

Needle Segmentation For Real-time Guidance of Minimally Invasive Procedures Using Handheld 2D Ultrasound Systems

Paul Mugume Okwija ¹, Joanitta Nabacwa ², Sylvia Imanirakiza ², Alvin Bagetuuma Kimbowa ², Cosmas Mwikirize ², Ilker Hacihaliloglu ², and Andrew Katumba ²

¹Makerere University

²Affiliation not available

October 30, 2023

Abstract

Background

Accurate needle placement is crucial during minimally invasive procedures such as biopsies, regional anesthesia, and fluid aspiration. 2D Ultrasound is widely used for needle guidance during such procedures, however, it has a limited field-of-view and poor needle visibility for steep insertion angles.

Methods

In this work, we propose a novel machine learning (ML)-based method for real-time needle segmentation in handheld 2D ultrasound systems. The proposed method involves a fast and simple annotation technique allowing for the labeling of large datasets. It then utilizes the U-Net architecture which is modified to allow for easy integration into a handheld ultrasound system. Two datasets were used in this work, one consisting of B-mode ultrasound videos obtained from human tissue and the other consisting of videos and frames from chicken, porcine and bovine tissue. The model is trained on 1262 frames and evaluated on 209 frames.

Results

This approach achieves an Intersection Over Union (IoU) of 0.75 and a dice coefficient of 0.851 on frames obtained from human tissue. The model is integrated into the processing pipeline of a portable ultrasound system and achieves an overall processing speed of about 8 frames per second. The proposed approach outperforms state-of-the-art methods for needle segmentation while achieving real-time integration. This work is a step forward towards real-time needle guidance using machine learning-based algorithms in handheld ultrasound systems.

Needle Segmentation For Real-time Guidance of Minimally Invasive Procedures Using Handheld 2D Ultrasound Systems

Paul Okwija Mugume¹, Joanitta Nabacwa¹, Sylvia Imanirakiza¹, Alvin Bagetuuma Kimbowa¹,

Cosmas Mwikirize¹, Ilker Hacihaliloglu^{2,3}, Andrew Katumba¹

Department of Electrical and Computer Engineering, Makerere University, Kampala Uganda ¹

Department of Radiology, University of British Columbia, BC Canada V6T 1Z2 ²

Department of Medicine, University of British Columbia, BC Canada V6T 1Z2 ³

Abstract—Accurate needle placement is crucial during minimally invasive procedures such as biopsies, regional anesthesia, and fluid aspiration. 2D Ultrasound is widely used for needle guidance during such procedures, however, it has a limited field-of-view and poor needle visibility for steep insertion angles. In this work, we propose a novel machine learning (ML)-based method for real-time needle segmentation in handheld 2D ultrasound systems. The proposed method involves a fast and simple annotation technique allowing for the labeling of large datasets. It then utilizes the U-Net architecture which is modified to allow for easy integration into a handheld ultrasound system. Two datasets were used in this work, one consisting of B-mode ultrasound videos obtained from human tissue and the other consisting of videos and frames from chicken, porcine and bovine tissue. The model is trained on 1262 frames and evaluated on 209 frames. This approach achieves an Intersection Over Union (IoU) of 0.75 and a dice coefficient of 0.851 on frames obtained from human tissue. The model is integrated into the processing pipeline of a portable ultrasound system and achieves an overall processing speed of about 8 frames per second. The proposed approach outperforms state-of-the-art methods for needle segmentation while achieving real-time integration. This work is a step forward towards real-time needle guidance using machine learning-based algorithms in handheld ultrasound systems.

Index Terms—ultrasound, minimally invasive procedures, segmentation, U-Net, real-time, integration

I. INTRODUCTION

Minimally invasive procedures such as biopsies, regional anesthesia, and fluid aspiration involve percutaneous needle insertion which requires tracking so that the needle safely reaches its intended target [1] [2] [3]. Ultrasound is the gold standard for visualization of needle progress vis-à-vis patient anatomy. This is because it is real-time, does not emit ionizing radiation, and is widely available [4] [5]. Traditional cart-based ultrasound systems generally have good image quality but are expensive, bulky, and high power consuming. [6]. Recently, handheld ultrasound systems, which are significantly more affordable, less power-consuming, and allow for point-of-care applications, have been developed. [7].

