# Hybrid Deep Neural Network for Brachial Plexus Nerve Segmentation in Ultrasound Images

Juul van Boxtel<sup>‡</sup> Department of Biomedical Engineering Eindhoven University of Technology Eindhoven, the Netherlands j.p.a.v.boxtel@student.tue.nl

Josien Pluim Department of Biomedical Engineering Eindhoven University of Technology Eindhoven, the Netherlands J.Pluim@tue.nl

Abstract-Ultrasound-guided regional anesthesia (UGRA) can replace general anesthesia (GA), improving pain control and recovery time. This method can be applied on the brachial plexus (BP) after clavicular surgeries. However, identification of the BP from ultrasound (US) images is difficult, even for trained professionals. To address this problem, convolutional neural networks (CNNs) and more advanced deep neural networks (DNNs) can be used for identification and segmentation of the BP nerve region. In this paper, we propose a hybrid model consisting of a classification model followed by a segmentation model to segment BP nerve regions in ultrasound images. A CNN model is employed as a classifier to precisely select the images with the BP region. Then, a U-net or M-net model is used for the segmentation. Our experimental results indicate that the proposed hybrid model significantly improves the segmentation performance over a single segmentation model.

*Index Terms*—medical imaging, brachial plexus, deep learning, convolutional neural network, segmentation

## I. INTRODUCTION

The brachial plexus (BP) is a part of the cervical nerves which originates from the spinal cord. It is partially located in the neck and partially in the axilla. It is an important part of the nervous system, as it innervates the upper limb. The BP contains cervical nerves C5-C8 and most of the thoracic T<sub>1</sub> nerves, often along with fibers from C<sub>4</sub> and/or T<sub>2</sub> [1]. The segmentation of BP regions is an important consideration in treatment planning for lung or head-and-neck cancer patients, as radiotherapy used to treat cancer patients can induce brachial plexopathy when the BP is overburdened with radiation energy, causing severe and irreversible effects [2]. Furthermore, the BP is used more for regional anesthesia (RA) in clavicular surgeries, in the form of interscalene BP blocks. It can replace general anesthesia, allowing better pain control, decreasing opioid consumption, and reducing recovery time [3], [4]. The classic RA procedure is to inject the anesthesia into the target nerve region rather blindly. This method poses

Vincent Vousten<sup>‡</sup> Department of Biomedical Engineering Eindhoven University of Technology Eindhoven, the Netherlands v.r.j.vousten@student.tue.nl

Nastaran Mohammadian Rad Department of Biomedical Engineering Eindhoven University of Technology Eindhoven, the Netherlands n.mohammadian.rad@tue.nl

risks as block failure, nerve trauma, and local anesthesia toxicity [5]. Ultrasound-guided RA (UGRA) has become a popular method to visualize this procedure. UGRA is often faster, requires fewer attempts, and in some cases provides a better sensory block compared to other RA techniques [6].

In general, tissue segmentation in ultrasound (US) images is a challenging task due to low contrast between background and the tissue, compared to other modalities like MRI and CT scans. Especially the segmentation of nerve tissue can be challenging due to speckled noise coming with the US modality and the fact that the nerve region does not form a salient structure in the images [7]. The need for trained experts, capable of recognizing the region of interest (ROI), limits the applicability of the UGRA procedure. Deep learning algorithms could provide a solution by automating the recognition of the ROI. This would make UGRA better available for more extensive use.

Deep learning methods have shown promising results in image segmentation applications [8]. Deep learning algorithms, and more specifically convolutional neural networks (CNNs) and U-net are widely used in medical applications such as tumor region identification and metastasis detection [9], [10]. CNNs are suitable to be applied to data that have a grid-like topology such as time-series and images [11]. A typical CNN has a hierarchical architecture that alternates a convolution layer, a rectified linear unit (ReLU) activation function, and a pooling layer to summarize the large input spaces into a lower-dimensional feature space. CNN solutions are among the best-performing systems on pattern recognition systems [11].

U-net [12] is the leading model architecture for medical image segmentation. The model consists of an encoder and a decoder part. The encoder part, the left side of the U-shape, resembles a normal CNN. It consists of subsequent convolution, activation, and pooling layers. The decoder part, the right side of the U-shape, is symmetrical to the encoder part.

