Scatterer Reconstruction and Parametrization of Homogeneous Tissue for Ultrasound Image Simulation

Oliver Mattausch[1] and Orcun Goksel[1]

Abstract—Numerical simulation of ultrasound images can facilitate the training of sonographers. A realistic appearance of simulated ultrasonic speckle is essential for a plausible ultrasound simulation. An efficient and realistic model for ultrasonic speckle is the convolution of the ultrasound point-spread function with a parametrized distribution of point scatterers. Nevertheless, for a given arbitrary tissue, such scatterer distributions that would generate a realistic image are not known a priori, and currently there is no principled method to extract such scatterer patterns for given target tissues to be simulated. In this paper we propose to solve the inverse problem, in which an underlying scatterer map for a given sample ultrasound image is estimated. From such scatterer maps, it is also shown that a parametrization distribution model can be built, using which other instances of the same tissue can be simulated by feeding into a standard speckle generation method. This enables us to synthesize images of different tissue types from actual ultrasound images to be used in simulations with arbitrary view angles and transducer settings. We show in numerical and physical tissue-mimicking phantoms and actual physical tissue that the appearance of the synthesized images closely match the real images.

I. Introduction

Ultrasound (US) simulation is essential in the context of training simulators for medical students and doctors. *Interpolative* methods use pre-recorded patient data, through which the ultrasound plane can be sliced in real-time with [5] or without [1] deformation effects. Despite the photo-realistic images, this constrains the simulations drastically by only allowing simulation at voxel where data is available and also limiting image settings and viewing angle to those used in acquisition. *Wave-based* simulation, such as Field II [7], is an alternative that is accurate but slow and thus typically used only for offline simulation. In contrast, fast ray-based methods [12], [11], [3] are suitable for interactive US generation. While ray-based methods successfully model large-scale structures and wave interactions, they are less suited to simulate sub-wavelength interactions, at the scale of which wave-like properties become predominant and create the typical noise patterns known as ultrasonic speckles. These are caused by the interference of diffuse reflections from countless microscopic structures in the tissue. These structures, which are too small to be observed directly but nevertheless scatter the ultrasound, are commonly called *scatterers*.

For simulation, spatial scatterer distributions are often created statistically [8], [7], from which convolution tech-
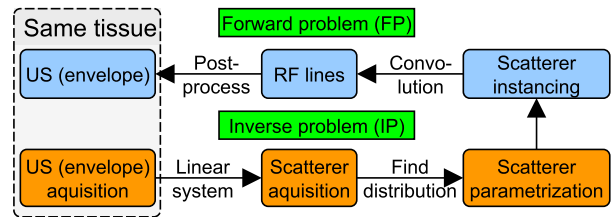
Fig. 1. Overview of the problem definition. Using an input image of ultrasonic speckles and assuming an underlying scatterer model, we first reconstruct the scatterer distribution by solving an inverse problem and then statistically parametrize this distribution. Subsequently, new tissue instances can be generated from this parametrization with the forward pipeline generating realistic speckle patterns resembling the input tissue.

niques for speckle generation can generate realistic speckle patterns [2], [4]. In contrast to earlier works modeling such scatterer locations as continuous spatial variables, a recent ray-based simulation [3] suggested the use of discrete scatterer approximations in order to facilitate the GPU rendering pipeline where scatterers are treated as 3D texture patterns. In that work and others, it has been also shown that for a *fine-tuned* parametrization of a spatial scatterer distribution, realistic looking tissue appearance can be generated. Nevertheless, finding such a parametrization for a particular tissue type currently would require a tedious process of manual trial-and-error by observing generated images, as the final appearance relates to the parametrization in a non-trivial fashion. To our knowledge, there is no automatic method available yet to find such a parametrization. We propose an inverse-problem approach to extract scatterers based on speckle observations, such that images of this tissue region can be simulated from arbitrary viewing angles and with different imaging parameters. Additionally, we propose to estimate parametric models for those scatterer distributions in order to simulate other instances of the same tissue type, e.g. to simulate a whole region from an exemplary patch.

