Deep Network for Scatterer Distribution Estimation for Ultrasound Image Simulation

Journal Article

Author(s):

Zhang, Lin; Vishnevskiy, Valery; Goksel, Orcun

Publication date:

2020-12

Permanent link:

<https://doi.org/10.3929/ethz-b-000448976>

Rights / license:

In Copyright - Non-Commercial Use Permitted

Originally published in:

IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control 67(12), <https://doi.org/10.1109/tuffc.2020.3018424>

Deep Network for Scatterer Distribution Estimation for Ultrasound Image Simulation

Lin Zhang, Valery Vishnevskiy, and Orcun Goksel

Abstract—Simulation-based ultrasound training can be an essential educational tool. Realistic ultrasound image appearance with typical speckle texture can be modeled as convolution of a point spread function with point scatterers representing tissue microstructure. Such scatterer distribution, however, is in general not known and its estimation for a given tissue type is fundamentally an ill-posed inverse problem. In this paper, we demonstrate a convolutional neural network approach for probabilistic scatterer estimation from observed ultrasound data. We herein propose to impose a known statistical distribution on scatterers and learn the mapping between ultrasound image and distribution parameter map by training a convolutional neural network on synthetic images. In comparison with several existing approaches, we demonstrate in numerical simulations and with in-vivo images that the synthesized images from scatterer representations estimated with our approach closely match the observations with varying acquisition parameters such as compression and rotation of the imaged domain.

Index Terms—Deep learning, image synthesis, deconvolution, medical training

I. INTRODUCTION

ULTRASOUND (US) is a low-cost, non-invasive and portable imaging modality, and it has been widely used in the clinical routine, especially for obstetrics and gynaecology as it does not involve ionizing radiation. However, ultrasound scanners produce images suffering from limited spatial resolution, signal-to-noise ratio and tissue contrast, which make the interpretation of images very difficult. Clinical ultrasound practice relies exclusively on sonographer expertise in navigating a hand-held probe and visual inspection of the acquired images; and an extensive training is required for being able to conduct such clinical examinations. The current education is mostly based on training on real patients, under the guidance of expert specialists. This form of education is inefficient due to time consumption of both specialists and patients and involve difficulties in finding volunteers especially when rare pathologies are concerned. Indeed, after completing the one year education students will only have seen up to 80% of possible pathologies [1]. Learning from real ultrasound images offline for training is an alternative option that does not require volunteers. However, this approach does not allow interactive training and is highly restricted to the available clinical images. Computer-assisted ultrasound simulation can aid training, especially important for rare pathologies, which are still vital to identify. New scenarios of any given case can be simulated with transformations on pre-collected data, and

different imaging parameters and conditions can be simulated from the same data. Furthermore, computerized simulations would allow for training in a real-time virtual-reality environment with complex anatomical scenarios and various pathologies, from which the medical students can obtain in-depth knowledge of possible clinical scenarios during training.

Wave propagation, convolution and ray-based methods are the three major techniques for US simulation. Wave-based approaches [2], [3] model ultrasound propagation in tissue by solving complex acoustic wave equations. Therefore, these approaches are relatively slow and not suitable for real-time applications. Convolution-based [4] methods approximate the ultrasound interactions as a spatial impulse response (point-spread function, PSF), which is then convolved with a representation of sub-wavelength tissue structures, known as *scatterers*, under the assumption that the acoustic field is linear. This can realistically reproduce typical ultrasound noisy texture, known as speckle, caused by the constructive and destructive interference of echoes scattered by countless tissue scatterers. On the one hand, speckle can be seen as noise that degrades tissue contrast [5], but on the other hand, it can help distinguish tissues and identify pathologies [6]. Restoration of this granular pattern is not only important for visual realism in a simulation scenario, but also for preserving structural and diagnostic information about the tissue. Ray-based methods [7], [8] simulate the propagation of ultrasonic wavefront as rays using computed graphics techniques, which allows to simulate interactions such as refractions and reflections, while simulating speckle using a PSF convolution with a texture representing scatterers. Using stochastic and sophisticated interaction models with Monte-Carlo ray sampling, this was shown in [8] to lead to impressively realistic US images even at real-time framerates.

There are recent works that use deep learning, and in particular generative adversarial networks, for ultrasound image synthesis [9], [10]. Nevertheless, these works aim to synthesize isolated individual images, without any intermediate (physical space) representation or parametrization, they are not designed for simulating images with speckle motion coherent with underlying physical motion, and therefore may not be applicable in a real-time, interactive simulation framework where temporal continuity of simulated frames is utmost important. In other words, for an infinitesimal movement of the probe to one side should result in the image moving infinitesimally to the other side. If every image is generated somewhat independently, this behaviour and the visual fluidity cannot be guaranteed. Whereas, using a fixed scatterer representation, the above condition can be satisfied. Such tissue representations

L. Zhang, V. Vishnevskiy, and O. Goksel are with the Computer-assisted Applications in Medicine group, ETH Zurich, Switzerland.

Funding provided by the Swiss Innovation Agency Innosuisse #25540PFLS.

can also allow scene editing operations such as copying and adding anatomy in the scatterer domain avoids any image artifacts, as shown in [11].

The problem with using scatterers in simulations is that such tissue representations are not known a priori. Assuming it can be modeled as a PSF convolution, finding a representation from observed US data can then be posed as a blind deconvolution problem. There has been several approaches and approximations to this problem, including inverse problem based solutions [12], variational methods [13], [14] and filtering techniques [15], [16]. Different from our motivation of ultrasound simulation, however, the above methods aim for image denoising / restoration to achieve higher contrast or quality for better diagnostic information. An iterative solution to the regularized deconvolution problem was proposed in [11] by solving the inverse problem jointly for multiple acquisitions of the same tissue, which are obtained efficiently with electronic beam steering. This allows to overconstrain the problem with observations from multiple PSFs. Discrete scatterer reconstructions were performed on a fine Cartesian grid in order to approximate sub-wavelength particles, where sparsity was enforced by using an ℓ_1 -norm regularization. However, solving such inverse problem is computationally and memory-wise only possible for small patches, which were tiled in [11] over the imaging field, and still required hours to days to solve for a single image. Accordingly, such a method is not efficient and scalable for scans of many images or to apply in 3D.

Scatterers can alternatively be estimated as statistical distributions of random variables. A pipeline for generating synthetic echocardiographic ultrasound sequences was presented in [17]. This was later extended to simulate vendor specific US images for speckle tracking algorithms [18]. For this purpose, a scatterer space was populated with randomly sampled 3D cloud points, whose amplitudes were assigned according to the template B-Mode images after the compensation for the log-compression. Simulated speckle statistics were reported to be in good agreement with the known fitting distributions. This approach however does not take the point spread function, and therefore the constructive and destructive interference between scatterers into account, and sampling from B-mode (similarly from RF envelope) assumes scatterers all contributing non-negatively, which is not the case for modulated RF nature of typical US PSF. A simple Gaussian-parametrized model was fit to the inverse-problem reconstructed scatterers in [19]. This method was demonstrated for homogeneous tissues that other instances of the same tissue can be obtained with the found model for simulating new images, which were however reported as lack of visual variety of real tissues in [8].

