Robust Tracking of Small Displacements with a Bayesian Estimator

Douglas M. Dumont[1] and Brett C. Byram[1]
[1]Department of Biomedical Engineering, Vanderbilt University, Nashville, TN, USA

Abstract— **Radiation-force-based elasticity imaging describes a group of techniques that use acoustic radiation force (ARF) to displace tissue in order to obtain qualitative or quantitative measurements of tissue properties. Because ARF-induced displacements are on the order of micrometers, tracking these displacements *in vivo* can be challenging. Previously, it has been shown that Bayesian-based estimation can overcome some of the limitations of a traditional displacement estimator like normalized cross-correlation (NCC). In this work, we describe a Bayesian framework that combines a generalized Gaussian-Markov random field (GGMRF) prior with an automated method for selecting the prior's width. We then evaluate its performance in the context of tracking the micrometer-order displacements encountered in an ARF-based method like acoustic radiation force impulse (ARFI) imaging. The results show that bias, variance, and mean-square error performance vary with prior shape and width, and that an almost one order-of-magnitude reduction in mean-square error can be achieved by the estimator at the automatically-selected prior width. Lesion simulations show that the proposed estimator has a higher contrast-to-noise ratio but lower contrast than NCC, median-filtered NCC, and the previous Bayesian estimator, with a non-Gaussian prior shape having better lesion-edge resolution than a Gaussian prior. *In vivo* results from a cardiac, radiofrequency ablation ARFI imaging dataset show quantitative improvements in lesion contrast-to-noise ratio over NCC as well as the previous Bayesian estimator.**

I. Introduction

Ultrasound has become a ubiquitous tool for the diagnostic imaging of soft-tissue in part due to its low-cost, ease-of-use, real-time imaging capabilities, and ability to non-invasively characterize motion within the human body. This last feature has led to the routine use of ultrasound-derived motion information in a number of applications, including Duplex ultrasonography, the registration of free-hand, volumetric data [1], tissue-Doppler echocardiography [2], thermal-strain imaging [3], clutter reduction [4], vector-velocity imaging [5], and ultrasound-based, elasticity imaging [6]–[10].

These techniques depend on accurately measuring the relative displacement between sequentially-received, pulse-echo signals. In practice, displacement estimation quality can be degraded by a number of factors, some dependent on the specific-type of algorithm chosen, while others are related to the characteristics of the transducer and the imaging target. For example, traditional phase-shift time-delay estimators (TDEs) can be subject to aliasing artifacts while time-shift TDEs can be degraded by correlation-peak hopping. Both types of estimators are performance-limited by finite transducer bandwidths, finite signal window-lengths, thermal noise, and signal decorrelation [11]. The effect these parameters have on estimator performance has been described previously and is given by the Cramér-Rao Lower Bound (CRLB), which places a bound on the minimum estimation error variance that can be achieved in a specific, imaging situation by an unbiased estimator [11].

A new class of displacement estimation algorithms have been developed that are capable of achieving significant improvement relative to a CRLB-limited estimator [12]–[14]. These methods—termed Bayesian speckle tracking or Bayesian regularization—use prior knowledge of the estimation task in order to improve the current estimate. Byram et al. showed that Bayesian speckle tracking could produce displacement estimates with a lower mean-square error relative to a CRLB-limited estimator for bulk displacement, strain-based elastography, and radiation-force based elasticity imaging [13]. McCormick et al. proposed an iterative, Bayesian regularization method and showed that significant improvements in estimate quality could be achieved in very few iterations for ultrasound strain images with strains greater than 5 % [14]. While both approaches demonstrate the reduction in estimate error that can be realized with these techniques, there is room for improvement. For example, Byram et al. proposed a directionally-dependent prior scheme, and hypothesized that the falsely-imposed causality limited its performance [13]. McCormick et al. reported a reduction in image quality for strain-fields smaller than 1 % [14], making it unclear how suitable

their approach will be for radiation-force based techniques such as Acoustic Radiation Force Impulse (ARFI) imaging or Shear Wave Elasticity Imaging (SWEI) [8], [10].

To address these limitations, we recently proposed a Bayesian estimator that uses a generalized-Gaussian Markov Random Field (GGMRF) prior, one that has no direction-dependence and allows for adjustments in prior shape [15]. Our initial results showed improvement in mean-square error compared to normalized cross-correlation, but the error analysis was limited to tracking ARF-induced displacements in simulated, homogeneous data with little noise. It was unclear in our preliminary work on how well the estimator performed in visualizing non-homogeneous structures in the presence of increased noise, how to reliably select the prior width from the data or what advantages the new prior offered in comparison to the previous approach [12], [13], [15].

In this paper, we review the algorithm's implementation and propose an automatic approach for selecting the prior's width directly from the data. Next, we compare the algorithm's performance to the previous Bayesian formulation suggested by Byram et al. as well as normalized cross-correlation [12], [13], and extend our analysis to quantifying estimator accuracy as well as lesion visualization over a larger parameter and noise space. Finally, we measure execution time and discuss directions for improvement.

II. METHODS

A. Previous Formulation

The estimator proposed by Byram et al. is formulated as follows [12], [13]. Briefly, Bayes' Theorem is written as

$$P_m(\tau_0|x) = \frac{P_m(x|\tau_0)P_m(\tau_0)}{\int P_m(x|\tau_0)P_m(\tau_0)d\tau + 0}, \quad (1)$$

which expresses the posterior probability density function (PDF) of the displacement estimate τ_0 given the observed, normalized cross-correlation function x computed between two time-shifted RF signals [13]. $P_m(x|\tau_0)$ is the likelihood function, $P_m(\tau_0)$ is the prior PDF , $P_m(x)$ is the marginal likelihood PDF, while m indexes the radio-frequency data axially. The likelihood function $P_m(x|\tau_0)$

$$P_m(x|\tau_0) \propto \exp\left[\frac{SNR_\rho}{\alpha}\sum_{n=\tau_0/\Delta}^{\tau_0/\Delta+N-1}\psi\right] \quad (2)$$

$$\psi = \frac{s_1(n\Delta)s_2(n\Delta - \tau_0)}{\sqrt{\sigma_{s_1}^2 \sigma_{s_2}^2}},$$

scales the normalized, cross-correlation function–calculated between RF signals s_1 and s_2 with sampling period Δ, kernel length N, and sample index n–by an empirically-determined, scaling term α and the signal-to-noise ratio SNR_ρ [12], [13]. The SNR_ρ term is a peak-correlation-coefficient estimate of the SNR [13], [16] and is given as

$$SNR_\rho = \frac{\rho}{1 - \rho}. \quad (3)$$

where ρ is peak correlation-coefficent computed between the time-shifted signals.

$P_m(\tau_0)$ represents the prior distribution of the estimate, and describes how additional information will be incorporated into the posterior PDF. Byram et al. proposed using the posterior distribution of the previous depth's estimate for the current depth's prior distribution, while also imposing a lower bound on the prior's distribution (denoted here as σ) to prevent the prior from becoming too narrow relative to the likelihood [13]. For distributions that fall below this bound, a Gaussian distribution with mean τ_{0m-1} (i.e. the previous depth's estimate) and standard deviation σ is used to represent the prior information for the current depth [13]. The displacement estimate is found using the maximum *a posteriori* (MAP) principle,

$$\hat{\tau_0} = \underset{\tau_0}{argmax}(P_m(\tau_0|x)) \quad (4)$$

which finds the displacement estimate for kernel m that maximizes the posterior probability given the previous estimate's prior and the current estimates's likelihood function [13].

B. Proposed Estimator

The previously-described estimator has several disadvantages which limit its performance. First, the posterior PDF is only influenced by the previous estimate. As will be demonstrated, this false-causal implementation can bias the current depth's estimate towards the previous depth's estimate, distorting the final displacement image and increasing the estimation error. It also constrains the performance of the estimator. Because the posterior PDF only considers the previous depth's distribution, the current prior must be kept artificially broad in order to ensure that the next's estimate's prior has not drifted too far away from the true estimate. This limits the reduction in measurement variance that can be achieved by the estimator. Second, the implementation only considers Gaussian approximations for the prior shape. Other prior shape functions can be attractive for their edge-preservation and noise-reduction properties [17]. Finally, the posterior PDF is only maximized locally based on a

limited amount of information in one iteration. It is likely that performance could be increased if the posterior PDF was continuously refined with additional information, until every estimate's posterior PDF is maximized globally over the entire image.

To overcome these limitations, we proposed an iterative approach that combines an edge-preserving generalized Gaussian-Markov random field (GGMRF) prior with a reformulated likelihood function [15]. We review the algorithm here. First, the likelihood $P_m(x|\tau_0)$ is rewritten explicitly as a function of the two time-shifted RF signals instead of the normalized cross-correlation function,

$$
\begin{aligned}
P_m(x|\tau_0) &= \prod_{n=0}^{N-1} (2\pi\sigma_n^2)^{-\frac{1}{2}} \exp\left[-\frac{1}{\sigma_n^2}(\zeta)^2\right] \\
&= (2\pi\sigma_n^2)^{-\frac{N}{2}} \exp\left[-\frac{1}{\sigma_n^2}\sum_{n=0}^{N-1}(\zeta)^2\right] \\
\zeta &= s_1[n] - s_2[n; -\tau_0],
\end{aligned}
\tag{5}
$$

where the data x is not taken as the cross-correlation function, but rather the RF signal $s_1[n]$. N is the kernel length, τ_0 is the displacement estimate for a kernel m containing n samples, and σ_n^2 is a noise term (described later) that quantifies the uncertainty in the probability distribution of the data [15], [18].[1] The function described in (5) expresses the likelihood of observing $s_1[n]$ given delayed signal $s_2[n; -\tau_0]$ that has been un-delayed by $-\tau_0$. The advantage of utilizing (5) over (2) is that it avoids the need to use either an extremely high RF upsampling rate (to obtain the resolution necessary to track micron-order displacements), or sub-sample estimation (i.e. parabolic)–which can be problematic when utilizing a non-Gaussian shaped prior.

Next, the false-causality is removed by reformulating the prior. Like Byram et al. [12], [13], we assume that the $P_m(\tau_0)$ can be modeled using adjacent information, but rather than restricting the prior distribution of τ_0 to information from the previous depth, we instead consider a localized neighborhood of estimates centered around the current estimate m. The influence of these adjacent estimates can be represented by a GGMRF prior [17],

$$
P_m(\tau_0) = \frac{1}{Z}\exp\left[-\frac{1}{p\lambda^p}\sum_{j\in B} w_j|\tau_0 - \tau_j|^p\right],
\tag{6}
$$

where p describes the shape of the distribution and controls the cost of spatial discontinuities (i.e. edges). The term λ scales the influence of the prior on the final posterior PDF. Z is a normalizing constant, τ_0 is the current estimate for kernel m, and w_j weights the influence of adjacent estimates τ_j within the neighborhood B [15].

The prior described by (6) has several properties that are attractive for estimating displacement fields. First, there are only two tuning parameters (p and λ). Second, the parameter p allows for a high-degree of control in terms of preserving or smoothing image features. When $p = 2$, the prior distribution is a Gaussian, and the estimated displacement profile will have smooth features. As p is reduced, the prior distribution becomes increasingly non-Gaussian and will not penalize sharper edges as heavily [15], [17]. Effectively, the GGMRF prior will assign a larger cost to neighborhoods composed of discontinuous estimates, and a smaller cost to neighborhoods that have a continuous displacement field.

