

Automatic brachial plexus segmentation to facilitate targeted echo block using convolutional neural networks

Nicolas Mauricio Ramírez-Triana^a, Oscar Hernán Ramírez-Agudelo^b, and Lola Xiomara Bautista-Rozo^c

^{a,c}Escuela de Ingeniería de Sistemas, Universidad Industrial Santander, Cra 27 Calle 9 Ciudad Universitaria, Bucaramanga, Santander, Colombia

^bGerman Aerospace Center (DLR), Institute for AI Safety and Security, Rathausallee 12, 53757 Sankt Augustin, Germany

ABSTRACT

The brachial plexus is a network of nerves located in the shoulder. It carries movement and sensory signals from the spinal cord to the arms and hands. The brachial plexus block is the most widely used anesthetic method in surgical procedures of the upper limbs. By using ultrasound imaging, trained professionals proceed to carry out echo-directed brachial plexus block procedure. However, the identification and location of the brachial plexus is difficult, which in turn may lead to health complications to the patient. Convolutional neural networks could assist this procedure by providing an automatic identification and segmentation of the region of interest. In this paper, by using the dataset provided by the competition *Ultrasound Nerve Segmentation*, an U-Net model is trained. Our implementation automatically locates and segments the brachial plexus in the aforementioned dataset. In our experiment, the model achieves a dice coefficient (DSC) score of $DSC = 0.87$. It is concluded that the proposed method satisfactorily locates and segments the brachial plexus. Our findings are in line with other deep learning works, suggesting that convolutional neural networks (CNN) could facilitate the echo brachial plexus block procedure.

Keywords: Medical imaging, brachial plexus, deep learning, convolutional neural networks, segmentation

1. INTRODUCTION

The brachial plexus is a network of nerves in the shoulder region that transmits motor and sensory signals from the spinal cord to the upper limbs. Injuries to this nerve complex, often caused by trauma, may result in pain, numbness, or muscle weakness in the arm and hand.

The brachial plexus block is a regional anesthesia technique widely used in upper limb surgeries. It involves the administration of a local anesthetic near the brachial plexus, enabling anesthesia from the shoulder to the fingertips. This procedure has a high clinical application rate, reported to range between 70% and 97% of all upper limb anesthesia cases.¹

However, accurate identification of the brachial plexus during the procedure remains challenging. The anatomical complexity of the region and the proximity to vital structures such as the subclavian artery and lung increase the risk of complications. These include intravascular injection, systemic local anesthetic toxicity,² seizures or central nervous system depression,³ and pneumothorax.⁴ As such, tools that can assist clinicians in reliably locating the brachial plexus are of high practical value.

In recent years, convolutional neural networks (CNN) have shown great promise in medical image analysis, particularly in ultrasound imaging. These models can learn to detect and segment anatomical features directly from imaging data. Prior work has demonstrated the successful use of CNN for identifying peripheral nerves in ultrasound scans, including the median nerve⁵ and the brachial plexus.⁶

Further author information: (Send correspondence to N.M.R.-T.)

N.M.R.-T.: E-mail: nicolas1ramireztr@gmail.com, Telephone: +57 3138234505

O.H.R.-A.: E-mail: oscar.ramirezagudelo@dlr.de

In this study, a CNN-based approach is proposed for the automatic segmentation of the brachial plexus in ultrasound images. A U-Net architecture⁷ is implemented and trained using the publicly available dataset provided by the *Ultrasound Nerve Segmentation* challenge.⁸ The objective is to support echo-guided nerve block procedures by enabling accurate and efficient identification of the brachial plexus. The remainder of this paper is organized as follows. Section 2 provides the relevant anatomical and procedural background. Section 3 details the methodology. The results are presented in Section 4, followed by a discussion in Section 5. Conclusions are drawn in Section 6.