<sup>&</sup>lt;sup>‡</sup>These authors contributed equally to this work.



(b) Segmentation model (U-net/M-net)

Fig. 1: The architecture of (a) CNN for the classification and (b) U-net and M-net for the segmentation.

It consists of upsampling layers, a concatenating layer which adds the feature map of the encoder layer, and subsequent convolutional layers. The U-net architecture has proven to be able to accurately localize ROIs, even when only trained on relatively small training sets [12]. Over the years, a variety of adaptations have been made to the U-net structure in attempts to improve its performance on specific datasets. A promising architecture is M-net [13]. It uses multiple scaled inputs and multiple outputs coming from the different layers, in order to densely supervise extracted feature maps.

The goal of this paper is to investigate the effect of implementing a hybrid model in segmenting the BP nerve region from US images. We hypothesize that a combination of CNN and a segmentation model provides more accurate BP segmentation compared to using only a segmentation model.

The rest of this paper is structured as follows: Section II presents our hybrid deep neural network architecture. The experimental materials and the procedures are described in Section III. The results of experiments are described and discussed in Sections IV and V. Section VI concludes this paper by summarizing our achievements.

#### II. METHODS

In this paper, we used a CNN architecture comprised of four convolutional layers. The architecture is shown in Figure 1a. Each convolutional layer consists of two sliding windows with a 3x3 kernel size and ReLU activation, followed by a 2x2 max-pooling layer. In the first layer, 32 filters were used, and this number was multiplied by a factor two every layer down. The output was then flattened and followed by two fully-connected layers. After the first fully-connected layer, a dropout layer with a dropout fraction of 50% was used.

The first segmentation model used was an adaptation of the U-net architecture [12]. The encoder part of the model consists of five convolutional layers, each consisting of two 5x5 sliding windows with ReLU activation, each followed by a GroupNormalization layer with group size 8. In the first layer, the number of filters is 32. This was multiplied by a factor two every layer down. A dropout layer with a dropout fraction of 40% was used inside the convolutional block, followed by an average-pooling layer. The decoder uses the same convolutional blocks, however without the dropout layer. Between the blocks, up-sampling is done by using transposed convolutions. The final output is given through a 1x1 convolution with Sigmoid activation.

The other model used for segmentation is the M-net architecture [7]. M-net is an extension of the U-net architecture, as displayed in Figure 1b. The M-net structure uses multiple inputs, which are down-sampled versions of the original input by factor two, four, and eight. These inputs are sent in the

TABLE I: The value of parameters for CNN, U-net, and M-net architectures.

Parameter	CNN	U-net	M-net	_
Optimizer	Adam	Adam	Adam	_
Batch Size	64	16	16	
Learning rate	10 <sup>-4</sup> (1)	10-4 (2)	10 <sup>-4</sup> (2)	
No. epochs	25	50	50	
Start no. filters	32	32	32	
Kernel size	3x3	5x5	3x3	
Dropout	0.5	0.4	0.2	

<sup>1</sup> The CNN optimization algorithm used the 'Reduce Learning Rate on Plateau' callback in the Keras API; when the validation accuracy would not improve for 3 consecutive epochs, the learning rate would be lowered by a factor 0.2. The value in the table represents the starting learning rate.

 $^2$  These models used learning rate decay. The LR of U-net was divided by ten at epoch ten, and then after each fifth epoch. The same goes for M-net, but then for every tenth epoch. The value in the table represents the starting learning rate.

TABLE II: Number of images in BP and no-BP classes before data augmentation.

	Non-filtered	Filtered	
BP	2322	1454	
No-BP	3313	2452	

subsequent layers from the encoder part of the U-net structure, by concatenating them with the feature map of the previous layer. In the decoder part, each layer creates an output image by sending the feature map through a 1x1 convolution with Sigmoid activation, and then up-sampling that output to match the shape of the original input. All four output images are then averaged to form a single final segmentation map. The inner (U) structure is similar to U-net, starting with 32 filters in the first layer. However, a dropout fraction of 20% is used and the normalization layers are left out. In the decoder, up-sampling is done by bilinear up-sampling. The hyperparameters used are given in Table I.

