

BREAST LESION SEGMENTATION IN ULTRASOUND IMAGES WITH LIMITED ANNOTATED DATA

Bahareh Behboodi^{*†} Mina Amiri^{*†} Rupert Brooks^{*‡} Hassan Rivaz^{*†}

^{*}Department of Electrical and Computer Engineering, Concordia University, Canada

[†]PERFORM Center, Concordia University, Canada

[‡]Nuance Communications

ABSTRACT

Ultrasound (US) is one of the most commonly used imaging modalities in both diagnosis and surgical interventions due to its low-cost, safety, and non-invasive characteristic. US image segmentation is currently a unique challenge because of the presence of speckle noise. As manual segmentation requires considerable efforts and time, the development of automatic segmentation algorithms has attracted researchers' attention. Although recent methodologies based on convolutional neural networks have shown promising performances, their success relies on the availability of a large number of training data, which is prohibitively difficult for many applications. Therefore, in this study we propose the use of simulated US images and natural images as auxiliary datasets in order to pre-train our segmentation network, and then to fine-tune with limited *in vivo* data. We show that with as little as 19 *in vivo* images, fine-tuning the pre-trained network improves the dice score by 21% compared to training from scratch. We also demonstrate that if same number of natural and simulation US images is available, pre-training on simulation data is preferable.

Index Terms— Segmentation, simulation, U-Net, fine-tuning

1. INTRODUCTION

Breast cancer has been reported as one of the leading causes of death among women worldwide. Although, digital mammography is an effective modality in breast cancer detection, it has limitations in detecting dense lesions which are similar to dense tissues [1], and further uses ionizing radiation. Therefore, ultrasound (US) imaging as a safe and versatile screening and diagnostic modality plays an important role in this regard. However, due to contamination of the US images with speckle noise, US images have low resolution and poor contrast between the target tissue and background; thus, their segmentation is currently a challenging task [2]. Researchers have utilized recent state-of-the-art deep learning techniques in order to overcome limitations in manual segmentation. Despite the success of deep learning techniques in computer vision tasks, their performance depends on the size of input data

which is limited specially in medical US images. The collection and annotation of US images require considerable effort and time which attain the need to a deep learning-based strategy that can be trained on as few annotated data as possible.

The U-Net architecture [3], as one of the most well-known networks for segmentation purposes, is built upon fully convolutional network. It involves several convolutional, max-pooling, and up-sampling layers. To cope with limited input data for training U-Net, researches have proposed various strategies based on data augmentation and transfer learning [2, 4, 5]. Data augmentation cannot truly capture the characteristics of the real data when very limited data is available. To this end, we propose a methodology based on transfer learning which utilizes US simulation data and natural images as an auxiliary dataset. The goal is to enhance the segmentation results while only few images are available. In our work, first, we pre-train the U-Net with US simulated and natural images separately, and then fine-tune the network with only 15% of available *in vivo* images. We demonstrate 21% improvement in segmentation results when small number of images are available.

2. METHOD

In deep learning approaches, the improvement in results depends on the number of training data. Therefore, such techniques perform better if they have larger amount of training data. In medical images, especially in US images, annotating enough number of training data is expensive, and thus, we take advantage of simulated US data as well as natural images as the auxiliary datasets for pre-training U-Net in our proposed workflow. To that end, our proposed workflow consists of three avenues as shown in Fig. 1. In the first avenue, U-Net is trained using only 15% of the *in vivo* dataset. In the second avenue, U-Net is first pre-trained on the simulated data, and then fine-tuned using the same 15% of the *in vivo* dataset which was used in the first avenue. And the last avenue is similar to the second avenue with the difference that natural images were used for pre-training. Section 2.5 will clarify each avenue in more details.

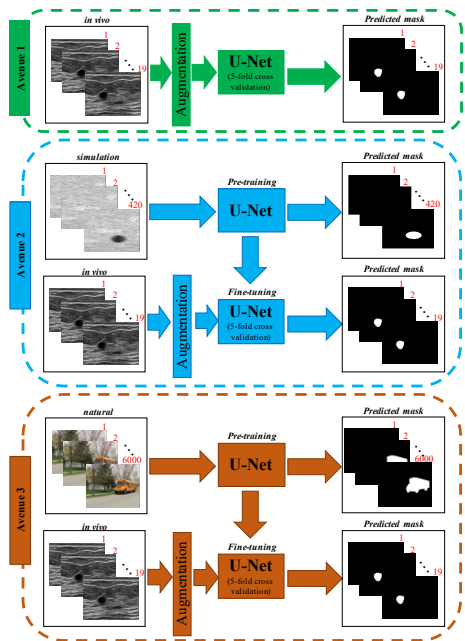


Fig. 1: Proposed workflow for training U-Net when limited annotated data is available.