Common hand-held devices on the market produce 2D ultrasound images. However, 2D ultrasound has a limited field-of-view since it can only display a cross-section of a 3D scene. Additionally, needle visibility in 2D ultrasound is poor for steep insertion angles which often result in minimal shaft information and a low-intensity needle tip [8]. Needle visibility also deteriorates with thin and bending needles which may become indistinguishable from body and motion artifacts. Finally, 2D ultrasound requires coordination between the surgeon, who performs the needle insertion and transducer placement, and the echo tech, who manipulates the ultrasound system. Overall, these limitations reduce the efficacy of minimally invasive procedures and could lead to injury, and post-procedure complications.

Various hardware-based methods have been developed to enhance needle visualization in 2D ultrasound. These include mechanical needle guides, which have predetermined angles of approach but this limits fine needle trajectory adjustments [9]. Another method involves the use of echogenic coatings that improve needle visibility [5]. With this approach, visualization becomes suboptimal for steeper insertion angles and more echogenic backgrounds. Electromagnetic and optical tracking systems have also been proposed, but they are susceptible to inaccuracies emanating from the influence of metallic objects in the operating environment [10]. Generally, hardware-based methods increase the cost of ultrasound systems and require additional hardware which disrupts normal workflow.

Alternatively, software-based techniques involving image processing have been proposed to facilitate automatic needle detection. These aim at either localizing or segmenting the needle within an ultrasound image. Of the two tasks, needle segmentation is key since it derives a more precise result with each pixel in the image assigned one of two categories i.e needle or background [2]. Deep learning-based methods have gained prominence in segmentation tasks in tandem with the progress made in the field of computer vision [2]. Several approaches have been proposed for needle segmenta-

tion in 3D ultrasound, for instance, Pourtaherian et al [11] proposed an approach that accurately detects short needles in 3D ultrasound. Younes et al [12] proposed a method that can automatically segment flexible needles in 3D ultrasound images. However, 3D ultrasound is computationally expensive and this hinders the integration of these methods for real-time needle detection. It is also not widely available and thus 2D ultrasound remains the gold standard for needle visualization.

Various methods for needle segmentation in 2D ultrasound have been developed. For instance, Lee et al proposed an approach that combines LinkNet, with the Spatial and Channel "Squeeze and Excitation" (scSE) network for needle detection and trajectory prediction [13]. Their approach utilizes a light deep learning architecture which is desirable for real-time integration. However, this approach achieves a low performance with an IoU of 0.41 and dice accuracy of 56.65%. Wijata et al proposed a fully automatic Convolutional Neural Network (CNN) architecture to segment the needle and a Radon Transform to determine the trajectory of the needle during core-needle biopsy [14]. Their approach achieves a relatively high IoU of up to 0.655 outperforming other state-of-the-art methods. However, this method is sensitive to superimposed text, arrows, lines, and other high-intensity artifacts. Overall, these approaches have not been evaluated in a real-time ultrasound system.

In this work, we focused on developing an accurate machine learning-based method for needle segmentation in 2D ultrasound and integrating it into a handheld 2D ultrasound system. We propose a robust and accurate machine learning algorithm based on the U-Net architecture [15] which we modify to reduce computational complexity without negatively impacting segmentation accuracy. We demonstrate that the proposed approach outperforms state-of-the-art methods even in presence of superimposed text and high-intensity artifacts in the ultrasound images. This makes it suitable for integration in computer-assisted interventional systems for needle segmentation. The main contributions of this paper are; 1) a novel approach, for needle segmentation in 2D ultrasound, consisting of a modified U-Net architecture; (2) real-time integration of the algorithm into a handheld 2D ultrasound system and its evaluation for real-time needle visualization; and (3) a simple and fast needle annotation technique for labeling 2D ultrasound images used in developing deep learning models.

