Challenge on Ultrasound Beamforming with Deep Learning (CUBDL)

Muyinatu A. Lediju Bell,[1,2,3] Jiaqi Huang,[2] Dongwoon Hyun,[4]
Yonina C. Eldar,[5] Ruud van Sloun,[6] Massimo Mischi[6]
[1]Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, MD
[2]Department of Biomedical Engineering, Johns Hopkins University, Baltimore, MD
[3]Department of Computer Science, Johns Hopkins University, Baltimore, MD
[4]Department of Radiology, Stanford University, Palo Alto, CA, USA
5Department of Math and Computer Science, Weizmann Institute of Science, Rehovot, Israel
[6]Department of Electrical Engineering, Eindhoven University of Technology, Eindhoven, Netherlands

Abstract*—Recent developments in deep learning have created immense potential for improving ultrasound beamforming. We organized a Challenge on Ultrasound Beamforming with Deep Learning (CUBDL) to benchmark methodologies in this space with two transmission data types: plane wave and focused transmissions. Plane wave ultrasound transmissions have created new opportunities for ultrafast ultrasound imaging, while focused ultrasound transmissions are more traditional and are widely used in most clinical ultrasound systems available today. For both transmission types, we challenged participants to obtain the best image quality under the fastest possible frame rates. CUBDL organizers solicited datasets from several leading ultrasound groups around the world and received a total of 106 data sequences including *in vivo*, *ex vivo*, simulated, and experimental phantom datasets. These submissions formed our test datasets, which were not released to participants while the challenge was open. The challenge was composed of three optional tasks (one including two subtasks) that were evaluated using the test datasets. Participants had the option to provide their results for a minimum of one up to a maximum of four tasks or subtasks: (1) beamforming with deep learning after a single plane wave transmission, which had two subtasks to either (a) match or (b) exceed traditional image quality metrics obtained with multiple plane wave transmissions; (2) beamforming with deep learning after a few plane wave transmissions; (3) beamforming with deep learning to achieve dynamic transmit focusing from datasets acquired with a single transmit focus. Evaluation included image quality metrics as well as network complexity metrics. A challenge website was created to provide information and updates: https://cubdl.jhu.edu/.

I. Introduction

Ultrafast imaging is often achieved by transmitting plane waves that span a wide region of interest (as opposed to transmitting focused beams line-by-line) [1], [2]. Plane wave imaging increases frame rates by over 100-fold, when compared to focused transmissions, thereby enabling applications, such as real-time brain activity monitoring, Doppler imaging, or shear wave elastography [1], [3], [4]. Increasing frame rates for future tasks requires using only a few plane waves for beamforming and as a result, suffers from image quality degradation. While previous work presented during the Plane-wave imaging challenge in medical ultrasound (PICMUS) [2] explored different beamforming methods for plane-wave images, including minimum variance beamforming [5], short-lag spatial coherence beamforming [6], and iterative maximum a posteriori beamforming [7] (among others), these methods are known to suffer from high computational complexity and complicate real-time applications. In direct comparison to plane wave imaging, focused transmissions are typically used in most commercial ultrasound imaging scanners [3], [4]. One downside of focused transmissions is that only one focus is allowed per image acquisition. While synthetic aperture beamforming [8] can be used to dynamically focus ultrasound transmissions, dynamic focusing generally involves a high computational load.

Recent developments in deep learning have created immense potential for ultrasound imaging research [9]. The Challenge on Ultrasound Beamforming with Deep Learning (CUBDL) was designed to explore the benefits of using deep learning to create images after: (1) plane wave ultrasound transmission, which has the potential to create new opportunities for ultrafast ultrasound imaging and (2) traditional focused ultrasound transmission, which is widely used in most clinical ultrasound systems available today. This goal was centered around three independent tasks associated with CUBDL. The first and second tasks challenged participants to balance both image quality and frame rates with the application of novel deep learning approaches to plane wave imaging. The third task challenged participants to explore deep learning approaches to achieve dynamic transmit focusing with the attractive possibility of faster and simpler computations than existing methods.

