Deep-Learning Based Adaptive Ultrasound Imaging From Sub-Nyquist Channel Data

Alon Mamistvalov, Ariel Amar, Naama Kessler, and Yonina C. Eldar, *Fellow, IEEE*

Abstract—**Traditional beamforming of medical ultrasound images relies on sampling rates significantly higher than the actual Nyquist rate of the received signals. This results in large amounts of data to store and process, imposing hardware and software challenges on the development of ultrasound machinery and algorithms, and impacting the resulting performance. In light of the capabilities demonstrated by deep learning methods over the past years across a variety of fields, including medical imaging, it is natural to consider their ability to recover high-quality ultrasound images from partial data. Here, we propose an approach for deep-learning-based reconstruction of B-mode images from temporally and spatially sub-sampled channel data. We begin by considering sub-Nyquist sampled data, time-aligned in the frequency domain and transformed back to the time domain. The data are further sampled spatially so that only a subset of the received signals is acquired. The partial data is used to train an encoder-decoder convolutional neural network (CNN), using as targets minimum-variance (MV) beamformed signals that were generated from the original, fully-sampled data. Our approach yields high-quality B-mode images, with up to two times higher resolution than previously proposed reconstruction approaches (NESTA) from compressed data as well as delay-and-sum (DAS) beamforming of the fully-sampled data. In terms of contrast-to-noise ratio (CNR), our results are comparable to MV beamforming of the fully-sampled data, and provide up to 2 dB higher CNR values than DAS and NESTA, thus enabling better and more efficient imaging than what is used in clinical practice today.**

Index Terms—**Beamforming, deep-learning, sub-Nyquist reconstruction, ultrasound imaging.**

I. Introduction

OVER the past decades, ultrasound has become a preferred scanning modality in a variety of clinical scenarios due to its non-ionizing and noninvasive nature, high availability, and relatively low cost. The scanning process consists of transmitting acoustic pulses to the tissue, receiving their echos and digitally compensating their arrival time due to the geometry, and finally summing over the channel. This process is also referred to as beamforming.

Performing beamforming in the time domain necessitates high sampling rates of the received signals. This requirement originates from the time-alignment step, in which sufficient delay resolution is obtained through oversampling and interpolation. In practice, signals are sampled at rates 4–10 times higher than the transducer central frequency, leading to sampling rates much higher than the Nyquist rate [1]–[3], which is considered as the minimum sampling rate to allow full reconstruction of the continuous-time signal before sampling [2], [4]. This leads to vast amounts of samples to transmit and process in order to produce the final image.

Data volume is crucial in particular in receive beamforming. At this stage, averaging the signals across the array is performed using either a pre-defined apodization, as in delay-and-sum (DAS) beamforming [5], or a per-pixel data-adaptive apodization, as in Wiener beamforming and coherence factor [6] or minimum-variance (MV) beamforming [7]. While the latter allows for better trade-off between the main lobe's width and side-lobes intensity, translating to improved resolution and contrast in the final image, its computational cost is high and increases with data size. Therefore, it is more expensive and harder to implement MV beamforming in real-time clinical applications, making DAS the method of choice, resulting in degraded image quality [8].

To circumvent the long processing time and the high computational cost, a variety of techniques have emerged which enable reconstruction of the beamformed signal from partial data. In [3] and [9]–[11], the authors shift the process of time-alignment to the frequency domain by drawing a connection between the set of Fourier coefficients of the received signals pre-alignment and the set of Fourier coefficients of the beamformed signal, allowing to sample the former at their effective Nyquist rate. A reduction to a sub-Nyquist rate is also considered, by sampling only a subset of the Fourier coefficients of the received signals. Reconstruction is then performed using compressed sensing (CS) methods [2], [12]–[14], relying on the finite-rate-of-innovation structure of the beamformed signal [2], [12], [15].

In [16], a method is proposed to reduce the data and hardware burden by using sparse arrays, namely, only a subset

of the receive elements are activated. Processing is then performed by convolutional beamforming, which is shown to preserve the array beampattern under appropriate conditions on the chosen array [17]. More specifically, the achieved beampattern is equivalent to that of a virtual array given by the sum co-array of the sparse array. Thus, using a sparse array whose sum co-array includes a full uniform linear array, yields enhanced resolution and contrast from fewer transmitting elements. This method is integrated with Fourier domain beamforming in [18] and [19], allowing sub-sampling in both space and time.

The above-mentioned works perform recovery from the partial samples by solving a minimization problem in an iterative manner. Since the process is repeated for every acquisition angle in the frame, this results in long processing times. Moreover, fixed weighting is applied on the signals prior to their summation, which results in degraded image resolution in comparison to adaptive beamforming methods [7], [8]. To improve on this, one would prefer to replace the fixed weighting with adaptive weights. However, those are calculated per scanned depth and transmission angle, making the process computationally expensive.

Inspired by the notable performance of deep learning over the past years across a wide range of fields and tasks [20]–[22], medical imaging included [23], [24], different uses of deep learning in ultrasound reconstruction have been investigated [25], [26]. In [8], a deep-learning-based MV beamformer is proposed, implemented with a fully-connected neural network over fully-sampled and spatially sub-sampled channel data. The objective is to ease the heavy computational burden of adaptive beamforming, as well as to improve performance by learning from samples. The works [27]–[29] target the same problem using encoder-decoder architectures, while [30] considers the combination of deep neural networks and MV beamforming for contrast enhancement. It suggests an ensemble of networks operating in the frequency domain over frequency sub-bands, either before or after computation of the adaptive weights. In [31], the authors expand the problem to sub-sampling of transducer elements and transmission angles. They theoretically justify their approach by drawing a connection between the encoder-decoder architecture and a low-rank Hankel matrix decomposition which models the problem. Other works that consider reconstruction from a partial set of plane-wave transmissions use fully convolutional networks [32], [33], encoder-decoder networks [34]–[37] and generative adversarial networks (GANs) [38], [39].

The aforementioned approaches propose fixed sampling schemes that do not depend on the transmitted pulses or the task at hand. This aspect is addressed in [40], where two concatenated models are proposed. The first model learns to sub-sample the data, while the second, whose architecture depends on the desired task, learns the recovery. Both are jointly trained in an end-to-end fashion. The method is tested in several different sub-sampling tasks, including temporal sub-sampling of partial Fourier measurements and recovery of the original signal from them. However, this particular task is tested *in-silico* only over simulated random signal vectors; time-alignment of the sub-sampled signals, which is

required for ultrasound image recovery, is not addressed, and neither is beamforming. Since traditional time-alignment is not possible over a sub-sampled grid, the application to recovery of temporally sub-sampled ultrasound data remains unexplored.

To the best of our knowledge, no other deep-learning-based method has been proposed for the recovery of temporally partial ultrasound channel data. Here, we address this issue based on the results in [3] and [16]. The input to our model are partial Fourier measurements of the received signals, from either the full array or from a subset of channels, emulating a sparse array. Those can be obtained by sub-Nyquist sampling implemented in hardware, as shown in [15] and [41]. Using a convolutional neural network (CNN) trained separately on each reduction factor with MV-beamformed targets, our model learns a transformation from the sub-sampled grid to a high-quality B-mode image. Despite the significant reduction in data volume, the CNN outperforms fully-sampled DAS in terms of resolution and is comparable to fully-sampled MV-beamforming in terms of contrast-to-noise ratio (CNR). As such, it offers improved image quality compared to the DAS beamformer used in clinical systems today, at a much lower cost in terms of hardware and software.

The remainder of this article is organized as follows. Section II shortly reviews DAS and Fourier domain beamforming. Section III introduces our proposed method, which is verified in Sections IV and V. Results are discussed in Section VI.