II. METHODS

A. Dataset

Our dataset comprises *ex vivo* animal ultrasound images and *in vivo* human ultrasound images. The *ex vivo* data were collected using three different imaging systems: SonixGPS (Analogic Corporation, Peabody, MA, USA) with a hand-held C5-2/60 curvilinear probe, at 30 frames-per-second (fps) and two 2D hand-held wireless US (Clarius C3 and Clarius L7, Clarius Mobile Health Corporation, Burnaby, BC, Canada) at 24 fps. Three freshly excised tissue types were used: porcine (thigh) and chicken (breast and thigh) and bovine. Furthermore, four needle types were used: a 17G SonixGPS vascular

TABLE I: Summary of the dataset. The human data were collected for different anatomical structures. Animal data were collected from chicken, porcine and bovine tissue.

Image Type		No of Images
Human Data	Femoral artery	223
	Axillary vein	127
	Poplitealartery	62
	Shoulder	59
	Saphenous vein	50
	Gluteal muscle	41
Animal Data	Porcine	927
	Bovine	389
	Chicken	1,724
Total		3,624

access needle (Analogic Corporation, Peabody, MA, USA), a 17G Tuohy epidural needle (Arrow International, Reading, PA, USA), a 21G clinical needle and a 22G spinal Quincke-type needle (Becton, Dickinson and Company, Franklin Lakes, NJ, USA). The needles were inserted in-plane with varying angles of insertion (10 – 40 degrees). In all experiments, data were collected mimicking clinical situations by introducing motion via large probe pressure while concurrently rotating the transducer. The human data were collected from online open-source resources.

A total of 3,624 images were collected and their distribution is summarized in Table I.

B. Preprocessing

The ultrasound images originally contained unwanted information, such as annotations and text, on the outer margins. MATLAB was used to crop out this information as shown in Fig. 1. All the frames were thereafter resized to 256x256. A simple and fast labeling technique was adopted where a linear region of interest (ROI) is drawn starting from the entry point of the needle and ending at the needle tip. This yields a pair of




Fig. 1: Illustration of frame cropping. The part of the image within the red box is what is retained and the rest is discarded. The unwanted metadata contains information about the imaging device used and its settings.




Fig. 2: Illustration showing the original image, labeled image, and corresponding mask generated. The generated mask shows a white line for the needle against a black background.

(x,y) coordinates corresponding to the start of the entry point of the needle and the needle tip for each image.

The coordinates are then used to automatically generate the corresponding mask by drawing white lines passing through the coordinates on a black background. The thickness of the line can then be treated as a hyperparameter during model training. For our scenario, the optimal line width was empirically determined as 7 pixels. Labeling was done using MATLAB’s VideoLabeler and the masks were generated using OpenCV as shown in Fig. 2. From the collected dataset, frames that were indistinguishable and those that didn’t contain the needle were discarded. The final labeled dataset consisted of a total of 2107 images summarized in Table II.

We randomly split the dataset into three sets; 60% for train (1262 images), 30% for validation (616 images) and 10% for testing (209 images). In order to eliminate sampling bias, the stratified sampling technique was used such that each category of the images has the same distribution throughout all the datasets as in the original dataset. The image datasets were normalized using L2 norm.

C. Model architecture

Our proposed framework, the Needle Segmentation U-Net (NSeg-UNet), shown in Fig. 3, is a modification of the U-Net architecture [15]. The U-Net architecture was chosen due to its competitive performance on medical image segmentation tasks requiring few images and yielding precise segmentations. The model receives as input, an image of size 256x256x1, and passes it through the contracting path (encoder) which extracts a feature map that captures the information contained within the image. The feature map is then passed through the expanding path (decoder) which propagates context information upwards to higher resolution layers. The model then outputs a 256x256x1 segmentation map containing a pixel-wise classification of a needle.

For each step in the contracting path, we apply two 3x3 convolutions each followed by a rectified linear unit (ReLU). We then perform downsampling at the end of each step by applying a 2x2 max pooling operation with a stride of 2. At the end of each step, the image size is halved and the number of channels is doubled. Each step in the expansive path starts with a 2x2 up-convolution followed by a concatenation with the corresponding feature map from the contracting path. We

TABLE II: The labeled dataset was split using the stratified sampling technique such that the data distribution is maintained across the splits.