The remainder of this paper is organized as follows. Section II provides more details on specific challenge tasks, evaluation datasets that were sourced from multiple ultrasound research labs around the world, example images from these datasets, and the evaluation metrics for each task. We recommend these metrics as the baseline standard for evaluation of any future ultrasound beamforming with deep learning work. Section III summarizes some key details about the submissions we received while the challenge was open. Section IV includes links to multiple resources that were developed to support this challenge, including the challenge website (which will remain available to provide future updates for the community) [10], the IEEE DataPort site that was used to accept submissions [11], and the website that contains evaluation code that we recommend for future use by anyone evaluating data in this research area, as well as example code to input network weights for beamforming evaluation [12]. Section V provides concluding thoughts.

II. Challenge Overview

A. Tasks

CUBDL was composed of three optional tasks. Participants had the option to provide their results for a minimum of one up to a maximum of four tasks or subtasks.

- **Task 1: Beamforming with deep learning after a single plane wave transmission.** Task 1 has two optional subtasks.
    1) **Task 1a** was explicitly focused on creating a high-quality image from a single plane wave to match a higher quality image created from multiple plane waves.
    2) **Task 1b** gave more freedom to create an image that will be benchmarked against the highest SNR, CNR, gCNR, and contrast. These values can be better than those obtained with multiple plane wave transmissions.
- **Task 2: Beamforming with deep learning after a few plane wave transmissions.** Task 2 imposed a maximum of 10 plane waves but lets participants choose from provided angles to create the best image quality possible.
- **Task 3: Beamforming with deep learning to achieve dynamic transmit focusing.** Task 3 enabled participants to compare the results of a deep learning dynamic transmit focusing implementation that will be useful with current transmit beamforming techniques implemented on most clinical systems today.

B. Datasets

No training data were included because a review of current literature on the topic of deep learning for ultrasound beamforming reveals that there are many different training approaches. For example, in [23] the authors do not use ultrasound images for training, but still arrived at good results when starting the training process with digital photographs. Another approach uses high-quality images constructed from synthetic aperture measurements during training [24]. Other techniques train with channel data input prior to applying receive delays [25], use sub-aperture beamformed data as the network input [23], [26], or only replace portions of the beamforming process [17], [27]–[29]. Additional methods are summarized in a recent review on this topic [9].

Given these multiple training approaches, the organizers decided to keep the training open-ended and focused on challenging participants to produce a network that achieved specific tasks and met specified requirements. We also recognize that generalization of trained networks to multiple ultrasound systems is a key requirement for advancement of this emerging research field. Toward this end, the CUBDL organizers solicited channel data from multiple ultrasound

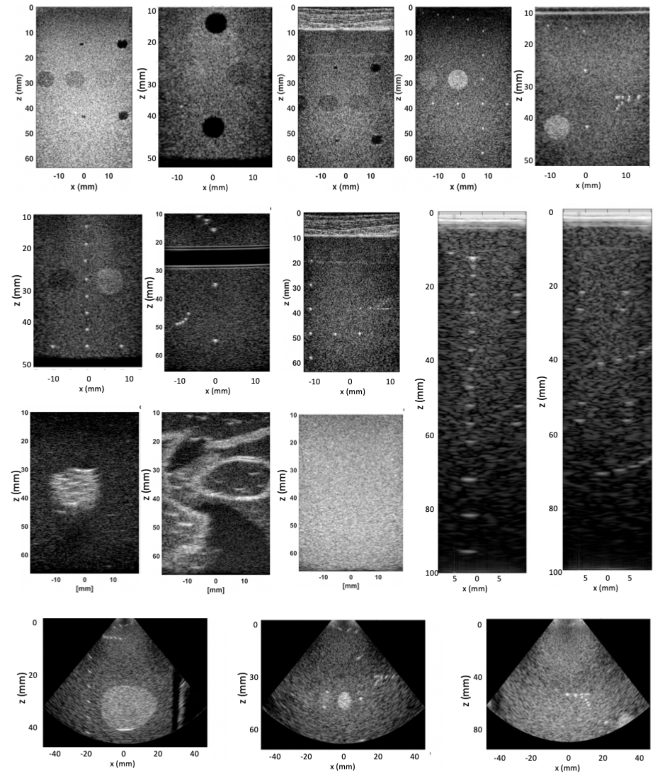

Fig. 1. Example phantom images.