For a given statistical model for scatterer distribution, we propose herein to estimate distribution parameter maps directly from observed ultrasound images. This can be used to instantiate new scatterer maps that would reproduce the original images when input to a convolution-based simulation. As the estimated parameter maps would represent a physical tissue space, they could be used to simulate new images faithfully with varying imaging conditions, viewing directions and other imaging parameter variations. Due to the power of deep neural

networks in learning patterns of visual inputs, we propose to learn the mapping between simulated images and parameter maps by training a convolutional neural network.

II. BACKGROUND

A. Forward Problem of Ultrasound Simulation

Based on the first order Born approximation (weak scattering) for soft tissues [20], the interaction between ultrasound field and tissue scatterers can be formulated as a 2D convolution model in discrete domain:

$$\mathbf{I}[l, a] = \mathbf{g}[l, a] * \mathbf{h}[l, a] + \gamma[l, a], \quad (1)$$

with radio-frequency (RF) US image intensity $\mathbf{I}[l, a]$, scatterer intensity $\mathbf{g}[l, a]$, spatial variant point spread function (PSF) $\mathbf{h}[l, a]$ and noise term $\gamma[l, a]$. $[l, a]$ are the lateral and axial coordinates with respect to the probe origin, elevational thickness is ignored here.

Ultrasound point spread function can be approximated with a two dimensional Gaussian kernel modulated by a cosine function in the axial direction [7], i.e.:

$$\mathbf{h}[l, a] = e^{-\frac{l^2}{\sigma_l^2} - \frac{a^2}{\sigma_a^2}} \cos(2\pi f_c a), \quad (2)$$

where f_c is the transducer center frequency, σ_l and σ_a determines the Gaussian shape along the lateral and axial direction. Due to nonuniform focusing and aperture, the PSF in US imaging is often assumed to be spatial-variant mainly along the axial direction [21], [11], [6]. Some recent works model PSF as continuously non-stationary blurring, such as based on semigroup theory [22] and with the diffraction effects during wave propagation [23]. For deconvolution tasks, PSF has often been assumed to be patchwise invariant, an approach we also adopt in this paper.

B. Inverse Problem of Scatterer Reconstruction

Eq. (1) can be equivalently written in a matrix-vector form as $\mathbf{Ax} + \mathbf{n} = \mathbf{b}$ with the convolutional matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$ associated with PSF, a vector of scatterer amplitudes $\mathbf{x} \in \mathbb{R}^N$, RF image intensities $\mathbf{b} \in \mathbb{R}^M$ and the acquisition noise term $\mathbf{n} \in \mathbb{R}^M$. Assuming that $N = M$ and imposing no constraints on $\mathbf{x}$, a solution of this system of linear equations is referred to as the tissue reflectivity function (TRF) [16], [13]. To mitigate ill-posed nature of this deconvolution problem, regularization is typically imposed to introduce a prior knowledge about the solution. Wiener filter [24] is a common choice for efficient image deconvolution, which solves the inverse problem based on ℓ_2 -norm regularization of the solution magnitude. Other regularizations have been also widely explored, such as a ℓ_1 or ℓ_p -norm based on the assumption of Laplacian [15] and generalized Gaussian distribution [6] for the TRF $\mathbf{x}$. Several computationally efficient methods have been proposed for TRF deconvolution with sophisticated forward models, such as axially varying kernels [25], physical model accounting for diffraction effects [23]. However, such TRF representation is difficult to attribute to a physical quantity and without any constraints on $\mathbf{x}$, the Wiener filter solution may overfit to the

Fig. 1: (a) Illustration of the proposed pipeline, note that PSF and scatterer maps are sampled for each training pair. (b) ScatParam CNN: we use stride of 2 to downsample axial dimension with image size indicated on the left for each layer, while the number of filters is shown on top of each layer.

observation $\mathbf{b}$, e.g. a slight change in acquisition parameters may yield largely different TRF estimates.

For image simulation purpose, estimated discrete scatterer map $\mathbf{x}$ needs to have sufficiently fine resolution to approximate the underlying continuum of scatterers, hence $N \gg M$. Mauttausch et al. [11] used the ℓ_1 norm, favoring sparse scatterer map. Assuming a Laplacian noise distribution for n , the objective is formulated as follows:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_1 + \lambda \|\mathbf{x}\|_1, \quad \text{s.t. } \mathbf{x} \geq 0, \quad (3)$$

with a regularization parameter λ . This formulation, known as regularized least absolute deviations (RLAD), allows more robust solution with respect to outliers in the error function, which could be caused by wave interactions other than scattering, e.g. directional reflections.

III. METHODS

For US simulation, we aim to find a tissue representation from observed image, which can be used to simulate the same tissue with varying imaging conditions. Rather than estimating a deterministic scatterer locations and amplitudes by solving a large-scale inverse problem, we propose to impose a statistical model on the scatterer distribution and infer corresponding parameters from the observation. We learn the mapping from a single US image to its parameter map in a supervised manner. The required paired data are generated by simulation, since no ground truth of tissue scatterer is available. An overview of our proposed pipeline, referred as ScatParam, can be seen in Fig. 1(a).

A. Statistical Model of Scatterer Distribution

We assume that each tissue type can be parametrized by a model with three parameters $(\rho_s, \mu_s, \sigma_s)$, the parameter $\rho_s \in [0, 1]$ for scatterer density, the mean μ_s and standard deviation σ_s for scatterer amplitude modeled to be normally distributed. Scatterer maps are sampled as follows: for each pixel, a Bernoulli distributed random variable is sampled, where the pixel takes the value one with probability ρ_s and

zero with $1 - \rho_s$. For non-zero pixels, their amplitudes are sampled from a Gaussian distribution $\mathcal{N}(\mu_s, \sigma_s)$. Tissue scattering strength is controlled by the mean, while σ_s models random fluctuations around that mean. Since we only consider fully developed speckles, we set ρ_s to a fixed value satisfying the Rayleigh criterion; herein set to a minimum of 100 scatterers per mm^2 for a fully-developed speckle pattern [26].