Keras library [14] was used in our implementation. The model was trained on a dedicated GPU server, containing an Intel Core i7-5930K CPU, 2 NVIDIA GeForce GTX Titan X (GM200) and 1 NVIDIA GeForce GTX Titan Xp (GP102) GPU and 62 GB of RAM.

### **III. EXPERIMENTS**

# A. Data

We use the data presented in [15]. This dataset consisted of US images of the neck with their corresponding masks, as can be seen in Figure 2. The images originate from 47 different subjects, with 119-120 images per subject. All images had a resolution of 420x580 pixels. Trained volunteers annotated the binary masks indicating the BP.

The dataset was found to contain contradictory annotations of close matching images. An example of this is shown in Figure 2. A second dataset, here we call it filtered data, was constructed, in which the close matching samples without an annotated BP were removed. Both the filtered and nonfiltered data will be tested in the experiments. The dataset



(a) Ultrasound Image of the neck



(c) Similar Ultrasound Image





(d) Incoherent Mask

Fig. 2: (a,b) An example of an ultrasound image from the neck containing the brachial plexus nerve region with the corresponding annotated ground truth highlighting this region. (c,d) An example of similar ultrasound images with a complete different annotated brachial plexus nerve region.

is unbalanced, with the no-BP class over-represented. The number of images in each class for both datasets are shown in Table II.

The data is then downsampled (128x128 in the CNN and 96x96 in the segmentation models). The images were meancentered with a standard deviation of 1. In all experiments, 5-fold cross validation was used for the model evaluation. For each fold, the train data was randomly divided into a train and validation set by a 80%-20% split. To increase the number of samples used for training the CNN, we augmented training data by 2500 extra samples using affine transformations (scaling, shearing, rotating and reflecting) on randomly selected samples.

## B. Experimental setup

To investigate the effect of employing the hybrid model on the performance of segmenting BP region, three experiments were conducted:

*Experiment 1: No classification:* In the first experiment, the segmentation models were trained on the data without any prior classification model. In this experiment, training data contains US images with and without BP regions.

*Experiment 2: hybrid model:* In the second experiment, the CNN architecture was used to identify images with BP. Then, the segmentation models (U-net and M-net) were trained on positive images only, and tested on the images that were classified by the CNN as BP.

*Experiment 3: perfect classification model:* In the final experiment, a perfect classification model was mimicked by manually discarding all images without an annotated mask from data, and the segmentation models were trained and

TABLE III: Experiment results (DSC) on filtered and non-filtered data.

	Non-filtered data		Filtered data	
	U-net	M-net	U-net	M-net
No classification	$0.53 \pm 0.06$	$0.59 \pm 0.07$	$0.60 \pm 0.08$	$0.65 \pm 0.08$
Hybrid model	$0.72 \pm 0.01$	$0.67 \pm 0.05$	$0.83 \pm 0.02$	$0.79 \pm 0.05$
Perfect classification	$0.92 \pm 0.00$	$0.83 \pm 0.10$	$0.93 \pm 0.00$	$0.86 \pm 0.09$



Fig. 3: Results of the different experiments visualized in boxplots.

tested to these sets. This experiment was done to be able to discuss the potential of a hybrid model with a better performing classification model.

# C. Evaluation

For evaluation of the classification results, two metrics were used in this research: the F1-score and the accuracy metric (see Equations 1 and 2). The accuracy metric is more suited for balanced datasets, whereas the F1-score is more suited for imbalanced datasets [16]. Since the dataset was somewhat imbalanced, as discussed in Section III-A, both metrics were evaluated.

For evaluation of the segmentation results, the dice similarity coefficient (DSC) was used (see Equation 3). DSC makes a pixel-wise comparison between the ground truth and the corresponding prediction. The DSC ranges between 0 and 1, where 0 indicates there is no overlap, and 1 represents a perfect overlap.