The posterior PDF $P_m(\tau_0|x)$ describes the updated, posterior PDF after considering the data and the additional information. Because our goal is to find the vector of displacement estimates τ_0 that maximizes the posterior PDF rather than the exact, normalized value of the posterior PDF itself, it can be advantageous computationally to rewrite (1) in the log-domain. Rewriting both (5) and (6) in the log-domain and dropping constant and normalizing terms, the log-posterior PDF can be written as

$$
\begin{aligned}
ln(P_m(\tau_0|x)) &\propto ln(P_m(x|\tau_0)) + ln(P_m(\tau_0)) \\
&\propto -\frac{1}{4\sigma_n^2}\sum_{n=0}^{N-1}(s_1[n] - s_2[n; -\tau_0])^2... \\
&\quad -\frac{1}{p\lambda^p}\sum_{j\in B} w_j|\tau_0 - \tau_j|^p.
\end{aligned}
\tag{7}
$$

Note that in (7), the log-likelihood ($ln(P_m(x|\tau_0))$) is a sum-squared difference (SSD) calculation between the RF signals, and the noise term σ_n^2 has been doubled based on the work by Walker [19]. The noise term σ_n^2, which describes the uncertainty in the log-likelihood

---

[1]Note that in Carlson et al.'s formulation [18], the likelihood is computed by delaying signal s_1 by τ_0 and then computing the sum-square difference between s_1 and s_2 when finding the maximum likelihood estimate of the time shift. Here, the initial time-shift and signal segments defining the kernels for s_1 and s_2 are already established by an initial normalized cross-correlation calculation. For convenience, we choose to then evaluate the likelihood for the Bayesian estimator by undelaying the shifted signal s_2 by $-\tau_0$, and computing the sum-square difference.

function, is computed as

$$\sigma_d^2 = \sigma_s^2 + \sigma_n^2$$

$$SNR_\rho = \frac{\sigma_s^2}{\sigma_n^2} = \frac{\sigma_d^2 - \sigma_n^2}{\sigma_n^2} \qquad (8)$$

$$\sigma_n^2 = \frac{\sigma_d^2}{SNR_\rho + 1},$$

where σ_d^2 is the power of the RF data, σ_s^2 is the power of the signal, and SNR_ρ is given by (3) [15].

The displacement estimates for the entire dataset are then found by

$$\hat{\vec{\tau}} = \underset{\vec{\tau}}{argmax} \sum_{m=0}^{M-1} ln(P_m(\tau_0|x)) \qquad (9)$$

which describes the vector of M displacements that produces the largest global log-posterior over the region-of-interest [15].
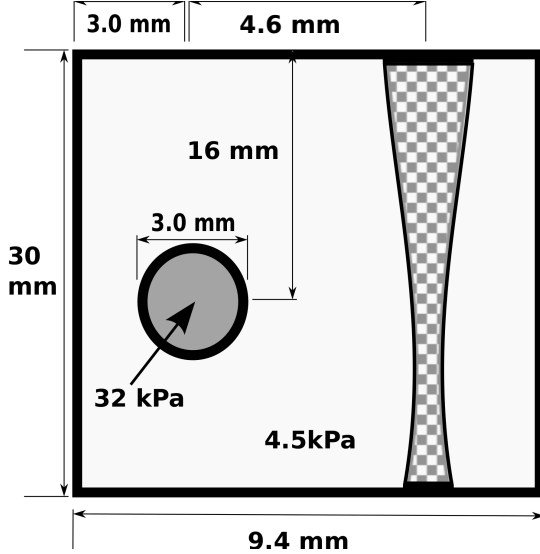

Fig. 1: Diagram showing the field-of-view for a simulated 2D ARFI scan of a 4.5 kPa phantom (off-white) with a 3 mm, 32 kPa embedded lesion (dark gray). Also shown is the lateral location of the ARF excitation and tracking beam (checkerboard pattern) used to characterize estimator performance when tracking a homogeneous ARF response. Figure is not to scale.

C. Implementation

Eqn. (9) is solved as a recursive, non-linear minimization problem, where the negative of (9) is minimized per iteration [15]. The algorithm was implemented in MATLAB (The MathWorks Inc., Natick, MA) using the unconstrained minimization routine fminunc and a quasi-Newton, unguided line-search. The algorithm is as follows:

1) For each kernel m, calculate the initial displacement estimate and cross-correlation coefficient, using a tracking kernel of length N, a search region of length S, and a normalized cross-correlation estimator with parabolic, sub-sample estimation. This estimate is also used to initialize the log-prior.
2) Estimate the noise power at each kernel m using (8) and the geometric mean of the two signal powers (computed over a kernel of length N centered on kernel m) as the estimate for σ_d^2
3) Estimate the log-likelihood by computing the SSD between the two RF signals after shifting the tracked signal s_2 by $-\tau_0$ back towards the reference signal s_1, using Fourier-domain interpolation to compute the sub-sample shift. Each SSD is computed over a kernel of length N.
4) Estimate the log-prior for kernel m using spatially-adjacent displacements (τ_j). The absolute value term, $|\tau_0 - \tau_j|^p$, is approximated by the smooth function($\sqrt{((\tau_0 - \tau_j)^2 + \epsilon)})^p$, where ϵ is a small number [20].
5) Update the log-posterior and sum the negative log-posterior over all M kernels.
6) Repeat steps (3) - (5) for a given number of iterations or until reaching a convergence tolerance on either the cost function (defined as TolFun in fminunc) or the change in parameter values (defined as TolX in fminunc).

D. Numerical Simulations

Estimator performance was evaluated in simulations modeling the response of a 4.5 kPa, linearly-elastic, isotropic phantom with a 3 mm embedded 32 kPa lesion to an acoustic-radiation force excitation[2]. The simulations were created using the simulation package Field II and finite-element methods described previously by Palmeri et al. for modeling the tracking of radiation-force induced displacements within an elastic medium [26]–[29]. Table I lists the parameters used in the simulations.

Two types of experiments were performed. To characterize estimator performance when tracking displacements in a homogeneous medium, 100 simulations with independent speckle realizations were generated with a f/2, 6.67 MHz, 67 μs ARF excitation beam focused at a point 1.8 cm into the phantom and offset 4.6 mm laterally from the lesion center. To characterize estimator performance for lesion visualization, a 2D ARFI image

---

[2]The lesion contrast for the simulation is not meant to be representative of a specific clinical imaging scenario. A wide range of lesion contrasts exist clinically, some of which are smaller or larger than the contrast selected here [21]–[25].

was created by translating the co-aligned excitation and tracking beams across a 9.4 mm imaging field-of-view to synthesize a 48-line ARFI displacement image (n = 25 independent speckle realizations). The simulated RF data were downsampled from the Field II simulation frequency of 1 GHz to 40 MHz, which is a typical scanner sampling frequency. Thermal noise was modeled by superimposing additive, white Gaussian noise (AWGN) on the RF signals to achieve a SNR (denoted hereafter as SNR_T to represent the thermal SNR) of either 10 dB, 15 dB or 30 dB (4.5 kPa homogeneous data, relative to the noiseless RF data). For the simulated ARFI lesion data, the RF data were corrupted to an SNR_T of either 10 dB or 30 dB. Fig. 1 shows the geometry of the 2D ARFI scan as well as the location of the ARF excitation and tracking beam used to simulate the ARF response from a homogeneous, 4.5 kPa region of the phantom.

Table II summarizes the parameters for each estimator (NCC, old Bayesian estimator, and proposed Bayesian estimator). The simulated RF data were first upsampled from 40 MHz to a base sampling frequency of 120 MHz. For each estimator, the kernel length was defined as the ratio of the sampling frequency ($f_s = 120$ MHz) to the center frequency ($f_c = 7$ MHz), and then multiplied by a constant (i.e. $2.73 f_s/f_c$). For the NCC estimator, the RF data and kernel length were upsampled by a factor of three to 360 MHz to compute the normalized cross-correlation function; the coarse-lag estimate was refined using sub-sample estimation. For the old Bayesian estimator, the RF data and kernel length were upsampled by a factor of nine to 1080 MHz to compute the noise term SNR_ρ and to ensure the likelihood function was sampled finely-enough to accurately construct the posterior PDF. Parabolic sub-sampling was then used to refine the coarse-lag, MAP estimate. For the proposed Bayesian estimator, the RF data and kernel length were kept at the base sampling frequency of 120 MHz to reduce the computational time when computing the likelihood. To further reduce the computational time for the Bayesian estimators, the axial field-of-view was restricted from approximately 0.5 cm to 2.5 cm, and the kernel-overlap was reduced to approximately 74%. The NCC result was cropped to the same field-of-view to facilitate comparison.

The proposed log-prior was computed using a two-point neighborhood B that considers information from adjacent kernels (i.e. the preceding and the subsequent axial kernel, or the two nearest kernels for kernels located at either edge), with a uniform weighting factor (w_j) of 0.5 [15]. Table II summarizes the parameter space for p and λ. Unless otherwise indicated, the optimization was run until reaching a user-specified number of iterations (1000 iterations) or reaching a user-specified termination tolerance (TolX, TolFun = 10^{-6}).

TABLE I: ARF-Simulation Parameters

Parameter	Value
Attenuation Coefficient	0.7 dB /MHz /cm
Probe bandwidth	50 %
Probe Focus	1.8 cm
Samp. Freq. (Down-sampled)	40 MHz
Samp. Freq. (Simulation)	1000 MHz
Speed-of-sound	1540 m/s
Tracking Frequency (f_c)	7 MHz
Tracking PRF	10 kHz

TABLE II: Estimator Parameters

Parameter	Value
NCC	
Kernel Length	$2.73\ f_s/f_c$
Kernel Overlap	96 %
Samp. Freq. (Up-sampled)	360 MHz
Bayesian(Old)	
α	4
Likelihood Kernel Length	$2.73\ f_s/f_c$
Likelihood Kernel Overlap	74 %
Minimum Prior σ	$10^{-5} \le \sigma \le 10^5\ \mu m$
Samp. Freq (Up-sampled)	1080 MHz
Bayesian(Proposed)	
ϵ	10^{-6}
Likelihood Kernel Length	$2.73\ f_s/f_c$
Likelihood Kernel Overlap	74 %
p	$1.05 \le p \le 2$
Prior λ	$10^{-5} \le \lambda \le 10^5\ \mu m$
Samp. Freq (Up-sampled, f_s)	120 MHz

E. Selection of λ and σ

The majority of the results presented in this manuscript characterize the parameter space of p and λ (or σ in the case of the old estimator) in varying noise environments. The parameters λ and σ in particular are important as they establish confidence in using additional information when constructing the log-posterior (or posterior) probability.