II. Methods

Typical convolution-based ultrasound simulation takes a scatterer distribution and convolves it with a point-spread function to generate simulated ultrasound radio-frequency (RF) data, which is then post-processed (envelope detection, dynamic-range adjustment, etc) to generate an image; called the *forward-problem* (FP) in the top row of Fig. 1. Accordingly, we define the *inverse problem* (IP) as recovering the scatterer parameterization from given ultrasound images.
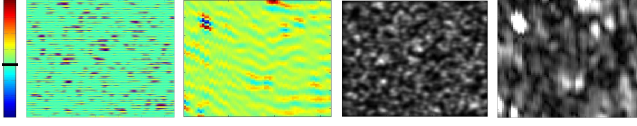
Fig. 2. RF data from synthetic and real phantoms (left) and after envelope detection (right). Envelope images are the input to our method.

A. Forward problem

We assume a well-established model for the ultrasound speckle and its simulation [2], [9], [8]. The methods are presented below in 2D ignoring the finite ultrasound beam-thickness, although they could extend to 3D without loss of generality. The model computes ultrasonic speckle intensity $I(x,y)$ by convolving point-like scatterers in the tissue $T(x,y)$ with the ultrasonic impulse response $H(x,y)$ from each of those, called the point-spread function (PSF), i.e.:

$$I(x,y) = T(x,y) * H(x,y) . \qquad (1)$$

$H(x,y)$ approximates the ideal $sinc$ kernel, a Gaussian modulated with a cosine in the axial direction y, i.e.:

$$H(x,y) = e^{\frac{x}{\sigma_x} + \frac{y}{\sigma_y}} \cos(2\pi f y) . \qquad (2)$$

T is traditionally modeled as a collection of scatterers with varying amplitudes and spatial positions [7]. Assuming to linearity, a speckle pattern is generated by the superposition (convolution) of the individual impulse responses of all such scatterers. The final B-mode image is then obtained after envelope detection, time-gain compensation, and dynamic-range (log) compression.

To utilize GPU pipelines, Bürger et al. [3] used a discretized version of this model where scatterers are represented on a discretized texture grid, i.e. $T(x,y)$. They also introduced a 3-parameter approximation to model tissue-specific sparse scatterer patterns. This uses a normal distribution $\mathcal{N}_{\mu,\sigma}$, which has two parameters, and a scatterer sparsity parameter r, which is the *ratio* of texels populated with a scatterer. We adopt a 2D version of this model to formulate the forward problem. The goal is then solving the inverse problem of extracting μ, σ, and r from a sample US image.

The steps for our proposed IP approach can be seen in the top row of Fig. 1, and are detailed below.

B. US (envelope) acquisition

The first input to our method is the US image of the tissue region that we wish to simulate. Using our method, a scatterer map of that region can then be recovered. Nevertheless, the goal of modeling the parametrization of the scatterers for a particular tissue type implies the constraint that the input image should be from a homogeneous tissue region and does not have large structures that can not be modeled statistically. Accordingly, a small homogeneous US image patch is selected from the tissue targeted for simulation. We acquire RF data and use its *envelope* (see Fig. 2) as input to our method, since (i.) the phase information in RF makes IP formulations ill-conditioned, and (ii.) the envelope contains
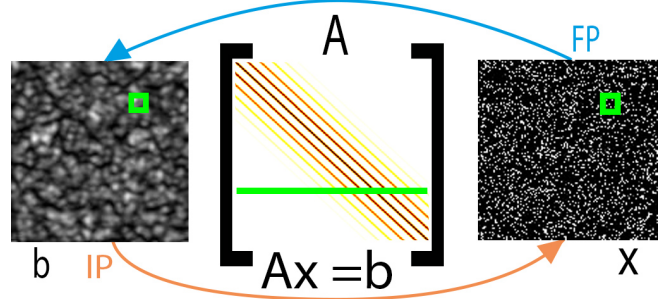

Fig. 3. Since convolution with a PSF is a linear operation and can be expressed with a sparse convolution matrix A acting on the scatterers x to yield image pixel intensities b, its effect can hence be reversed by solving the inverse problem using convex optimization.

sufficient information for convolution-based representation of US texture. We do not apply further processing such as log-compression to avoid nonlinear transformations to the signal. Precisely, we apply Hilbert transform to raw RF data to obtain the envelope signal without the carrier wave.