All acquired data was tested for normality with the Shapiro-Wilk test [17]. The models were then tested for different means with a two sided T-test.

$$Accuracy = \frac{TP + TN}{TP + FP + TN + FN}$$
(1)

$$F1 = \frac{2TP}{2TP + FP + FN} \tag{2}$$

$$DSC = \frac{2 \cdot |M \cap GT|}{|M| + |GT|} \tag{3}$$

## **IV. RESULTS**

The CNN achieved an average F1-score of  $0.72 \pm 0.01$  and an accuracy score of  $0.77 \pm 0.01$  on the non-filtered data, whereas on the filtered data, an average F1-score of  $0.82 \pm$ 0.02 and an accuracy score of  $0.87 \pm 0.01$  were found.

The DSCs of the experiments are listed in Table III. In the 'no classification' experiment, M-net outperforms U-net in both the non-filtered and the filtered data (p < .001, p = .002 respectively). M-net and U-net both perform better on the filtered data, compared to the non-filtered data (both p < .001).

From experiment 2, it was found that both hybrid architectures performed better in both datasets compared to the first experiment, as can be seen in Figure 3. Both models still perform better on the filtered data, compared to the non-filtered data. The U-net variant has shown the most improvement, and now performs better than M-net in both datasets (all p < .001). In U-net, the variance has dropped as results became more stable.

The third experiment shows an even better performance for both models in both datasets. Again, U-net has improved the most. The variance of U-net has decreased even more. This is not the case for M-net, where the variance has risen. Figure 3 shows the outliers (determined with the 1.5xIQR rule) in the M-net performance, explaining the large variance. Both models still perform better on the filtered data, however the discrepancy in performance between the two data sets is less compared to the first experiment.

#### V. DISCUSSION

Our experimental results revealed the potential of using hybrid models. The 'hybrid model' experiment showed an increase in performance of all models when prior classification is added, even when the classification is not perfect. The 'perfect classification' experiment showed that promising results can be achieved when a perfect classification method can be approached or reached. It also showed that results will keep increasing as the classification becomes more accurate.

U-net seems to have benefited more from the hybrid solution than M-net. This indicates that the M-net variant can better handle the large number of images without an annotated BP region, which can be attributed to the more densely supervised nature of the architecture [7].

The difference between the performance on the non-filtered data and the filtered data increased for the hybrid variant of both U-net and M-net. The incoherent data may have caused errors in the learning process of the models. In the hybrid model, this would happen in both the classification and segmentation models, which would have caused the error to propagate. This propagating error could have been the cause of the increased difference in performance between filtered and non-filtered data. This difference decreased in the third experiment. Since a perfect classification model was simulated, no images with a fully negative ground truth were sent to the segmentation models. This means that no incoherent data was being fed to the segmentation models when using the nonfiltered dataset, since only incoherent negatives were removed in the filtered dataset. This can explain the reduction in performance difference between filtered and non-filtered data in the second experiment.

For the U-net model, the variance decreased when using the hybrid model, indicating a more stable model. Since most images without an annotated mask are filtered out by the CNN, more homogeneous data is fed to the segmentation model. This increases the likelihood for the segmentation model to recognize the BP regions.

The results of M-net in our first experiment do not correspond with the results presented in [7], with an average DSC of 0.59 in this study, compared to the reported 0.88 in [7]. Discrepancies between results can be expected since there are differences in methods. For one, the images used for testing are unknown, since each study created an own test set from the original data. Furthermore, the authors did not provide information about their test methods, making it harder to compare results.

In future research, improving the classification model performance should be considered as an important step, since the perfect classification experiment showed the potential for much better segmentation performance. Moreover, merging the models into a single 'chain' could improve time performance.

## VI. CONCLUSION

The goal of this research was to improve image segmentation of the BP nerve structure from US images, which was done by employing a hybrid deep neural network model. A CNN was used as the classification model, and U-net and M-net were employed as segmentation models. Three experiments were conducted on two datasets with filtered and nonfiltered data. The first experiment did not use any prior classification, the second experiment used the CNN for classification, and in the third experiment, a perfect classification model was simulated by removing all negative samples from data. From these experiments, it could be concluded that performance significantly improved when using a hybrid model, and even more when using a 'perfect classification' model, thus showing the promising possibilities of using hybrid model for accurate BP segmentation.