II. EXISTING BEAMFORMING METHODS

A. Delay-and-Sum (DAS) Beamforming

Consider a phased-array transducer of M elements aligned along the x-axis, where m_0 denotes the central element, and δ_m denotes the distance to the mth element.

The imaging cycle begins at time $t = 0$, when a short pulse is transmitted from the array in direction θ. Denote by $(x, z) = (ct \sin \theta, ct \cos \theta)$ the coordinates of the pulse at time t, as it propagates through the tissue at speed c. Assume that a point reflector positioned at this location scatters the energy such that an echo is received by all array elements, at a time depending on their location. Beamforming is the operation of averaging the reflections while compensating for these differences in arrival time.

Let $\varphi_m(t)$ be the signal received by the mth element, and

$$\hat{\tau}(t; \theta) = t + \frac{d_m(t; \theta)}{c} \tag{1}$$

be its time of arrival, where

$$d_m(t; \theta) = \sqrt{(ct \cos \theta)^2 + (\delta_m - ct \sin \theta)^2} \tag{2}$$

is the distance traveled by the reflection to the element. Applying an appropriate delay to the mth signal results in its alignment to the origin m_0

$$\hat{\varphi}_m(t; \theta) = \varphi_m(\tau_m(t; \theta)) \tag{3}$$

where

$$\tau_m(t; \theta) = \frac{1}{2}\left(t + \sqrt{t^2 - 4(\delta_m/c)t \sin \theta + 4(\delta_m/c)^2}\right). \tag{4}$$

The final beam is derived by averaging the aligned signals received by the entire array

$$\Phi(t;\theta) = \frac{1}{M}\sum_{m=1}^{M}\hat{\varphi}_m(t;\theta). \tag{5}$$

B. Fourier Domain Beamforming

Applying appropriate time delays necessitates high sampling rates, 4–10 times higher than the probe's central frequency. To address this problem, it was shown in [3] that beamforming can be implemented equivalently in the frequency domain, bypassing the need for oversampling as no shifts are actually performed.

Let $c[k]$ be the kth Fourier series coefficient of the beam $\Phi(t;\theta)$, which can be expressed as

$$c[k] = \frac{1}{M}\sum_{m=1}^{M}\hat{c}_m[k] \tag{6}$$

where $\hat{c}_m[k]$ is given by

$$\hat{c}_m[k] = \frac{1}{T}\int_0^T I_{[0,T_B(\theta)]}(t)\hat{\varphi}_m(t;\theta)e^{-i(2\pi/T)kt}dt. \tag{7}$$

Here, $I_{[0,T_B(\theta)]}$ is an indicator function for the beam's support and T is defined by the penetration depth of the transmitted pulse. Following the derivation in [42] and [3], $\hat{c}_m[k]$ can be written as

$$\hat{c}_m[k] = \sum_n c_m[k-n]Q_{k,m,\theta}[n] \tag{8}$$

where $c_m[k]$ are the Fourier coefficients of the signal received in the mth element with no time-alignment applied to it, and $Q_{k,m,\theta}[n]$ are the Fourier coefficients of a distortion function $q_{k,m}(t;\theta)$, that effectively transfers the beamforming delays defined in (4) to the frequency domain. The function $q_{k,m}(t;\theta)$ depends on the geometry of the array alone, and therefore, its Fourier coefficients can be computed offline and stored in memory.

Since most of the energy of $\{Q_{k,m,\theta}[n]\}$ is concentrated around the dc component, the infinite sum in (8) can be approximated sufficiently with the finite sum

$$\hat{c}_m[k] \cong \sum_{n\in\nu(k)} c_m[k-n]Q_{k,m,\theta}[n] \tag{9}$$

where $\nu(k)$ depends on the decay properties of $\{Q_{k,m,\theta}[n]\}$. Substituting (9) into (6) yields the beamformed signal in the frequency domain

$$c[k] \cong \frac{1}{M}\sum_{m=1}^{M}\sum_{n\in\nu(k)} c_m[k-n]Q_{k,m,\theta}[n]. \tag{10}$$

Applying an inverse Fourier transform on $\{c[k]\}$ results in the beamformed signal in time. The relationship (10), proves that the Fourier coefficients of the beam can be obtained as a linear combination of the Fourier coefficients of the non-delayed received signals. Therefore, it is possible to transfer the process of beamforming to the frequency domain while yielding similar results.
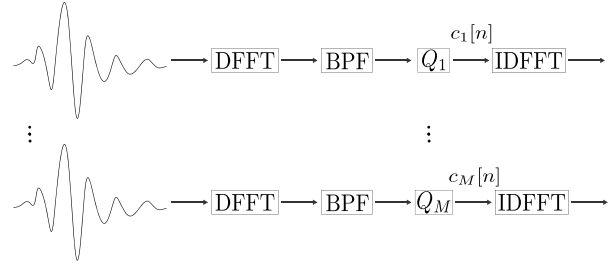

Fig. 1. Temporal sub-sampling scheme per one transmission angle. Each fully-sampled signal is digitally filtered using a bandpass filter of the desired width around the central transmission frequency. The acquired subset of Fourier coefficients is then multiplied by the corresponding Fourier coefficients $\{Q_{k,m,\theta}[n]\}$ (9), effectively applying TOF-correction in the frequency domain [3]. The resulting signals are zero-padded and transformed back to the time domain.

The required set of Fourier coefficients of the received signals can be obtained in hardware using low-rate sampling, significantly lower than the rate required for time-domain beamforming [15], [41]. Further reduction to a sub-Nyquist rate is achieved by obtaining only a subset of the coefficients, resulting in a subset of the beam's coefficients. In this scheme, however, the inverse Fourier transform does not sufficiently recover the beamformed signal in time, and additional methods are required for full recovery.

As indicated in (10), Fourier domain beamforming incorporates DAS beamforming of the partial data and thus yields similar resolution; however, how to integrate adaptive beamforming is not clear. To overcome this challenge, we introduce a deep-learning-based approach, substituting adaptive beamforming on the sub-sampled grid by learning a direct transformation to a high-quality B-mode image.

III. METHOD

A. Data Sub-Sampling and Pre-Processing

Our pre-processing pipeline consists of digitally emulating temporal sub-sampling with two sampling factors, each constituting a distinct dataset, and applying time-of-flight (TOF) correction in the frequency domain. The two datasets are denoted in the following as $D_{\text{temp}_{\times a}}$, $D_{\text{temp}_{\times b}}$, where a, b are the sampling factors. Specifically, each dataset is created as follows: The fully-sampled signals are first transformed to the Fourier domain; then, they are filtered using a rectangular bandpass filter which corresponds to the desired frequency band around the transmission's central frequency. The resulting coefficients are multiplied with the Fourier coefficients $\{Q_{k,m,\theta}[n]\}$ (9), following the method in [3]. From this point on we assume that we are given only the temporally sub-sampled data. Then, the data is transformed back to the time domain by restoring the negative spectrum of each signal (i.e., its transposed-conjugate, since the signals are real-valued), padding with an appropriate-sized vector of zeros to maintain the original resolution in time, and performing an inverse discrete Fourier transform. The process is described in Fig. 1.

