ORIGINAL ARTICLE

Time-aware deep neural networks for needle tip localization in 2D ultrasound

Cosmas Mwikirize[1] · Alvin B. Kimbowa[1] · Sylvia Imanirakiza[1] · Andrew Katumba[1] · John L. Nosher[2] · Ilker Hacihaliloglu[2,3]

Abstract

Purpose Accurate placement of the needle is critical in interventions like biopsies and regional anesthesia, during which incorrect needle insertion can lead to procedure failure and complications. Therefore, ultrasound guidance is widely used to improve needle placement accuracy. However, at steep and deep insertions, the visibility of the needle is lost. Computational methods for automatic needle tip localization could improve the clinical success rate in these scenarios.

Methods We propose a novel algorithm for needle tip localization during challenging ultrasound-guided insertions when the shaft may be invisible, and the tip has a low intensity. There are two key steps in our approach. First, we enhance the needle tip features in consecutive ultrasound frames using a detection scheme which recognizes subtle intensity variations caused by needle tip movement. We then employ a hybrid deep neural network comprising a convolutional neural network and long short-term memory recurrent units. The input to the network is a consecutive plurality of fused enhanced frames and the corresponding original B-mode frames, and this spatiotemporal information is used to predict the needle tip location.

Results We evaluate our approach on an ex vivo dataset collected with in-plane and out-of-plane insertion of 17G and 22G needles in bovine, porcine, and chicken tissue, acquired using two different ultrasound systems. We train the model with 5000 frames from 42 video sequences. Evaluation on 600 frames from 30 sequences yields a tip localization error of 0.52 ± 0.06 mm and an overall inference time of 0.064 s (15 fps). Comparison against prior art on challenging datasets reveals a 30% improvement in tip localization accuracy.

Conclusion The proposed method automatically models temporal dynamics associated with needle tip motion and is more accurate than state-of-the-art methods. Therefore, it has the potential for improving needle tip localization in challenging ultrasound-guided interventions.

Keywords Needle tip localization · Ultrasound · LSTM · Minimally invasive procedures

Introduction

This paper focuses on the clinical challenge of needle tip localization during minimally invasive ultrasound-guided procedures such as regional anesthesia and biopsies [1–3]. These procedures are typically performed using the

✉ Cosmas Mwikirize
  cosmas.mwikirize@mak.ac.ug

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

2 Department of Radiology, Rutgers Robert Wood Johnson Medical School, New Brunswick, NJ 08901, USA

3 Department of Biomedical Engineering, Rutgers University, Piscataway, NJ 08854, USA

conventional 2D B-mode ultrasound, and the needle may be inserted using one of two techniques: in-plane, where the needle is inserted parallel to the ultrasound beam, or out-of-plane, where the needle insertion plane and the ultrasound beam are perpendicular. Although in-plane insertion should ideally produce a conspicuous needle shaft and tip, it is common for the needle to veer away from the narrow field of view, producing no shaft and/or a low-intensity tip. Out-of-plane insertion, on the other hand, usually produces no shaft information and a low-intensity tip. In either case, the interventional radiologist must rely on recognizing low-intensity features associated with tip motion while concurrently manipulating the ultrasound transducer and the needle, a challenging task, which is exacerbated by motion artifacts, noise, and high-intensity anatomical artifacts. Therefore,

accurate and consistent visualization of the needle tip is often difficult to achieve. Consequently, it is common for a non-experienced radiologist to miss the anatomical targets, and this could lead to injury, increased hospital stays, and reduced efficacy of procedures.

To address this challenge, several methods have been proposed, and these can broadly be categorized as hardware or software based. On the hardware front, mechanical needle guides, which are designed to keep the needle aligned with the ultrasound beam, are prominent [4]. Some needle guides have predetermined angles of approach, while others permit minor adjustments, but overall, needle guides are inefficient in procedures where fine trajectory adjustments are required, or out-of-plane insertion is desired. Another method involves the integration of sensors at the needle tip [5, 6], but this makes the needles more expensive. 3D/4D US gives a wider field of view, overcoming the limitations of 2D ultrasound [7], but current technology has poor resolution and a low frame rate, making it unsuitable for real-time applications. Electromagnetic/optical tracking systems [8–13] have been proposed, but they necessitate specialized needles and probes, thus adding a huge cost to the basic ultrasound system. Furthermore, electromagnetic systems are susceptible to interference from metallic objects in the operating environment. Lastly, robotic systems facilitate autonomous or semi-autonomous needle insertion [14, 15], but they are expensive and not practical for routine procedures.

