Extending CohereNet to Retain Physical Features when Classifying Benign or Malignant Breast Masses

Alycen Wiacek*, Najim Dehak*, and Muyinatu A. Lediju Bell*†‡
*Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, MD
†Department of Biomedical Engineering, Johns Hopkins University, Baltimore, MD
‡Department of Computer Science, Johns Hopkins University, Baltimore, MD

***Abstract*—Breast ultrasound is often used as a supplement to mammography. However, standard breast ultrasound imaging often has a high false positive rate which limits its use as a screening tool. The spatial coherence of ultrasound signals has particularly impactful implications for breast mass diagnoses when considering both quantitative ultrasound and coherence-based ultrasound beamforming techniques. However, coherence can be computationally expensive to compute. Our recent work demonstrated a novel deep neural network architecture, named CohereNet, which has the ability to estimate coherence features by leveraging the universal approximation properties of deep neural networks. The work in this paper extends the CohereNet architecture to perform binary classification and extract unique features from raw ultrasound data that differentiate benign from malignant breast masses. This extended network classified breast masses as benign or malignant with 83% classification accuracy when tested with data that was not used during training. With the CohereNet architecture as the backbone of this extended network, we leverage coherence features during training and ultimately display network-modified coherence functions. Overall, our extended network architecture and training strategy are promising for a new class of deep learning methods that simultaneously extract features from raw ultrasound data, retain physical interpretations, and diagnose breast masses as benign or malignant.**

I. Introduction

Ultrasound is often used as a diagnostic tool to supplement mammography. However, the false positive rate of breast ultrasound can be as high as 93% depending on the type of mass under investigation [1]–[3]. Alternatives to improve this false positive rate often involve examining the features encoded in backscattered ultrasound waves. For example, quantitative ultrasound (QUS) relies on features within raw backscattered ultrasound data to differentiate benign from malignant breast masses using techniques such as the backscatter coefficient (BSC), effective scatterer diameter (ESD), and effective acoustic concentration (EAC), as well as envelope statistics through the fitting of models such as the homodyned-K distribution [4]. In particular, the homodyned-K distribution includes the estimation of a parameter, k, which is a measure of the ratio of the coherent to diffuse signal energy [4]. In one study investigating quantitative ultrasound parameters of *in vivo* breast lesions, this k parameter revealed a statistically significant difference between benign and malignant breast lesions [5].

There are additional coherence-based parameters that can be estimated through the analysis of the spatial coherence of backscattered ultrasound waves to benefit breast mass diagnosis. For example, the short-lag spatial coherence (SLSC) beamformer creates images based on spatial coherence information [6] with one goal of improving the distinction between solid and fluid-filled breast masses [7], [8]. This possibility was further refined by developing a custom clinical overlay [9]. In addition, our group later demonstrated that the distinction between solid and fluid-filled masses can be achieved without requiring image formation or reader input, which was possible with a quantitative coherence-based feature known as lag-one coherence [10].

The combination of previous studies demonstrating multiple uses of spatial coherence information in breast mass diagnosis highlights the overall diagnostic power of this physical feature contained within backscattered ultrasound waves. However, the calculation of spatial coherence from backscattered ultrasound data can be computationally intensive. Therefore, previous work leveraged the universal approximation properties of deep neural networks (DNNs) to improve the computational complexity of computing the coherence function. Specifically, a network architecture named CohereNet was developed to estimate spatial coherence functions while retaining the integrity and benefits of SLSC beamforming [11], [12]. Considering that we successfully developed this deep learning framework to efficiently estimate spatial coherence functions, we believe that there is additional potential to extract useful diagnostic information that combines spatial coherence information, deep learning, and breast mass classification to determine whether a suspicious breast mass is benign or malignant.

The objective of this paper is to extend our deep learning framework to analyze raw backscattered ultrasound data and extract useful features that distinguish benign from malignant breast masses. This new framework leverages coherence features by building on the CohereNet architecture. We then visualize the impact of class labels (i.e., benign or malignant) on what we term a *network-modified coherence function*, which can be thought of as a spatial coherence function that encodes both spatial coherence and the specific diagnosis of the breast mass.

2021 IEEE International Ultrasonics Symposium (IUS) | 978-1-6654-0355-9/21/$31.00 ©2021 IEEE | DOI: 10.1109/IUS52206.2021.9593710

II. Methods

A. Datasets

Seventeen patients with suspicious hypoechoic breast masses were enrolled in our study after informed consent and approval from the Johns Hopkins Medicine Institutional Review Board (IRB). Raw radio-frequency (RF) ultrasound channel data were acquired immediately prior to a scheduled core-needle biopsy using an Alpinion ECUBE12R research ultrasound scanner connected to either an Alpinion L8-17 or Alpinion L3-8 linear ultrasound transducer. These data included two transducer locations per patient (i.e., one radial and one anti-radial) and 20 scans per transducer location.