Since no additional processing steps are performed on the under-sampled data to account for the loss of frequencies outside the selected bandwidth, transforming it back to the
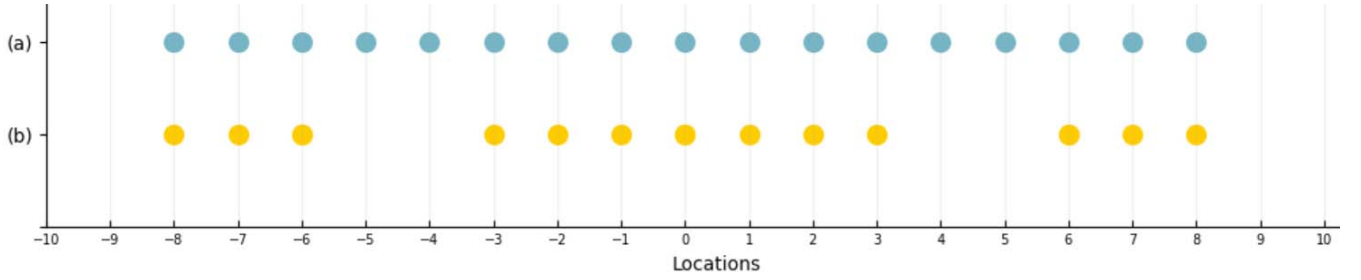
Fig. 2. Elements position of (a) full array (b) sparse array. In this example, $M/2 = 9$, $A = 3$, and $B = 3$.
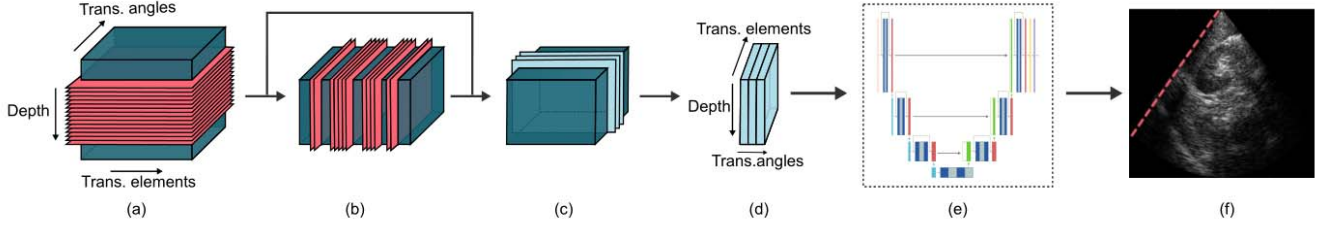

Fig. 3. Overview of the proposed method's pipeline. (a) Temporal sub-sampling and time-of-flight correction of the fully-sampled data in the frequency domain, as detailed in Fig. 1. (b) Spatial sub-sampling, which digitally emulates reception from a sparse array [16] (optional). (c) Slicing the sub-sampled data to samples, each consisting of data from three consecutive transmission angles. (d) Permuting the second and third dimensions of each sample, such that the signals received across the array constitute the different channels of input to the network. (e) Feeding the sample to the neural network. (f) Prediction of the network, which consists of the beamformed radio frequency signal in one transmission angle.

time domain introduces aliasing artifacts as well as degraded resolution. However, working in this domain rather than staying in the frequency domain, allows us to model the recovery of the sub-sampled signals and their beamforming as a standard computer-vision problem, somewhat similar to image enhancement or artifact removal. Moreover, it simplifies the task since the difficulties of designing a neural network that operates in the frequency domain and works with complex numbers are avoided.

To further reduce data volume, a third dataset $D_{\text{spatio}-\text{temp}}$ is generated by sampling $D_{\text{temp}_{\times a}}$ spatially, following the sparse arrays approach presented in [16]. The sampling pattern is obtained as follows. Let M be the full array's width, and $A, B \in N^+$ be a factorization of $M/2$ such that $AB = M/2$. Define the arrays U_A, U_B, and U_C as

$$U_A = \{-(A-1), \ldots, 0, \ldots, A-1\}$$
$$U_B = \{mA \colon m = -(B-1), \ldots, 0, \ldots, B-1\}$$
$$U_C = \{m \colon |m| = M - A, \ldots, M - 1\}. \quad (11)$$

Then the indices of half of the desired array are given by

$$\hat{U} = U_A \cup U_B \cup U_C \quad (12)$$

and indices for the entire sparse array are obtained by symmetrically concatenating those around the central element

$$U = (M/2 - \hat{U}) \cup (M/2 + \hat{U}). \quad (13)$$

As an example, for $M/2 = 9$, $A = 3$, and $B = 3$, the set $\hat{U}$ has only 13 elements out of 17 in the full array, as depicted in Fig. 2.

In contrast to [16], no convolution is performed on the filtered data before it is introduced to the network since

such pre-processing resulted empirically in degraded results. This can be explained by the fact that the network does not integrate components inspired by the concept of convolutional beamforming, and moreover, is required to learn different, more complex apodization weights. Since those are computed originally from the signals themselves rather than from the convolved signals, feeding the latter as input to the network might resolve in a harder learning task.

Each sub-sampled dataset is used to train and test the network separately. It is introduced to the network in 3-D data cubes in order to exploit correlation in all three dimensions, each input sample consisting of all scanned depths per three consecutive transmission angles θ_{j-1}, θ_j, θ_{j+1}, and all channels of either the original or the sparse array. The network's output is the beamformed signal at transmission angle θ_j (see Fig. 3). Correspondingly, each target consists of the MV-beamformed signal at θ_j.

B. Network Architecture

The desired output of our deep-beamformer is an enhanced B-mode image. Hence, an exact recovery of each sub-sampled signal prior to beamforming is not necessarily required. In other words, we would rather let the network learn a non-exact recovery from the partial samples of each signal to the original signal, if such is able to produce a better beamformed image. To allow that relaxation, we tackle the recovery from partial samples and beamforming jointly in a single model, instead of handling each stage separately.

Our learning model is a variant of UNet, an encoder-decoder CNN which was originally developed for segmentation tasks with limited amount of training data [43]. It consists of three
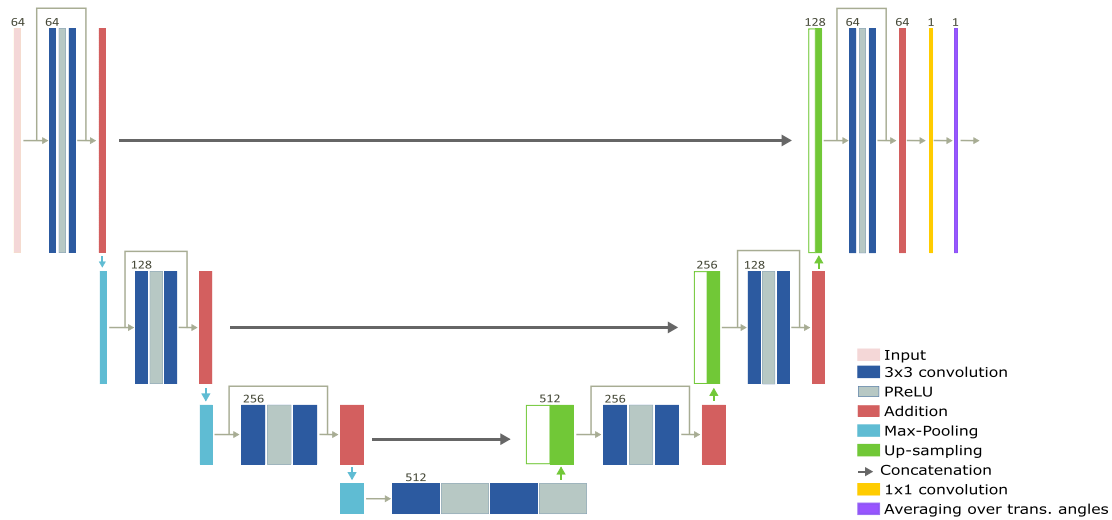
Fig. 4. Proposed architecture. The number of channels in each block is indicated above the block's first box; the light-gray arrow in each block indicates an additional residual connection, which enables the summation of the block's input with its output. The final averaging of beamformed signals from three transmission angles to one is indicated in the last box.

blocks in a contractive path, three blocks in an expansive path and one bottleneck block between them; corresponding blocks are connected with skip connections, i.e., concatenation along the third dimension. Each block of the model consists of two 3×3 convolutions, followed by either a max-pooling layer for a contractive block or an upsampling layer for an expansive block, both operating along the first dimension.