2.1. *In vivo* Data

In vivo dataset includes 163 breast B-mode US images with lesions and the mean image size of 760×570 . The images as well as their delineation of lesions are publicly available upon request [1]. The breast lesions of interest are generally hypoechoic (i.e. tissues with lower echogenicity), that is, darker than surrounding tissue. Only 15% of the total number of *in vivo* images were used as training and validation datasets and the remaining 85% were set as the testing datasets. The size of training dataset was selected 4 times larger than the size of validation images yielding 19, 5, and 139 images for training, validation, and testing datasets, respectively.

2.2. Simulation Data

To simulate B-mode images, a MATLAB-based publicly available US simulation software, Field_II [6] was used. Number of RF lines, centre frequency, sampling frequency, and speed of the sound were respectively set to 50, $3.5MHz$, $100MHz$, and $1540m/s$. In our simulation phantom, the surface started at $30mm$ from the transducer surface and the axial, lateral, and elevational distances were initiated as $60mm$, $40mm$, and $10mm$, respectively. The scatterers were randomly distributed in our virtual phantom such that each mm^3 of phantom had in average 4 scatterers, to allow for fast ultrasound simulation. In our simulated images we considered each image to randomly have either hyperechoic (i.e. tissues with higher echogenicity), that is brighter than surrounding tissue, or hypoechoic lesions, or both at the same

time in order to let our network learn better the various possible textures of the US images. The intensities for hyperechoic lesions were set k times higher than the background where k was an integer in range of $1 - 10$, however, for hypoechoic lesions the intensities were set l times the background where l was a random variable between 0 and 1. The location of the lesions was randomly selected with circle or ellipsoid shapes. A total of 700 images were simulated and separated to training, validation, and testing sets with splitting factors of 60%, 15%, and 25% of total number of images, yielding 420, 105, and 175 images, respectively. It worth mentioning that as the *in vivo* data consisted of hypoechoic lesions, in the masks of simulated data only the pixels inside of the hypoechoic lesions were set to 1, and the remaining pixels were set to 0. Therefore, there were some simulated images with zero segmented lesions in their masks.

2.3. Natural Data

The natural images are publicly available at [7]. The dataset consists of 10000 images of salient objects with their annotations. In our work, the dataset was split to training, validation, and testing sets with splitting factors of 60%, 15%, and 25% of total number of images, yielding 6000, 2500, and 1500 images, respectively.

2.4. U-Net Architecture

The U-Net structure previously proposed by [3] utilizes several conv-block, max-pooling, up-sampling, and skip connection layers as illustrated in Fig. 2. Each conv-block consists of repetition of two convolution layers while in the contraction and expansion paths, followed by max-pooling and up-sampling layers, respectively. In this work, the kernel sizes of convolution, max-pooling, and up-sampling layers were set to 3×3 , 2×2 , and 2×2 , respectively. As a pre-processing step, all the images were resized to 388×388 , mirrored with the mirroring factor of 92 pixels, yielding images with size 572×572 , and normalized to the range of $[0, 1]$. Thus, the size of the input and output data was $(batch, 572, 572, 1)$ and $(batch, 388, 388, 2)$, respectively, where *batch* indicates the number of images in each batch. The activation and loss functions, optimizer, learning rate, number of epochs, batch size, weight initializer, and kernel regularizer were initialized as stated in Table 1. The Dice score is defined as $DSC = \frac{2|G \cap P|}{|G| + |P|}$, where G and P is ground truth and predicted masks, respectively.

2.5. Experiments

As previously mentioned, in this work we propose three avenues to study the impact of simulated and natural images as the auxiliary datasets for US segmentation (see Fig. 1). In the following paragraphs, we will explain each avenue in detail:

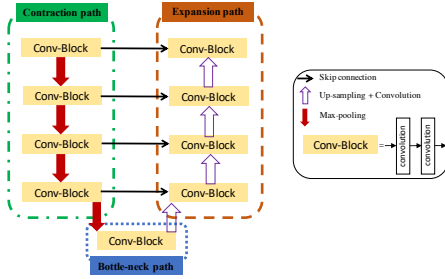


Fig. 2: U-Net structure with its contraction, bottle-neck, and expansion paths.