Similar to the pre-processing performed for the original CohereNet architecture [12], these data were pre-processed by computing the coherence function (i.e., spatial coherence as a function of element spacing):

$$\hat{R}(m) = \frac{1}{N-m} \sum_{i=1}^{N-m} \frac{\sum_k s_{k,i}(n) s_{k,i+m}(n)}{\sqrt{\sum_k s_{k,i}^2(n) \sum_k s_{k,i+m}^2(n)}}, \quad (1)$$

where m is the lag in number of elements, N is the number of receive elements in the transducer, $s_{k,i}(n)$ is a time-delayed, zero-mean kernel consisting of k axial samples, each received at element i and centered at depth n. In addition, due to the sliding window of selected active receive subapertures of the Alpinion linear array, the training data were filtered to only include the center scanlines which contained $N = 64$ active receive elements for each computation. These coherence functions were computed over an axial kernel $k = 7$ samples, resulting in the input to the network, $s_k(n) = \{s_{k,i}(n) \; \forall \; i = 1, ..., N\}$.

To ensure inclusion of pixels from regions associated with the core-needle biopsy diagnosis, a region of interest (ROI) was defined at the center of each solid mass. This ROI definition was performed by visualizing both the reconstructed B-mode image and the annotated clinical screeshot to ensure only pixels within the solid mass were selected. Each pixel in this ROI corresponds to one axial kernel of raw channel data and represents one input to the neural network, $s_k(n)$. The output of the neural network was the binary classification (i.e., benign or malignant), which was defined based on the result of the core-needle biopsy. Each ROI contains approximately 1,000 pixels, resulting in 1,000 input/output pairs per scan. Considering the 17 patients, 20 scans per patients, and 1000 input/output pairs per scan, the complete dataset consisted of approximately 340,000 input/output pairs.

This dataset of input kernels and output binary labels was then divided into training, validation, and testing set individually by patient. Eleven patients were assigned to the training set, 3 patients were assigned to the validation set, and the remaining 3 patients were assigned to the testing set. Separating the dataset by patient ensured that the features learned by the network were not specific to the patients in the training set.

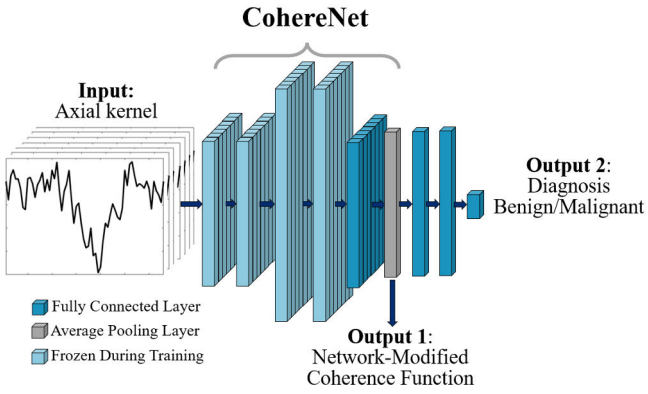

Fig. 1. Extended CohereNet architecture. The gray block indicates the output layer of the original CohereNet architecture which is now the network-modified coherence function.

TABLE I
Extended CohereNet architecture

Layer	Type	Size	Activ.
Input	-	7 x 64	-
1	Fully Connected	7 x 64	ReLU
2	Fully Connected	7 x 128	ReLU
3	Fully Connected	7 x 128	ReLU
4	Fully Connected	7 x 64	Tanh
Output 1	Average Pool	1 x 64	-
5	Fully Connected	1 x 64	ReLU
6	Fully Connected	1 x 64	ReLU
Output 2	Fully Connected	1 x 2	Softmax

B. Network architecture

Extending from the previously demonstrated CohereNet architecture [12], the extended network architecture adds three fully connected layers which offer 8,450 additional trainable parameters to perform binary classification. The complete architecture is shown in Fig. 1, where "Output 1" is the original output of CohereNet and represents the estimated coherence function. "Output 2" represents the additional classification task and is a vector of length two representing the two-class binary classification (i.e., benign or malignant). To prevent overfitting, weights within layers 5 and 6 were randomly set to 0 with a probability of 50% during the training phase.

The network was trained using Keras [13] with a Tensorflow backend [14]. Fine tuning was performed by initializing the network with the optimized weights from CohereNet and freezing the first three fully connected layers (represented by the lighter blue color in Fig 1). The remaining layers were allowed to update through the training process. The network was trained with a batch size of 64, an Adam optimizer [15], and a learning rate of 0.001 over twenty epochs. Training was performed on a system with an Intel Xeon E7 processor and four Tesla P40 GPUs with 24 GB of memory.

III. Results & Discussion

Given an input axial kernel of channel data, the network was able to classify the kernel of raw ultrasound data as benign

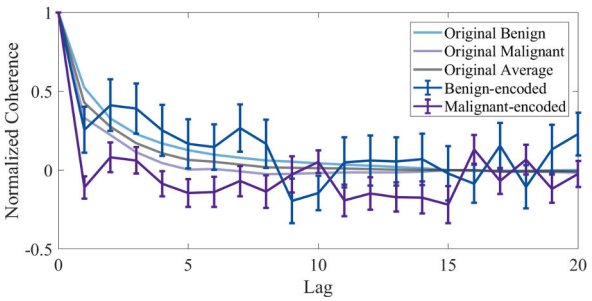

Fig. 2. Original coherence functions (averaged over multiple masses) and mean $\pm$ one standard deviation of network-modified coherence functions for benign (blue) and malignant (purple) masses.

or malignant with an overall classification accuracy of 83% at Output 2. This accuracy is calculated from a test set containing 60 regions of interest within 3 different masses not included during training or validation.