The Rayleigh criterion and distribution statistics for estimating interference with isolated scatterers [26] has only been studied for scatterers distributed randomly at continuum spatial locations. Typical convolution based simulation packages, such as Filed II [4] and SIMUS [27], accordingly use scatterer representations with floating-point locations in continuum domain. Reconstructing these on a discrete map therefore necessitate a sufficiently high grid resolution to approximate the continuum. Following [11] we choose to use an isotropic grid spacing (i.e. Cartesian grid) with the native axial resolution of the raw RF data, which inherently also satisfies the Nyquist criterion for lateral sampling. For instance, for a sampling frequency of 40 MHz and speed-of-sound of 1540 m/s, we use a scatterer map with a resolution of roughly $20 \mu\text{m}$ and with 5% of the pixels populated with scatterers, resulting in a scatterer density of 130 per mm^2 .

B. Training Set Generation

Since scatterers are an abstract tissue representation, no point-wise ground truth exists, thus we create network training data by means of simulation. For parameter map generation, we use random synthetic shapes by overlapping irregular geometric shapes, with a procedure similar to [28], where random coarse gray-scale patterns are interpolated at a finer resolution and finally thresholding them to create random shapes, as exemplified in Fig. 2(a). These aim to represent a rich variety of potential tissue structures without assuming particular anatomical priors, both to be invariant to any anatomical assumptions and region-of-interest as well as to allow the network for better generalization. We assume uniform distribution for Gaussian mean: $\mu_s \sim \mathcal{U}(0, 1)$ and a fixed Gaussian standard deviation $\sigma_s = 0.05$. We assign one sampled mean value to each region, assuming the scatterers in

Fig. 2: Generation of training images: (a) scatterer parameter maps, (b) envelope images simulated with scatterers sampled from the corresponding parameter maps, and (c) the corresponding parameter maps estimated by our method.

each tissue region following the same distribution. The scatterer maps are sampled according to the procedure described in Section III(A). RF images are generated by convolving a sampled scatterer map with a PSF. We assume spatially invariant PSFs for image patches, allowing very fast US image generation. PSFs are sampled from the analytic expression in Eq. (2), with the transducer center frequency $f_c = 6$ MHz, the sampling frequency $f_s = 40$ MHz and normalized by its ℓ_2 norm. The vertical and axial spreading are uniformly sampled from $\sigma_t^2 \in [0.2, 1]$ mm² and $\sigma_a^2 \in [0.02, 0.05]$ mm². The images are then corrupted by additive Gaussian noise. The noise level is uniformly sampled in the interval $[2, 20]\%$ of the average signal value. The envelope images are used as the input to the neural network, which are taken as the absolute value of Hilbert transform of RF images, illustrated in Fig. 2(b). For training, 4000 parameter maps of 64×128 pixels with the random geometric shapes above were first generated offline. Then on-the-fly during each training batch, random scatterer maps were spatially sampled from a random subset of the parameter maps, following which PSF convolution and envelope detection (Hilbert transform) were also carried out on-the-fly for each training image.

C. Network Architecture and Training

The general network architecture for ScatParam is illustrated in Fig. 1(b). An encoder-decoder network is used to extract features from the input US image and estimate its corresponding parameter map, used for scatterer map sampling. The model comprises an encoder and a decoder part along the axial direction, with skip connections between the corresponding layers. This design choice is due to very low lateral resolution in US imaging being 10–20 times lower than the axial. Parameter maps are estimated at a coarser axial resolution than input envelope images, assuming spatially smooth tissue content. This facilitates a more efficient utilization of network weights and hence inference power. Our preliminary experiments with equally high resolution in encoder and decoder did not indicate results substantially superior to our presented architecture. We use strided convolution to perform layer pooling and upsampling. Exponential linear unit activation [29] is used at each layer except the output layer, which is linear. The network is trained using Adam optimizer [30] with a learning rate of 10^{-4} , minimizing the ℓ_1 -norm based loss function between the

true $\mathbf{x}$ and estimated parameter maps $\hat{\mathbf{x}}$ given the acquisition $\mathbf{y}$ as the input:

$$L(\Theta) = \mathbb{E} \|\mathbf{x} - \hat{\mathbf{x}}(\mathbf{y}; \Theta)\|_1, \quad (4)$$

with the network parameters Θ , the empirical average $\mathbb{E}$ over the training sampling procedure. The batch size is set to 16. The network is trained for 20000 iterations.

The proposed pipeline shown in Fig. 1 is summarized as follows:

- 1) Data generation: simulation of synthetic ultrasound images involves scatterer map sampling and convolution with point spread functions;
- 2) Offline training: a convolutional neural network is trained with the simulated paired data for parameter map estimation;
- 3) Scatterer distribution estimation and sampling: for each observation, a scatterer map is sampled from the estimated distribution parameter map for synthesizing new images;
- 4) Convolution-based simulation: the sampled scatterer map is fed into the convolution based simulator, which generates images with desired imaging parameters in real time.

IV. EXPERIMENTS AND RESULTS

We study our proposed method comparatively to its alternatives with experiments conducted on numerical simulations of synthetic phantoms, as well as on actual data acquired from a gelatin phantom and in vivo tissue. Given the background above, we consider three alternatives to compare our method against:

- Sampling scatterers from envelope image (**SampleEnv**): This is an adaptation of the method proposed in [17] for sampling continuous scatterers from log-uncompressed B-mode images. We herein adapt this to sample on discrete scatterer maps and, for better accuracy, we use the original envelope images instead, as we have access to them.
- Tissue reflectivity function (**TRF**): This is the simple Wiener filter estimation to the deconvolution inverse problem [31], [16]. We use a spatially constant filter kernel, as the PSF computed or estimated at the center of the imaged field of view.
- Iterative scatterer reconstruction (**ScatRec**) [11]: This is an inverse problem based approach is referred here as ScatRec [11], which reconstructs scatterer map based on a single observation.
- Deep Learning based estimation of parametric scatterer maps (**ScatParam**): This is our proposed method trained on simulated images, as detailed in the previous section.

For deconvolution based methods, TRF and ScatRec, in simulated data we used the known PSF from the simulations and for acquired data, we used a cepstrum-domain PSF estimation method described in [32], followed by least square fitting to the known parametric form in Eq. (2) to project them on our PSF model manifold, which our trained network is better conditioned on.

Fig. 3: 1D signal example with simulated RF and envelope signals by TRF (left) and envelope signals by SampleEnv and ScatParam (right). Black star depicts the scatterer locations.

Since different scatterer representations cannot be compared directly and no ground truth is available, we evaluate the performance of scatterer estimation on the envelope images simulated from the estimated scatterer maps. For resynthesizing images, we used the same forward simulation for all the methods, namely a discrete image-space convolution of the scatterers estimated by any particular method with the same depth-dependent PSFs estimated for deconvolution based methods. The convolution was implemented in Matlab to operate separately for each PSF on rows of image pixels.