Table 1: U-Net parameters

| Parameter | Value |
|---|----------------|
| Activation function (except last layer) | ReLU [8] |
| Activation function (last layer) | Softmax |
| Loss function | Dice score |
| Optimizer | Adam [9] |
| Learning rate | 0.00001 |
| No. of epochs | 150 |
| Batch size | 8 |
| Weight initializer | He-normal [10] |
| Kernel-regularizer | L2-norm |

2.5.1. Avenue 1: Train U-Net on *in vivo* data from scratch

In the first avenue, the U-Net structure with above-mentioned parameters, was trained on *in vivo* images from scratch using 19 and 5 images as training and validation sets, respectively, and was tested on 139 images. We call this trained network as *Pt_in vivo*. Due to small number of training data, we used 5-fold cross-validation to prevent variation in performance. Prior to each optimization iteration, we performed "on-the-fly" augmentation by applying random height-shift, width-shift, and zooming.

2.5.2. Avenue 2: Train U-Net with US simulated images and fine-tune with *in vivo* data

In this avenue, U-Net was first trained using 420 and 105 simulation images as its training and validation sets, respectively. Similar to the first avenue, the U-Net was initialized using parameters mentioned in Table 1. For the simplicity, we refer the trained U-Net with simulated data as *Pt_sim*. Afterwards, the contraction path of *Pt_sim* was fine-tuned on *in vivo* training and validation sets based on parameters in Table 1 except that weights were initialized using the *Pt_sim* weights. We call the fine-tuned network as *Ft_sim_in vivo* which was tested on *in vivo* test set. 5-fold cross-validation and "on-the-fly" augmentation was used for fine-tuning our *Ft_sim_in vivo* network.

2.5.3. Avenue 3: Train U-Net with natural images and fine-tune with *in vivo* data

In this step, similar to the *Avenue 2* described above, U-Net was first pre-trained and then fine-tuned on *in vivo*. However, for pre-training the network we used 6000 and 2500 natural images as training and validation sets, respectively. For simplicity, the pre-trained U-Net with natural images is referred as *Pt_nat* and the fine-tuned network using *Pt_nat* is referred as *Ft_nat_in vivo*. 5-fold cross-validation and "on-the-fly" augmentation was used in the fine-tuning step.

3. RESULTS

3.1. Evaluation Criteria

In this work, we used Dice Similarity Coefficient (*DSC*) as our evaluation criteria. It is worth noticing that we also used *DSC* as our loss function, however, in the evaluation step the predicted masks which were the output of the last layer (i.e. Softmax layer) were first binarized using *argmax* function, and then compared with the ground truth masks. As our dataset was unbalanced (i.e. number of background pixels were higher than the lesion pixels), we only report the *DSC* scores of the foreground (i.e. lesions) masks ignoring the *DSC* score of the background.

3.2. Experimental Results

Table 2 presents the *DSC* scores of the predicted masks derived from *Pt_in vivo*, *Pt_sim*, *Ft_sim_in vivo*, *Pt_nat*, and *Ft_nat_in vivo* networks for both training and testing *in vivo* sets. The *DSC* score for test set increases when we fine-tune the pre-trained network no matter what type of images were used during pre-training. Therefore, pre-training the network performs better than training from scratch with limited training data. It is worth mentioning that we used 6000 number of natural images and 420 number of simulated images during pre-training. However, when we decreased the number of natural images in *Avenue 3* from 6000 to 420 in order to be equal to the number of simulated images, the *DSC* score was reduced from 0.56 to 0.38 as shown in Table 2 (*Pt_nat420*, and *Ft_nat420_in vivo* are referred as repetition of *Avenue 3* using 420 natural images). As a result, pre-training the network using simulated data is preferable as the auxiliary dataset than using natural images when same number of images from both datasets is available. Figure 3 demonstrates examples of the predicted masks with their *DSC* scores.

We had 6000 natural images in which 29 hours was needed to train the *Ft_nat_in vivo* network. For training on simulation, 2 hours, and for training/fine-tuning on *in vivo*, 5 minutes was needed. As more annotations become available, although the U-Net is better trained, more time is needed for the pre-training step.

