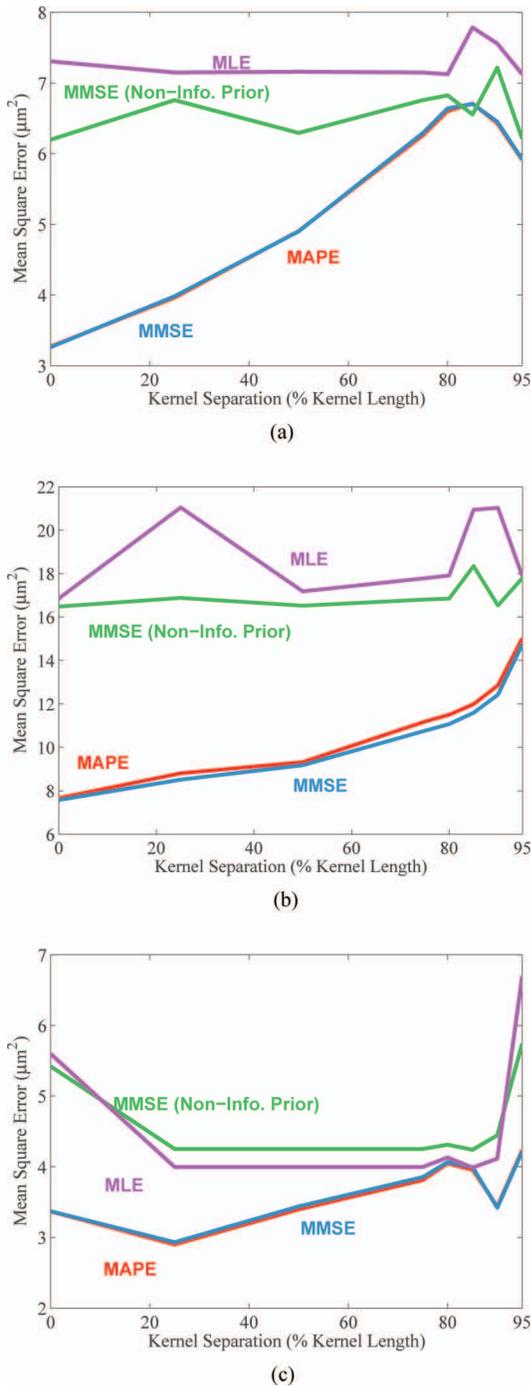


Fig. 12. Acoustic radiation force imaging (ARFI) displacement estimation mean-square error (MSE) displayed as a function of kernel separation. Estimation mean-square error (MSE) is shown for (a) the leading edge, (b) the peak, and (c) the trailing edge of the ARFI displacement as shown in Fig. 11(a). The results for all 4 estimators are shown in each figure.

may work slightly better, particularly for environments with high thermal noise. The MMSE can also be shown to be computationally more efficient if subsample estimators are not used in conjunction with the MAPE or MLE.

Finally, biased displacement estimation has only been applied to 1-D displacement scenarios. Two- and three-dimensional realizations of biased displacement estimation