Moreover, Parametric ReLU (PReLU) [44] is used as an activation function rather then ReLU in order to assist the learning process. The model outputs three lines of the beamformed image, which corresponds to the data slice of three consecutive transmission angles we entered as input. We then average over the three angles to get a single line of the beamformed image.

Choosing UNet may seem non-trivial due to the large receptive field created in the deeper layers of the model, toward the bottleneck. It promotes both local and more global features, the latter being typically used in tasks which consider large objects, such as segmentation and classification. Nevertheless, UNet's compatibility to some classes of inverse problems was addressed in [45], and examples of it being applied to different reconstruction problems, ultrasound reconstruction included, were published over the last years [35], [46]. In essence, the class of inverse problems consists of tasks in which a signal is retrieved from a set of measurements, obtained by some forward model. They can be solved by iterative processes which aim to minimize the distance between the actual measurements and the approximated signal after it passes through the forward model, with some regularization derived from prior knowledge about the signal. Such solutions include a multiplication of the operator associated with the forward model with its adjoint and an inversion of the result. Therefore, when the support of the filters associated with these operations is not compact, large receptive field may be a desirable feature in a network that learns to output similar results [45]. An overview of the chosen architecture is given in Fig. 4.

C. Loss Function

Ultrasound channel data is characterized by a large dynamic range, and is typically compressed after beamforming to obtain the final B-mode image. Therefore, to promote visual similarity between the output of a learning model and its corresponding beamformed target, we follow [8] and apply compression within the loss function used in training, in a variant of mean-squared-error (MSE). The function, named signed-mean-squared-logarithmic-error (SMSLE), is defined as

$$L_{\text{SMSLE}} = 0.5 \cdot \left\| \log_{10}(B_{\text{Pred}}^+) - \log_{10}(B_{\text{MV}}^+) \right\|_2^2$$
$$+ 0.5 \cdot \left\| \log_{10}(B_{\text{Pred}}^-) - \log_{10}(B_{\text{MV}}^-) \right\|_2^2 \quad (14)$$

where B^+, B^- are the positive and negative parts of the radio frequency beamformed data, B_{Pred} is the model's prediction and B_{MV} is the MV-beamformed target [8].

To further promote perception-based similarity between the beamformed signals, a second term, 1-D variant of structural similarity index (SSIM) [47], was added to the loss function. This term, which employs similar compression, is given by

$$L_{\text{SSIM}} = 0.5 \cdot \left(1 - \text{SSIM}_{1D}\left(\log_{10}(B_{\text{Pred}}^+), \log_{10}(B_{\text{MV}}^+)\right)\right)$$
$$+ 0.5 \cdot \left(1 - \text{SSIM}_{1D}\left(\log_{10}(B_{\text{Pred}}^-), \log_{10}(B_{\text{MV}}^-)\right)\right)$$
$$(15)$$

where SSIM_{1D} is defined for two signals x, y as

$$\text{SSIM}_{1D}(x, y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)}. \quad (16)$$

Here, $\mu_x, \mu_y, \sigma_x, \sigma_y$, and σ_{xy} are the means, standard deviations and cross correlation of the two images, calculated over an 11-pixel 1-D sliding Gaussian window ($\sigma = 1.5$), and C_1, C_2 are constants meant to stabilize the division in image regions where the local means or standard deviations are close to zero. We choose the stabilization constants to be $C_1 = (k_1 \cdot L)^2, C_2 = (k_2 \cdot L)^2$ where L is the dynamic range

of pixel values and $k_1 = 0.01, k_2 = 0.03$ as suggested by Wang *et al.* [48].

Overall, the model was optimized using the following cost function:

$$L = 0.5 \cdot L_{\text{SMSLE}} + 0.5 \cdot L_{\text{SSIM}}. \tag{17}$$

D. Metrics for Evaluation

To the best of our knowledge, no other approach was proposed for deep-learning-based reconstruction from temporally sub-sampled data. Therefore, we cannot directly compare our model to other deep-leaning based state-of-the-art approaches. Instead, we concentrate on quantitative evaluation to demonstrate the benefits of our model. CNR is used to evaluate contrast, and full-width-at-half-maxima (FWHM), calculated over simulated point scatterers and *in vivo* data, is used to evaluate axial and lateral resolution. Overall similarity to the target image is evaluated using SSIM.

CNR and gCNR are evaluated over beamformed phantom scans after envelope detection and logarithmic compression. It is calculated from two regions with different intensities in each image, namely, an anechoic cyst phantom scan and its background, and is given by

$$\text{CNR} = 20 \cdot \log_{10}\left(\frac{|\mu_c - \mu_b|}{\sqrt{\sigma_c^2 + \sigma_b^2}}\right) \tag{18}$$

where μ_c, μ_b, σ_c, and σ_b are the means and standard deviations of the anechoic cyst phantom and the background, respectively [49]. Similarly, we also used the generalized CNR (gCNR) metric for evaluation over beamformed phantom scans. This metric is less sensitive to large dynamic ranges as in ultrasound scans. The gCNR is calculated using

$$\text{gCNR} = 1 - \int_{-\infty}^{\infty} \min_{x}\{p_o(x), p_i(x)\}dx \tag{19}$$

where p_i and p_o are prior probabilities assigned based on the size of the regions of interest, p_i is inside the region of interest, and p_o is outside the region of interest.

FWHM is calculated by first computing the peaks per column and row in each image frame, and then calculating the width for each peak, and averaging the results over all axial and lateral cuts in the frame, respectively. By averaging over all possible cuts, we obtain a reliable resolution measurement that does not prefer a specific image direction.

SSIM is evaluated per dataset over the test set of each of the three folds, in reference to the fully-sampled MV target images. It is defined for two images x, y as in (16), with the exception of using a 2-D sliding window as in the original definition of the function [48].

IV. EXPERIMENTAL SETUP

In vivo data for training and testing was acquired by scanning three healthy volunteers, using a P4-2v Verasonics phased-array transducer with 64 elements.

The scanning procedure was not a predefined protocol but an in-house scan of healthy volunteers and a visiting doctor that performed the scanning. Approval of the institutional review board (IRB) of Weizmann Institute of Science and informed consent were obtained for collecting the *in vivo* data. The dataset consists of organs of the abdominal cavity— liver, gallbladder, bladder, kidneys, and the Aorta. The carrier frequency was 2.7 MHz and the sampling rate was 10.9 MHz, which is twice the Nyquist rate, resulting in 1918 samples per image line. Two temporally sub-sampled datasets were generated from it using the described scheme. In the first, $D_{\text{temp}\times 5}$, 400 samples per image line were sampled, resulting in sampling rate of 0.42 of the Nyquist rate and fivefold reduction in comparison to the original data volume. In the second, $D_{\text{temp}\times 9}$, 220 samples per image line were sampled, resulting in a sampling rate of 0.23 of the Nyquist rate and ninefold reduction in comparison to the original data volume. As described in Section III-A, a third dataset, $D_{\text{spatio-temp}}$, was generated by spatial sampling of $D_{\text{temp}\times 5}$, omitting 37 out of the 64 transmitting elements. This results in 11-fold reduction in comparison to the original data volume.