As a second input, we need the PSF, $H(x,y)$, of ultrasound imaging within this patch. By using US envelope data, the PSF in our method effectively simplifies to:

$$H'(x,y) = e^{\frac{x}{\sigma_x} + \frac{y}{\sigma_y}} . \qquad (3)$$

Given σ_x and σ_y, the PSF can then be approximated. In simulations, such PSF can be known precisely from the FP. For acquired data, the PSF is currently estimated empirically from the size of the speckles, although it could be determined via simulations from accurate models of the transducer [10], e.g. Field II, or experimentally by imaging sub-wavelength synthetic features, e.g. wires, in degassed water.

C. Scatterer reconstruction

1) Setting up the inverse problem: For extracting the scatters from an image, one has to solve the scatterers $T(x,y)$ based on Eq. 1 for given speckle image intensities $I(x,y)=T(x,y) * H'(x,y)$. Convolution can be written as a linear operation $Ax = b$ as shown in Fig. 3. A is then the convolution matrix, each row of which contains all PSF values at columns corresponding to image indices that this PSF kernel would operate on for that row. x is the column vector of all scatterers discretized on a grid, b is the resulting column vector of image intensities. For a scatterer texture resolution $n_s \times m_s$ and a speckle image resolution $n_b \times m_b$, x has then $n_s m_s$ elements, b has $n_b m_b$ elements, and A has $n_s m_s \times n_b m_b$ elements.

Since $H'(x,y)$ is separable, the convolution matrix can be written as a product $A = CD$ of Toeplitz matrices, where C and D represent the convolution in lateral and axial directions, respectively. In our implementation, the PSF convolution kernel is cutoff at 4 standard deviations where the energy becomes negligible. For example, $\sigma_x = 2.5$ and $\sigma_y = 1.5$ pixels (in image scale – not RF) then results in a window size of 20 and 12 pixels for the lateral and axial convolution kernels, and A has accordingly at most 240 nonzero entries per row.

Since the scatterers is an abstract definition merely for the approximation of subwavelength particles, there is indeed no clear definition for their true resolution, i.e. Field II uses a continuous representation, while others use a discretization at an arbitrary resolution [3]. It is thus rather their representation power of the actual image content and, for our case, speckle appearance that matters. The resolution must clearly be high enough to simulate visual complexity of a realistic image, which hence should be at least as fine as the image resolution. Empirical tests with a scatterer resolution equal to the image resolution and up to 16 times the image resolution gave us satisfactory results, also with reasonable computation times of seconds to minutes for scatterer reconstruction.

2) Solving the inverse problem: If and when the number of measurements is equal or less than the number of unknowns, e.g. while using a single image to reconstruct the scatterers at a higher resolution, then the linear system is underdetermined and hence regularization or additional constraints are required to get viable solutions. We use total-variation regularization using the L_1 norm in order to obtain sparse scatterer reconstructions by solving:

$$\hat{x} = \arg\min_{x} \quad ||Ax - b||_2 + \delta||x||_1 \quad \text{s.t. } x \geq 0 , \quad (4)$$

which favors positive sparse x with small scatterer amplitudes. The constraint ensures the scatterer responses to be positive, modeling the actual physics while also increasing solution robustness. δ is empirically fixed to 10^{-5} in all our examples. This optimization problem was then solved using the CVX software package [6]. The result $\hat{x}$ is then a scatterer texture as seen in Fig. 3, which would produce the input image when convolved with the PSF.