Software-based methods, on the other hand, rely on image analysis methods applied to the B-mode ultrasound images, to facilitate automatic needle recognition. This review will focus on machine learning-based methods, which have been shown to outperform classical computer vision methods. Hatt and colleagues suggested a method for needle localization, utilizing an Adaboost classifier and beam-steered ultrasound images [16]. Their approach requires a visible needle shaft, which is easier to obtain on ultrasound systems with beam steering capability and difficult otherwise. Moreover, this method would not work for out-of-plane needles. Beigi et al. presented a learning-based method for segmentation of imperceptible needle motion, relying on optical flow and support vector machines [17], but the method is computationally expensive (1.18 s per frame). Pouraherian and colleagues have proposed a framework for needle detection in 3D ultrasound, using orthogonal-plane convolutional networks [18]. As earlier noted, 3D ultrasound is not widely available, and 2D ultrasound remains the standard of care.

In our previous work, we demonstrated deep learning approaches for needle shaft and tip localization, based on convolutional neural networks (CNNs). One instance [19] focuses on needle tip localization for in-plane needles, in individual frames, when the shaft is at least partially visible. Recently, we presented other methods targeting challenging

procedures where the needle shaft may not be visible [20, 21]. This latter work employed a novel foreground detection scheme, in which the needle tip feature is extracted from consecutive frames, using dynamic background information. The enhanced needle frames are then fed to CNNs, one at a time, for needle tip localization. Although the methods in [20, 21] achieved good tip localization accuracy and high computational efficiency, tip localization was affected by motion artifacts in the clinical setting, for example, those arising from physiological activity such as breathing or pulsation.

In this paper, we build on our previous work and propose a more robust and accurate needle tip localization strategy, suitable for localization of both in-plane and out-of-plane needles under 2D ultrasound guidance, in which there is no shaft information. In the new approach, we enhance the needle tip using the foreground detection scheme introduced in [20, 21]. However, instead of using individual enhanced frames as input to a neural network, we feed a consecutive fused image sequence, derived from fusion of the enhanced frames and the corresponding B-mode frames, to a time-aware neural network which consists of a unified convolutional neural network (CNN) and a long short-term memory (LSTM) recurrent neural network. The CNN acts as a feature extractor, with stacked convolutional layers which progressively create a hierarchy of more abstract features. The LSTM on the other hand models temporal dependencies in time-series data. The combination of CNNs and LSTMs is thus able to capture time dependencies on features extracted by convolutional operations, thus supporting sequence prediction. In our case, the network learns spatiotemporal features associated with needle tip movement, for example, needle tip appearance and trajectory information, and successfully localizes the needle tip in the presence of abrupt intensity changes and motion artifacts.

The main contribution of this paper is a novel CNN–LSTM learning approach, optimized for learning temporal relationships emanating from needle tip motion events. Since the proposed framework does not rely on needle shaft visibility, it is appropriate for the localization of thin needles and both in-plane and out-of-plane trajectories. We demonstrate that the new approach has a significant edge over the prior art, thus making it a good candidate for integration in a computer-assisted interventional system for needle tip localization.

Methods

The proposed method is designed for hand-held 2D US probes during in-plane and out-of-plane needle insertion. We split the problem of motion-based needle localization into two parts: (a) motion detection in each frame and (b)

spatial–temporal feature extraction. The two components of the proposed method are illustrated in Fig. 1: (1) Similar to our previous work in [20, 21], we extract needle tip features caused by otherwise imperceptible scene changes arising from needle motion in the 2D ultrasound image ("Needle tip feature extraction from ultrasound frame sequences" section). This is achieved by logical subtraction of the current frame (foreground) from the previous frame, which acts as a dynamic reference (background). We achieve an enhanced needle frame without requiring a priori information about the needle trajectory, (2) we fuse the enhanced tip images and the corresponding B-mode images and feed multiple consecutive fused images to a novel CNN–LSTM framework which localizes the needle tip in the last frame of the sequence ("Needle tip localization" section). In the next sections, we describe the main aspects of these methods.