Threefold cross-validation was used in training, resulting in three trained models per dataset. Each model was trained and validated using a different subset of two patients in an 80%–20% split, setting the third patient aside for testing. Each prediction of the model corresponds to the beamformed signal in one transmission angle; as each frame consists from 128 transmissions, this results in $4 \cdot 10^4$ training samples and $2 \cdot 10^4$ testing samples in average per fold. No organ-wise division was performed during training or testing, meaning that both stages operate on data from multiple organs. Targets for training were generated from MV beamforming of the fully-sampled data, time-aligned in the time domain.

For quantitative evaluation of CNR, tissue-mimicking phantoms Gammex 403GSLE and 404GSLE were scanned by the 64-element phased array transducer P4-2v with similar transmission specifications as the *in vivo* dataset. For calculating the FWHM we created a dataset of point scatters simulations using the Verasonics simulator for the P4-2V phased array US probe, using the same setup that was used for the *in vivo* and phantom scans.

The network was implemented with Keras, using Tensorflow backend. It was trained separately on each dataset for up to 100 epochs using Adam optimizer, with an initial learning rate of $3 \cdot 10^{-5}$ for the temporal sub-sampling and $3 \cdot 10^{-6}$ for the spatial-temporal sub-sampling. Weights were initialized using He Normal initialization [44].

V. RESULTS

In this section we evaluate the suggested method; we present quantitative and qualitative evaluations followed by a discussion of the results.

A. Numerical Study

Table I presents a quantitative evaluation of the lateral and axial resolution of the proposed method, averaged over the three models trained per reduction factor. The results were
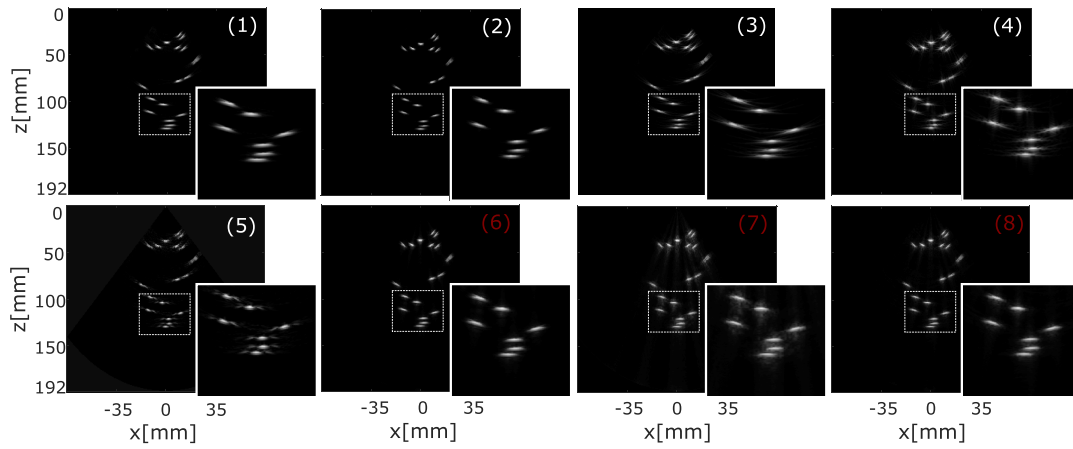
Fig. 5. Simulated point scatterers, obtained by (1) DAS (fully-sampled), (2) MV (fully-sampled), (3) NESTA ($D_{\text{temp}_{\times 5}}$), (4) NESTA ($D_{\text{temp}_{\times 9}}$), (5) NESTA ($D_{\text{spatio-temp}}$), (6) proposed method ($D_{\text{temp}_{\times 5}}$), (7) proposed method ($D_{\text{temp}_{\times 9}}$), and (8) proposed method ($D_{\text{spatio-temp}}$). In (5), NESTA is applied after convolving the data according to the convolutional beamforming framework. Predictions of the proposed method are averaged over the three models trained per reduction factor. While the proposed method demonstrates high axial resolution over the simulated frames, lateral resolution seems slightly degraded, and in frames which were sub-sampled by the ×9 sampling factor, axial halos are visually detected.
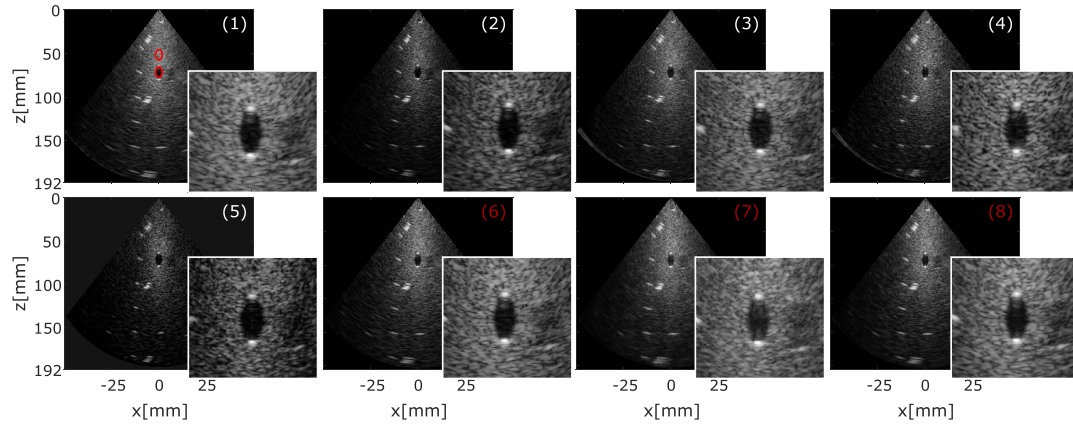

Fig. 6. Anechoic cyst phantom scan, obtained by (1) DAS (fully-sampled), (2) MV (fully-sampled), (3) NESTA ($D_{\text{temp}_{\times 5}}$), (4) NESTA ($D_{\text{temp}_{\times 9}}$), (5) NESTA ($D_{\text{spatio-temp}}$), (6) proposed method ($D_{\text{temp}_{\times 5}}$), (7) proposed method ($D_{\text{temp}_{\times 9}}$) and (8) proposed method ($D_{\text{spatio-temp}}$). In (5), NESTA is applied after convolving the data according to the convolutional beamforming framework. Red circles indicate the regions used for computing the CNRs and gCNRs.

TABLE I
RESOLUTION EVALUATION WITH FULL WIDTH AT HALF
MAXIMA (FWHM) PARAMETER

	Simulated Scatterers		*In-Vivo* images	
	lateral	axial	lateral	axial
DAS fully-sampled	0.38	0.28	0.93	0.46
MV fully-sampled	0.3	0.28	0.68	0.32
NESTA x5 reduction	0.51	0.24	1.01	0.55
NESTA x9 reduction	0.67	0.35	1.07	0.67
NESTA x11 reduction *	0.37	0.26	0.72	0.39
Proposed x5 reduction	0.4	0.17	0.76	0.37
Proposed x9 reduction	0.45	0.23	0.57	0.28
Proposed x11 reduction	0.41	0.2	0.71	0.35

* NESTA$_{\text{x11 reduction}}$ is applied after convolving the sub-sampled data according to the convolutional beamforming (COBA) framework, as opposed to the proposed method.
** Results are in mm.

TABLE II
CNR, STRUCTURAL SIMILARITY INDEX (SSIM)
EVALUATION AND gCNR

	CNR (dB)	SSIM	gCNR
DAS fully-sampled	9.99	0.75	0.89
MV fully-sampled	9.31	1	0.86
NESTA x5 reduction	10.37	0.87	0.88
NESTA x9 reduction	8.04	0.78	0.78
NESTA x11 reduction *	7.29	0.47	0.80
Proposed x5 reduction	9.82	0.75	0.77
Proposed x9 reduction	8.61	0.80	0.74
Proposed x11 reduction	9.47	0.77	0.78

network's outputs to the MV targets, assessed on the phantom scans and *in vivo* test scans, respectively. Again, results were averaged over the three models trained per reduction factor.