To investigate the relationship between λ, p, and SNR_ρ for the proposed estimator and σ and SNR_ρ for the old estimator, fifty independent speckle realizations were chosen from the 4.5 kPa data. The reference data and the tracked data from the peak-displacement time-step (0.3 ms after ARF excitation) were then corrupted
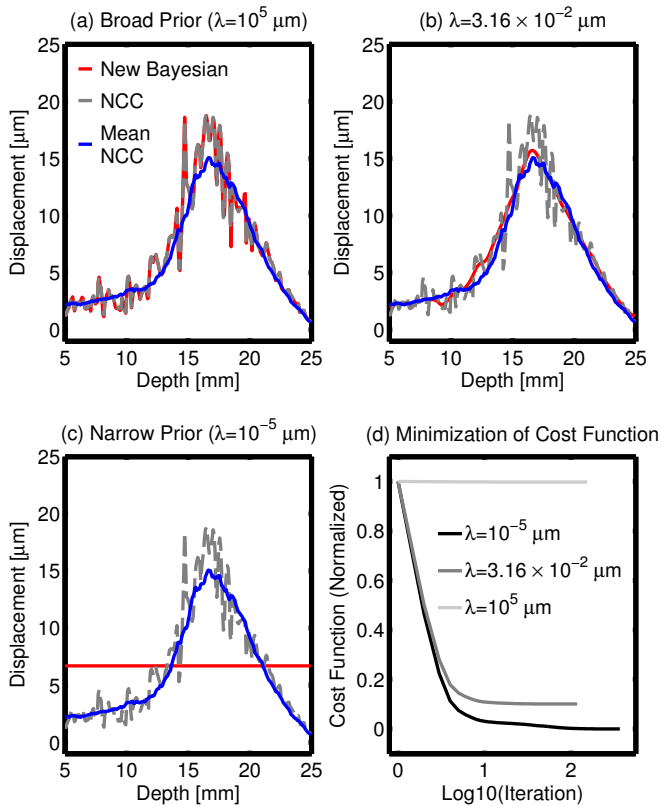
Fig. 2: ARFI displacement profiles (units are in μm) as a function of axial depth from one realization, for the NCC estimator (light-gray, dashed) and the proposed Bayesian estimator ($p = 2$, red) for three prior widths: (a) $\lambda = 10^5$, (b) $\lambda = 3.16 \times 10^{-2}$, and (c) $\lambda = 10^{-5}\,\mu$m. The mean displacement profile of the NCC data averaged over the other 99 realizations is shown in blue. (d) The minimization of the cost function vs iteration (normalized by the initial value) for the results shown in a-c.

with AWGN to a final SNR_T of either 10, 12, 14, 16, 18, 20, 22, 25, 30, or 40 dB. Brent's minimization algorithm [14], [30] was then used to minimize the mean-square displacement error between the estimated Bayesian displacement profile and the true displacement profile (mean, noise-free NCC displacement profile averaged over all 100 realizations, not including the realization being analyzed) in order to find the optimal value of λ or σ for each realization[3]. Brent's algorithm was run using a tolerance of 10^{-5}. The following power-law model was then fit to the optimized λ or σ data using a non-linear regression with bi-square weighting

$$\lambda = ax^b + c, \qquad (10)$$

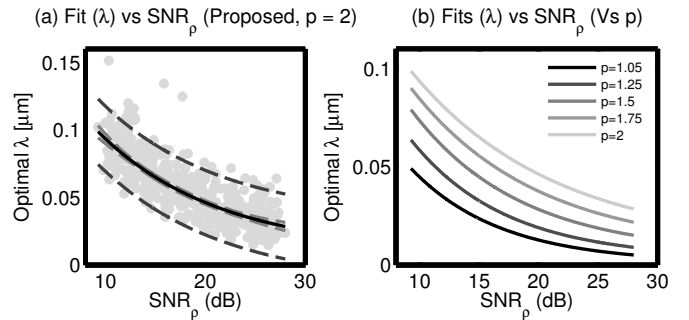[3]McCormick et al. also use Brent's algorithm to investigate their optimal parameter [14].

Fig. 3: (a) Optimal λ (gray dots, n = 50 realizations, 10 noise cases) plotted against the median axial SNR_ρ, for p = 2. The black line shows the regression fit for the power-law model described by (10). The black and gray dashed lines show the upper and lower 95 % prediction and confidence intervals for the fit, respectively. (b) Power-law fits as a function of p. Note that in (8) and (10), SNR_ρ is linear. Here, the data are shown with SNR_ρ defined on a decibel scale to aid visualization.

where x is the median axial SNR_ρ as defined by (3), and a, b, c are the parameters estimated by the model fit.

F. Baseline Error Metrics

Estimator performance was quantified by calculating the bias, variance and mean-square error using the homogeneous, 4.5 kPa data and the following equations

$$bias = E[\tau_0 - \tau_t]$$
$$variance = E[(\tau_0 - \tau_t)^2] - bias^2 \qquad (11)$$
$$MSE = variance + bias^2$$

where τ_o is the estimated displacement and τ_t is the true displacement. The finite-element model displacements were not used as the true values as ultrasonic displacement tracking of ARF-induced displacements typically underestimates the true value [27]. Instead, the mean NCC displacement of all realizations (not including the realization being analyzed) of the noise-free data was used as the true estimate. Error metrics were calculated for displacement profiles obtained 0.3 ms after ARF excitation (approximate time of the peak-displacement) over a range of λ and p. Error metrics were also computed as a function of time following the ARF excitation, using (10) to select λ for the proposed estimator and σ for the old estimator. The results are compared to the NCC data, as well as the NCC data filtered with a one-dimensional median filter. Although median-filtered is commonly performed in two dimensions, the NCC data is only median-filtered axially as the proposed Bayesian estimator is restricted to the axial dimension. The chosen kernel size (approximately 0.6 mm) is more
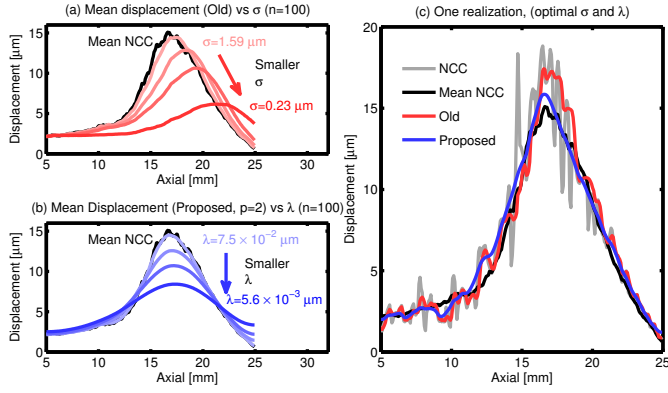
Fig. 4: (a) The mean, axial displacement profile for the previous Bayesian estimator (n =100 realizations) for varying σ (red). The results demonstrate how the false-causal implementation of the prior biases the estimate towards the previous depth, shifting the peak displacement deeper. (b) The same as (a), but for the proposed Bayesian estimator. Unlike (a), the peak displacement is not shifted axially. (c) Example results for one realization for NCC (gray), the old Bayesian estimator (red), the proposed Bayesian estimator (blue) and the mean NCC result (black). σ and λ are selected by (10)

than three times larger than the Bayesian neighborhood and is representative of median filter kernel lengths (axial dimension only) used previously in ARFI imaging (mean 0.43 mm, range (0.17 mm - 0.9 mm)) [31]–[35].

G. Lesion Image Quality

ARFI 2D lesion images were created with each estimator (NCC, NCC with median filtering, old estimator, and proposed estimator), using (10) to select λ or σ for the Bayesian estimators. Lesion contrast and contrast-to-noise were then calculated as

$$Contast = \frac{\mu_o - \mu_i}{\mu_o}, CNR = \frac{|\mu_o - \mu_i|}{\sqrt{\sigma_o^2 + \sigma_i^2}}, \quad (12)$$

where μ_i, μ_o, σ_i^2 and σ_o^2 are the displacement mean and variance inside and outside the lesion, respectively. A 3 mm diameter circle was used to define points inside the lesion, while a 3 mm circle centered at the same depth but offset laterally was used to define points outside the lesion. The same regions were used for all measurements.

Lesion-edge resolution was quantified using an approach similar to one proposed by Rouze et al. [36], where two sigmoid functions are fit to data along an axial line through the center of the lesion,

$$\tau(y) = (\tau_n - \tau_i) \left(\frac{1}{\exp \frac{(n-n_1)}{W_1} + 1} \right) + ...$$
$$(\tau_f - \tau_i) \left(\frac{1}{\exp \frac{-(n-n_2)}{W_2} + 1} \right) + \tau_i, \quad (13)$$

where y is the axial depth, τ_i is the displacement inside the lesion, τ_n and τ_f are the displacements outside the lesion at the near and far lesion boundaries, respectively, n_1 and n_2 are the midpoints for the background-to-lesion and lesion-to-background transitions, respectively, and W_1 and W_2 are windowing parameters that describe the width of each transition. The 80 %-to-20 % lesion-to-background transition distance is given as

$$T_{80,20} = 2ln(4)W_{1,2}, \quad (14)$$

which describes the distance for the sigmoid function to transition from 80 % to 20 % of the background displacement across a given boundary [36]. The width of the lesion is estimated as the difference between the two transition midpoints n_2 and n_1.

H. Execution Time

Execution time was measured by implementing each estimator in MATLAB, using a Dell Precision T5600 Intel Xeon 3.1 GHz computer with 16 GB memory, and recording the time it takes for each estimator to calculate the ARF-induced displacement between two RF lines. Execution time (per realization) was measured for each method over 100 realizations of the 10 dB and 30 dB SNR_T RF data. The parameters listed in Table II were used, except that the kernel overlap was decreased for the NCC estimator to the value used for the Bayesian estimators (i.e. all estimators used the same number of kernels). Equation (10) was used to automatically select λ or σ for the two Bayesian estimators. Execution time was also quantified for the proposed Bayesian estimator as a function of the number of kernels used to estimate the displacement between two RF lines (30 dB SNR_T, one phantom realization). For the experiment quantifying execution time as a function of the number of kernels, the maximum number of iterations was increased to 2000 to ensure convergence for datasets using more kernels.

I. In vivo Example

ARFI displacement images of a cardiac, radio-frequency (RF) ablation in a canine model were made with each estimator (NCC, median-filtered NCC, previous and proposed Bayesian) in order to evaluate in vivo
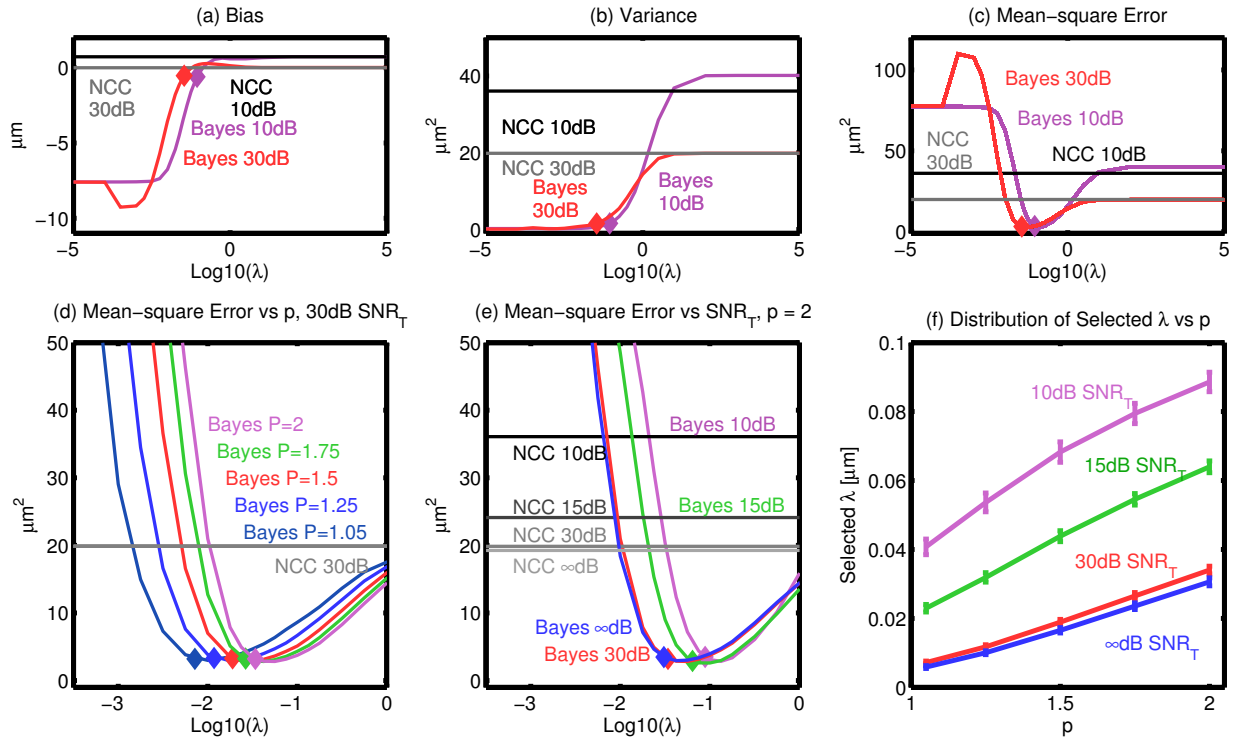
Fig. 5: (a) Bias, (b) variance, and (c) mean-square error (MSE) comparing the Bayesian estimator (red, purple lines, $p = 2$) against the NCC estimator (gray, black) for the 10 dB SNR_T and 30 dB SNR_T data. The plots show the peak bias, variance, and MSE chosen from a 1 mm axial region centered on the peak-axial displacement (determined from the mean noise-free NCC data). (d) MSE as a function of p and λ when SNR_T is held constant, and (e) MSE as a function of SNR_T and λ when p is held constant. Diamonds denote the error metrics when using (10) to select λ. (f) Distribution of λ selected by (10) for varying p and SNR_T.