Table 2: Mean and standard deviation of *DSC* scores for predicted masks of *in vivo* train and test sets over 5-fold cross-validation

| Network Name | | Train <i>in vivo</i> | Test <i>in vivo</i> |
|--------------------------|-----------------------|----------------------|-----------------------------------|
| <i>Avenue 1</i> | <i>Pt_in vivo</i> | 0.73 ± 0.03 | 0.37 ± 0.04 |
| <i>Avenue 2</i> | <i>Pt_sim</i> | 0.29 ± 0.22 | 0.27 ± 0.19 |
| | <i>Ft_sim_in vivo</i> | 0.79 ± 0.05 | 0.45 ± 0.03 |
| <i>Avenue 3</i> | <i>Pt_nat</i> | 0.13 ± 0.27 | 0.14 ± 0.26 |
| | <i>Ft_nat_in vivo</i> | 0.85 ± 0.04 | 0.57 ± 0.02 |
| <i>Ft_nat420_in vivo</i> | | 0.78 ± 0.15 | 0.40 ± 0.03 |

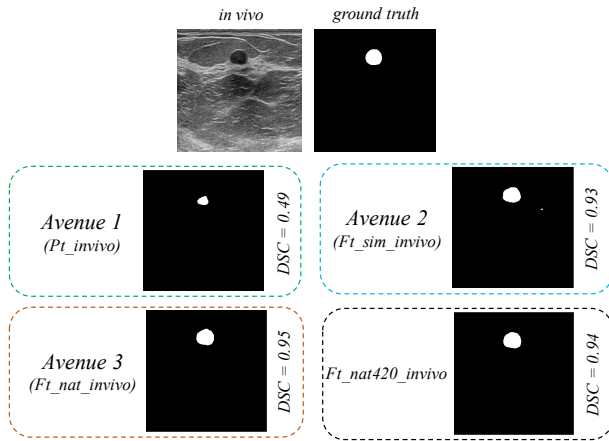


Fig. 3: Examples of segmentation results and their *DSC* scores derived from *Avenue 1*, *Avenue 2*, *Avenue 3*, and *Ft_nat420_in vivo*.

4. CONCLUSION AND FUTURE WORK

In this work, we showed that pre-training the network performs better than training the network from scratch especially when the number of annotations is limited. We proposed the use of simulated US images as the auxiliary dataset for pre-training the network. In addition, we confirmed that natural images can be also considered as the auxiliary dataset, however, thousands of them are required for optimum results which led to hours of pre-training. Therefore, we conclude that US simulation images are the preferred auxiliary dataset for pre-training the network. As our future work, we will validate our strategy for different type of *in vivo* datasets and segmentation applications such as prostate cancer and muscle segmentation.

5. ACKNOWLEDGEMENT

This research was funded by Richard and Edith Strauss Foundation and by NSERC Discovery Grant RGPIN 04136. The authors would like to thank NVIDIA for donating the GPU.

6. REFERENCES

- [1] MH Yap, G Pons, J Martí, S Ganau, M Sentís, R Zwigelaar, A K Davison, and R Martí, “Automated breast ultrasound lesions detection using convolutional neural networks,” *IEEE Journal of Biomedical and Health Informatics*, vol. 22, no. 4, pp. 1218–1226, 2017.
- [2] S Liu, Y Wang, X Yang, B Lei, L Liu, SX Li, D Ni, and T Wang, “Deep learning in medical ultrasound analysis: A review,” *Engineering*, 2019.
- [3] O Ronneberger, P Fischer, and T Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical Image Computing and Computer-assisted Intervention*. Springer, 2015, pp. 234–241.
- [4] S Pereira, A Pinto, V Alves, and CA Silva, “Brain tumor segmentation using convolutional neural networks in mri images,” *IEEE Transactions on Medical Imaging*, vol. 35, no. 5, pp. 1240–1251, 2016.
- [5] HC Shin, NA Tenenholtz, JK Rogers, CG Schwarz, ML Senjem, JL Gunter, KP Andriole, and M Michalski, “Medical image synthesis for data augmentation and anonymization using generative adversarial networks,” in *International Workshop on Simulation and Synthesis in Medical Imaging*. Springer, 2018, pp. 1–11.
- [6] JA Jensen, “Field: A program for simulating ultrasound systems,” in *10th Nordicbaltic Conference on Biomedical Imaging, Vol. 4, Supplement 1, Part 1: 351–353*. Citeseer, 1996.
- [7] C Xia, J Li, X Chen, A Zheng, and Y Zhang, “What is and what is not a salient object? learning salient object detector by ensembling linear exemplar regressors,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 4142–4150.
- [8] V Nair and GE Hinton, “Rectified linear units improve restricted boltzmann machines,” in *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, 2010, pp. 807–814.
- [9] DP Kingma and J Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [10] K He, X Zhang, S Ren, and J Sun, “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 1026–1034.