A comparison view of the beamformed simulated data and anechoic cyst phantom is presented in Figs. 5 and 6,

calculated for simulated point scatters and *in vivo* test data, as detailed in Section III-D. Table II presents a quantitative evaluation of CNR and gCNR, as well as similarity of the
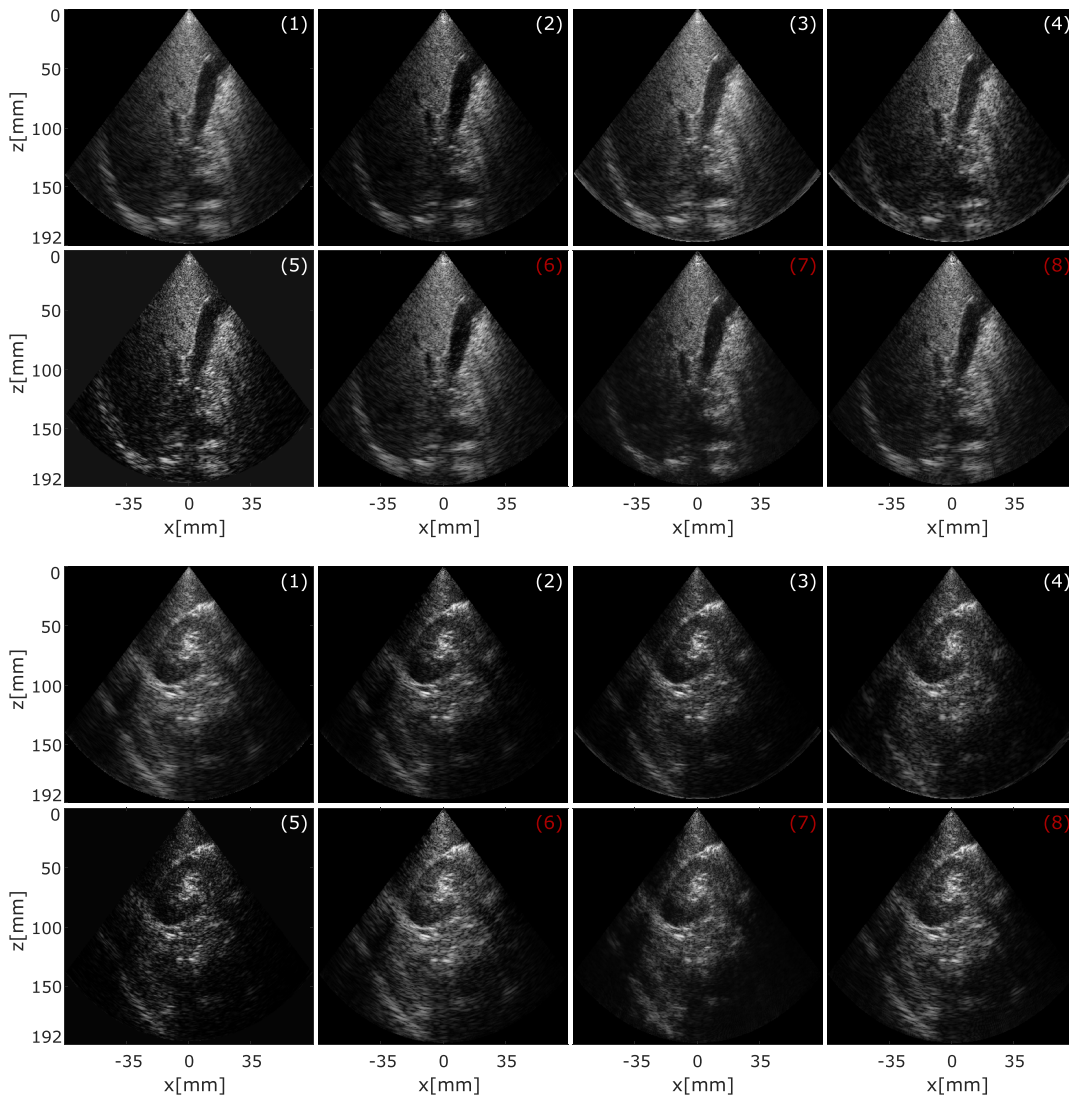
Fig. 7. Abdominal cavity images from testsets of two of the folds, obtained by (1) DAS (fully-sampled), (2) MV (fully-sampled), (3) NESTA ($D_{\text{temp}_{\times 5}}$), (4) NESTA ($D_{\text{temp}_{\times 9}}$), (5) NESTA ($D_{\text{spatio-temp}}$), (6) proposed method ($D_{\text{temp}_{\times 5}}$), (7) proposed method ($D_{\text{temp}_{\times 9}}$), and (8) proposed method ($D_{\text{spatio-temp}}$). In (5), NESTA is applied after convolving the data according to the convolutional beamforming framework.

respectively. Fig. 7 presents temporally and spatially sub-sampled *in vivo* data from testsets of the different reduction factors, beamformed by the proposed method and compared to the previously suggested approaches.

B. Discussion

An example frame from the simulated set, reconstructed by all compared methods, is presented in Fig. 5. Improved lateral resolution can be seen in the images reconstructed by the suggested method over all reduction factors, in comparison to the NESTA reconstructions. In particular, side beams which are present at the NESTA reconstruction of the ×11 reduction factor (Fig 5, image 5) are not present at all in the suggested method's reconstruction (Fig. 5, image 8). However, in the image which was sub-sampled by the ×9 sampling factor (Fig. 5, image 7), axial halos are visually detected. We believe

that these results stem from the fact that the networks were not trained on synthetic data, and therefore, had difficulty in compensating for input differences over the temporally-harder dataset.

In Fig. 6, improved contrast in comparison to DAS can be seen in the ×5 and ×11 reduction factors (images 6, 8), while degradation is apparent over the ×9 reduction factor (image 7).

Fig. 7 depicts that over $D_{\text{temp}_{\times 5}}$ and $D_{\text{spatio-temp}}$, our method yields results that are visually comparable to the target images in terms of both resolution and contrast. Moreover, there is good preservation of speckles and weaker reflections. Over $D_{\text{temp}_{\times 9}}$, however, slight suppression of speckles and mild pale artifacts are detected, especially in low SNR areas at the margins of the frame.

In terms of CNR, our method is comparable to DAS and MV over the ×5 sub-sampling factor, but is slightly degraded

over the $\times 9$ sampling factor; in comparison to NESTA, it is again slightly degraded over the $\times 5$ sub-sampling factor but is better over the $\times 9$ sampling factor. These last results suggest more robustness to the reduction in sampling rate. For the gCNR metric, the proposed method is comparable to NESTA although achieving slightly worse results in the first dataset. However, in the higher reduction cases, we get an effectively similar gCNR to NESTA. Therefore, these results combined with the CNR prove the high robustness of the proposed method to data size reduction in terms of contrast.

One of the notable advantages of the suggested technique is that unlike other adaptive beamforming techniques, such as MV, the proposed method, once learned can be used to beamform different types of scanned objects, and indeed it can be seen to recover high quality US images of a wide range of organs.

Finally, we consider two issues regarding the suggested method. First, a possible drawback of the proposed mechanism can be the fact that it does not produce an actual B-mode US image, but only a single beamformed image line. However, there are well-known steps of post-processing to visualize the image which can be added here as well. Second, using delayed data as target means our network should be able to produce both negative and positive values. Therefore, one limitation of the architecture is that we must use activation functions that do not truncate negative values (ReLU for example). Using ReLU in deep learning has been proven to make models faster in backpropagation time, the gradients tends to vanish less, and we get a sparser representation of the data; however, in our method we cannot use standard ReLU due to the reason we mentioned above.