Image Type		Number	Training	Validation	Test
Human data		476	285	144	47
Animal data	Chicken	748	448	226	74
	Bovine	370	222	111	37
	Porcine	513	307	155	51
Total		2107	1262	636	209

then apply two 3x3 convolutions each followed by a ReLU. At the end of each step, the image size is doubled and the number of channels is halved. In the final layer which consists of a 1x1 convolution, we use the sigmoid activation function.

One important modification in our architecture is that we use padded convolutions which maintain the size of the feature maps in every step in the contracting path. This eliminates the need to crop the feature maps during the skip connections which helps reduce the model inference time. We also use fewer filters, starting with 32 filters in the first step and 512 filters in the last step of the contracting path, so as to obtain a relatively small model size with good performance.

D. Training

The network was trained using the pre-processed B-mode ultrasound images as input and the corresponding generated masks as the target. During training, we aim to directly maximize the intersection over union (IoU), also known as Jaccard similarity. The IoU measures the overlap of the prediction with the ground truth. Floris et al [16] showed that training directly on IoU can significantly increase model performance as compared to training on the conventional binary cross entropy loss. Therefore, the IoU, at threshold of 0.5, was calculated for each prediction using equation 1.

$$IoU = \frac{(y_{pred} * y_{true}) + smooth}{(y_{pred} + y_{true}) - (y_{pred} * y_{true}) + smooth} \quad (1)$$

where y_{pred} is the predicted mask, y_{true} is the ground truth mask and smooth is a value that removes division by 0 error. We, therefore, train the model using the IoU loss given by equation 2.

$$IoU \text{ loss} = 1 - IoU \quad (2)$$

$$Dice = \frac{2 * (y_{pred} * y_{true}) + smooth}{(y_{pred} + y_{true}) + smooth} \quad (3)$$

Another metric we track during training is the Dice coefficient defined by equation 3. The dice coefficient is also a similarity metric which helps to measure the percentage of overlap between the predicted segmentation map and the ground truth. The optimizer used was the adaptive moment estimation (ADAM) with a learning rate of 0.001. Training was done on Google Colab, powered by a 12 GB Tesla K80 GPU with 16GB RAM.




Fig. 3: Model Architecture of the NSeg-UNet. The blue boxes represent feature maps. The number of channels for each feature map is denoted at the top of each box. The dimensions of each feature map are denoted at the lower left edge of each box.

III. EXPERIMENTS

The NSeg-UNet was trained for 70 epochs with a batch size of 32. For each epoch, the model was evaluated on the validation data and the best model was chosen as the one with the lowest validation IoU loss.

We also tried alternative approaches. We modified the NSeg-UNet to include attention gates as in [17] and trained it in a similar way as the NSeg-UNet. We also implemented the method proposed by Wijata [14] and trained it on our dataset. Finally, we implemented a transfer learning approach where we replaced the contracting path with pretrained models.




Fig. 4: Segmentation results for two B-mode frames obtained from chicken tissue with the needle inserted at shallow angles. The model performance is not sensitive to the direction in which the needle is inserted.




Fig. 5: Segmentation results for two B-mode frames obtained from human data. The first row shows a scenario where the needle was inserted at a steep angle. The second row shows a scenario where the image had superimposed text.




Fig. 6: Segmentation result for a frame, obtained from human tissue, where the needle has partially veered away from the field of view.

The models were frozen and training was done for only the expanding path. This allows the model to utilize the features learned by the pretrained models without distorting the weights. The pretrained models used were VGG16 and VGG19 with ImageNet weights.

IV. RESULTS AND DISCUSSION

In Fig. 4, we illustrate the needle segmentation results for two frames with different needle directions at shallow angles of insertion. It can be seen that the NSeg-UNet accurately segments the needle even when the needle is faintly visible. Fig. 5 illustrates the model performance at steeper angles of insertion. The model accurately segments the needle even in the presence of other high intensity artifacts. In very challenging situations where the needle veers away from the field view, as shown in Fig. 6, the model yields broken segmented predictions. This is not unexpected since the model utilizes only spatial information to make predictions.