If modeling only the given region is the goal, then this result achieves this, and can be used in ray-tracing and with different PSFs to generate different viewing angles and imaging parameters. However, if a generalization is required to inpaint larger models and generate arbitrary geometrical scenes, then the following step of statistically parameterizing this scatterer texture shall be employed.

D. Scatterer parametrization

We assume the scatterer distribution can be parameterized for a homogeneous tissue region with a statistical model [3], the parameters of which can be estimated from the reconstructed scatterer texture by a typical maximum-likelihood estimation step. Indeed, the 3 parameters $\{\mu, \sigma, \text{and } r\}$ can be estimated simply as follows. The number of non-zero scatterer texels using a threshold ϵ (set to 0.2 empirically) gives the ratio r. The mean and the standard deviation of these non-zero scatterers then yield μ and σ, respectively.

When instantiating a new scatterer texture, one should note that negative-amplitude scatterer, which are physically not possible, may be generated by the above normal distribution. Depending on how these scatterers with negative amplitudes are treated, the final scatterer statistics may indeed slightly change. We chose to clamp scatterer amplitudes to zero, hence slightly lowering the actual value of r in the synthesized texture.
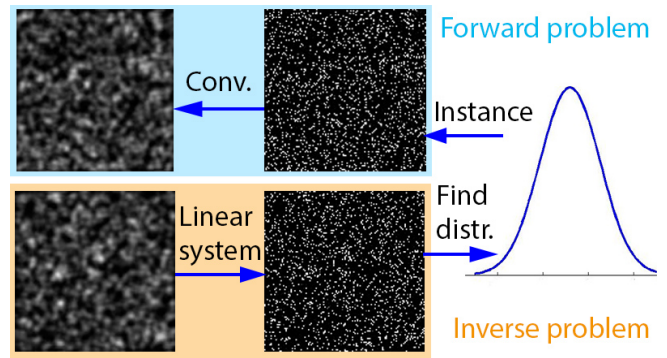

Fig. 4. Our algorithm using a synthesized input and a known PSF. Bottom row: From the input speckle image we acquire a texture of scatterers and their distribution. Top row: Using the forward pipeline, we instantiate new scatterers and create speckle images representing the same tissue.
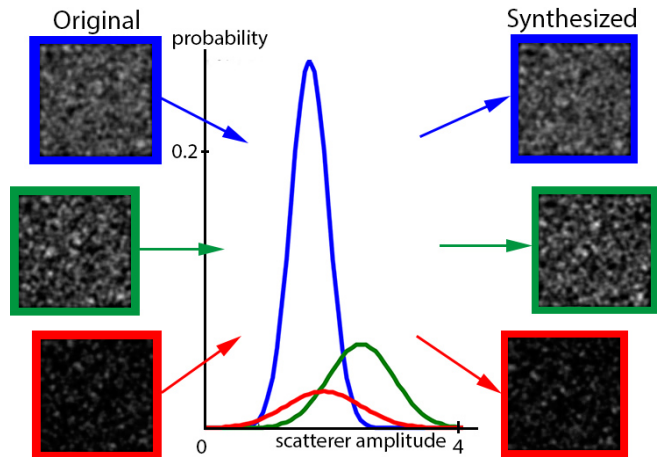

Fig. 5. Our method used on synthetic input images with varying parameter values (μ, σ, and r) for speckle generation, simulating different tissue appearances. The y-axis represents the probability of a texel having the particular scatterer amplitude denoted on the x-axis.