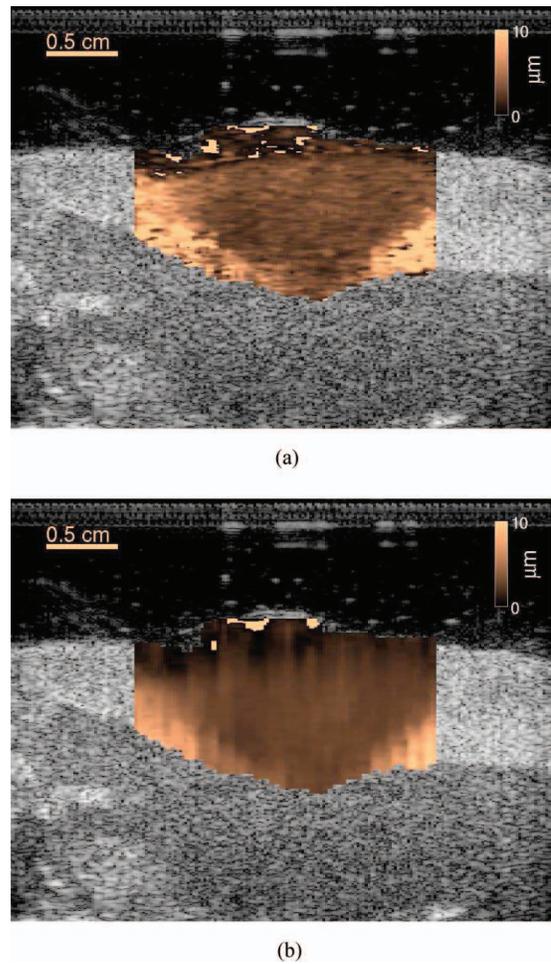


Fig. 13. An *in vivo* comparison of unbiased and biased acoustic radiation force imaging (ARFI) images of a cryoablation lesion (copper scale) embedded in a B-mode image (grayscale). (a) An image of the cryoablation lesion with displacements estimated using normalized cross-correlation. (b) The same data are used to create the image, but the displacements are estimated using the maximum *a posteriori* estimator (MAPE). The contrasts of the normalized cross-correlation and MAPE-derived displacement images are 0.4370 and 0.3988, respectively. The contrast-to-noise ratios (CNRs) of the normalized cross-correlation and MAPE-derived images are 0.2534 and 0.6236, respectively. The MAPE image on the right shows subtle vertical displacement streaks resulting from the parallel tracking sequence in conjunction with the early time step after the initial push used to form this image. Because the image is created at an early time after excitation, there has not been sufficient time for the shear wave to propagate away from the initial excitation position and a banding artifact results. This effect is also present in the image on the left, but the artifact is masked by the larger variance of the unbiased estimator.

can easily be accomplished by extending the normalized cross-correlation function in (6) to two or three dimensions. (The associated α value would also change based on the kernel size.) This has been applied successfully to *in vivo* cardiac speckle tracking [25].

V. CONCLUSION

Biased estimators were devised using several simple prior schemes, the non-informative prior, and a dynamic prior PDF. These prior schemes were coupled with two

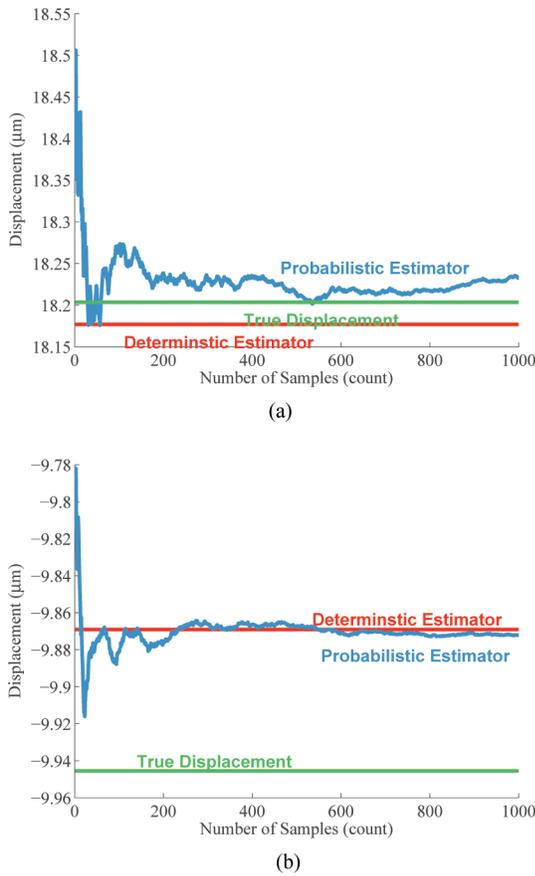


Fig. 14. This figure shows two example results comparing a deterministic implementation [direct multiplication of probability density functions (PDFs)] of the proposed estimator and a true Bayesian implementation (randomly drawing from the prior distribution) for the simple case implemented here. (a) A relatively shallow depth (0.4 mm) with a wide prior probability on τ_0 . (b) A still shallow but deeper depth (5 mm), which has a more refined prior.

methods of extracting displacement estimates from posterior distributions, the MMSE and the MAPE. Using the new biased estimators, it was shown that with relatively small amounts of additional information the CRLB can be surpassed. Although the biased estimators obviously contained bias for appropriate selection of the prior PDF, the dominant noise mechanism was still the variance of the error.