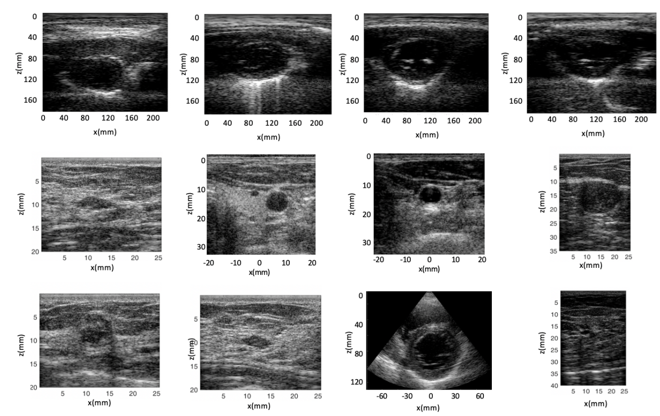

Fig. 2. Example *in vivo* images. A subset of the *in vivo* images in the test set were previously published in [13]–[15], [15]–[20].

groups worldwide to be used as test data for the challenge. The compiled database of test images includes a total of 106 datasets with the following breakdown: (1) 49 experimental phantom data sequences acquired with plane wave transmissions; (2) 25 *in vivo* data sequences of the heart of thirteen patients, the carotid of two healthy volunteers, and the brachioradialis of a healthy volunteer, each acquired with plane or diverging wave transmission; (3) 6 experimental phantom data sequences acquired with focused transmissions; (4) 24 *in vivo* data sequences of the breast of ten patients, the carotid of a healthy volunteer, and the heart of a healthy volunteer,

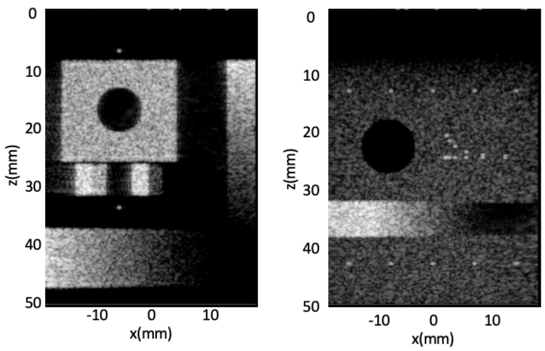

Fig. 3. Example simulation images [21], [22]

each acquired with focused transmissions; and (5) 2 Field II [30], [31] simulations.

The phantom data consisted a total of 12 different phantoms from 6 manufacturers including: (1) CIRS models 040, 049, 054GS, 050 and 059; (2) GAMMEX models 404GSLE, 403TM and 410 SCG; (3) NPL Thermal Test Phantom; (4) CAR Blue Phantom Elastography Breast Model; (5) True Phantom Solutions Brain Phantom; and (6) Dansk Phantom Service Model 453. Three of the phantom acquisitions included a layer of *ex vivo* porcine abdominal tissue to introduce acoustic clutter.

This wide range of channel data was acquired with 4 ultrasound scanners and 8 ultrasound transducers. The acquisition center frequencies ranged from 2.5 MHz to 12.5 MHz. The sampling frequencies ranged from 10 MHz to 78.125 MHz. The ultrasound transducers consisted of linear and phased arrays. Participants were provided with the Data Guide available on the CUBDL website [10], which listed the breakdown of these parameters for each provided dataset.