2. BACKGROUND

2.1 Anatomy

The brachial plexus is a nerve structure located at the base of the neck that connects the spinal cord to all peripheral nerves in the arm. It is made up of five spinal roots, the ventral rami of the four lower cervical roots, and the first thoracic root (C5, C6, C7, C8, and T1), these roots join to form three primary trunks (UPPER: C5, C6; MIDDLE: C7; LOWER: C8, T1), subsequently, these trunks divide into divisions, fascicles, branches, and terminal nerves (See Figure 1 upper-panel). The brachial plexus is in charge of completely innervating (with the exception of some areas) the upper limb¹ (see also lower-panel of Fig. 1).




Figure 1. Anatomy of the Brachial Plexus and location in the human body.⁹

Due to the proximity of the lung to the brachial plexus at the level of the clavicle, the most frequent complication associated with this block is pneumothorax, with a risk of up to 6.1%.⁴ Other complications of supraclavicular block include puncture of the subclavian artery, and spread of local anesthetic to cause paralysis of the stellate ganglion, phrenic nerve, and recurrent laryngeal nerve.

2.2 Echo-directed brachial plexus block

In the realization of the block the patient first should be seated with an inclination between 30 and 45 degrees, with the head elevated and turned slightly away from the block (see Fig. 2).

Then, the trained professional anesthetizes the patient and by guiding the process with ultrasound imaging proceeds to insert the needle to the brachial plexus to the area known as 'Corner Pocket', i.e. between the artery subclavian and the first rib¹ (see also lower-panel Fig. 1).

2.2.1 Alternative approaches

Deep learning methods, such as CNN, have shown promising results in the field of medical ultrasound imaging.⁵ The automation of the location process of the Region of Interest (ROI) could help in the brachial plexus block procedure. The CNN model would identify and segment the brachial plexus, reducing the risks associated with such treatment e.g.⁵ and.⁶

In⁵ the authors implemented an ultrasound automatic segmentation of the median nerve by using CNN. They show satisfactory results, locating and segmenting the median nerve. They also created promising measurements for applications in the clinical diagnosis of carpal tunnel syndrome.



Figure 2. Position of the patient and instruments for echo-guided brachial plexus block.¹⁰

The work of⁶ proposes a hybrid deep learning model to help to improve the location and segmentation of the brachial plexus. It consists of a CNN, an U-Net,^{7,11} and a M-Net.¹² The CNN is used as a classifier. Then, the U-Net and M-Net are used for the segmentation.

Based on the aforementioned works, a CNN approach is proposed here. The method is introduced in the next section.

3. METHODOLOGY

This section outlines the dataset (Sect. 3.1), preprocessing pipeline (Sect. 3.2), model architecture, training configuration, and evaluation metrics used in this study (Sect. 3.3).

3.1 Dataset

The dataset used in this work was obtained from the publicly available *Ultrasound Nerve Segmentation* challenge on Kaggle.⁸ It consists of 16,778 ultrasound images collected from 47 different persons. Each image corresponds to a scan of the brachial plexus region and is optionally accompanied by a binary mask indicating the location of the nerve.

For this study, a subset of 11,270 images was used. Among these, 5,635 images are labeled with corresponding segmentation masks, while the remaining 5,635 are raw images without annotations. Images lacking annotations were excluded from model training and evaluation. Figure 3 shows a random example of the training set. Panel a) displays the ultrasound scan image. Panel b) corresponds to its respective mask.




Figure 3. **Training set case.** Panel a) displays a random ultrasound scan image. Panel b) corresponds to its respective mask.

3.2 Data preprocessing

To standardize the inputs for training, all images and masks were resized to 128×128 pixels. The dataset was split into training, validation, and test sets in a 72%, 20%, and 8% ratio, respectively:

- *Training set:* 4057 of brachial-plexus ultrasound scans and their respective annotated masks. For a total of 8114 images.
- *Validation set:* 1127 of brachial-plexus ultrasound scans and their respective annotated masks. For a total of 2254 images.
- *Testing set:* 451 of brachial-plexus ultrasound scans and their respective annotated masks. For a total of 902 images.

Data augmentation was applied on-the-fly to the training set using Keras' `ImageDataGenerator`. These augmentations introduce controlled variability in position, scale and orientation, thereby enhancing model generalization. Table 1 summarizes the parameters employed.