III. RESULTS

Figure 4 illustrates the algorithmic steps on a sample numerical phantom image of 64^2 pixels resolution, while using a 128^2 scatterer texture. The input image was generated using the 3-parameter model and a known PSF of σ_x=2.1 and σ_y=1.6. The scatterer reconstruction in this example takes 47 seconds in Matlab using CVX, and yields simulated speckles of 75.4 dB *peak signal-signal-to-noise ratio* compared to the original known speckles. Visually, the synthesized speckle image in the top row closely mimics the input speckle image. The acquired parametrization was $\mathcal{N}_{\mu=2.08, \sigma=1.01}$ and r=0.11, which closely matches the original parameters $\mathcal{N}_{\mu=2.10, \sigma=1.00}$ and $r = 0.10$.

Figure 5 shows the method applied using the same PSF as above but to tissue samples simulated with different parameters – covering a range of near-underdeveloped to fully-developed speckles. In this figure, a resolution of 128^2 was used for both speckle image and scatterer texture. Differences in the reconstructed parametrizations and the faithfulness in the appearance of simulated to input images can be observed for all combinations of parameters.
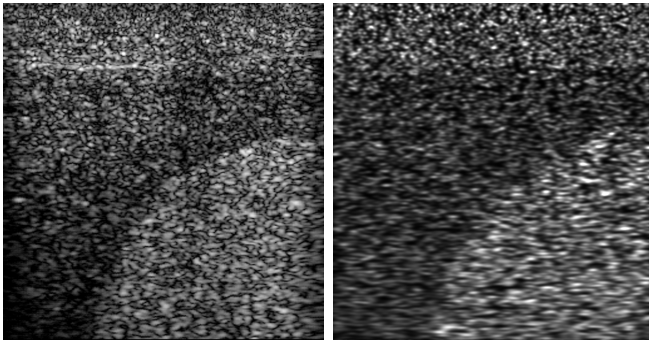
Fig. 6. (Left) Female pelvic phantom and (right) the homogeneous tissues synthesized with our method (using 3 input images). Bottom circular region is the corner of uterus (not imaged fully due to the low penetration of the used high-frequency linear probe).
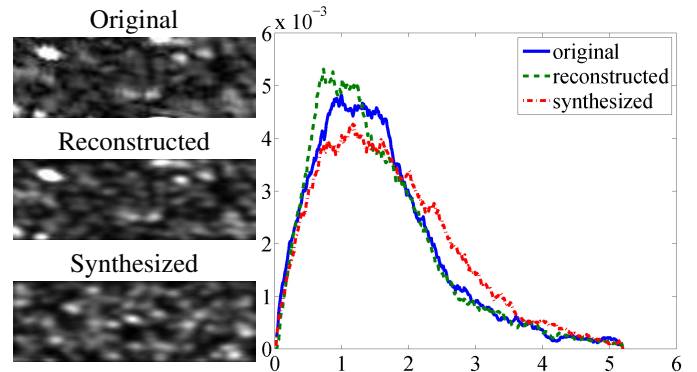

Fig. 7. Top: A sample from an actual tissue ultrasound image. Middle: The speckle image simulated from the scatterers reconstructed by the inverse problem. Bottom: An image simulated from synthesized speckles using the acquired tissue parametrization. The plot on the right compares the histograms of speckle intensity (normalized by the number of scatterers).

Figure 6 shows an image acquired from a pelvic ultrasound phantom using an Ultrasonix SonixTouch machine and a linear array transducer operating at $6.6\,\mathrm{MHz}$. The image on the right shows the homogeneous regions of the phantom simulated with our method, using three 64^2 input images to model tissue parametrizations.

Figure 7 shows an example image from actual human tissue (from the forearm) acquired at $10\,\mathrm{MHz}$. We took a sample window of 64^2 pixels from a sufficiently (but not fully) homogeneous region of an RF image with overall resolution 256×2496 (lateral/axial), which we converted into the signal envelope (as shown in Fig. 2). From the images, we estimated the PSF visually as $\sigma_x=1.25$ and $\sigma_y=2.5$, and used a 10×20 pixels convolution window (for a cutoff at 4σ). The scatterer texture reconstruction resolution was set to 128^2. The images shown are drawn in physical aspect ratios. Despite the uncertainties in the visually-estimated PSF, the input image (top) is reconstructed successfully, c.f. (center), by the inverse-problem solution. The main differences of the instantiated image (bottom) are seen to be due to the non-homogeneous tissue features, which cannot be captured by a statistical model. The histogram plot on the right shows that the statistics of the input, reconstructed, and re-instantiated speckles match closely. The remaining difference might be due to the approximated PSF.