Data were provided by 9 groups total, referenced hereafter by the short-hand 3-letter code provided in parentheses: (1) Department of Radiology, Mayo Clinic, US (MYO); (2) Microelectronic Systems Design Laboratory, University of Florence, Italy (UFL); (3) Signal Processing Systems group, Eindhoven University of Technology, Netherlands (EUT); (4) CREATIS, Insa Lyon, France (INS); (5) Research Group for Digital Signal Processing and Image Analysis, University of Oslo, Norway (OSL); (6) Ultrasound Elasticity and Imaging Laboratory, Columbia University, USA (COL); (7) Department of Biomedical Engineering, Tsinghua University, China (TSH); (8) Department of Biomedical Engineering, Lund University, Sweden (LUN); and (9) Photoacoustic and Ultrasonic Systems Engineering Lab, Johns Hopkins University, USA (JHU).

In order to provide participants with example images from the test set database, Figs. 1-3 show a sampling from 30% of the above-described data.

C. Evaluation Metrics

The following general metrics apply to Tasks 1-3:

- contrast
- contrast-to-noise ratio (CNR)
- generalized CNR (gCNR) [32], [33]
- speckle signal-to-noise ratio (SNR)
- resolution (measured as the axial and lateral full widths at half maximimum of line profiles through singular point targets)
- network complexity (measured as the total number of trainable parameters in the model).

With Task 1a, we have an additional opportunity to assess performance by matching the images achieved with a high number of multiple plane wave transmissions. Therefore, we additionally assessed the following image-to-image correlation metrics for this task:

- L1 Loss
- L2 Loss
- PSNR
- Normalized Cross-Correlation.

For Tasks 1a, 2, and 3, we were concerned with preserving speckle statistics. Therefore, SNR was measured and participants were ranked based on their ability to preserve the SNR of ground truth measurements. Given the more futuristic outlook of Task 1b, we allowed participants to obtain the highest possible SNR, regardless of speckle preservation.

For Task 3, we proposed to measure the general image quality evaluation metrics (i.e., contrast, CNR, gCNR, SNR, axial and lateral resolution), the preservation of speckle, and the speckle-based resolution both at and away from the transmit focus in order to assess the effectiveness of deep learning-based dynamic transmit focusing.

For each task or subtask, participants were rank ordered using each metric described above and received a rank for each metric. These rankings will be grouped into two categories: (1) image quality and (2) network complexity (because we are interested in balancing image quality with display frame rates). We then averaged the ranks of the metrics obtained by each participant within these two groups. The average rank from each group was summed. This scoring system is represented mathematically as follows:

$$\begin{aligned}\text{Final Score} = &\frac{\sum \text{Image Quality Metric Rankings}}{T_{IQ}} \\ &+ \frac{\sum \text{Network Complexity Metric Rankings}}{T_{NC}},\end{aligned}$$

where T_{IQ} and T_{NC} are the total numbers of image quality metric rankings and network complexity rankings, respectively. To win the challenge, a participant needed to have the lowest final score.

D. Example Code

Example code [12] was provided to show participants how the test data would be organized. Plane wave and focused transmit raw data were provided to the participants as Python classes. We also provided a pixel grid that contained the coordinates for each reconstructed pixel.

Challenge participants were asked to implement models that take the appropriate data object and grid as input and produce the final beamformed image as output (e.g., perform delay-and-sum (DAS) beamforming). As starter code, we provided PyTorch [34] and TensorFlow [35] implementations of DAS that converted the raw data from PICMUS [2] into beamformed images.

Participants were free to incorporate deep learning at any point in the image reconstruction pipeline (e.g., before, during, or after DAS, or replacing DAS entirely). The only requirement was that the final output of the model be image data (or IQ data that was converted into image data, as in the provided example).

We also provided functions for some of the image metrics, although it was up to the challenge participants to select the proper regions of interest (ROIs) in the image to be used with the metrics. Participants were also informed that during model evaluation, ROIs would be specified via the "grid" input.

III. Submission Summary

The CUBDL organizers received a total of four submissions. Three of the submitters participated in Task 1a and one participated in Task 1b, which reduced the number of test datasets for the evaluation of participant results. There were a total of 21 unique datasets forming a total 30 different test cases, submitted by 6 different universities (EUT, INS, MYO, OSL, UFL, TSH).