Table 1. Data augmentation parameters applied to the training set.

Parameter	Value
Rotation range	$\pm 5^\circ$
Width shift range	5% of image width
Height shift range	5% of image height
Zoom range	$\pm 10\%$
Horizontal flip	Enabled
Fill mode	reflect

3.3 Model architecture

A U-Net⁷ convolutional neural network was implemented for segmentation. The network comprises a contracting path (encoder) and an expansive path (decoder). The encoder includes four convolutional layers with Exponential Linear Unit (ELU) activation functions, each followed by Group Normalization (group size of 8), dropout (rate = 0.5), and max pooling. The number of filters starts at 32 and doubles at each downsampling step.

The decoder performs upsampling via transposed convolutions and includes skip connections from the encoder to preserve spatial information. A final 1×1 convolution with sigmoid activation produces the binary segmentation mask.

The architecture is visualized in Fig. 4 (rendered using PlotNeuralNet¹³).




Figure 4. The architecture of the U-Net (visualisation obtained by using¹³).

3.4 Training configuration

The model was implemented using TensorFlow 2.4¹⁴ and trained on a dedicated GPU server with the following hardware specifications:

- 2x Intel Xeon Gold 5118 (16 cores each)
- NVIDIA Titan V GPU (12 GB VRAM)
- 256 GB RAM
- 48 TB HDD (RAID configuration)

The training was conducted for a maximum of 100 epochs with early stopping enabled (patience = 10 epochs). The Adam optimizer was used with an initial learning rate of 10^{-3} . A batch size of 32 was used across all experiments. The selected hyperparameters were determined through preliminary exploration and are summarized in Table 2.

Table 2. The value of parameters for the UNet.

Parameter	Value
Optimizer	Adam
Dropout	0.5
Epochs	100
Batch size	32
Learning rate	10^{-3}
Number of filters	32
Activation function	ELU

3.5 Evaluation

Segmentation performance was quantified primarily using the Dice coefficient (DSC).¹⁵ Furthermore, the Intersection over Union (IoU)¹⁶ and the Area Under the Receiver Operating Characteristic Curve (AUC-ROC)¹⁷ were also calculated as complementary metrics to support the decision of the best result adopted. All three metrics range from 0 to 1, with higher values indicating better agreement between prediction and ground truth.

The Dice coefficient (DSC) measures the spatial overlap between the predicted region Y and the ground truth region X :

$$DSC = \frac{2|X \cap Y|}{|X| + |Y|} \quad (1)$$

It is optimized directly on the DSC via the corresponding loss:

$$Loss(DSC) = 1 - DSC \quad (2)$$

The Intersection over Union (IoU), also known as the Jaccard index, is defined as:

$$IoU = \frac{|X \cap Y|}{|X \cup Y|} \quad (3)$$

Finally, AUC-ROC evaluates the model's pixel-wise discrimination ability by computing the area under the curve of true positive rate versus false positive rate, providing a threshold-independent measure of classification performance.

In the fitting process, the model takes as input the training and validation datasets. In addition, early stopping has also been implemented with a tolerance of 10 epochs. Finally, once the model has been computed, it is evaluated in the testing data set.

4. RESULTS

This section presents the results following the methodology proposed in Sect. 3, Sect. 5 discusses the findings.

4.1 Dice coefficient (*DSC*) curve

Figure 5 presents the evolution of the validation learning curves for Dice coefficient (DSC), Intersection over Union (IoU) and Area Under the ROC Curve (AUC-ROC). During the first ten epochs, all three metrics exhibit a steep increase on the validation set (orange curves), with DSC, IoU and AUC-ROC reaching its maximum early at epoch 56 with values of 0.8709, 0.7850 and 0.8447 respectively. Beyond epoch 56, the curves enter a plateau, and training is terminated at epoch 66 by early stopping. Accordingly, the model checkpoint saved at epoch 56 optimally balances DSC, IoU and AUC-ROC performance.