A. 1D Example

In contrast to US deconvolution (TRF and the regularized version ScatRec) that aims to reproduce the RF and hence speckle image precisely on a pixel level, the goal of SampleEnv and ScatParam are rather to generate realistic speckle textures that match the speckle distribution and tissue contrast of a real image. We herein illustrate their differences using a simplified 1D example depicted in Fig. 3. We first simulated a ground truth 1D ultrasound RF signal (RF GT) of 40 MHz sampling, by convolving a Gaussian-modulated sinusoidal PSF for a 5 MHz pulse with randomly placed scatterer locations within $[20, 180]\mu\text{s}$ (Scats GT). Scatterer amplitudes were randomly sampled from a Gaussian distribution with a mean of 1 and standard deviation of 0.05. For comparison, for each estimated representation we resynthesized the RF signal by convolving with the known PSF and then plot the envelope of these in comparison to the ground truth envelope (Env GT).

For TRF, we deconvolved RF GT using Wiener filtering with the known PSF, which yields almost a perfect match to the original data for this ideal, no-noise scenario, cf. RF TRF and Env TRF in Fig. 3(left). For SampleEnv and ScatParam, we resampled random scatterer locations shown in Fig. 3(right), the amplitudes of which were set for SampleEnv using the envelope intensities and for ScatParam from the original scatterer distribution, with the assumption that such distribution can be correctly estimated by our method. ScatParam and SampleEnv both do not reproduce the RF data exactly, but generate an envelope intensity distribution similar to the original envelope.

B. Evaluation Metrics

Several evaluation metrics are utilized to assess the simulation performance and compare our method ScatParam

with SampleEnv, TRF and ScatRec. Three image-based metrics, mean image intensity (I), signal-to-noise ratio (SNR), contrast-to-noise ratio (CNR) and one histogram based metric, Kullback-Leibler (KL) divergence, are used to calculate the mismatch between ground truth and simulated envelope images.

Mean Image Intensity (I) can capture any global intensity shift in the simulated images, which can be caused by systematic biases in constructive or destructive interference when the scatterers are not distributed truly stochastically or in a view-dependent way. For instance, if the scatterers estimated from one direction align in a structured way, when the object is imaged from an oblique direction (e.g., with the half-wavelength projected on the rotation angle aliased with the structure), these scatterers may then interfere with each other mostly destructively, creating an artificial intensity drop in the image. The change in mean image intensity is calculated as follows

$$\Delta I = \frac{|I_t - I_s|}{I_t}, \quad \text{with } I = \frac{1}{N} \sum_{j=1 \dots N} s_j, \quad (5)$$

where s_j denotes the image intensity value at the j -th pixel, and N is the number of pixels. I_t and I_s are the mean image intensities of ground truth and simulated image, respectively. Hereafter, the subscript t refers to the ground truth and s to a simulated image.

Signal-to-noise Ratio (SNR) measures the global statistics of signal-and-noise ratio. For the simulation purpose, we aim at reproducing the images at the same SNR level as the ground truth. Any mismatch in SNR is then quantified as follows

$$\Delta \text{SNR} = \frac{|\text{SNR}_t - \text{SNR}_s|}{\text{SNR}_t}, \quad \text{with } \text{SNR} = \frac{\mu}{\sigma}, \quad (6)$$

where μ and σ denote the mean and standard deviation of envelope image intensities.

Contrast-to-noise ratio (CNR) mismatch is defined similarly as

$$\Delta \text{CNR} = \frac{|\text{CNR}_t - \text{CNR}_s|}{\text{CNR}_t}, \quad \text{with } \text{CNR} = \frac{|\mu_{s1} - \mu_{s2}|}{\sigma_{s1} + \sigma_{s2}}, \quad (7)$$

where μ_{s1} , μ_{s2} , σ_{s1} , and σ_{s2} denote the means and standard deviations of the image intensities within two contrasting regions. This metric is clinically relevant, as incorrect tissue contrast in the simulated images may lead to the learning of false diagnostic cues during medical training.

Kullback-Leibler (KL) divergence compares the statistics between the histograms of two images as follows:

$$\text{KL}(h_s || h_t) = \sum_{l=1 \dots D} h_s[l] \log \left(\frac{h_s[l]}{h_t[l]} \right), \quad (8)$$

where h_t and h_s are the normalized histograms of the true and simulated images respectively. The number of histogram bins D is set herein to 50. Histogram statistics are widely explored for tissue characterization [33], [34]. Hence, a large discrepancy in the histograms could indicate a mismatch in the speckle pattern appearance. Since computing histograms over the whole image could miss local speckle texture information, we calculate a KL divergence metric locally within patches

Fig. 4: Simulated images of the numerical phantom for views rotated by 15°, 30°, and 45°. Envelope images are shown.

(herein non-overlapping patches of $3 \times 3 \text{ mm}^2$ corresponding to 10λ per dimension) and report herein the metric mean over all patches.

For calculating ΔSNR , ΔCNR and KL divergence, simulated image (s) is normalized (or brightness equalized) with respect to ground truth image (t) by multiplying a factor $\frac{\sum_j t_j}{\sum_j s_j}$ similarly to [11], to eliminate effects in these metrics from any global intensity shift, which is captured separately by ΔI .

C. Synthetic Data

This experiment evaluates the invariance of the reconstructed scatterer maps to various US imaging parameters. To that end, we used simulated images from Field II with controllable imaging conditions. A numerical phantom of $15 \times 15 \text{ mm}^2$ with a 3 mm circular inclusion was simulated for imaging at 6.0 MHz center frequency, with a 128 element 40 mm linear transducer, a single transmit focus at the center, and dynamic receive focusing. The phantom was placed 15 mm away from the transducer to avoid near field effects.

We evaluate the simulation results for 1) phantom rotation, which emulates imaging a region of interest from different viewing directions; 2) phantom compression, which emulates image plausibility under potential tissue deformation, e.g. induced by probe compression.

1) *Rotation Experiment*: In this experiment, we evaluate the invariance of the reconstructed scatterer maps to phantom rotation. For this, the box phantom were rotated around the phantom center with varying angles. Fig. 4 illustrates the results for the rotated views. The images were cropped to the region of interest. The top row shows the ground truth envelope images simulated by Field II for varying rotation angles ranging from 0° to 45° with a 15° increment. For ScatRec, the image intensity drastically decreases with increasing rotation angle and the phantom almost disappears for 45° due to destructive interference, as reported in [11]. The images

Fig. 5: Performance with respect to rotation angle in terms of normalize mean image intensity difference (ΔI (%)), signal-to-noise ratio difference (ΔSNR (%)), contrast-to-noise ratio difference (ΔCNR (%)) and histogram difference (KL divergence).

simulated by SampleEnv appear dark, whereas the image mean intensity remains similar to the ground truth images for TRF and ScatParam during rotation. It can be well observed that the reconstruction of the proposed ScatParam is robust for different viewing angles.