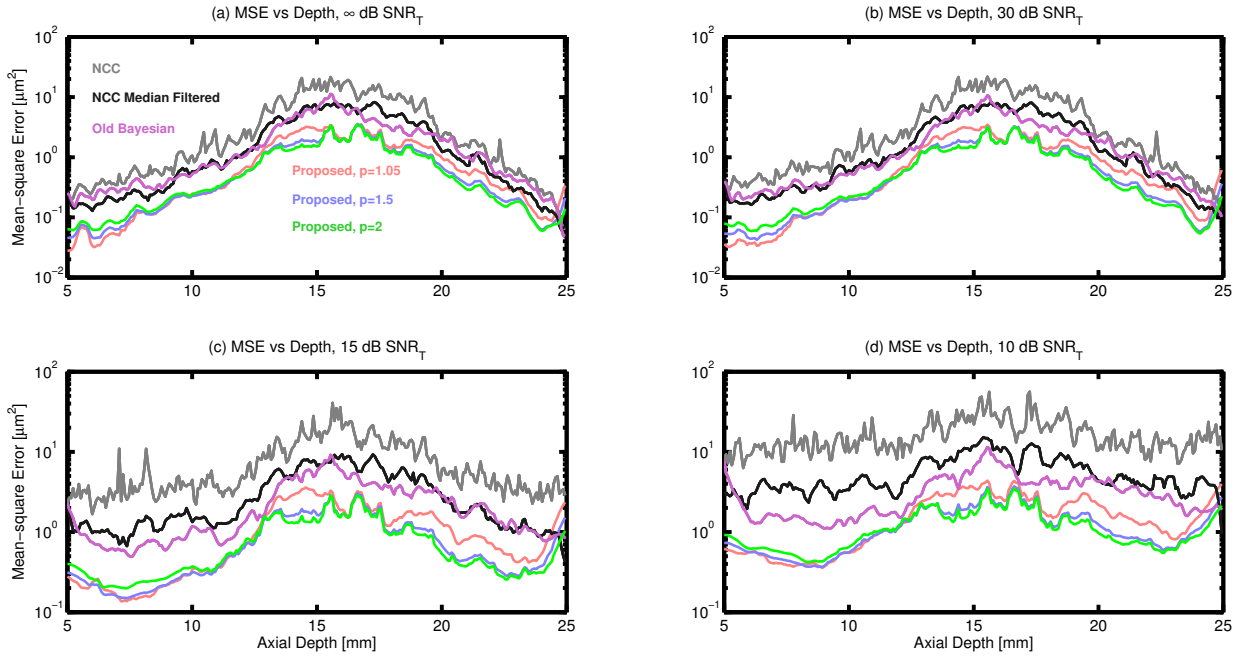
Fig. 6: MSE as a function of axial depth for (a) ∞ dB, (b) 30 dB, (c) 15 dB, and (d) 10 dB SNR_T, for the NCC estimator (gray), the NCC estimator after median filtering (black), the previous estimator (purple) and the proposed Bayesian estimator (red, blue and green) for varying p. For the Bayesian estimators, λ and σ are selected using (10). The results demonstrate the MSE improvements that can be realized with the proposed estimator and suggest that (10) is able to select λ appropriately with for varying p and noise.
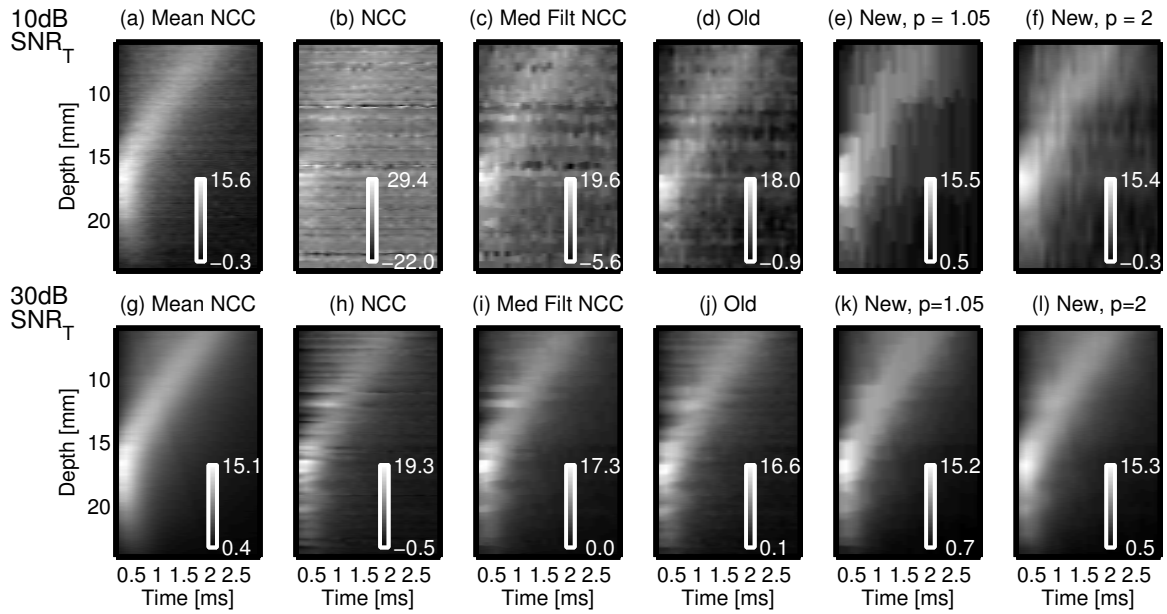
Fig. 7: Displacement vs time images for the (a) mean NCC, (b) NCC, (c) median-filtered NCC, (d) previous Bayesian estimator, (e) proposed non-Gaussian estimator (p=1.05) and (f) proposed Gaussian estimator (p=2) with $SNR_T = 10$ dB. (g)-(l) The same for $SNR_T = 30$ dB. Each image shows the entire dynamic range of the data to avoid saturating the image and to directly evaluate the ability of each technique to reduce the measurement error and recover the true displacement. The units are in μm. The proposed Bayesian results show a qualitative improvement and match favorably with the mean NCC result. The Gaussian results ($p = 2$) have a smoother axial profile.

image quality (imaging data courtesy of Dr. Stephanie Eyerly and Dr. Patrick Wolf, Duke University, Durham, NC). The data were acquired under an open-chest preparation as described in Eyerly et al. [37]. For the NCC and the Bayesian estimators, the same parameters described in Table II were used, except the kernel length and kernel overlap factor are slightly smaller due to the higher center-frequency (f_c=8 MHz). Eqn. (10) was used to automatically select λ or σ for each location, and p was set to either 1.05 or 2.

III. RESULTS

A. Simulation Results - 4.5 kPa Homogeneous Region

Fig. 2(a-c) shows characteristic ARFI displacement profiles (30 dB SNR_T, 0.3 ms after excitation, one realization) for the NCC estimator (gray), the Bayesian estimator when using an appropriately-selected λ (red, Fig. 2(b)), and the Bayesian estimator when λ is selected from either extreme (red, Fig. 2(a, c), overly-broad and overly-narrow λ). The mean NCC displacement profile averaged over the other 99 realizations is shown in blue. The results illustrate that for an overly-broad prior distribution (Fig. 2(a), very large λ)—where the prior has a nearly-uniform probability distribution and little relative

weight is assigned to prior information—the two estimators produce nearly identical results. For an overly-narrow prior distribution (Fig. 2(c), very small λ)—where little relative weight is placed on the data—the cost function becomes minimized when every estimate in the prior term has the same value, biasing all the estimates toward one value. In this example, the Bayesian result for a prior of $10^{-5}\,\mu$m ($6.683\pm3.7\times10^{-5}\,\mu$m, averaged axially) converges to the axial average of the NCC data (mean axial displacement $6.684\,\mu$m). When λ is appropriately-selected (Fig. 2(b)), the Bayesian estimates (red) have lower noise and match well with the mean displacement profile. Fig. 2(d) shows the iteration history for the cost function (normalized by the initial value) and suggests that most of the minimization of the cost function occurs within ten iterations.