IV. DISCUSSION AND CONCLUSIONS

Despite several ultrasound simulation approaches using scatterer models, to date there had been no methods to extract or parametrize such models for simulating a particular tissue type. We presented a novel approach to reconstruct scatterer representations of tissue for use in ultrasound image simulation, potentially from different viewing angles and with different imaging parameters. We also presented a parameterization of these representations for generating other samples of the same tissue, that can be used to inpaint different geometrical models of the same anatomy, e.g. to fill missing data.

We presented qualitative comparisons of simulated images, with quantitative distributions as an insight into their realism. Quantitative metrics for judging the realism of ultrasound

simulations are not readily available, and user studies with sonographers may aid us evaluate our techniques in the future. Methods for estimating the PSF experimentally or from simulations were not the focus of this work, and will be studied next. A limitation of our method is that the used parametrization model can only describe homogeneous and isotropic tissue which follows a simple normal distribution. In the future, more elaborate statistical models can be employed. Also, by extending our method to 3D, these scatterer models can be incorporated in ray-based simulations with ultrasound interactions such as reflection and refraction.

REFERENCES

[1] D. Aiger and D. Cohen-Or. Real-time ultrasound imaging simulation. *Real-Time Imaging*, 4(4):263–274, 1998.
[2] J C Bamber and R J Dickinson. Ultrasonic b-scanning: a computer simulation. *Physics in Medicine and Biology*, 25(3):463, 1980.
[3] B. Bürger, S. Bettinghausen, Radle M., and J. Hesser. Real-time gpu-based ultrasound simulation using deformable mesh models. *IEEE Trans. Med. Imaging*, 32(3):609–618, 2013.
[4] H. Gao, HF. Choi, P. Claus, S. Boonen, S. Jaecques, GH. Van Lenthe, G. Van der Perre, W. Lauriks, and J. D'hooge. A fast convolution-based methodology to simulate 2-d/3-d cardiac ultrasound images. *IEEE TUFFC*, 56(2):404–409, 2009.
[5] O. Goksel and Salcudean S. E. B-mode ultrasound image simulation in deformable 3-d medium. *IEEE Trans. Med. Imaging*, 28(11):1657–69, 2009.
[6] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, v2.1. http://cvxr.com/cvx, March 2014.
[7] J. Jensen. Simulation of advanced ultrasound systems using field ii. In *Proceedings of ISBI*, pages 636–639. IEEE, 2004.
[8] J. Meunier and M. Bertrand. Ultrasonic texture motion analysis: theory and simulation. *IEEE Trans. Med. Imaging*, 14(2):293–300, 1995.
[9] J. Meunier, M. Bertrand, and G. Mailloux. A model for dynamic texture analysis in two-dimensional echocardiograms of the myocardium. *SPIE*, 0768:193–200, 1987.
[10] J. Ng, R. Prager, N. Kingsbury, G. Treece, and A. Gee. Wavelet restoration of medical pulse-echo ultrasound images in an em framework. *IEEE TUFFC*, 54(3):550–568, 2007.
[11] T. Reichl, J. Passenger, O. Acosta, and O. Salvado. Ultrasound goes gpu: real-time simulation using cuda. *Proceedings of SPIE*, 7261:726116–726116–10, 2009.
[12] W. Wein, S. Brunke, A. Khamene, M.R. Callstrom, and N. Navab. Automatic ct-ultrasound registration for diagnostic imaging and image-guided intervention. *Medical Image Analysis*, 12(5):577–585, 2008.