APPENDIX:

STOCHASTIC VERSUS DETERMINISTIC IMPLEMENTATIONS

Bayesian estimators are appropriately implemented stochastically. For example, the estimators developed here should be implemented by drawing random values from the prior PDF that describes the *a priori* knowledge of τ_0 . However, deterministic approximations can in some cases be faster, or at least can be more readily implemented in existing code. It will be shown that appropriate stochastic implementations of the Bayesian estimators and deterministic implementations yield statistically indistinguishable results. The estimators will be implemented deterministically by directly multiplying the prior distribution over τ_0

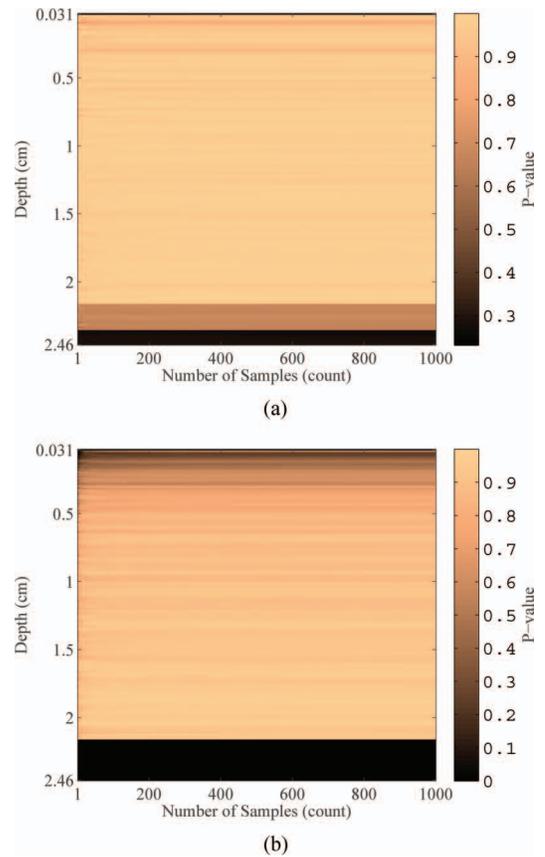


Fig. 15. This figure demonstrates the similarity between the deterministic and probabilistic implementations of the biased estimators: (a) mean comparison and (b) variance comparison. The two figures show the actual p -values calculated as a function of the number of samples from the prior and depth.

and the likelihood function. To make the sampling simple, the prior PDF for both cases will always be a normal distribution with mean and variance equivalent to the previous posterior distribution.

Two example results for specific data realizations are shown in Fig. 14. The results show how the random sampled implementation converges to a final estimate as more samples are drawn. The example shows that the deterministic and stochastic implementations for a given realization of data do not necessarily converge to the same result. Fig. 15 shows the results for simulating 1000 data realizations and comparing the displacement estimate error distribution between the two methods. The figure shows the resulting p -value from statistical tests to compare the mean and variance of the error distributions for the two methods as a function of depth and of number of drawn samples for the stochastic implementation. The results show that the error distributions vary only at very shallow depths (i.e., broad prior distribution) and a small number of draws from the prior distribution. The exception to this is at deeper depths (i.e., narrow priors) where a small portion of the estimates derived from the stochastic implementation yield large errors, which result in statistically different distributions.

Although the deterministic and stochastic approach are shown to be statistically equivalent for the simple case implemented in this paper, this similarity should not automatically be expected to hold for more complex problems.

ACKNOWLEDGMENTS

The authors thank N. Danieley for his computer assistance, M. Jakovljevic for helpful suggestions and D. Dumont and P. Wolf with their assistance in acquiring the *in vivo* data.