Table III summarizes the experimental results obtained for the different approaches when tested on 172 images. It can be seen that the NSeg-UNet outperforms all the other methods with an IoU of 0.75 and a dice coefficient of 0.851 on the human data. The Attention U-Net achieves a lower performance with an IoU of 0.725 and a dice coefficient of 0.829 on the human data. Attention gates proved to improve performance for tissue/organ segmentation and localization in [17] as they enable the model to focus on target structures which are usually of low intensity. However, for medical instrument segmentation, attention gates may not yield significant improvement since medical instruments, such as needles, already exhibit a high intensity under ultrasound. Similarly, the transfer learning approaches yielded lower performance, with VGG16 achieving an IoU of 0.680 and VGG19 achieving an

IoU of 0.685. This can be attributed to the fact that these pretrained models were trained on ImageNet, which consists of real-world images and no grayscale medical images.

The results indicate that model performance on a given dataset depends on the number of images available in that dataset. For instance, all models achieve highest performance on the chicken data, which had the largest number of images, and least performance on the bovine data, which had the least number of images. This is because a higher number of images allows the model to better learn the representation of the needle in a given dataset.

The inference time is a key metric since it is crucial in the real-time integration of any needle segmentation method. The models achieve similar inference time on Google Colab's 12 GB Tesla K80 GPU. The NSeg-UNet achieves an average inference time of about 54ms which translates to 18 frames per second. The original U-Net architecture [15] achieves an average inference time of 152ms on the same machine. It can be seen that the NSeg-UNet achieves over 64% improvement in inference time as compared to the original U-Net architecture. Further, the NSeg-UNet is less complex and smaller in terms of parameters, with about 7.7 million parameters as compared to the original U-Net architecture that consists of about 31 million parameters. Therefore, it achieves better computational efficiency when integrated.

A. Integration

The proposed model was integrated into the image processing pipeline of a hand-held 2D ultrasound system, the Clarius L7 [18]. The device connects to iOS and Android devices and is operated through an application installed on the devices. Clarius provides an application programming interface (API) that allows for real-time wireless streaming of ultrasound images from the probe to multiple devices including smartphones and computers. The API uses the transmission control protocol (TCP) to create a connection with the devices receiving the ultrasound images as shown in Fig. 7.

A general machine learning integration workflow mainly involves three steps; data input and preprocessing, model inference and post-processing [2]. The data input step was fully handled by the Clarius API. We then used a 3.6 GHz Intel (R) Core™ i5 Windows laptop computer, with an NVidia GTX 1050 4GB VRAM GPU, to preprocess the streamed image, make inference and also carry out post-processing. The preprocessing done involves resizing the image to 256x256x1, and normalizing it using L2 normalization.

TABLE III: Comparison of the NSeg-UNet, Attention U-Net, and the transfer learning approaches using VGG16 and VGG19 in terms of IoU, dice coefficient, number of parameters and average inference time on a single image.

Model	Human		Chicken		Porcine		Bovine		Model parameters	Inference time (s)
	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice		
NSeg-UNet	0.750	0.851	0.778	0.870	0.755	0.655	0.610	0.746	7,771,297	0.054
Attention U-Net	0.725	0.829	0.777	0.868	0.764	0.670	0.589	0.729	9,086,021	0.057
Wijata	0.707	0.819	0.756	0.855	0.719	0.603	0.515	0.659	511,425	0.052
Vgg16_encoder	0.680	0.795	0.762	0.859	0.735	0.625	0.541	0.681	18,549,761	0.077
Vgg19_encoder	0.685	0.799	0.744	0.846	0.734	0.620	0.491	0.633	21,499,649	0.079

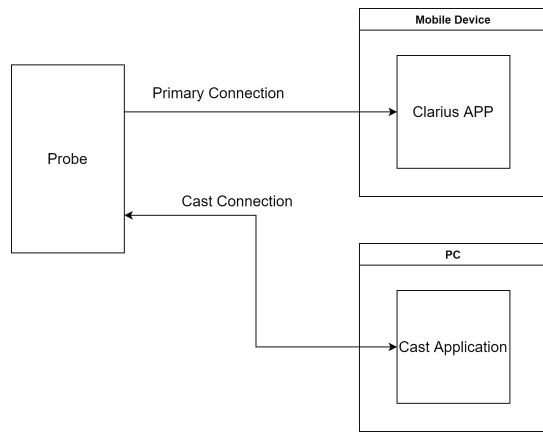


Fig. 7: Connections between the Clarius probe, Cast API and the Clarius App. All the connects are wireless.