Fig. 3 shows an example of how λ (or σ) can be determined empirically as a function of the noise term SNR_ρ. Fig. 3(a) shows the optimal λ (gray dots) that minimizes the mean-square, axial displacement error (relative to the mean, noiseless NCC data) for p=2 for each realization. The black line shows the estimated power-law fit to this data as a function of the median SNR_ρ. The dark gray and black dashed lines show the parameter confidence and prediction intervals for the fit,
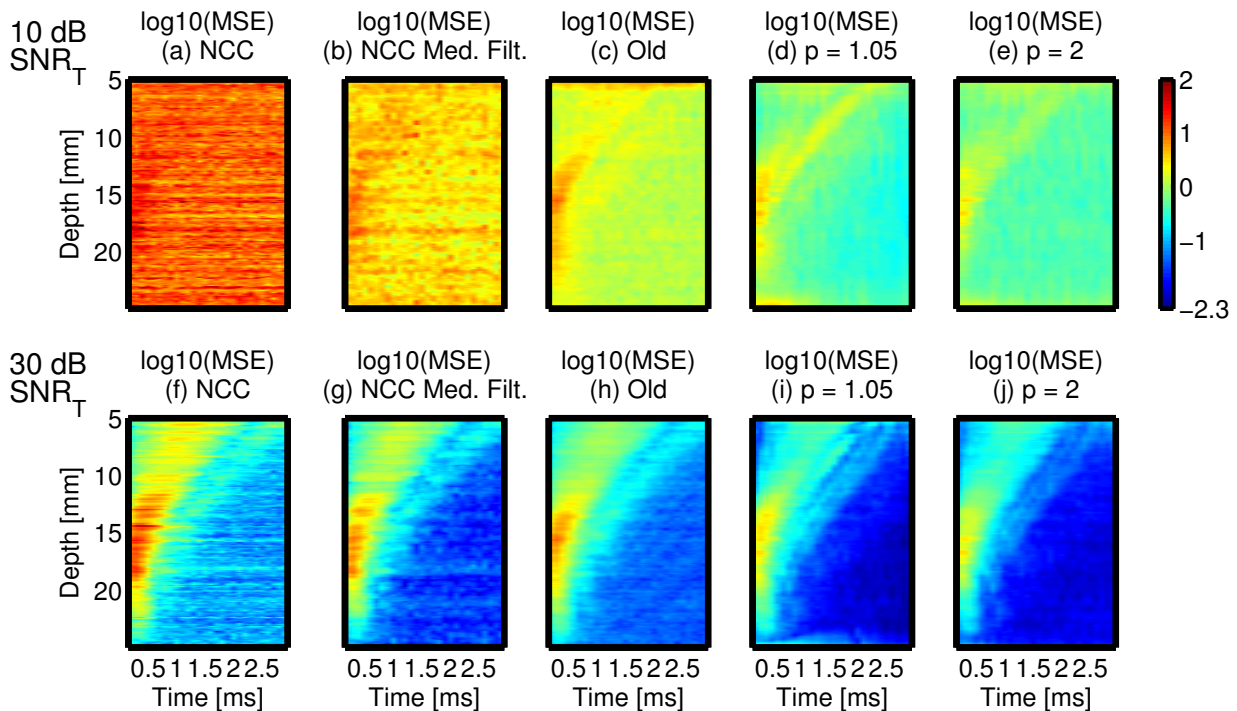
Fig. 8: Mean-square error (units are $log10(\mu m^2)$) for the (a,f) NCC estimator, (b,g) NCC estimator with median-filtering, (c,h) the previous Bayesian estimator, (d,i) the proposed estimator with $p = 1.05$, and (e,j) the proposed estimator with $p = 2$ as a function of time following the excitation. The results demonstrate that the proposed Bayesian estimators offer superior performance in terms of reduced MSE throughout the entire recovery period. The dynamic range is the same for all images and is selected from the minimum and maximum values of the data.

respectively. The parameter λ (or σ) is then estimated by measuring the median SNR_ρ for a given dataset, and using the power-law relationship (10) to determine λ (or σ). Fig. 3(b) shows the power-law fits for λ as a function of p for the proposed estimator. Table III lists the fit parameters to (10) for the previous and proposed Bayesian estimators, where x is the median SNR_ρ measured for a given dataset.

TABLE III: Power-Law Fit Parameters for Selecting λ or σ

Bayesian (Old)	$\sigma = ax^b + c$		
	a	b	c
	11.55878	-0.52161	0.79651
Bayesian (Proposed)	$\lambda = ax^b + c$		
	a	b	c
$p=1.05$	0.16377	-0.57473	0.00109
$p=1.25$	0.19321	-0.54472	0.00325
$p=1.5$	0.20297	-0.47737	0.00575
$p=1.75$	0.19788	-0.41176	0.00781
$p=2$	0.19195	-0.35983	0.00962

Fig. 4 motivates the selection of the proposed Bayesian estimator over the previous estimator. Fig. 4(a)

demonstrates how the false-causal implementation of the previous estimator biases the displacement towards the previous depth's estimate for increasingly-narrow priors, shifting the peak displacement axially and distorting the shape of the displacement profile. Fig. 4(b) demonstrates that by removing this false-causality and appropriately reformulating the problem, the location of the peak displacement is no-longer shifted axially, and narrower priors can be used without biasing the result away from the true displacement. Fig. 4(c) shows characteristic displacement profiles for both Bayesian estimators when using (10) to select the optimal λ or σ. Fig. 4(c) clearly demonstrates the superior performance of the proposed estimator over the previous estimator.

Fig. 5(a-c) shows the trends in peak bias, variance, and mean-square error (MSE) of the proposed estimator for data located within the region of peak displacement. The results are plotted against the parameter λ for two different values of SNR_T. The diamonds denote the bias, variance and MSE when using (10) to select λ automatically. The matched results for the NCC estimator are shown in gray or black. Fig. 5(d) shows the MSE (zoomed-in for visualization) as a function of p and λ for a fixed noise case (30 dB SNR_T) while Fig. 5(e)
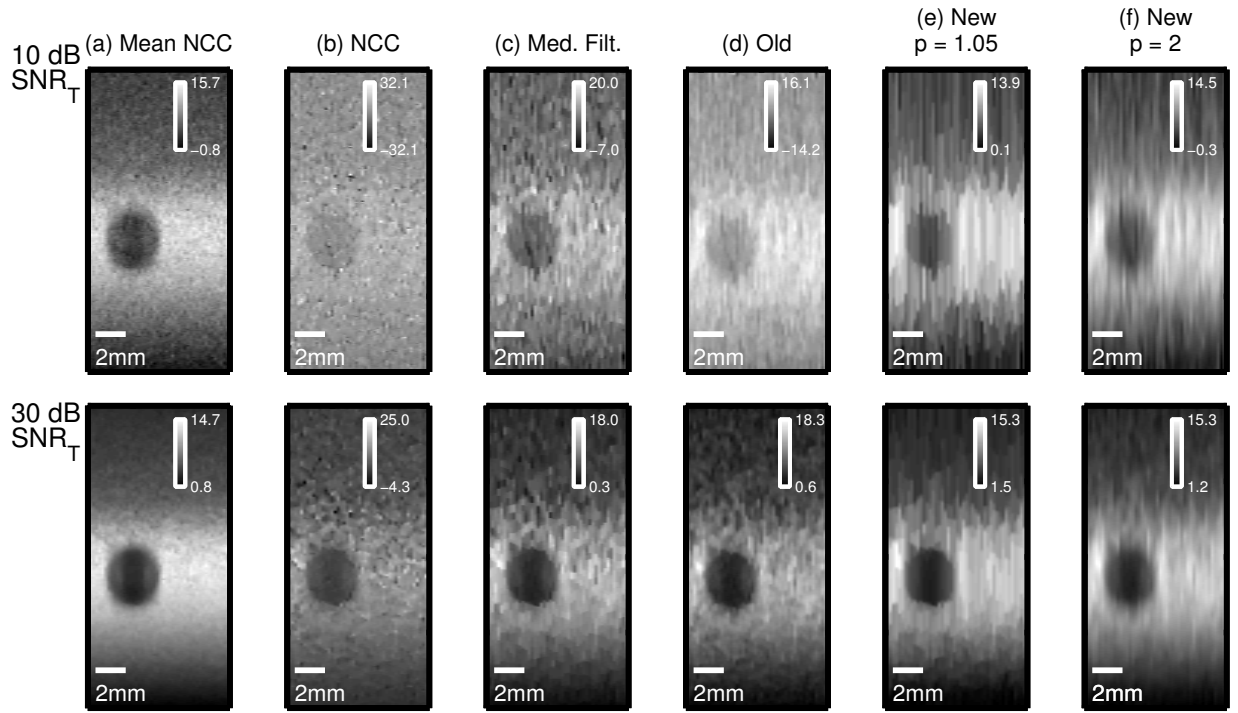
Fig. 9: Example 2D ARFI images (units are in μm) for (column b) NCC, (column c) NCC with median filtering, (column d) the previous Bayesian estimator, (column e) proposed estimator with $p = 1.05$, and (column f) proposed estimator with $p = 2$. (column a) The mean NCC result (n=25 realizations). Each image shows the entire dynamic range of the data to avoid artificially-saturating the image and represents the image that would be shown with no *a prior* knowledge of the true displacement range, which is typically not available clinically. Units are in μm. The proposed Bayesian images have lower variability compared to the NCC images, where the poor contrast-to-noise ratio reduces the apparent visibility of the lesion. Median-filtering or using the previous Bayesian estimator increases the visibility of the lesion compared to NCC, but does not reduce the jitter to the same level as the proposed method. For the 10 dB SNR_T dataset, the previous estimator (d) is unable to correct several of the negative false peaks at the starting depth of the dataset, due to a lack of prior information before those points, limiting its qualitative performance. The non-Gaussian ($p = 1.05$, column e) images compare favorably to the mean NCC result, while the Gaussian prior images ($p = 2$, column f) are blurred along the lesion edge.

shows the MSE as a function of λ and noise for a fixed prior shape (p = 2). The colored diamonds show the corresponding MSE when using (10) to select λ. Fig. 5(f) shows the distributions of the automatically-selected λ.

Fig. 5(a-c) demonstrates the dependence of estimator MSE on bias and variance. For a broad prior distribution (i.e. large λ), the Bayesian MSE is mostly determined by the variance, and for $SNR_T = 30$ dB, the error metrics are similar between the estimators. A larger bias is observed in the 10 dB SNR_T data for both estimators while a larger variance and MSE is observed for the Bayesian estimator, which is hypothesized to be due to the lower-performance of the implemented sum-square difference algorithm, compared to the NCC calculation. As the prior distribution is made narrower– placing more emphasis on prior information–the MSE of the Bayesian estimator decreases with the variance,

reaching an optimal value that is dependent on p and the noise. The larger MSE observed for the 10 dB SNR_T data when using an overly-broad prior is a small trade-off compared to an order-of-magnitude decrease in MSE when using an appropriately-weighted prior distribution (Fig. 5(c)).

Fig. 5(d,e,f) demonstrates the robustness of (10) for selecting λ appropriately in order to minimize the MSE for varying p and noise. The results show that an approximately one-order of magnitude improvement in MSE can be achieved when selecting λ with (10), and that this error reduction occurs with a large decrease in variance (Fig. 5(b)) accompanied by a small increase in bias (Fig. 5(a)). Fig. 5(d,e) demonstrates that there is an approximately an half order-of-magnitude range in λ that provides an order of magnitude in improvement in MSE, and at least an order-of-magnitude range in λ that reduces the MSE by a factor of two. Fig. 5(d) illustrates
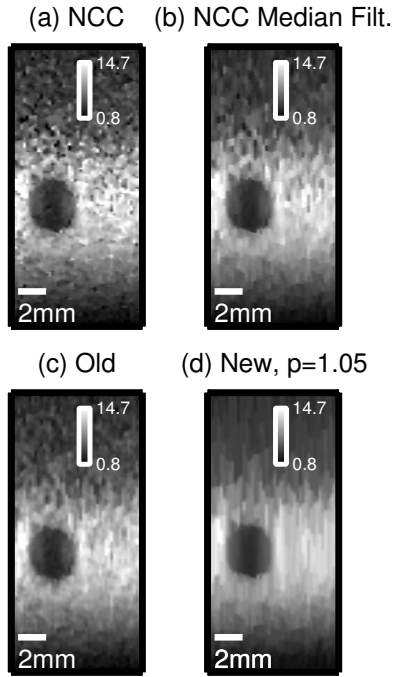
Fig. 10: Example 2D ARFI images for a) NCC, b) NCC with median-filtering, c) the previous estimator, and d) the proposed estimator with p=1.05 for the 30 dB SNR_T data. The images are shown with the same dynamic range as the mean, 30 dB SNR_T displacement data (Fig. 9(a)) in order to compare lesion features between estimators. Note that the data is now saturated. While the lesion is visible in all images, the proposed Bayesian image exhibits lower measurement variability throughout the field-of-view. The apparent contrast and edge resolution are similar in all the datasets. Subtle streaking is apparent in the proposed Bayesian image, due to the 1D nature of the algorithm.

the trend between performance and prior shape, where a lower value of p broadens the prior width, requiring a smaller value of λ to achieve equivalent performance compared to a more Gaussian-like p. Fig. 5(e) illustrates the trend between performance and noise, where narrower priors are required when optimizing MSE for cleaner data. Fig. 5(f) shows that the distribution in the selected λ is small. The results illustrate that it is not just the prior width that is important for minimizing MSE, but rather the prior width and shape in relation to the likelihood distribution of the data should both be considered in order to optimize performance.