Fig. 2 shows visualizations of the mean $\pm$ one standard deviation of network-modified coherence functions for correctly classified masses obtained at Output 1. These network modified coherence functions were compared to the standard computation for coherence functions (i.e., calculated using Eq. 1) and are referred to as "original coherence functions". These original coherence functions were obtained by averaging over the coherence functions of correctly classified benign and malignant kernels of channel data. In addition, the average of these two original coherence functions is plotted in gray in Fig. 2. In comparison to the original coherence functions, the network-modified coherence functions are distinctly different between benign and malignant masses, particularly at short lag values (i.e., $m < 8$). These network-modified coherence functions also deviate from the respective original coherence functions for both benign and malignant breast masses.

These results indicate that the proposed network architecture and training strategy are promising for the future use of deep learning with raw ultrasound data to retain physical interpretations while diagnosing benign or malignant breast masses. By initializing the network with CohereNet weights, we guided the network to leverage coherence information, which likely aided in differentiating benign from malignant masses. This approach enables us to visualize the network-modified coherence functions, which can be used as a classification tool. Future work will focus on deriving a theoretical framework for the observed differences in benign-encoded and malignant-encoded coherence functions.

IV. Conclusion

This paper is the first to leverage deep learning to extract unique features from raw backscattered ultrasound data with the primary goal of differentiating benign from malignant breast masses. As a secondary goal, we additionally visualized how learned breast mass classification features affect network-modified coherence functions. Overall, these results indicate that the proposed framework and training strategy are promising for the future use of deep learning with raw ultrasound data to retain physical interpretations while diagnosing benign or malignant breast masses.

V. Acknowledgements

Funding for this research was provided by NIH Trailblazer Award R21 EB025621 (awarded to M.A.L.B.). A.W. additionally acknowledges support from the Achievement Rewards for College Scientists (ARCS) Foundation as a 2021-2022 ARCS Scholar.

References

[1] W. A. Berg, J. D. Blume, J. B. Cormack, E. B. Mendelson, D. Lehrer, M. Böhm-Vélez, E. D. Pisano, R. A. Jong, W. P. Evans, M. J. Morton *et al.*, "Combined screening with ultrasound and mammography vs mammography alone in women at elevated risk of breast cancer," *Jama*, vol. 299, no. 18, pp. 2151–2163, 2008.

[2] W. A. Berg, A. I. Bandos, E. B. Mendelson, D. Lehrer, R. A. Jong, and E. D. Pisano, "Ultrasound as the primary screening test for breast cancer: analysis from acrin 6666," *JNCI: Journal of the National Cancer Institute*, vol. 108, no. 4, 2016.

[3] W. A. Berg, A. G. Sechtin, H. Marques, and Z. Zhang, "Cystic breast masses and the acrin 6666 experience," *Radiologic Clinics*, vol. 48, no. 5, pp. 931–987, 2010.

[4] M. L. Oelze and J. Mamou, "Review of quantitative ultrasound: Envelope statistics and backscatter coefficient imaging and contributions to diagnostic ultrasound," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 63, no. 2, pp. 336–351, 2016.

[5] I. Trop, F. Destrempes, M. El Khoury, A. Robidoux, L. Gaboury, L. Allard, B. Chayer, and G. Cloutier, "The added value of statistical modeling of backscatter properties in the management of breast lesions at us," *Radiology*, vol. 275, no. 3, pp. 666–674, 2015.

[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] A. Wiacek, O. M. H. Rindal, E. Falomo, K. Myers, K. Fabrega-Foster, S. Harvey, and M. A. L. 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, 2018.

[8] 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 & Biology*, vol. 46, no. 6, pp. 1380–1394, 2020.

[9] A. Wiacek, E. Falomo, K. Myers, O. M. H. 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," in *2018 IEEE International Ultrasonics Symposium (IUS)*. IEEE, 2018, pp. 1–4.

[10] A. Wiacek, E. Oluyemi, K. Myers, L. Mullen, and M. A. L. Bell, "Coherence-based beamforming improves the diagnostic certainty of breast ultrasound exams," in *2020 IEEE International Ultrasonics Symposium (IUS)*. IEEE, 2020, pp. 1–4.

[11] 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.

[12] 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*, vol. 67, no. 12, pp. 2574–2583, 2020.

[13] F. Chollet *et al.*, "Keras," https://keras.io, 2015.

[14] M. Abadi, A. Agarwal, P. Barham, E. Brevdo *et al.*, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from tensorflow.org. [Online]. Available: http://tensorflow.org/

[15] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.