We asked the first authors of each submission to self-identify their level of experience with beamforming prior to participation in the challenge on a scale of 1 to 5, with 1 being novice and 5 being expert. We received the responses reported in Fig. 4.

These challenge participants are scheduled to present their submitted networks at the virtual IEEE International Ultrasonics Symposium (IUS), through a series of live poster and recorded oral presentations. The winner (to be announced at the virtual symposium) will receive a cash prize (sponsored by Verasonics, Inc.).

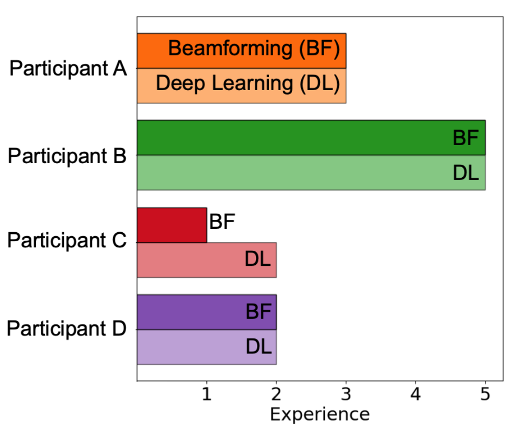

Fig. 4. Beamforming and deep learning experience levels of challenge participants prior to CUBDL (1=novice, 5=expert).

IV. CUBDL-Related Resources

In addition to evaluating and comparing submissions using the metrics described in Section II-C, one additional goal of CUBDL was to create a database that the ultrasound and deep learning community can use for future work. Toward this end, we share the following websites containing the challenge description, test data, and evaluation code:

- **CUBDL Website:** https://cubdl.jhu.edu/. This website contains the challenge description and is the official source of all resources for the challenge, including the following two links. Future updates to data or code availability will be advertised on this website. This website also contains a link to literature references that focus on applications of deep learning in ultrasound systems [10].
- **IEEE DataPort:** https://ieee-dataport.org/competitions/challenge-ultrasound-beamforming-deep-learning-cubdl. This website was used by data contributors to upload their datasets, by participants to submit results, and by the organizers to download and share datasets and submissions for evaluation [11].
- **CUBDL Code:** https://gitlab.com/dongwoon.hyun/cubdl. This website shares example code to input network weights for beamforming evaluation, as well as our evaluation code [12].

We also intend to release all unrestricted test datasets with the publication of a journal paper describing the top submissions, with the ultimate goal of providing useful reference benchmark datasets and networks for future follow-up work. These releases will be available through one or more of the above sites.

V. Conclusion

Deep learning is revolutionizing many fields, and we believe that it will similarly impact ultrasound imaging. CUBDL has therefore been set up as an inclusive challenge, with the aim of advancing the state of the art of ultrasound beamforming by deep learning – driven by our community – and facilitating the debate and discussion associated with it. The tools, resources, and tasks developed and designed for this challenge are recommended as a benchmark standard for both beginners and experts in this research area going forward.

Acknowledgments

We thank our sponsors (Verasonics, Inc., IEEE, and IUS) for their generous support of this challenge, as well as the many individuals and groups who contributed data to make this challenge possible. The following individuals provided data from their institutions: Alycen Wiacek, Jiaqi Huang, and Muyinatu Bell from Johns Hopkins University; Ping Gong and Shigao Chen from Mayo Clinic; Ben Luijten and Massimo Mischi from Eindhoven University of Technology; Xi Zhang and Jiawen Luo from Tsinghua University; Magnus Cinthio from Lund University; Vincent Perrot and Hervé Liebgott from Insa Lyon; Ole Marius Hoel Rindal from University of Oslo; Alessandro Ramalli and Piero Tortoli from University of

Florence; Julien Grondin and Elisa Konofagou from Columbia University. We also thank Olivier Bernard for creating data sequence OSL010 described in the CUBDL Data Guide available on the challenge website [10] and Hervé Liebgott for his helpful advice. M.A.L.B. additionally acknowledges NIH Trailblazer Award R21 EB025621 for partial support of her time on this project.