Fig. 6 compares the MSE for the proposed Bayesian estimator (using (10) to select λ) as a function of depth for varying p and SNR_T, against NCC, median-filtered NCC, and the previous Bayesian estimator (optimal σ). The results demonstrate that for any SNR_T, the proposed Bayesian estimator reduces the MSE for most

depths by approximately a half-order to one-order of magnitude compared to NCC. The exception is at the edges, where the Bayesian MSE is larger, and for some realizations (∞ dB and 30 dB SNR_T), exceeds the NCC MSE. The results also demonstrate the superior performance for the proposed Bayesian estimator to median-filtered NCC as well as the previous Bayesian estimator. For higher SNR_T, the results show that there is a strong spatial-dependence with depth for the MSE, with the largest MSEs observed within the region of peak axial displacement, which is consistent with a shearing-driven, decorrelation of the tracked RF data. For lower SNR_T, the spatial dependence of the MSE with depth is less apparent, suggesting that the dominant noise mechanism is from thermal noise rather than shearing.

Fig. 7 shows how the displacement profile changes as a function of elapsed time following the ARF excitation for each estimator, with the mean NCC result (averaged over 100 realizations) shown for comparison. Each image shows the entire dynamic range of the data to avoid artificially saturating the data and characterizes the full performance of each estimator. The image quality for the proposed estimator is qualitatively better for both noise realizations and prior shapes compared to NCC, median-filtered NCC, and the old estimator. The Gaussian prior images (Fig. 7(f, l)) appear to have a smoother axial profile and compare more favorably to the mean NCC result than the non-Gaussian prior images (Fig. 7(e, k)) images.

Fig. 8 shows the MSE as a function of axial depth and recovery-time, for two different noise realizations. The results demonstrate the robustness of the proposed Bayesian estimator to varying noise environments and changes in axial-displacement profiles, and show that the improvement in MSE can be maintained throughout the period of displacement recovery. For high SNR_T, the largest MSEs for the proposed Bayesian estimator are either localized at the image edges or follow the location of peak displacement (and thus the location of peak shearing) as it propagates toward the transducer. For a low SNR_T, the spatial-dependence of MSE on axial position is less clear for the NCC estimator and is smaller for the previous and proposed Bayesian estimators. Both results are consistent with displacement-estimate data degraded primarily by thermal noise.

B. Simulation Results - Lesion

Fig. 9 shows characteristic lesion images for NCC, median-filtered NCC, the previous Bayesian estimator, and the proposed Bayesian estimator for two prior shapes. Each image shows the entire dynamic range of
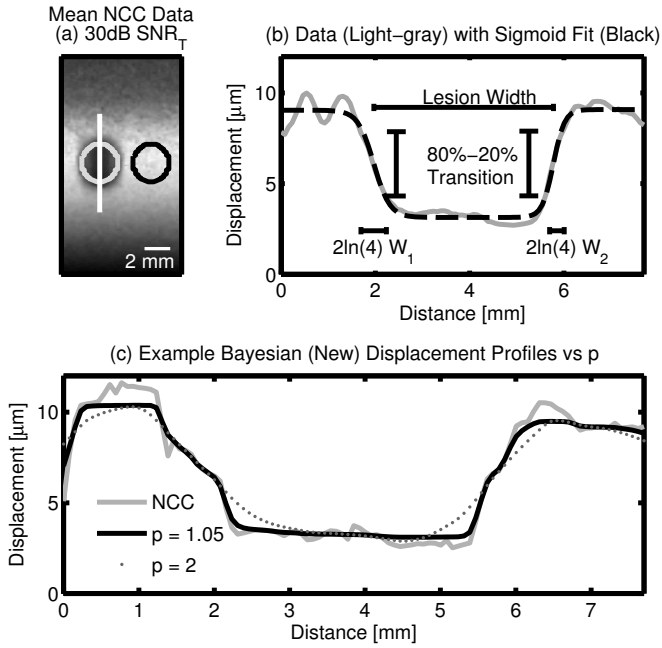
Fig. 11: (a) Example lesion image showing the regions used in the contrast and contrast-to-noise ratio analysis, as well as the axial line (white) used for quantifying edge performance. (b) Double-sigmoid (black dashed line) fit to the data (gray) for the image shown in (a). The lines illustrate the measurements extracted from the fit. (c) Displacement profiles through the center of the lesion (as denoted by the line in (a)) from one simulation dataset. The non-Gaussian prior shape ($p = 1.05$, solid black line) has better edge-preservation performance than the Gaussian prior shape ($p = 2$, gray dotted line),which blurs the edges.

each dataset to avoid artificially saturating the data. The figure shows the images that would be made with no *a prior* knowledge of the true displacement range, and allows for a qualitative comparison of the full performance of each estimator. The proposed Bayesian images have noticeably lower variability than the corresponding NCC images, where the high variability within the NCC image reduces the apparent visibility of the lesion for both noise realizations. Median-filtering the NCC data or using the previous Bayesian estimator improves the image quality considerably, but does not reduce the measurement variability as much as the proposed Bayesian estimators, particularly outside the lesion. The results also show the effect of p on feature preservation. For a Gaussian prior, the lesion margins appear to be more blurred than the corresponding non-Gaussian results, where the lesion-margin is sharper and better defined.

Fig. 10 shows similar data as in Fig. 9, but uses the same dynamic range for all images. Here, the dynamic

range is chosen based on the maximum and minimum values of the mean displacement data (30 dB SNR_T) in order to better compare specific lesion features between the proposed estimator, the previous estimator, NCC, and NCC with median filtering. Unlike in Fig. 9, the images in Fig. 10 are saturated. The figure highlights the ability of the proposed estimator—when using an appropriately-selected prior shape (p=1.05)—to reduce measurement variability without a significant loss of edge resolution. The figure also demonstrates the 1D nature of the algorithm, which introduces a subtle, vertical streaking in the image as the method only considers prior information in the axial dimension.

Fig. 11 show the regions used in the lesion data analysis, an example of the double-sigmoid fit used to evaluate lesion width and edge resolution, and a representative example of the displacement profile through the lesion for the proposed estimator using two different prior shapes. The results demonstrate that the proposed Gaussian prior ($p = 2$) blurs the lesion edge noticeably compared to the non-Gaussian, $p = 1.05$ prior. Fig. 12(a,b) shows that the proposed Bayesian estimators have lower contrast–especially at 10 dB SNR_T–but significantly higher CNR than the NCC, median-filtered NCC, and the previous Bayesian estimator.

Fig. 12(c,d,e) shows the edge-performance and estimated lesion width for each estimator. Only the 30 dB SNR_T data are shown. The fits did not model the 10 dB SNR_T data well enough to allow for the edge resolution to be quantified consistently. The results show that the proposed Bayesian estimator using a Gaussian prior ($p=2$) has much worse edge resolution compared to the other methods, while the proposed non-Gaussian estimator ($p=1.05$) preserves the lesion edges. Our initial expectation was that a Gaussian prior would outperform a non-Gaussian prior both in terms of contrast and CNR, but would have worse edge-preservation due to the quadratic difference penalty in the prior. Here, the data shows that the contrast and CNR performance are similar between the proposed priors for 30 dB SNR_T with a slight loss in contrast at 10 dB SNR_T, while the non-Gaussian prior has much better edge resolution. This is an important result because it suggests that high-spatial frequencies can be preserved by the non-Gaussian prior, without a significant loss in CNR performance for high SNR_T, and only a modest loss in contrast performance at a low SNR_T.

C. Execution Time

Table IV shows the execution time for each method, with the statistics showing the mean and standard deviation over 100 realizations. For the proposed Bayesian
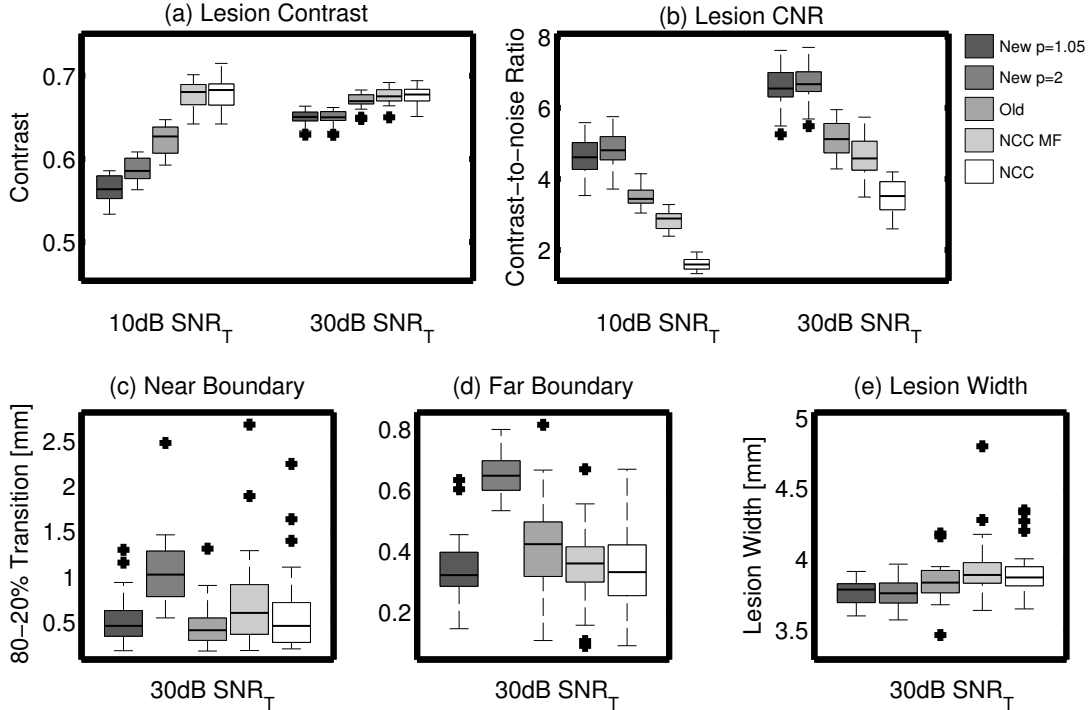
Fig. 12: Standard box-and-whisker plots showing the results for (a) lesion contrast, (b) lesion CNR, (c) near boundary, background-to-lesion transition distance, (d) far-boundary, lesion-to-background transition distance, and (e) lesion width for the proposed Bayesian estimator (p = 1.05, dark gray), the proposed Bayesian estimator (p = 2, medium gray), the previous Bayesian estimator (light, medium gray), median-filtered NCC (light gray), and NCC (white). Outliers are marked by crosses. For (c), (d), and (e), results are only shown for $SNR_T = 30$ dB. The sigmoid-fits were found to be unreliable for the 10 dB SNR_T data, particularly for NCC, and were excluded from analysis.

estimators, the number of iterations to reach convergence are also shown. The results show that the MATLAB execution time for the proposed Bayesian estimators are several orders of magnitude slower relative to the previous Bayesian estimator as well as NCC, which has been shown previously to be capable of achieving run-times on the order of milliseconds for a highly-optimized implementation in C [38]. The results also show that using a strongly non-Gaussian prior (p=1.05) is more computationally-expensive than a Gaussian prior (p=2).