VI. Conclusion

In this work we presented a deep-learning-based method for high-quality reconstruction of temporally and spatially sub-sampled channel data, obtained by the schemes presented in [3] and [16]. We have shown that an encoder-decoder CNN, trained on multiorgan scans using a loss function which incorporates domain knowledge, can be used to directly learn the production of high-quality B-mode images from the degraded data. Our method yields results with high resolution and contrast, which are comparable to model-based beamforming of the fully-sampled data. It performs particularly well on moderate temporal sampling, either with or without spatial sampling; yet, it seems to be more resilient to reduction in temporal sampling rate than the iterative reconstruction method which was previously used for the task in [3] and [16]. These results indicate that our method can be plugged into previously proposed schemes of sub-Nyquist and sparse-array ultrasound processing, to improve performance while alleviating the requirements of the data's transmission and processing.

Acknowledgment

References

[1] B. D. Steinberg, "Digital beamforming in ultrasound," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 39, no. 6, pp. 716–721, Nov. 1992.

[2] Y. C. Eldar, *Sampling Theory: Beyond Bandlimited Systems*. Cambridge, U.K.: Cambridge Univ. Press, 2015.

[3] T. Chernyakova and Y. Eldar, "Fourier-domain beamforming: The path to compressed ultrasound imaging," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 61, no. 8, pp. 1252–1267, Aug. 2014.

[4] C. E. Shannon, "Communication in the presence of noise," *Proc. IRE*, vol. 37, no. 1, pp. 10–21, Jan. 1949.

[5] K. E. Thomenius, "Evolution of ultrasound beamformers," in *Proc. IEEE Ultrason. Symp. (IUS)*, San Antonio, TX, USA, vol. 2, Nov. 1996, pp. 1615–1622.

[6] C.-I. C. Nilsen and S. Holm, "Wiener beamforming and the coherence factor in ultrasound imaging," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 57, no. 6, pp. 1329–1346, Jun. 2010.

[7] J. Capon, "High-resolution frequency-wavenumber spectrum analysis," *Proc. IEEE*, vol. 57, no. 8, pp. 1408–1418, Aug. 1969.

[8] B. Luijten *et al.*, "Adaptive ultrasound beamforming using deep learning," *IEEE Trans. Med. Imag.*, vol. 39, no. 12, pp. 3967–3978, Dec. 2020, doi: 10.1109/TMI.2020.3008537.

[9] A. G. J. Besson *et al.*, "A sparse reconstruction framework for Fourier-based plane-wave imaging," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 63, no. 12, pp. 2092–2106, Dec. 2016.

[10] J.-Y. Lu, "2D and 3D high frame rate imaging with limited diffraction beams," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 44, no. 4, pp. 839–856, Jul. 1997.

[11] J.-Y. Lu, J. Cheng, and J. Wang, "High frame rate imaging system for limited diffraction array beam imaging with square-wave aperture weightings high frame rate imaging system for limited diffraction array beam imaging with square-wave aperture weightings," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 53, no. 10, pp. 1796–1812, Oct. 2006.

[12] Y. C. Eldar and G. Kutyniok, *Compressed Sensing: Theory and Applications*. Cambridge, U.K.: Cambridge Univ. Press, 2012.

[13] N. Dobigeon, A. Basarab, D. Kouamé, and J.-Y. Tourneret, "Regularized Bayesian compressed sensing in ultrasound imaging," in *Proc. 20th Eur. Signal Process. Conf. (EUSIPCO)*, 2012, pp. 2600–2604.

[14] H. Liebgott, A. Basarab, D. Kouame, O. Bernard, and D. Friboulet, "Compressive sensing in medical ultrasound," in *Proc. IEEE Int. Ultrason. Symp.*, Oct. 2012, pp. 1–6.

[15] R. Tur, Y. C. Eldar, and Z. Friedman, "Innovation rate sampling of pulse streams with application to ultrasound imaging," *IEEE Trans. Signal Process.*, vol. 59, no. 4, pp. 1827–1842, Apr. 2011.

[16] R. Cohen and Y. C. Eldar, "Sparse convolutional beamforming for ultrasound imaging," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 65, no. 12, pp. 2390–2406, Dec. 2018.

[17] R. Cohen and Y. C. Eldar, "Sparse array design via fractal geometries," *IEEE Trans. Signal Process.*, vol. 68, pp. 4797–4812, 2020.

[18] A. Mamistvalov and Y. C. Eldar, "Sparse convolutional beamforming for wireless ultrasound," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2020, pp. 9254–9258.

[19] A. Mamistvalov and Y. C. Eldar, "Compressed Fourier-domain convolutional beamforming for wireless ultrasound imaging," 2020, *arXiv:2010.13171*.

[20] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.

[21] Y. Guo, Y. Liu, A. Oerlemans, S. Lao, S. Wu, and M. S. Lew, "Deep learning for visual understanding: A review," *Neurocomputing*, vol. 187, pp. 27–48, Apr. 2016.

[22] Z. Wang, J. Chen, and S. C. H. Hoi, "Deep learning for image super-resolution: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 10, pp. 3365–3387, Oct. 2021, doi: 10.1109/TPAMI.2020.2982166.

[23] H. Greenspan, B. V. Ginneken, and R. M. Summers, "Guest editorial deep learning in medical imaging: Overview and future promise of an exciting new technique," *IEEE Trans. Med. Imag.*, vol. 35, no. 5, pp. 1153–1159, Mar. 2016.

[24] B. Sahiner *et al.*, "Deep learning in medical imaging and radiation therapy," *Med. Phys.*, vol. 46, no. 1, pp. e1–e36, Jan. 2019.

[25] R. J. van Sloun, R. Cohen, and Y. C. Eldar, "Deep learning in ultrasound imaging," *Proc. IEEE*, vol. 108, no. 1, pp. 11–29, Aug. 2019.

[26] D. Hyun *et al.*, "Deep learning for ultrasound image formation: CUBDL evaluation framework and open datasets," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 68, no. 12, pp. 3466–3483, Dec. 2021.

[27] S. Khan, J. Huh, and J. C. Ye, "Universal deep beamformer for variable rate ultrasound imaging," 2019, *arXiv:1901.01706*.

[28] S. Khan, J. Huh, and J. C. Ye, "Adaptive and compressive beamforming using deep learning for medical ultrasound," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 67, no. 8, pp. 1558–1572, Aug. 2020.

[29] W. Simson *et al.*, "End-to-end learning-based ultrasound reconstruction," 2019, *arXiv:1904.04696*.

[30] R. Zhuang and J. Chen, "Deep learning based minimum variance beamforming for ultrasound imaging," in *Smart Ultrasound Imaging and Perinatal, Preterm and Paediatric Image Analysis*, vol. 11798. Cham, Switzerland: Springer, 2019, pp. 83–91.

[31] Y. H. Yoon, S. Khan, J. Huh, and J. C. Ye, "Efficient B-mode ultrasound image reconstruction from sub-sampled RF data using deep learning," *IEEE Trans. Med. Imag.*, vol. 38, no. 2, pp. 325–336, Feb. 2019.

[32] M. U. Ghani, F. C. Meral, F. Vignon, and J.-L. Robert, "High frame-rate ultrasound imaging using deep learning beamforming," in *Proc. IEEE Int. Ultrason. Symp. (IUS)*, Glasgow, U.K., Oct. 2019, pp. 295–298.

[33] M. Gasse, F. Millioz, E. Roux, D. Garcia, H. Liebgott, and D. Friboulet, "High-quality plane wave compounding using convolutional neural networks," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 64, no. 10, pp. 1637–1639, Oct. 2017.