Figure 5. Learning curve of the metrics (DSC, IOU, AUC-ROC) during training

4.2 Prediction

Once the best model has been obtained, it is visualized. Figure 6 shows the step-by-step process for a random case of the test set. Panel a) shows the ultrasound scan of a randomly selected test case. Panel b) displays the corresponding ground truth mask, annotated by⁸ Panel c) overlays the ultrasound image with the ground truth mask. Panel d) presents the ultrasound scan overlaid with the prediction generated by the best-performing model.

Table 3 summarizes the validation results for Dice coefficient (DSC), Intersection over Union (IoU), Area Under the ROC Curve (AUC-ROC).

Table 3. Validation performance: DSC, IoU, AUC-ROC for the epoch 56.

Metric	Value
DSC	0.8709
IOU	0.7850
AUC-ROC	0.8447




Figure 6. **Test set case.** Panel a) displays the ultrasound scan image of one random case. Panel b) corresponds to its respective mask annotated by.⁸ Panel c) is the combination of panels a) and b). Panel d) shows the ultrasound scan (i.e. panel a) overlying the prediction obtained by our best model.

5. DISCUSSION

5.1 Comparison to Related Work

Van Boxtel et al.⁶ proposed a hybrid deep neural network combining a CNN classifier with U-Net and M-Net segmentation modules, achieving a DSC of 0.92 for brachial plexus segmentation. Although their architecture outperforms our simpler U-Net model in terms of accuracy, our method maintains the advantages of lower computational complexity and a reduced number of hyperparameters, making it more accessible for deployment in real-time clinical settings.

5.2 Strengths and Practical Implications

A key strength of this work is the use of a standard U-Net architecture without any ensemble or hybrid design. Despite its simplicity, the model achieves competitive segmentation accuracy and operates efficiently during both training and inference. This makes it a suitable candidate for integration into ultrasound machines or decision-support tools aimed at assisting anesthesiologists.

Moreover, the pipeline is reproducible, leveraging only publicly available data and widely used software libraries. The low-latency prediction suggests potential for real-time clinical use, though further validation is necessary.

5.3 Limitations and Error Analysis

Several limitations must be acknowledged. First, the dataset was partitioned using a fixed train-validation-test split. This may limit the generalization ability of the model to new patients or imaging settings. Future work should consider cross-validation or subject-wise splits to more thoroughly assess robustness.

Second, mild data augmentation was applied during training—comprising small random rotations, horizontal and vertical shifts, slight zooms and flips—to increase variability in probe angle and anatomical presentation. However, these transformations were intentionally conservative; employing more aggressive augmentation strategies (e.g. larger rotation ranges, elastic deformations or intensity perturbations) could further improve resilience to imaging artifacts and patient heterogeneity.

Third, all images were resized to 128×128 pixels, which may have caused the loss of fine-grained anatomical detail. While downscaling was done to reduce computational costs, higher-resolution inputs could improve segmentation precision, especially in small or complex nerve structures.

Qualitative review of failed predictions suggests that the model sometimes confuses bright surrounding structures, such as vessels or bone interfaces, with nerve tissue. These false positives could be mitigated by incorporating attention mechanisms,¹⁸ anatomical priors, or transformer-based modules that better capture context.

5.4 Opportunities for Future Work

Approximately 5,500 ultrasound images in the dataset remain unlabeled. These could be leveraged using semi-supervised learning or pseudo-labeling methods to enhance model performance without requiring new annotations. This approach would be particularly useful in clinical settings where annotated data is scarce.

Additionally, future experiments could explore model ensembles, multi-scale architectures, and integration of spatial context (e.g., Doppler or anatomical landmarks). Attention-based modules such as Attention U-Net¹⁸ could also improve feature localization by highlighting relevant regions in the image.

Finally, clinical evaluation is needed to determine the real-world utility of the model. Collaborations with anesthesiologists could assess interpretability, ease of use, and effectiveness in reducing procedure time and complication rates. Deploying the model in embedded systems or point-of-care devices may also be explored to support decision-making in operating rooms.