References

[1] O. Couture, M. Fink, and M. Tanter, "Ultrasound contrast plane wave imaging," *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 59, no. 12, pp. 2676–2683, 2012.

[2] H. Liebgott, A. Rodriguez-Molares, F. Cervenansky, J. A. Jensen, and O. Bernard, "Plane-wave imaging challenge in medical ultrasound," in *2016 IEEE International Ultrasonics Symposium (IUS)*. IEEE, 2016, pp. 1–4.

[3] J.-Y. Lu, H. Zou, and J. F. Greenleaf, "Biomedical ultrasound beam forming," *Ultrasound in Medicine & Biology*, vol. 20, no. 5, pp. 403–428, 1994.

[4] K. E. Thomenius, "Evolution of ultrasound beamformers," in *1996 IEEE Ultrasonics Symposium. Proceedings*, vol. 2. IEEE, 1996, pp. 1615–1622.

[5] J.-F. Synnevag, A. Austeng, and S. Holm, "Benefits of minimum-variance beamforming in medical ultrasound imaging," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 56, no. 9, pp. 1868–1879, 2009.

[6] M. A. Lediju, G. E. Trahey, B. C. Byram, and J. J. Dahl, "Short-lag spatial coherence of backscattered echoes: Imaging characteristics," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 58, no. 7, pp. 1377–1388, 2011.

[7] T. Chernyakova, D. Cohen, M. Shoham, and Y. C. Eldar, "iMAP Beamforming for High-Quality High Frame Rate Imaging," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 66, no. 12, pp. 1830–1844, 2019.

[8] J. A. Jensen, S. I. Nikolov, K. L. Gammelmark, and M. H. Pedersen, "Synthetic aperture ultrasound imaging," *Ultrasonics*, vol. 44, pp. e5–e15, 2006.

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

[10] "Challenge on Ultrasound Beamforming with Deep Learning (CUBDL)." [Online]. Available: https://cubdl.jhu.edu/

[11] "Challenge on Ultrasound Beamforming with Deep Learning (CUBDL)," 2019. [Online]. Available: http://dx.doi.org/10.21227/f0hn-8f92

[12] [Online]. Available: https://gitlab.com/dongwoon.hyun/cubdl

[13] J. Grondin, V. Sayseng, and E. E. Konofagou, "Cardiac strain imaging with coherent compounding of diverging waves," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 64, no. 8, pp. 1212–1222, 2017.

[14] A. Wiacek, O. M. H. Rindal, E. Falomo, K. Myers, K. Fabrega-Foster, S. Harvey, and M. A. Lediju Bell, "Robust short-lag spatial coherence imaging of breast ultrasound data: Initial clinical results," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 66, no. 3, pp. 527–540, March 2019.

[15] A. Wiacek, E. Falomo, K. Myers, O. Hoel Rindal, K. Fabrega-Foster, S. Harvey, and M. A. L. Bell, "Clinical feasibility of coherence-based beamforming to distinguish solid from fluid hypoechoic breast masses," *IEEE International Ultrasonics Symposium (IUS)*, 2018.

[16] A. Wiacek, E. Oluyemi, K. Myers, L. Mullen, and M. A. L. Bell, "Coherence-based beamforming increases the diagnostic certainty of distinguishing fluid from solid masses in breast ultrasound exams," *Ultrasound in Medicine and Biology*, vol. 46, no. 6, pp. 1380 – 1394, 2020.

[17] A. Wiacek, E. González, and M. A. L. Bell, "CohereNet: A Deep Learning Architecture for Ultrasound Spatial Correlation Estimation and Coherence-Based Beamforming," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 2020.

[18] O. M. H. Rindal, S. Aakhus, S. Holm, and A. Austeng, "Hypothesis of improved visualization of microstructures in the interventricular septum with ultrasound and adaptive beamforming," *Ultrasound in Medicine & Biology*, vol. 43, no. 10, pp. 2494 – 2499, 2017.