We evaluate the simulation performance quantitatively for the above experiment by investigating the error for rotations with 1° increments, with the results plotted in Fig. 5. ScatRec is close to the ground truth near 0° given any metric, but it deviates largely from the ground truth at larger rotation angles, which corroborates the visual observations above. The errors for SampleEnv are consistently high irrespective of rotation, indicating not a successful scatterer reconstruction. In general, the error metrics of TRF vary without any pattern, remaining relatively similar across the angles, except a large variance in CNR. This is in agreement with the observation in Fig. 4 that the tissue contrast in the 15° rotated view of TRF is diminished. The metrics for ScatParam exhibit minor fluctuations and remain overall relatively low and thus superior compared to the other methods.

Tab I. summarizes the mean, median and maximum error across the 1° increment results in Fig. 5. ScatParam is seen to achieve the lowest error overall, nearly 30% lower than the second best method in ΔCNR and KL divergence, demonstrating the representativeness and robustness of the estimated scatterer map with respect to probe rotation.

2) *Compression Experiment*: With this, we investigate the invariance of reconstructed scatterer maps with respect to physical deformation. An axial phantom compression was simulated by interpolating the estimated scatterer maps on grids deformed by varying levels of axial strain e . Simulated images with all methods at $e = \{10, 30, 50\}\%$ compression are shown in Fig. 6. The images simulated by TRF are corrupted

TABLE I: Mean (mean), median (med), and maximum (max) errors for rotations between 0° and 45° with 1° increments. Bold number indicates the smallest value per column.

Metric	ΔI (%)			ΔSNR (%)			ΔCNR (%)			$KL (\times 10^{-2})$		
	mean	med	max	mean	med	max	mean	med	max	mean	med	max
SampleEnv	50.1	50.2	50.8	11.9	11.7	17.7	16.3	16.6	22.4	20.4	20.7	25.5
TRF	9.1	8.6	18.1	4.0	3.7	8.4	9.4	8.0	31.0	16.7	16.5	24.1
ScatRec	59.6	68.0	86.9	3.4	2.8	13.1	27.2	22.7	57.5	22.2	18.2	44.2
ScatParam	5.7	7.4	11.0	2.3	1.8	6.9	6.6	5.2	19.5	12.5	12.3	19.3

Fig. 6: Simulated images of the numerical phantom for axial compressions of 10%, 30%, and 50% strain. Envelope images are shown.

by aliasing artifacts. The simulation results of SampleEnv and ScatRec look similar with reduced intensity and degraded speckle pattern for large compression strain. The simulated images by ScatParam are visually closest to the ground truth images in terms of speckle appearance and contrast, but slightly hyperechoic for large compression.

We investigate the simulation errors for compression ranging from $e = 10\%$ to $e = 50\%$ with 1% increments in Fig. 7. Similarly to the rotation experiment, the error metrics of ScatRec increase proportionally with increasing compression, whereas our method performs consistently superior. SampleEnv and TRF yield large errors, especially in ΔCNR and KL divergence metrics.

The observations are supported by numerical results shown in Tab. II, reporting the mean, minimum, and maximum errors for the above plots. Our proposed method ScatParam is seen to be superior in terms of ΔSNR , ΔCNR , and KL divergence. It achieves approximately 78% lower error in the mean ΔSNR , 68% in the mean ΔCNR and 28% in the mean KL divergence compared to the second best method. TRF achieves 6% lower error than ScatParam in the mean ΔI here.

Fig. 7: Performance evolution with increasing axial strain e in terms of difference in image mean intensity (ΔI (%)), signal-to-noise ratio (ΔSNR (%)), contrast-to-noise ratio (ΔCNR (%)) and histogram difference (KL).

TABLE II: Mean (mean), median (med) and maximum (max) errors across 10% to 50% strain with 1% increments. Bold number indicates the smallest value per column.

Metric	ΔI (%)			ΔSNR (%)			ΔCNR (%)			$KL (\times 10^{-2})$		
	mean	med	max	mean	med	max	mean	med	max	mean	med	max
SampleEnv	38.5	39.4	46.3	12.7	13.3	16.2	17.6	17.2	24.5	16.6	16.5	21.4
TRF	24.4	24.1	36.1	8.3	9.3	13.3	21.5	22.2	26.9	20.3	20.4	26.6
ScatRec	25.6	27.1	38.3	12.9	13.1	22.1	14.6	14.6	24.9	18.9	19.9	32.3
ScatParam	26.0	25.1	41.7	1.8	1.6	8.0	4.7	4.1	13.9	11.9	11.7	16.6

D. Gelatin Phantom

Next we investigate the performance of our method for a real ultrasound scan of a gelatin phantom with corn starch as the scattering medium. A circular inclusion was made by adding twice as high as the concentration of starch in the background. The beamformed RF images were collected by a Fukuda Denshi UF-760AG ultrasound machine with a linear probe FUT-LA385-12P. The results in the top row of Fig. 8(a) show an excellent agreement of ScatRec with the ground truth image, as expected from an overconstrained optimization, while the SampleEnv and TRF results exhibiting speckle textures different than the ground truth. Our proposed method ScatParam generates images almost indistinguishable from the observation in terms of speckle texture and tissue contrast. The experiment is further evaluated with quantitative metrics in Tab. III. ScatRec achieves the closest match to the ground

TABLE III: Quantitative evaluation metrics for the gelatin phantom. Bold number indicates the smallest error per column.

Method	ΔI (%)	ΔSNR (%)	ΔCNR (%)	$KL (\times 10^{-2})$
SampleEnv	44.0	19.2	13.3	25.5
TRF	5.0	10.0	24.0	21.2
ScatRec	1.7	0.6	1.3	2.7
ScatParam	7.4	0.9	8.1	17.0

Fig. 8: (a) Envelope images with scatterers estimated from the gelatin phantom, images simulated at 0° (top) and 45° (top). (b) Histogram of the ground truth and rotated views inside the area shown by a red rectangle in one sample image. Quantifying the difference between the ground truth and simulated image histograms using KL divergence indicates the following errors: 0.101 for SampleEnv, 0.593 for TRF, 0.588 for ScatRec, and 0.032 for ScatParam.

truth image, as its scatterer estimation overfits to the observed phantom image. Comparing the remaining three, ScatParam yields CNR, SNR, and KL divergence metrics lower than TRF and SampleEnv, indicating that our method better preserves contrast and intensity distribution.

The bottom row in Fig. 8(a) shows the simulated 45° rotated views of the phantom. For an isotropic scattering phantom, the mean echo and speckle texture statistics would be invariant to the viewing direction and hence, after rotation, a speckle appearance similar to the initial view is expected. Such coherence pattern may naturally not stay pixel-wise constant after, e.g., rotation, therefore no pixel-wise error metric was employed. Similarly to the rotation experiment of Field II, the image simulated by ScatRec appears hypoechoic for 45° angle, while the speckle pattern of TRF is severely distorted.