Fig. 13(a) shows how the execution time changes as a function of the number of Bayesian kernels used to compute the final displacement estimate. Fig. 13(b,c) shows the final estimated displacement for the largest (n=1039 kernels) and smallest kernel overlap (n=87 kernels). The NCC results are also shown for comparison. The results suggest that the execution time for the proposed estimator grows at an order of approximately n^2 with respect to the number of kernels, and that the speed penalty can be mitigated to a degree by reducing the number of kernels used to compute the estimate. The results also show that execution times below 10 seconds

TABLE IV: Execution Time

	10 dB Data (n=100)	
Estimator	Execution Time (s)	Iterations
NCC	0.061±0.005	-
Old	0.218±0.003	-
New, p = 1.05	211.2±13.9	340±22
New, p = 2	44.1±2.7	162±10
	30 dB Data (n=100)	
Estimator	Execution Time (s)	Iterations
NCC	0.062±0.006	-
Old	0.217±0.005	-
New, p = 1.05	148.9±11.6	327±26
New, p = 2	34±2.2	123±8

are possible for a MATLAB-only implementation and a Gaussian prior.

D. In vivo Example

Fig. 14 shows *in vivo* ARFI images of a radio-frequency ablation performed in a canine model. In each ARFI image, the ablation lesion is visible as a circular
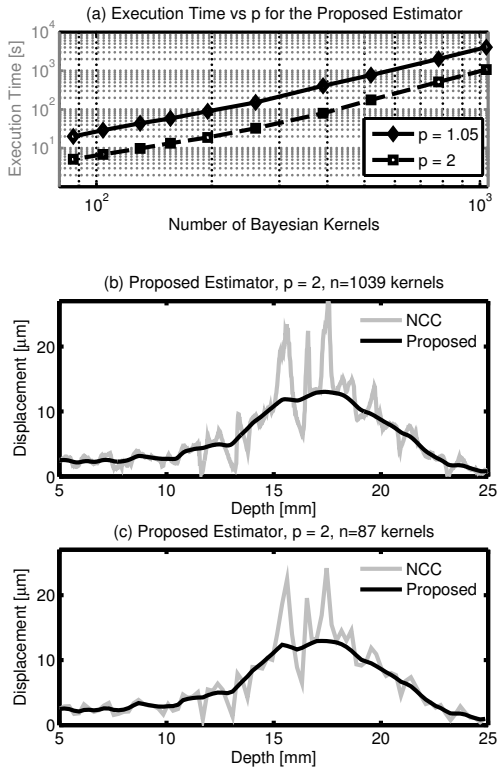
Fig. 13: (a) Execution time for the proposed Bayesian estimator ($p = 1.05, 2$) as a function of the number of Bayesian kernels used to estimate the displacement between two RF lines with an SNR_T of $30\,\mathrm{dB}$. The results show that the execution time increases approximately on the order of n^2 with increasing number of kernels. (b,c) The displacement result when using (b, black) 1039 Bayesian kernels and (c, black) 87 kernels. The NCC result for the same number of kernels is shown in light-gray.

region of lower displacements on the left side of the image, while the non-ablated, right-ventricle appears as a region of higher-displacements to the right of the lesion. The proposed Bayesian images appear to have less axial variability than the corresponding NCC, median-filtered NCC, and previous Bayesian images, particularly outside the lesion.[4] The proposed Bayesian image when using a Gaussian prior ($p = 2$) shows slightly more blurring along the lower-right lesion margin than the corresponding non-Gaussian ($p = 1.05$) prior image. All five images show an alternating, vertical banding artifact due to the parallel-RX tracking scheme used to acquire the data [39]; this artifact is more apparent in the two

[4]Here, the dynamic range is chosen from data within the regions used for the contrast and CNR analysis to ensure that only displacements within the ventricle—and thus actual tissue—are used to scale each image and allow for a qualitative assessment of the full performance of each estimator.

Bayesian images, where the banding is enhanced by the lower measurement variability of the proposed Bayesian estimator. The contrast of the lesion for the two proposed Bayesian estimators (0.8464 for non-Gaussian prior, and 0.8463 for the Gaussian prior) is comparable to the contrast obtained with NCC (0.8460), median-filtered NCC (0.8461), and the previous Bayesian estimator (0.8449), while the contrast-to-noise ratio—4.759, 5.643, 5.426. 6.250, and 6.466 for the NCC, median-filtered NCC, previous Bayesian estimator, non-Gaussian proposed, and Gaussian-proposed prior images—is higher.

IV. DISCUSSION

The results demonstrate the improvements in estimate quality that can be realized with the proposed estimator, and show the algorithm's robustness to a range of noise environments and imaging scenarios. Overall, the data show that the proposed Bayesian approach using (10) to select the prior width is superior in terms of mean-square error (MSE) and contrast-to-noise ratio to the previous Bayesian estimator, NCC, and median-filtered NCC, with comparable performance in edge-resolution when using a non-Gaussian, $p = 1.05$ prior. Compared to previous Bayesian approaches, the results demonstrate that the proposed framework is appropriate for small-displacement, elasticity imaging, [14] and show that significant improvements can be realized be eliminating the false-causal limitation of the previous prior, and replacing the Bayesian framework with an adaptive, iterative approach [13].

While Fig. 5 shows that a wide range of p and λ exist in which improvement can be realized, the optimal choice of the estimator parameters are likely to be application-specific. For imaging applications such as ARFI imaging in which the preservation of high-frequency imaging features is crucial, a non-Gaussian prior is an attractive choice due its ability to preserve edges with comparable contrast and CNR performance to a Gaussian prior. For applications such as SWEI imaging in homogeneous tissue where the main goal is to characterize the propagation velocity and dispersion of the excited shear-wave, a near-Gaussian or Gaussian shape may be a more appropriate choice for the prior, given the faster computational time of the algorithm and its tendency to more accurately preserve the axial displacement profile of a homogeneous region (see Fig. 7).

For any constrained estimator, the ability to select the estimator parameters automatically is advantageous, both for ease-of-implementation as well as robustness to a range of imaging scenarios. The lack of insight into how to select λ based on the underlying data was a
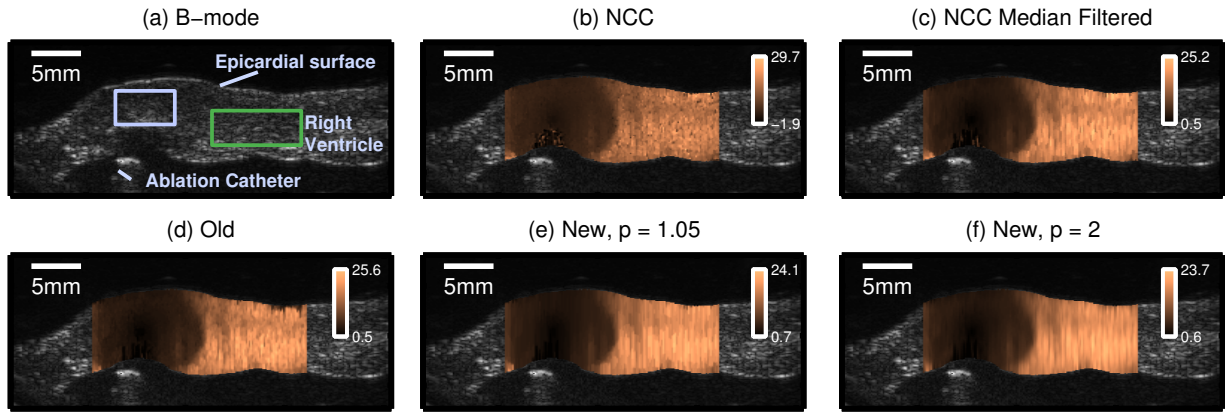
Fig. 14: (a) B-mode image marking anatomical features and the regions (blue, green) used in the contrast and CNR analysis. Example ARFI images for (b) NCC, (c) median-filtered NCC, (d) previous Bayesian estimator, (e) proposed Bayesian estimator ($p = 1.05$), and (f) the proposed Bayesian estimator ($p=2$). The ARFI imaging data is overlayed on the B-mode data. The prior widths for the Bayesian estimators were selected using (10). The dynamic range for each ARFI image is chosen based on the maximum and minimum displacements in the regions shown in (a). Displacement units are given in μm.

major limitation in our initial investigation, making the method tedious to use [15]. Figs. 5, 7, 8, and 9 suggest that selecting λ with (10) provides a reasonable result that is robust to noise as well as changes in the initial prior shape. Fig. 5 does suggest that the selected λ will introduce a small amount of bias into the estimate within the region of peak displacement, which is not surprising given that λ is optimized and selected based on the median noise over the entire dataset. For applications in which no amount of bias can be tolerated, Fig. 5 gives guidance on how λ can be shifted appropriately to minimize the estimator bias while maintaining a comparable mean-square error.

The data in several figures (Figs. 7, 9, and 14) were all presented using the full dynamic range resulting from each estimator, in order to show the complete performance of each estimator. In some cases, it can be more convenient to compare the lesions on a constant dynamic range, which is shown in Fig. 10. In this case, the four images were all displayed scaled between the minimum and maximum of the mean of the 30 dB SNR_T NCC data sets. Unfortunately, this information is not available clinically and setting an appropriate dynamic range for a qualitative ARFI image is challenging. One of the benefits of the proposed algorithm as shown in Figs. 7, 9, and 14 is that images can be made using the full dynamic range of the data without worrying about suppressing information or washing out the image. This characteristic is important for clinical workflow because our results suggest compelling images can be created with no *a prior* knowledge of the displacement dynamic

range, which likely varies with imaging application and between patients.

Further improvements are possible. Our expectation was that the MSE of the proposed Bayesian estimator would have little spatial-dependence with depth. Fig. 6 shows that while the algorithm can reduce the MSE significantly, Fig. 5 and Fig. 8 suggests that the method is unable to reduce the peak-excitation MSE below several μm^2—independent of thermal noise—and that the largest MSEs are localized either at the image-edge or in areas of peak-shearing. Ideally, the noise model should fully account for both signal decorrelation as well as thermal noise; the data suggests that the estimator can only recover estimates degraded by shearing-induced decorrelation up to a certain point. One weakness of the proposed method is that the noise term does not consider correlated, colored-noise. Increased performance can likely be realized with a different noise model that better accounts for correlated-noise within the displacement estimate data, or with a prior scheme that locally-adapts the prior width based on the local noise, rather than using a globally-determined, median estimate. Additionally, because the algorithm only considers information in the axial dimension, the algorithm's performance can vary in the lateral dimension, resulting in a subtle streaking pattern observed in the low SNR images in Fig. 9(e) and (f). It is possible a more optimal result could be obtained with a 2D implementation that considers prior information from both spatial dimensions.