6. CONCLUSION

By using the dataset,⁸ a U-Net architecture is implemented to automatically locate and segment the brachial plexus in ultrasound images. The resulting model achieves a good performance with a dice coefficient score of $DSC = 0.87$. Our result is comparable to other deep learning works presented in the literature. It is concluded that CNN models would facilitate the identification of the brachial plexus, reducing the risk of complications in the brachial plexus block procedure.

REFERENCES

- [1] Sandeep Kusre, Andrew McEwen, G. M., "Ultrasound-guided supraclavicular brachial plexus block," *Regional Anaesthesia* (2018).
- [2] Mather, L. E., Copeland, S. E., and Ladd, L. A., "Acute toxicity of local anesthetics: Underlying pharmacokinetic and pharmacodynamic concepts," *Regional Anesthesia & Pain Medicine* **30**(6), 553–566 (2005).
- [3] Crews, J. C. and Roth, T. E., "Seizure after levobupivacaine for interscalene brachial plexus block," *Anesthesia & Analgesia* **96**(4), 1188–1190 (2003).
- [4] Macfarlane AJ, Prasad GA, C. V. B. R., "Does regional anesthesia improve outcome after total knee arthroplasty?," *Clin Orthop Relat Res.* **467**, 2379–2402 (Sep 2009).
- [5] Horng, M.-H., Yang, C.-W., Sun, Y.-N., and Yang, T.-H., "Deepnerve: A new convolutional neural network for localization and segmentation of the median nerve in ultrasound image sequences," *Ultrasound in Medicine & Biology* **46**(9), 2439–2452 (2020).
- [6] Van Boxtel, J., Vousten, V., Pluim, J., and Rad, N. M., "Hybrid deep neural network for brachial plexus nerve segmentation in ultrasound images," *2021 29th European Signal Processing Conference (EUSIPCO)* (Aug 2021).
- [7] Ronneberger, O., Fischer, P., and Brox, T., "U-net: Convolutional networks for biomedical image segmentation," *CoRR* **abs/1505.04597** (2015).
- [8] Kaggle, "Ultrasound Nerve Segmentation Dataset." <https://www.kaggle.com/c/ultrasound-nerve-segmentation> (2016). Accessed: 2025-04-16.
- [9] "Resumen del plexo braquial." <https://danielrestituyo.medium.com/resum%C3%A9n-del-plexo-braquial-5c406bd9eb75>. Accessed: 2022-02-15.
- [10] "In plane technique." <https://temena.com/wp-content/uploads/2016/05/In-Plane-Technique-IP-Ultrasound-Temena-Group.jpg>. Accessed: 2021-10-15.
- [11] Abraham, N., Illanko, K., Khan, N., and Androustos, D., "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)]*, 85–89 (2019).

- [12] Mehta, R. and Sivaswamy, J., “M-net: A convolutional neural network for deep brain structure segmentation,” in *[2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)]*, 437–440 (2017).
- [13] “Plotneuralnet.” <https://github.com/HarisIqbal88/PlotNeuralNet>. Accessed: 2022-02-23.
- [14] Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X., “TensorFlow: Large-scale machine learning on heterogeneous systems,” (2015). Software available from tensorflow.org.
- [15] Harisha, L., “Dice coefficient! what is it?.” <https://lathashreeh.medium.com/dice-coefficient-what-is-it-ff090ec97bda> (feb 2023). Medium; Accessed: 2025-07-16.
- [16] Rosebrock, A., “Intersection over union (iou) for object detection.” <https://pyimagesearch.com/2016/11/07/intersection-over-union-iou-for-object-detection/> (nov 2016). Last updated April302022; Accessed: 2025-07-16.
- [17] Bhalla, D., “A complete guide to area under curve (auc).” <https://www.listendata.com/2014/08/learn-area-under-curve-auc.html> (aug 2014). Accessed: 2025-07-16.
- [18] Oktay, O., Schlemper, J., Folgoc, L. L., Lee, M., Heinrich, M., Misawa, K., Mori, K., McDonagh, S., Hammerla, N. Y., Kainz, B., Glocker, B., and Rueckert, D., “Attention u-net: Learning where to look for the pancreas,” *arXiv preprint* (2018). Preprint.