The ground truth rotated views are not obtained here. Nevertheless, we evaluate the results by comparing the histogram of the observed 0° envelope image to the histograms of the rotated envelope images (after brightness equalization) for the homogeneous region inside the red rectangle depicted in Fig. 8(a). Here Rayleigh statistics serve as an important criterion for assessing ultrasound speckle texture, as the envelope intensity should follow a Rayleigh distribution for fully developed speckles [35]. The histogram of SampleEnv closely follows the observed ground truth histogram, i.e. an ideal Rayleigh distribution, consistent with the observation reported in [17]. Nevertheless, our approach ScatParam is seen to have an even better agreement with the observed pattern; indeed, with an over 3-folds lower KL divergence score as listed in the figure caption.

E. In Vivo Experiment

For this purpose, beamformed RF data from an in-vivo scan of the liver was collected. Scatterer representations were

estimated for the imaged region using all four methods presented. Fig. 9 depicts the results simulated in the acquisition configuration of 0° and simulating a rotation of 45° given a fixed dynamic range computed at min and max of the original 0° view, considering the fixed imaging settings during real-life probe and scene manipulations. To highlight details imperceivable due to any intensity reduction, the last row in Fig. 9 shows the same simulated 45° views, dynamic range adjusted individually. At 0° , all methods perform somewhat similarly, with SampleEnv slightly hypoechoic, TRF misrepresenting speckle texture and ScatRec reconstructing an exact replica, as expected. Noticeably, given the dynamic compression for B-mode images presented for the in-vivo experiment, the loss of brightness in the image generated by SampleEnv is less prominent compared to the previous experiments showing envelope images. At 45° , however, one can observe a global intensity drop in SampleEnv and ScatRec and intensity increase in TRF with the distorted speckle patterns – indicating that these three all did not estimate robust scatterer representations. After compensating for the global intensity shift in the last figure row, the speckle textures for TRF and ScatRec still look unrealistic, whereas SampleEnv do not faithfully preserve structural details, e.g. at the bottom image region, with increased echogenicity in the upper region. In contrast, our method ScatParam emulates the image rotation of scattering regions realistically and preserves mean intensity as well as structural detail and contrast, after rotation.

Assuming a homogeneous isotropic structure for the liver tissue (marked with a red rectangle in Fig. 9), a histogram comparison is also performed to compare post-rotation speckle appearance to the original image. Given the clear difference in mean intensities, a raw B-mode histogram comparison indicates the superiority of ScatParam, by a large margin. For a comparison of the speckle texture alone, we therefore first brightness equalized the images (cf. Fig. 10(a)) and then

Fig. 9: Images simulated in the original 0° configuration from scatterers estimated from an in vivo liver image (1. row), estimated scatterer maps, with only a random 8% shown for sake of visualization (2. row), the images simulated from the same scatterers after a rotation of 45° (3. row) and the same rotated views, dynamic range adjusted individually (4. row). B-mode images are shown here, in contrast to all other images in this paper showing envelope images.

performed the histogram comparison. As seen in Fig. 10(b), ScatParam histogram even after after rotation closely matches to that of original ground truth image histogram, following an ideal Rayleigh distribution. KL divergence results reported in Fig. 10 caption corroborate with this observation.

V. DISCUSSION

In this work, we demonstrated a learning based approach for probabilistic scatterer estimation in the context of realistic ultrasound image simulation. The proposed framework ScatParam involves sparse scatterer model with Gaussian distribution for scatterer amplitude and estimation of the Gaussian mean directly from US envelope images by neural network. Similar isotropic scatterer distributions were used in several

earlier works, e.g. [4], [6], [7], [11], [17], [36]. In our preliminary experiments for estimating parameter maps of both mean and standard deviation, we have found that different such combinations may generate similar image outputs, making such double parameter estimation ill-posed. Therefore, in this work we fixed the standard deviation and estimated only the mean value. In addition, we assumed a fixed density for our model, since it is shown that in case of fully developed speckles the tissue characteristics are not affected by density [37]. Estimating scatterer density in addition to the amplitude mean could allow us to distinguish between partially and fully developed speckles. Nevertheless, we observed comparable performance with only mean estimation for both numerical simulations and in-vivo experiments presented herein.

For complex in-vivo scatterer distributions more sophisti-

Fig. 10: (a) Image patches in the red marked region in Fig. 9(a) for the 0° view (top) and the 45° view (bottom). Envelope images after brightness equalization are shown here. (b) Histogram of the ground truth patch compared with the simulated rotated patches, showing a best match with ScatParam. Quantifying the difference between the ground truth and simulated image histograms using KL divergence indicates to the following errors: 0.040 for SampleEnv, 0.059 for TRF, 0.018 for ScatRec, and 0.004 for ScatParam.

cated models may be required. For example, for muscle fibers, anisotropic parametrizations such as with tensor, wavelet, or frequency-domain representations may be more suitable. Indeed, instead of hand-crafted parametric models, (arbitrary) distributions could potentially be parametrized using a neural network to be inferred from observed images.

In the paper, we have conducted several experiments evaluating the invariance of tissue properties with geometric transformations, which is of great importance for ensuring plausible simulation of the same tissue content under different imaging conditions. Rotation and axial compression are chosen, since they well represent the clinical examinations such as transducer in-plane tilting and compression. Furthermore, these experimental scenarios model potential variation in speckle appearance from directional changes (isotropy) and concentration changes (axial strain) in scatterer configurations.

In comparison to other deconvolution algorithms, our method does not require PSF estimation in advance. It was reported in [11] that accurate PSF estimation is needed as input to an inverse problem based method, analogously important for the Wiener filter and other deconvolution algorithms. Our trained CNN is able to capture the PSF information in the form of different speckle textures in input images, hence estimated parameter maps would be independent of imaging system. One can produce visually plausible images without considering PSF as in the method SampleEnv. However, the

simulated speckle statistics with SampleEnv herein are not fully in agreement with the observations, since by dismissing the modulated nature of PSF and thereby any destructive interference, SampleEnv cannot fully model the interference between scatterers. Setting scatterer amplitudes directly using envelope intensities causes the incorrect translation of speckle variations into scatterer maps, i.e. higher amplitude scatterers lumped around the peaks of speckles. This leads to brighter hyperechoic and darker hypoechoic regions after PSF convolution. Such granular appearance is less visible after dynamic range compression in B-mode images, e.g. in Fig. 9, whereas these and resulting overall intensity reduction become apparent in the envelope images shown for the rest of the experimental results. Note that we present envelope images, as they allow easier interpretation of speckle patterns and imaging physics, where effects of methodological choices on the results are not masked by any graymap transform.