The image is then passed as input to the trained NSeg-UNet model which yields a segmentation map of the needle. In order to reduce latency in the segmentation, inference should be synced with the ultrasound device streaming rate. This can be achieved by sampling the input image stream at appropriate intervals and only making predictions on these images.

In our scenario, the Clarius ultrasound system has a frame rate of 24 fps, the NSeg-UNet model has an inference time of about 0.1s on the laptop, and an overall integration time of 0.13s. To sync the model’s inference with the ultrasound display rate, we chose a sampling interval of 8 frames and this was determined empirically starting from an interval of 1




Fig. 8: Post processing results. The first column shows the original B-mode images. The second column shows the results of the post processing step that would be displayed to the radiologist on a graphical user interface (GUI).

frame. Sampling the input image stream ensures that the user is always seeing a segmentation of the current image.

In the post-processing step, a threshold is applied to the segmentation map as required. A higher threshold results into sharper predictions but could result in broken contours in the prediction. We use a threshold of 0.5 as the default and then obtain the contour of the region containing the needle. The contour is then overlaid onto the streamed image using OpenCV. Fig. 8 shows the final image with the contour in red. Finally, the image is displayed on a graphical user interface (GUI) on the laptop.

V. CONCLUSION

This paper presents a novel approach for real-time needle segmentation in 2D ultrasound for handheld ultrasound systems. The proposed method was tested across multiple datasets consisting of human, chicken, porcine and bovine images ultrasound images. The model achieves an intersection over union (IoU) of 0.75 and a dice coefficient of 0.851 outperforming state of the art methods on the human dataset. The method can accurately segment needles for both shallow and steep angles of insertion and in the presence of high intensity artifacts. The proposed method was integrated into the processing pipeline of an ultrasound system, the Clarius L7, and qualitatively evaluated in real-time. The method achieves accurate real-time needle segmentation with about 8 fps.

Our proposed approach could be used to augment clinicians and thus improve target detection rate, reduce procedure time, and improve success rate of minimally invasive procedures such as biopsies, nerve blocks and regional anesthesia. Future work will investigate utilizing spatial-temporal information to improve performance on needles that veer out of the field of view. We will also investigate formulation of the sampling rate so as to find a general optimal approach on the number of frames to skip. Further, we shall also undertake clinical validation of the integrated system in in-vivo animal studies.

SUPPLEMENTARY MATERIAL

To further demonstrate the performance of the proposed method in real-time, we recorded evaluation sessions on *ex vivo* chicken tissue. In addition, more qualitative results are provided for each dataset category. This will be available at <http://ieeexplore.ieee.org>.

ACKNOWLEDGMENT

This work was funded by the Government of Uganda through the Makerere University Research and Innovations Fund. The authors would like to acknowledge Rutgers University for providing access to the aggregated human image datasets and part of the *ex vivo* data.

REFERENCES

- [1] D. L. Brown, “Chapter 1: Local Asthetics and Regional Anesthesia Equipment,” in *Atlas of Regional Anesthesia*, 4th ed. Elsevier Health Sciences, Jan. 2010, google-Books-ID: w92VFBqCtPYC.
- [2] H. Yang, C. Shan, A. F. Kolen, and P. H. N. de With, “Medical instrument detection in ultrasound-guided interventions: A review.” [Online]. Available: <http://arxiv.org/abs/2007.04807>