The presented study primarily focuses on the parameters p and λ due to their importance in scaling the weight of additional information relative to the data. While it is

possible to study the effect of other estimator parameters such as Bayesian neighborhood size and the spatial-separation between Bayesian kernels, our experience so far has been that many of these parameters are co-dependent and appear to resample a scaled or shifted version of the prior shape and prior width parameter space. For example, the Bayesian framework presented here assumes independence between neighboring pixels and there is likely some degree of statistical dependence between Bayesian kernels for the presented data, given a kernel-separation distance of approximately $80\,\mu$m and an expected speckle-correlation length of several hundred μm. If the distance between adjacent Bayesian kernels is decreased, the increased statistical dependence between adjacent kernels can be accounted for when selecting λ. This is illustrated in Fig. 13(b-c), where the number of Bayesian kernels used for the estimation in Fig. 13(b) are approximately twelve times greater than the number of kernels used to make the image in Fig. 13(c). In this case, nearly-identical results are obtained between the two datasets by simply reducing the λ used in Fig. 13(b) by approximately a factor of 12.

While the proposed method is appropriate for any displacement estimation task, the high computational-time of the optimization in its current form likely restricts the method to off-line processing rather than real-time displacement estimation. In particular, applications that involve off-line, model-based reconstruction of the viscoelastic or elastic parameters from the raw, displacement data—where displacement accuracy likely impacts reconstruction performance (i.e. viscoelastic SWEI imaging)—or are SNR-limited (i.e. deep abdominal ARFI and SWEI imaging) may benefit from the technique. While Fig. 13 shows that the computational time can be mitigated by simply reducing the number of the kernels used in the estimation, it is likely that this choice will be application specific. A greater number of kernels may be more desirable for ARFI imaging than group-velocity SWEI estimation, where the focus is on tracking the time-of-flight of the propagating wave, rather than imaging detailed structures. Future work will focus on increasing the computational efficiency of the algorithm. The current, non-parallel, MATLAB implementation is computationally-slow and it is likely that a multi-threaded, C/OpenCL or GPU-based implementation could help offset the computational penalty and make the method more suitable for real-time processing.

V. Conclusions

We have presented a new framework for Bayesian motion estimation that combines a generalized-Gaussian Markov Random Field prior with an empirical parameter-selection process to estimate radiation-force induced small-displacements. The method allows for fine-tuning of the prior shape to accommodate a range of elasticity-imaging applications (i.e. ARFI imaging vs SWEI velocity estimation). The results demonstrate that significant improvements in estimate quality can be realized for small-displacement, elasticity imaging applications.

References

[1] A. H. Gee, R. James Housden, P. Hassenpflug, G. M. Treece, and R. W. Prager, "Sensorless freehand 3d ultrasound in real tissue: speckle decorrelation without fully developed speckle." *Medical image analysis*, vol. 10, no. 2, pp. 137–149, Apr 2006.

[2] J. Gorcsan, "Tissue doppler echocardiography." *Current opinion in cardiology*, vol. 15, no. 5, pp. 323–329, Sep 2000.

[3] S.-W. Huang, K. Kim, R. S. Witte, R. Olafsson, and M. O'Donnell, "Inducing and imaging thermal strain using a single ultrasound linear array." *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 54, no. 9, pp. 1718–1720, Sep 2007.

[4] M. a. Lediju, M. J. Pihl, S. J. Hsu, J. J. Dahl, C. M. Gallippi, and G. E. Trahey, "A motion-based approach to abdominal clutter reduction." *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 56, no. 11, pp. 2437–49, Nov. 2009.

[5] J. Udesen and J. A. Jensen, "Investigation of transverse oscillation method." *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 53, no. 5, pp. 959–971, May 2006.

[6] S. Chen, M. W. Urban, C. Pislaru, R. Kinnick, Y. Zheng, A. Yao, and J. F. Greenleaf, "Shearwave dispersion ultrasound vibrometry (sduv) for measuring tissue elasticity and viscosity." *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 56, no. 1, pp. 55–62, Jan 2009.

[7] M. Fatemi and J. F. Greenleaf, "Ultrasound-stimulated vibro-acoustic spectrography." *Science (New York, N.Y.)*, vol. 280, no. 5360, pp. 82–85, Apr 1998.

[8] K. R. Nightingale, M. L. Palmeri, R. W. Nightingale, and G. E. Trahey, "On the feasibility of remote palpation using acoustic radiation force." *The Journal of the Acoustical Society of America*, vol. 110, no. 1, pp. 625–634, Jul 2001.

[9] J. Ophir, I. Céspedes, H. Ponnekanti, Y. Yazdi, and X. Li, "Elastography: a quantitative method for imaging the elasticity of biological tissues." *Ultrasonic imaging*, vol. 13, no. 2, pp. 111–134, Apr 1991.

[10] A. P. Sarvazyan, O. V. Rudenko, S. D. Swanson, J. B. Fowlkes, and S. Y. Emelianov, "Shear wave elasticity imaging: a new ultrasonic technology of medical diagnostics." *Ultrasound Med. Biol.*, vol. 24, no. 9, pp. 1419–1435, Nov 1998.

[11] W. Walker and G. Trahey, "A fundamental limit on the performance of correlation based phase correction and flow estimation techniques," *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, vol. 41, no. 5, pp. 644–654, Sept 1994.

[12] B. Byram, G. Trahey, and M. Palmeri, "Bayesian speckle tracking. part i: an implementable perturbation to the likelihood function for ultrasound displacement estimation." *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 60, no. 1, pp. 132–143, Jan 2013.

[13] ——, "Bayesian speckle tracking. part ii: biased ultrasound displacement estimation," *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, vol. 60, no. 1, pp. 144–157, January 2013.

[14] M. McCormick, N. Rubert, and T. Varghese, "Bayesian regularization applied to ultrasound strain imaging." *IEEE transactions on bio-medical engineering*, vol. 58, no. 6, pp. 1612–1620, Jun 2011.

[15] D. Dumont, M. Palmeri, S. Eyerly, P. Wolf, and B. Byram, "Feasibility of using a generalized-gaussian markov random field prior for bayesian speckle tracking of small displacements," in *Ultrasonics Symposium (IUS), 2014 IEEE International*, Sept 2014, pp. 1845–1848.

[16] E. I. Céspedes, J. Ophir, and S. Alam, "The combined effect of signal decorrelation and random noise on the variance of time delay estimation," *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, vol. 44, no. 1, pp. 220–225, Jan 1997.

[17] C. Bouman and K. Sauer, "A generalized gaussian image model for edge-preserving map estimation," *Image Processing, IEEE Transactions on*, vol. 2, no. 3, pp. 296–310, Jul 1993.

[18] J. Carlson and F. Sjoberg, "Simultaneous maximum likelihood estimation of time delay and time scaling," in *Proceedings of the 6th Nordic Signal Processing Symposium - NORSIG 2004*, June 2004, pp. 260–263.

[19] W. Walker, "Adaptive ultrasonic imaging performance for near-field aberrating layers," Ph.D. dissertation, Duke University, 1995.

[20] S.-I. Lee, H. Lee, P. Abbeel, and A. Y. Ng, "Efficient l1 logistic regression," in *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI-06)*, July 2006.

[21] V. Egorov and A. Sarvazyan, "Mechanical imaging of the breast," *Medical Imaging, IEEE Transactions on*, vol. 27, no. 9, pp. 1275–1287, Sept 2008.

[22] A. Guibal, C. Boularan, M. Bruce, M. Vallin, F. Pilleul, T. Walter, J. Y. Scoazec, N. Boublay, J. Dumortier, and T. Lefort, "Evaluation of shearwave elastography for the characterisation of focal liver lesions on ultrasound." *European radiology*, vol. 23, no. 4, pp. 1138–1149, Apr 2013.

[23] M. Muller, J.-L. Gennisson, T. Deffieux, M. Tanter, and M. Fink, "Quantitative viscoelasticity mapping of human liver using supersonic shear imaging: preliminary in vivo feasibility study." *Ultrasound in medicine & biology*, vol. 35, no. 2, pp. 219–229, Feb 2009.

[24] A. Samani, J. Zubovits, and D. Plewes, "Elastic moduli of normal and pathological human breast tissues: an inversion-technique-based investigation of 169 samples." *Physics in medicine and biology*, vol. 52, no. 6, pp. 1565–1576, Mar 2007.

[25] H. Xie, S. Zhou, V. Shamdasani, Y. Shi, J. Robert, J. Fraser, S. Chen, and J. Greenleaf, "Quantitative assessment of thermal lesion stiffness in the liver: Initial ex vivo results," in *Ultrasonics Symposium (IUS), 2011 IEEE International*, Oct 2011, pp. 2098–2101.

[26] 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 transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 52, no. 10, pp. 1699–1712, Oct 2005.

[27] 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 transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 53, no. 7, pp. 1300–1313, Jul 2006.

[28] J. A. Jensen, "Field: A program for simulating ultrasound systems," in *Medical and Biological Engineering and Computing, Vol. 4, Supplement 1, Part 1*, 1996, pp. 351–353.

[29] J. A. Jensen and N. B. Svendsen, "Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 39, pp. 262–267, 1992.

[30] R. Brent, *Algorithms for Minimization without Derivatives*. Englewood Cliffs, NewJersey: Prentice-Hall, 1973.

[31] P. Hollender, S. Rosenzweig, K. Nightingale, and G. Trahey, "Micro-elasticity (μ-e): Cnr and resolution of acoustic radiation force impulse imaging and single- and multiple track location shear wave elasticity imaging for visualizing small targets," in *Ultrasonics Symposium (IUS), 2014 IEEE International*, Sept 2014, pp. 703–706.

[32] S. J. Hsu, R. R. Bouchard, D. M. Dumont, P. D. Wolf, and G. E. Trahey, "In vivo assessment of myocardial stiffness with acoustic radiation force impulse imaging." *Ultrasound in medicine & biology*, vol. 33, no. 11, pp. 1706–1719, Nov 2007.

[33] S. Hsu, B. Fahey, D. Dumont, P. Wolf, and G. Trahey, "Challenges and implementation of radiation-force imaging with an intracardiac ultrasound transducer," *Ultrasonics, Ferroelectrics, and Frequency Control, IEEE Transactions on*, vol. 54, no. 5, pp. 996–1009, May 2007.

[34] S. J. Hsu, R. R. Bouchard, D. M. Dumont, C. W. Ong, P. D. Wolf, and G. E. Trahey, "Novel acoustic radiation force impulse imaging methods for visualization of rapidly moving tissue." *Ultrasonic imaging*, vol. 31, no. 3, pp. 183–200, Jul 2009.

[35] S. Rosenzweig, M. Palmeri, and K. Nightingale, "Analysis of rapid multi-focal-zone arfi imaging," *Ultrasonics, Ferroelectrics, and Frequency Control, IEEE Transactions on*, vol. 62, no. 2, pp. 280–289, Feb 2015.

[36] N. C. Rouze, M. H. Wang, M. L. Palmeri, and K. R. Nightingale, "Parameters affecting the resolution and accuracy of 2-d quantitative shear wave images." *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 59, no. 8, pp. 1729–1740, Aug 2012.

[37] S. A. Eyerly, M. Vejdani-Jahromi, D. M. Dumont, G. E. Trahey, and P. D. Wolf, "The evolution of tissue stiffness at radiofrequency ablation sites during lesion formation and in the peri-ablation period," *Journal of Cardiovascular Electrophysiology*, 2015, to be published.

[38] G. F. Pinton, J. J. Dahl, and G. E. Trahey, "Rapid tracking of small displacements with ultrasound." *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 53, no. 6, pp. 1103–1117, Jun 2006.

[39] J. J. Dahl, G. F. Pinton, M. L. Palmeri, V. Agrawal, K. R. Nightingale, and G. E. Trahey, "A parallel tracking method for acoustic radiation force impulse imaging." *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 54, no. 2, pp. 301–312, Feb 2007.