We herein chose Wiener filtering as a basic deconvolution baseline for TRF, with low computation and memory requirements. The results with artifacts and reduced contrast of TRF indicate that the use of the same low resolution of the input RF image in the output scatterers hinder interpolation in the scatterer domain after transformations, such as compressions. Although there are more recent forms of TRF, e.g. [25], [23], ScatRec [11] was chosen herein as a state-of-the-art baseline performing a sophisticated deconvolution approach as an optimization of an inverse-problem definition; with a suitable model, i.e. norms and regularizers; with depth-dependent PSF; non-negative scatterer constraint; and with a higher scatterer resolution than the RF image domain. Regardless of baselines, our results indicate that our learning based solution ScatParam as a fast-implementable network solution performs satisfactorily for image simulation and without requiring complex convolution modeling, PSF estimation, iterative optimization, or any other complex processing steps. Note that using only Gaussian noise model for the training samples, ScatParam is able to successfully estimate scatterer parametrizations for no-noise numerical experiments as well as unknown-noise phantom and in-vivo examples, potentially indicating the robustness of the proposed parametrization and the respectively trained NN to an assumed noise model.

Underconstrained inverse problem based approaches can be improved by using multiple measurements of the same tissue with different imaging parameters. A successful example was illustrated in [11], where ScatRec with multiple observations from beam steering was shown to be more robust to viewing angle changes. However, the aligned beam-steered images are in general not available from clinical scanners, let alone raw RF data. Our method takes envelope images as input and can thus accept clinical B-Mode images with slight modification of network training. For a fair comparison, we herein used the single view version of inverse problem, i.e. ScatRec1 in [11].

In this work any ultrasound image appearance is attributed solely to isotropic scattering, where attenuation variations and coherent reflections are not considered. For instance, all the methods compared in Fig. 8(a) attribute the slight attenuation in the original image behind the gelatin inclusion (after a depth of 40 mm) to some form of lower scattering

amplitude, thereby resulting in the attenuation not correctly reproduced in the rotated images, e.g. extending diagonally rather than vertically. This demonstrates the need to take directional attenuation and reflections into account during any scatterer estimation process, which should be a focus of future studies. Furthermore, any potential reflections at anatomical boundaries, i.e. between supra-wavelength structures, would also be attributed to result from the reconstructed scatterers, which may then incorrectly reproduce the tissue from different viewing angles. In all our numerical examples, any such reflections were thus avoided. The gelatin phantom was made with isotropic scatterers and for the in-vivo example the liver was chosen for its relatively homogeneous speckle appearance. Additionally we conducted evaluations within small regions-of-interest selected in line with our assumptions. Accordingly, with the experiments we evaluated the methods in terms of speckle appearance, as opposed to realism of simulated attenuation or reflection effects.

Since the network is trained with isotropically distributed scatterers, any estimation for non-isotropic regions may be sub-optimal or likely be applicable only when viewed from a similar direction. Nevertheless, it may be possible to separate the direction-dependent image content prior to capturing only scattering effects. One can then simulate such directional wave interactions at a later time using ray tracing techniques [7], [8]. For instance, reflection boundaries could be detected and removed using a simple phase symmetry (PS) algorithm [38], which is designed to estimate directional reflections at tissue boundaries such as bone surfaces. In a way similar to estimating and compensating for reflections, acoustic attenuation can indeed also be reconstructed a-priori [39] in order to spatially normalize incident acoustic energy to decouple its effect from our reconstructed scatterer amplitudes.

Our network training takes approximately 6 hours on Nvidia Titan XP GPU. Once trained, our network can estimate scatterer maps in milliseconds at inference time, for arbitrary input image size (as being a fully convolutional network architecture). In contrast, the inverse problem based approach takes approximately 2 hours for one image from a single view. For multiple (beam-steered) observations, the computation time would increase exponentially, quickly making this method infeasible for large images and 3D volumes.

VI. CONCLUSION

We have demonstrated a learning-based technique to efficiently estimate the distribution of tissue scatterer representation, which can then be fed directly into convolution- or ray-tracing-based simulation techniques [7], [8] to simulate realistic images for sonographer training. The proposed network is trained only with synthetic images generated with random shapes and spatial invariant convolution. In comparison to the state-of-the-art methods, we demonstrate with numerical simulations the proposed estimation pipeline being robust for simulating images at different viewing angles and tissue deformations. The method is further evaluated on a tissue-mimicking gelatin phantom and an in-vivo liver image, demonstrating the generalization ability of our network to real ultrasound scans.

REFERENCES

- [1] S. Köhn, R. H. Van Lengen, G. Reis, M. Bertram, and H. Hagen, "Ves: virtual echocardiography system," in *IASTED International Conference on Visualization, Imaging, and Image Processing*, 2004.
- [2] M. Verweij, B. Treeby, K. Dongen, van, and L. Demi, *Simulation of ultrasound fields*. Elsevier, 2014, pp. 465–499.
- [3] B. E. Treeby, J. Jaros, A. P. Rendell, and B. Cox, "Modeling non-linear ultrasound propagation in heterogeneous media with power law absorption using a k-space pseudospectral method," *The Journal of the Acoustical Society of America*, vol. 131, no. 6, pp. 4324–4336, 2012.
- [4] J. A. Jensen, "Field: A program for simulating ultrasound systems," in *Nordic-Baltic Conference on Biomedical Imaging*. Citeseer, 1996.
- [5] P. C. Tay, C. D. Garson, S. T. Acton, and J. A. Hossack, "Ultrasound despeckling for contrast enhancement," *IEEE Transactions on Image Processing*, vol. 19, no. 7, pp. 1847–1860, 2010.
- [6] M. Alessandrini, S. Maggio, J. Porée, L. De Marchi, N. Speciale, E. Franceschini, O. Bernard, and O. Basset, "A restoration framework for ultrasonic tissue characterization," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 58, no. 11, pp. 2344–2360, 2011.
- [7] B. Burger, S. Bettinghausen, M. Radle, and J. Hesser, "Real-time GPU-based ultrasound simulation using deformable mesh models," *IEEE Transactions on Medical Imaging*, vol. 32, no. 3, pp. 609–618, 2012.
- [8] O. Mattausch, M. Makhinya, and O. Goksel, "Realistic ultrasound simulation of complex surface models using interactive monte-carlo path tracing," in *Computer Graphics Forum*, vol. 37, no. 1. Wiley, 2018, pp. 202–213.
- [9] Y. Hu, E. Gibson, L.-L. Lee, W. Xie, D. C. Barratt, T. Vercauteren, and J. A. Noble, "Freehand ultrasound image simulation with spatially-conditioned generative adversarial networks," in *Molecular imaging, reconstruction and analysis of moving body organs, and stroke imaging and treatment*. Springer, 2017, pp. 105–115.
- [10] F. Tom and D. Sheet, "Simulating patho-realistic ultrasound images using deep generative networks with adversarial learning," in *IEEE International Symposium on Biomedical Imaging*. IEEE, 2018, pp. 1174–1177.
- [11] O. Mattausch and O. Goksel, "Image-based reconstruction of tissue scatterers using beam steering for ultrasound simulation," *IEEE Transactions on Medical Imaging*, vol. 37, no. 3, pp. 767–780, 2017.
- [12] Z. Chen, A. Basarab, and D. Kouamé, "Compressive deconvolution in medical ultrasound imaging," *IEEE Transactions on Medical Imaging*, vol. 35, no. 3, pp. 728–737, 2015.
- [13] N. Zhao, A. Basarab, D. Kouamé, and J.-Y. Tournet, "Joint segmentation and deconvolution of ultrasound images using a hierarchical bayesian model based on generalized gaussian priors," *IEEE Transactions on Image Processing*, vol. 25, no. 8, pp. 3736–3750, 2016.
- [14] M. Alessandrini, A. Palladini, L. De Marchi, and N. Speciale, "Expectation maximization for joint deconvolution and statistics estimation," in *Acoustical Imaging*. Springer, 2011, pp. 335–343.
- [15] O. Michailovich and A. Tannenbaum, "Blind deconvolution of medical ultrasound images: A parametric inverse filtering approach," *IEEE Transactions on Image Processing*, vol. 16, no. 12, pp. 3005–3019, 2007.
- [16] T. Taxt, "Restoration of medical ultrasound images using two-dimensional homomorphic deconvolution," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 42, no. 4, pp. 543–554, 1995.
- [17] M. Alessandrini, M. De Craene, O. Bernard, S. Giffard-Roisin, P. Allain, I. Waechter-Stehle, J. Weese, E. Saloux, H. Delingette, M. Sermesant *et al.*, "A pipeline for the generation of realistic 3D synthetic echocardiographic sequences: Methodology and open-access database," *IEEE Transactions on Medical Imaging*, vol. 34, no. 7, pp. 1436–1451, 2015.
- [18] M. Alessandrini, B. Chakraborty, B. Heyde, O. Bernard, M. De Craene, M. Sermesant, and J. D'Hooge, "Realistic vendor-specific synthetic ultrasound data for quality assurance of 2-d speckle tracking echocardiography: Simulation pipeline and open access database," *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 65, no. 3, pp. 411–422, 2017.
- [19] O. Mattausch and O. Goksel, "Scatterer reconstruction and parametrization of homogeneous tissue for ultrasound image simulation," in *International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 2015, pp. 6350–6353.
- [20] J. A. Jensen, J. Mathorne, T. Gravesen, and B. Stage, "Deconvolution of in-vivo ultrasound B-mode images," *Ultrasonic Imaging*, vol. 15, no. 2, pp. 122–133, 1993.