#### REFERENCES

- [1] E. N. Marieb and K. Hoehn, *Human Anatomy & Physiology*, 10th ed. London: Pearson Education, 2015, ch. 13.
- [2] J. Yang, A. Amini, R. Williamson, L. Zhang, Y. Zhang, R. Komaki, Z. Liao, J. Cox, J. Welsh, L. Court, and L. Dong, "Automatic contouring of brachial plexus using a multi-atlas approach for lung cancer radiation therapy," *Practical Radiation Oncology*, vol. 3, no. 4, pp. 139–147, oct 2013, doi:10.1016/j.prro.2013.01.002.

- [3] E. Pincus, "Regional Anesthesia: An Overview," AORN Journal, vol. 110, no. 3, pp. 263–272, 2019, doi:https://doi.org/10.1002/aorn.12781.
- [4] S. Banerjee, R. Acharya, and B. Sriramka, "Ultrasound-guided interscalene brachial plexus block with superficial cervical plexus block compared with general anesthesia in patients undergoing clavicular surgery: A comparative analysis," *Anesthesia: Essays and Researches*, vol. 13, no. 1, p. 149, 2019, doi:10.4103/aer.aer\_185\_18.
- [5] A. Kakade and J. Dumbali, "Identification of nerve in ultrasound images using U-net architecture," in 2018 International Conference on Communication information and Computing Technology (ICCICT), 2018, pp. 1–6, doi:10.1109/ICCICT.2018.8325894.
- [6] S. S. Liu, J. E. Ngeow, and J. T. YaDeau, "Ultrasound-Guided Regional Anesthesia and Analgesia: A Qualitative Systematic Review," *Regional Anesthesia & Pain Medicine*, vol. 34, no. 1, pp. 47–59–47–59, 2009, doi:10.1097/AAP.0b013e3181933ec3.
- [7] N. Abraham, K. Illanko, N. Khan, and D. Androutsos, "Deep Learning for Semantic Segmentation of Brachial Plexus Nerves in Ultrasound Images Using U-Net and M-Net," in 2019 3rd International Conference on Imaging, Signal Processing and Communication (ICISPC), 2019, pp. 85–89, doi:10.1109/ICISPC.2019.8935668.
- [8] R. Wang, H. Shen, and M. Zhou, "Ultrasound Nerve Segmentation of Brachial Plexus Based on Optimized ResU-Net," in 2019 IEEE International Conference on Imaging Systems and Techniques (IST), 2019, pp. 1–6, doi:10.1109/IST48021.2019.9010317.
- [9] J. Seetha and S. S. Raja, "Brain tumor classification using convolutional neural networks," *Biomedical and Pharmacology Journal*, vol. 11, no. 3, pp. 1457–1461, Sep. 2018, doi:10.13005/bpj/1511.
- [10] B. Kong, X. Wang, Z. Li, Q. Song, and S. Zhang, "Cancer metastasis detection via spatially structured deep network," in *Lecture Notes* in *Computer Science*. Springer International Publishing, 2017, pp. 236–248, doi:10.1007/978-3-319-59050-9\_19.
- [11] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, http://www.deeplearningbook.org.
- [12] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," *CoRR*, vol. abs/1505.0, pp. 234–241, 2015, doi:10.1007/978-3-319-24574-4\_28.
- [13] R. Mehta and J. Sivaswamy, "M-net: A Convolutional Neural Network for deep brain structure segmentation," in 2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017), 2017, pp. 437–440, doi:10.1109/ISBI.2017.7950555.
- [14] F. Chollet et al., "Keras," https://keras.io, 2015.
- [15] "Ultrasound nerve segmentation dataset," 2016, data retrieved from Kaggle, https://www.kaggle.com/c/ultrasound-nerve-segmentation/.
- [16] A. Tharwat, "Classification assessment methods," *Applied Computing and Informatics*, vol. ahead-of-p, no. ahead-of-print, aug 2020, doi:10.1016/j.aci.2018.08.003.
- [17] S. S. Shapiro and M. B. Wilk, "An analysis of variance test for normality (complete samples)," *Biometrika*, vol. 52, no. 3/4, pp. 591–611, 1965.