- [3] M. J. Barrington and R. Kluger, "Ultrasound guidance reduces the risk of local anesthetic systemic toxicity following peripheral nerve blockade," vol. 38, no. 4, pp. 289–299.
- [4] J. Samantha and C. Barbara, "Healthcare technology basics," in *Introduction to Clinical Engineering*, ser. 1st Edition, p. 280.
- [5] J. S. Hopstaken, L. de Jong, and J. J. Fütterer, "Increased needle visibility in ultrasound-guided percutaneous liver biopsy by an echogenic sheath: A proof of concept study in a human cadaver," vol. 44, no. 6, pp. 959–967. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8172504/>
- [6] A. L. Falkowski, J. A. Jacobson, M. T. Freehill, and V. Kalia, "Hand-held portable versus conventional cart-based ultrasound in musculoskeletal imaging," vol. 8, no. 2, p. 2325967119901017. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7016312/>
- [7] *Portable Ultrasound Devices in the Pre-Hospital Setting: A Review of Clinical and Cost-Effectiveness and Guidelines*, ser. CADTH Rapid Response Reports. Canadian Agency for Drugs and Technologies in Health. [Online]. Available: <http://www.ncbi.nlm.nih.gov/books/NBK349868/>
- [8] T. L. de Jong, N. J. van de Berg, L. Tas, A. Moelker, J. Dankelman, and J. J. van den Dobbelsteen, "Needle placement errors: do we need steerable needles in interventional radiology?" vol. 11, pp. 259–265. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6080661/>
- [9] H. Elsharkawy, R. Babazade, S. Kolli, H. Kalagara, and M. L. Soliman, "The infinity plus ultrasound needle guidance system improves needle visualization during the placement of spinal anesthesia," vol. 69, no. 4, pp. 417–419. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4967642/>
- [10] K. Umbarje, R. Tang, R. Randhawa, A. Sawka, and H. Vaghadia, "Out-of-plane brachial plexus block with a novel SonixGPS(TM) needle tracking system," vol. 68, no. 4, pp. 433–434.
- [11] A. Pourtaherian, F. Ghazvinian Zanjani, S. Zinger, N. Mihajlovic, G. C. Ng, H. H. M. Korsten, and P. H. N. de With, "Robust and semantic needle detection in 3d ultrasound using orthogonal-plane convolutional neural networks," vol. 13, no. 9, pp. 1321–1333.
- [12] H. Younes, S. Voros, and J. Troccaz, "Automatic needle localization in 3D ultrasound images for brachytherapy," in *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, Apr. 2018, pp. 1203–1207, iSSN: 1945-8452.
- [13] J. Y. Lee, M. Islam, J. R. Woh, T. S. M. Washeem, L. Y. C. Ngoh, W. K. Wong, and H. Ren, "Ultrasound needle segmentation and trajectory prediction using excitation network," vol. 15, no. 3, pp. 437–443. [Online]. Available: <https://doi.org/10.1007/s11548-019-02113-x>
- [14] A. Wijata, J. Andrzejewski, and B. Pyciński, "An automatic biopsy needle detection and segmentation on ultrasound images using a convolutional neural network," vol. 43, no. 5, pp. 262–272.
- [15] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation." [Online]. Available: <http://arxiv.org/abs/1505.04597>
- [16] F. van Beers, A. Lindström, E. Okafor, and M. Wiering, "Deep Neural Networks with Intersection over Union Loss for Binary Image Segmentation:," in *Proceedings of the 8th International Conference on Pattern Recognition Applications and Methods*. Prague, Czech Republic: SCITEPRESS - Science and Technology Publications, 2019, pp. 438–445. [Online]. Available: <http://www.scitepress.org/DigitalLibrary/Link.aspx?doi=10.5220/0007347504380445>
- [17] O. Oktay, J. Schlemper, L. L. Folgoc, M. Lee, M. Heinrich, K. Misawa, K. Mori, S. McDonagh, N. Y. Hammerla, B. Kainz, B. Glocker, and D. Rueckert, "Attention U-Net: Learning Where to Look for the Pancreas," May 2018, arXiv:1804.03999 [cs]. [Online]. Available: <http://arxiv.org/abs/1804.03999>
- [18] Clarius 17 HD3. [Online]. Available: <https://store.clarius.com/products/17-hd3-handheld-ultrasound-scanner>