- [21] J. G. Nagy and D. P. O’Leary, “Restoring images degraded by spatially variant blur,” *SIAM Journal on Scientific Computing*, vol. 19, no. 4, pp. 1063–1082, 1998.
- [22] O. V. Michailovich, “Non-stationary blind deconvolution of medical ultrasound scans,” in *Medical Imaging 2017: Ultrasonic Imaging and Tomography*, vol. 10139. International Society for Optics and Photonics, 2017, p. 101391C.
- [23] A. Besson, L. Roquette, D. Perdios, M. Simeoni, M. Arditì, P. Hurley, Y. Wiaux, and J.-P. Thiran, “A physical model of nonstationary blur in ultrasound imaging,” *IEEE Transactions on Computational Imaging*, vol. 5, no. 3, pp. 381–394, 2019.
- [24] E. E. Hundt and E. A. Trautenberg, “Digital processing of ultrasonic data by deconvolution,” *IEEE Transactions on Sonics and Ultrasonics*, vol. 27, no. 5, pp. 249–252, 1980.
- [25] M. I. Florea, A. Basarab, D. Kouamé, and S. A. Vorobyov, “An axially variant kernel imaging model applied to ultrasound image reconstruction,” *IEEE Signal Processing Letters*, vol. 25, no. 7, pp. 961–965, 2018.
- [26] B. Oosterveld, J. Thijssen, and W. Verhoef, “Texture of B-mode echograms: 3D simulations and experiments of the effects of diffraction and scatterer density,” *Ultrasonic Imaging*, vol. 7, no. 2, pp. 142–160, 1985.
- [27] S. Shahriari and D. Garcia, “Meshfree simulations of ultrasound vector flow imaging using smoothed particle hydrodynamics,” *Physics in Medicine & Biology*, vol. 63, no. 20, p. 205011, 2018.
- [28] V. Vishnevskiy, R. Rau, and O. Goksel, “Deep variational networks with exponential weighting for learning computed tomography,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2019, pp. 310–318.
- [29] D.-A. Clevert, T. Unterthiner, and S. Hochreiter, “Fast and accurate deep network learning by exponential linear units (ELUs),” *arXiv preprint arXiv:1511.07289*, 2015.
- [30] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [31] J. A. Jensen and S. Leeman, “Nonparametric estimation of ultrasound pulses,” *IEEE Transactions on Biomedical Engineering*, vol. 41, no. 10, pp. 929–936, 1994.
- [32] O. Mattausch and O. Goksel, “Image-based PSF estimation for ultrasound training simulation,” in *International Workshop on Simulation and Synthesis in Medical Imaging*. Springer, 2016, pp. 23–33.
- [33] P. M. Shankar, J. M. Reid, H. Ortega, C. W. Piccoli, and B. B. Goldberg, “Use of non-rayleigh statistics for the identification of tumors in ultrasonic B-scans of the breast,” *IEEE Transactions on Medical Imaging*, vol. 12, no. 4, pp. 687–692, 1993.
- [34] P.-H. Tsui, C.-K. Yeh, C.-C. Chang, and Y.-Y. Liao, “Classification of breast masses by ultrasonic nakagami imaging: a feasibility study,” *Physics in Medicine & Biology*, vol. 53, no. 21, p. 6027, 2008.
- [35] J. W. Goodman, “Statistical properties of laser speckle patterns,” in *Laser speckle and related phenomena*. Springer, 1975, pp. 9–75.
- [36] J. Bamber and R. Dickinson, “Ultrasonic B-scanning: a computer simulation,” *Physics in Medicine & Biology*, vol. 25, no. 3, p. 463, 1980.
- [37] J. M. Thijssen and B. J. Oosterveld, “Texture in tissue echograms. speckle or information?” *Journal of Ultrasound in Medicine*, vol. 9, no. 4, pp. 215–229, 1990.
- [38] I. Hacıhaliloglu, R. Abugharbieh, A. J. Hodgson, and R. N. Rohling, “Bone surface localization in ultrasound using image phase-based features,” *Ultrasound in medicine & biology*, vol. 35, no. 9, pp. 1475–1487, 2009.
- [39] R. Rau, O. Unal, D. Schweizer, V. Vishnevskiy, and O. Goksel, “Attenuation imaging with pulse-echo ultrasound based on an acoustic reflector,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2019, pp. 601–609.