[34] W. Simson, M. Paschali, N. Navab, and G. Zahnd, "Deep learning beamforming for sub-sampled ultrasound data," in *Proc. IEEE Int. Ultrason. Symp. (IUS)*, Kobe, Japan, Oct. 2018, pp. 1–4.

[35] O. Senouf *et al.*, "High frame-rate cardiac ultrasound imaging with deep learning," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, Granada, Spain, 2018, pp. 126–134.

[36] S. Vedula *et al.*, "High quality ultrasonic multi-line transmission through deep learning," in *Proc. Int. Workshop Mach. Learn. Med. Image Reconstruction*, Granada, Spain, 2018, pp. 147–155.

[37] D. Perdios, M. Vonlanthen, A. Besson, F. Martinez, M. Arditi, and J.-P. Thiran, "Deep convolutional neural network for ultrasound image enhancement," in *Proc. IEEE Int. Ultrason. Symp. (IUS)*, Kobe, Japan, Oct. 2018, pp. 1–4.

[38] A. A. Nair, T. D. Tran, A. Reiter, and M. A. L. Bell, "A generative adversarial neural network for beamforming ultrasound images: Invited presentation," in *Proc. 53rd Annu. Conf. Inf. Sci. Syst. (CISS)*, Baltimore, MD, USA, Mar. 2019, pp. 1–6.

[39] W. Choi, M. Kim, J. HakLee, J. Kim, and J. BeomRa, "Deep CNN-based ultrasound super-resolution for high-speed high-resolution B-mode imaging," in *Proc. IEEE Int. Ultrason. Symp. (IUS)*, Kobe, Japan, Oct. 2018, pp. 1–4.

[40] I. A. M. Huijben, B. S. Veeling, K. Janse, M. Mischi, and R. J. G. Van Sloun, "Learning sub-sampling and signal recovery with applications in ultrasound imaging," *IEEE Trans. Med. Imag.*, vol. 30, no. 12, pp. 3955–3966, Dec. 2020, doi: 10.1109/TMI.2020.3008501.

[41] E. Baransky, G. Itzhak, N. Wagner, I. Shmuel, E. Shoshan, and Y. Eldar, "Sub-Nyquist radar prototype: Hardware and algorithm," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 50, no. 2, pp. 809–822, Apr. 2014.

[42] N. Wagner, Y. C. Eldar, and Z. Friedman, "Compressed beamforming in ultrasound imaging," *IEEE Trans. Signal Process.*, vol. 60, no. 9, pp. 4343–4657, Sep. 2012.

[43] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, Munich, Germany, 2015, pp. 234–241.

[44] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Santiago, Chile, Dec. 2015, pp. 1026–1034.

[45] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser, "Deep convolutional neural network for inverse problems in imaging," *IEEE Trans. Image Process.*, vol. 26, no. 9, pp. 4509–4522, Sep. 2016.

[46] D. Perdios, M. Vonlanthen, F. Martinez, M. Arditi, and J.-P. Thiran, "CNN-based image reconstruction method for ultrafast ultrasound imaging," 2020, *arXiv:2008.12750*.

[47] D. Breakey and C. Meskell, "Comparison of metrics for the evaluation of similarity in acoustic pressure signals," *J. Sound Vib.*, vol. 332, no. 15, pp. 3605–3609, Jul. 2013.

[48] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.

[49] M. S. Patterson and F. S. Foster, "The improvement and quantitative assessment of B-mode images produced by an annular array/cone hybrid," *Ultrason. Imag.*, vol. 5, no. 3, pp. 195–213, 1983, doi: 10.1177/016173468300500301.

Alon Mamistvalov received the B.Sc. degree in physics and the B.Sc. degree in electrical engineering from the Technion—Israel Institute of Technology, Haifa, Israel, in 2017. He is currently pursuing the M.Sc. degree in mathematics and computer science with the Weizmann Institute of Science, Rehovot, Israel.

His research interests include signal processing, medical imaging, advanced signal processing methods for ultrasonic imaging and deep learning for ultrasound imaging.

Ariel Amar received the B.Sc. degree in electrical and computer engineering from The Hebrew University of Jerusalem, Jerusalem, Israel, in 2018. He is currently pursuing the M.Sc. degree in mathematics and computer science with the Weizmann Institute of Science, Rehovot, Israel.

His research interests include deep learning, medical imaging, signal processing, and deep learning for ultrasonic imaging.

Naama Kessler received the B.Sc. degree in computer science from The Hebrew University of Jerusalem, Jerusalem, Israel, in 2019.

Then, she worked as an Algorithm Developer at the Weizmann Institute of Science, Rehovot, Israel, where she is currently working under the same role in the industry. Her professional interests include computer vision, machine learning for computer vision, image processing, and medical imaging.

Yonina C. Eldar (Fellow, IEEE) received the B.Sc. degree in physics and the B.Sc. degree in electrical engineering from Tel-Aviv University (TAU), Tel-Aviv, Israel, in 1995 and 1996, respectively, and the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2002.

She is currently a Professor with the Department of Mathematics and Computer Science, Weizmann Institute of Science, Rehovot, Israel. Previously, she was a Professor with the Department of Electrical

Engineering, Technion, where she held the Edwards Chair in engineering. She is also a Visiting Professor with MIT; a Visiting Scientist with the Broad Institute; and an Adjunct Professor with Duke University, Durham, NC, USA; and was a Visiting Professor at Stanford University, Stanford, CA, USA. Her research interests include in the broad areas of statistical signal processing, sampling theory and compressed sensing, learning and optimization methods, and their applications to biology and optics.

Dr. Eldar is a member of the Israel Academy of Sciences and Humanities (elected 2017) and a EURASIP Fellow. She was a member of the Young Israel Academy of Science and Humanities and the Israel Committee for Higher Education. She was a Horev Fellow of the Leaders in Science and Technology Program at the Technion and an Alon Fellow. She has received many awards for excellence in research and teaching, including the IEEE Signal Processing Society Technical Achievement Award in 2013, the IEEE/AESS Fred Nathanson Memorial Radar Award in 2014, and the IEEE Kiyo Tomiyasu Award in 2016. She received the Michael Bruno Memorial Award from the Rothschild Foundation, the Weizmann Prize for Exact Sciences, the Wolf Foundation Krill Prize for Excellence in Scientific Research, the Henry Taub Prize for Excellence in Research (twice), the Hershel Rich Innovation Award (three times), the Award for Women with Distinguished Contributions, the Andre and Bella Meyer Lectureship, the Career Development Chair at the Technion, the Muriel & David Jacknow Award for Excellence in Teaching, and the Technion's Award for Excellence in Teaching (two times). She received several best paper awards and best demo awards together with her research students and colleagues, including the SIAM Outstanding Paper Prize, the UFFC Outstanding Paper Award, the Signal Processing Society Best Paper Award, and the *IET Circuits, Devices & Systems* Premium Award, and was selected as one of the 50 most influential women in Israel. She is the Editor-in-Chief of *Foundations and Trends in Signal Processing*, a member of the IEEE Sensor Array and Multichannel Technical Committee and serves on several other IEEE committees. In the past, she was a Signal Processing Society Distinguished Lecturer, a member of the IEEE Signal Processing Theory and Methods and Bio Imaging Signal Processing technical committees. She was the co-chair and the technical co-chair of several international conferences and workshops. She served as an Associate Editor for IEEE TRANSACTIONS ON SIGNAL PROCESSING, the *EURASIP Journal on Advances in Signal Processing*, the *SIAM Journal on Matrix Analysis and Applications*, and the *SIAM Journal on Imaging Sciences*.