REFERENCES

- [1] B. C. Byram, G. E. Trahey, and M. Palmeri, "Bayesian speckle tracking. Part I: An implementable perturbation to the likelihood function for ultrasound displacement estimation," *IEEE Trans. Ultrason. Ferroelect. Freq. Control.*, vol. 60, no. 1, pp. 132–143, 2013.
- [2] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Upper Saddle River, NJ: Prentice-Hall, 1993.
- [3] 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.
- [4] I. Cespedes, M. Insana, and J. Ophir, "Theoretical bounds on strain estimation in elastography," *IEEE Trans. Ultrason. Ferroelect. Freq. Control.*, vol. 42, no. 5, pp. 969–972, Sep. 1995.
- [5] W. Walker and G. Trahey, "A fundamental limit on the performance of correlation-based phase correction and flow estimation techniques," *IEEE Trans. Ultrason. Ferroelect. Freq. Control.*, vol. 41, no. 5, pp. 644–654, 1994.
- [6] W. Walker and G. Trahey, "A fundamental limit on delay estimation using partially correlated speckle signals," *IEEE Trans. Ultrason. Ferroelect. Freq. Control.*, vol. 42, no. 2, pp. 301–308, 1995.
- [7] M. A. Lubinski, S. Y. Emelianov, K. R. Raghavan, A. E. Yagle, A. R. Skovoroda, and M. O'Donnell, "Lateral displacement estimation using tissue incompressibility," *IEEE Trans. Ultrason. Ferroelect. Freq. Control.*, vol. 43, no. 2, pp. 247–256, 1996.
- [8] E. T. Jaynes, *Probability Theory: The Logic of Science*, vol. 1, New York, NY: Cambridge University Press, 2003.
- [9] R. Zahir-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.
- [10] 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.
- [11] 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. Ferroelect. Freq. Control.*, vol. 57, no. 4, pp. 866–882, Apr. 2010.
- [12] M. Schlaikjer and J. A. Jensen, "Maximum likelihood blood velocity estimator incorporating properties of flow physics," *IEEE Trans. Ultrason. Ferroelect. Freq. Control.*, vol. 51, no. 1, pp. 80–92, 2004.
- [13] R. F. Wagner, M. F. Insana, and S. W. Smith, "Fundamental correlation lengths of coherent speckle in medical ultrasonic images," *IEEE Trans. Ultrason. Ferroelect. Freq. Control.*, vol. 35, no. 1, pp. 34–44, Jan. 1988.
- [14] D.S.P. Committee, *Programs for Digital Signal Processing*. Piscataway, NJ: IEEE Press, 1979.
- [15] T. Varghese and J. Ophir, "A theoretical framework for performance characterization of elastography: The strain filter," *IEEE Trans. Ultrason. Ferroelect. Freq. Control.*, vol. 44, no. 1, pp. 164–172, 1997.
- [16] 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. Ferroelect. Freq. Control.*, vol. 52, no. 10, pp. 1699–1712, Oct. 2005.
- [17] 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. Ferroelect. Freq. Control.*, vol. 53, no. 7, pp. 1300–1313, Jul. 2006.
- [18] J. A. Jensen and N. B. Svendsen, "Calculation of pressure fields from arbitrarily shaped, apodized and excited ultrasound transducers," *IEEE Trans. Ultrason. Ferroelect. Freq. Control.*, vol. 39, no. 2, pp. 262–267, 1992.
- [19] J. A. Jensen, "Field: A program for simulating ultrasound systems," *Med. Biol. Eng. Comput.*, vol. 34, suppl. 1, pt. 1, pp. 351–353, 1996.
- [20] W. Debski, "Probabilistic inverse theory," in *Advances in Geophysics (Advances in Geophysics, vol. 52)*. San Diego, CA: Elsevier Academic Press, 2010, pp. 1–102.
- [21] L. R. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," in *Proc. IEEE*, vol. 77, no. 2, pp. 257–286, 1989.
- [22] 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.
- [23] 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.
- [24] Y. Petrank, L. Huang, and M. O'Donnell, "Reduced peak-hopping artifacts in ultrasonic strain estimation using the Viterbi algorithm," *IEEE Trans. Ultrason. Ferroelect. Freq. Control.*, vol. 56, no. 7, pp. 1359–1367, 2009.
- [25] 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.