[19] T. Erlöv, R. Larsson, E. Boni, A. Ramalli, R. Ahlgren, and M. Cinthio, "Improved tracking performance in high frame rate imaging using iterative phase tracking," in *2019 IEEE International Ultrasonics Symposium (IUS)*, Oct 2019, pp. 2158–2161.

[20] X. Zhang, J. Li, Q. He, H. Zhang, and J. Luo, "High-quality reconstruction of plane-wave imaging using generative adversarial network," in *2018 IEEE International Ultrasonics Symposium (IUS)*, Oct 2018, pp. 1–4.

[21] O. M. H. Rindal, A. Austeng, A. Fatemi, and A. Rodriguez-Molares, "The effect of dynamic range alterations in the estimation of contrast," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 66, no. 7, pp. 1198–1208, July 2019.

[22] PICMUS, "Plane-wave imaging evaluation framwork for medical ultrasound," 2020. [Online]. Available: https://www.creatis.insa-lyon.fr/EvaluationPlatform/picmus/

[23] D. Hyun, L. L. Brickson, K. T. Looby, and J. J. Dahl, "Beamforming and speckle reduction using neural networks," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 66, no. 5, pp. 898–910, 2019.

[24] D. Perdios, A. Besson, F. Martinez, M. Vonlanthen, M. Arditi, and J.-P. Thiran, "On problem formulation, efficient modeling and deep neural networks for high-quality ultrasound imaging: Invited presentation," in *2019 53rd Annual Conference on Information Sciences and Systems (CISS)*. IEEE, 2019, pp. 1–4.

[25] A. A. Nair, T. D. Tran, A. Reiter, and M. A. L. Bell, "A deep learning based alternative to beamforming ultrasound images," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 3359–3363.

[26] A. A. Nair, K. N. Washington, T. D. Tran, A. Reiter, and M. A. L. Bell, "Deep learning to obtain simultaneous image and segmentation outputs from a single input of raw ultrasound channel data," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 2020.

[27] A. C. Luchies and B. C. Byram, "Deep neural networks for ultrasound beamforming," *IEEE Transactions on Medical Imaging*, vol. 37, no. 9, pp. 2010–2021, 2018.

[28] B. Luijten, R. Cohen, F. J. De Bruijn, H. A. Schmeitz, M. Mischi, Y. C. Eldar, and R. J. Van Sloun, "Adaptive ultrasound beamforming using deep learning," *IEEE Transactions on Medical Imaging*, 2020.

[29] A. Wiacek, E. Gonzalez, N. Dehak, and M. A. L. Bell, "CohereNet: A deep learning approach to coherence-based beamforming," in *2019 IEEE International Ultrasonics Symposium (IUS)*. IEEE, 2019, pp. 287–290.

[30] J. Jensen, "Field: A program for simulating ultrasound systems," *Medical & Biological Engineering & Computing*, vol. 34, no. sup. 1, pp. 351–353, 1997.

[31] J. A. Jensen and N. B. Svendsen, "Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 39, no. 2, pp. 262–267, March 1992.

[32] A. Rodriguez-Molares, O. M. H. Rindal, J. D'hooge, S.-E. Måsøy, A. Austeng, and H. Torp, "The generalized contrast-to-noise ratio," in *2018 IEEE International Ultrasonics Symposium (IUS)*. IEEE, 2018, pp. 1–4.

[33] A. Rodriguez-Molares, O. M. H. Rindal, J. D'hooge, S.-E. Måsøy, A. Austeng, M. A. L. Bell, and H. Torp, "The generalized contrast-to-noise ratio: a formal definition for lesion detectability," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 67, no. 4, pp. 745–759, 2019.

[34] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds. Curran Associates, Inc., 2019, pp. 8024–8035. [Online]. Available: http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf

[35] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser *et al.*, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from tensorflow.org. [Online]. Available: http://tensorflow.org/