Dataset overview

The 2D B-mode US data used in this work were collected using two imaging systems: SonixGPS (Analogic Corporation, Peabody, MA, USA) with a hand-held C5-2/60 curvilinear probe, at 30 frames-per-second (fps) and 2D hand-held wireless US (Clarius C3, Clarius Mobile Health Corporation, Burnaby, BC, Canada) at 24 fps. Three needle types were used in our experiments: a 17G SonixGPS vascular access needle (Analogic Corporation, Peabody, MA, USA),

a 17G Tuohy epidural needle (Arrow International, Reading, PA, USA), and a 22G spinal Quincke-type needle (Becton, Dickinson and Company, Franklin Lakes, NJ, USA). The needles were inserted in freshly excised bovine, porcine, and chicken tissue, with the chicken tissue overlaid on a lumbosacral spine phantom, in-plane (25° to 60°) and out-of-plane up to a depth of 70 mm. For experiments conducted with the SonixGPS needle, we collected tip localization data from the electromagnetic (EM) tracking system (Ascension Technology Corporation, Shelburne, VT, USA). The data were collected by a clinician who introduced motion seen in clinical situations, via large probe pressure while concurrently rotating the transducer. We collected 80 video sequences (45 in-plane, 35 out-of-plane: 40 with SonixGPS system and 40 with Clarius C3 system), with each video sequence having more than 300 frames. The experiment particulars are shown in Table 1. Data for training and validation were extracted from 42 sequences. Test experiments were conducted on 600 frames extracted from 30 left out sequences. The test data were chosen to focus on sequences with large motion artifacts.

Needle tip feature extraction from ultrasound frame sequences

We consider a temporal sequence of ultrasound frames, with each frame denoted by the spatiotemporal function
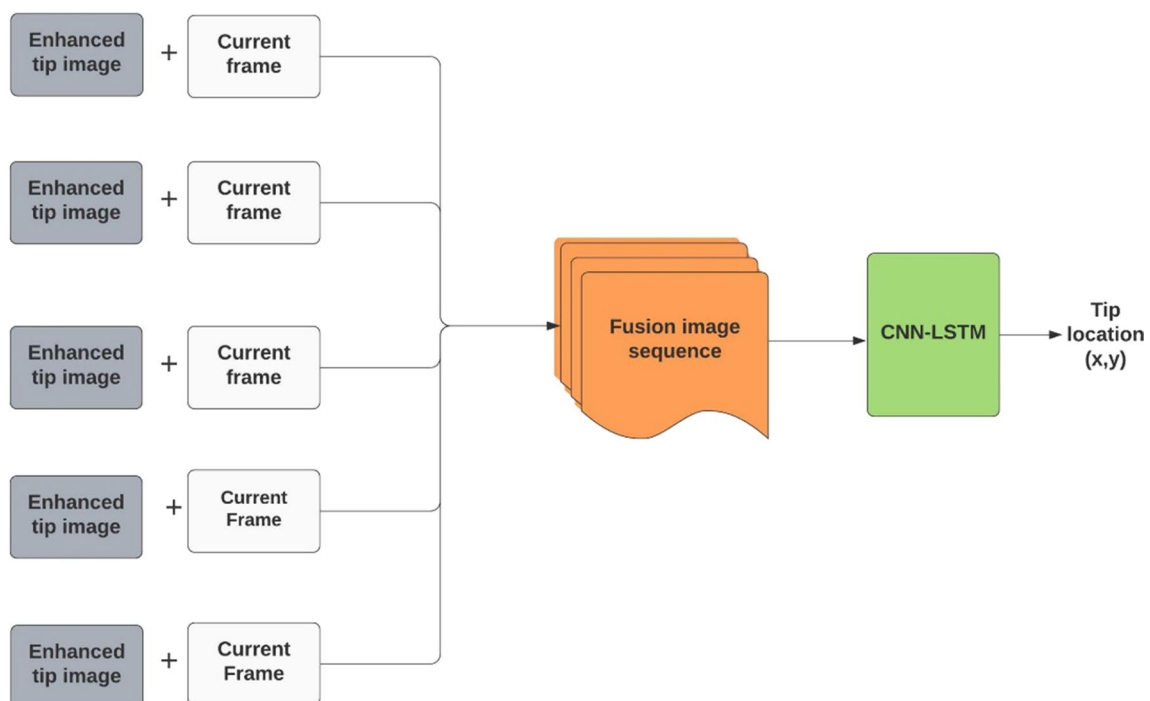

Fig. 1 Block diagram of the proposed framework for needle tip localization from spatiotemporal information. The input to the neural network consists of a consecutive sequence of five fused images derived from enhanced tip images and the corresponding B-mode images

Table 1 Experimental details for 2D US data collection

Imaging system	Needle type, dimensions, insertion profile, tissue	# of videos	Pixel size (mm)
SonixGPS	17G SonixGPS (1.5 mm, 70 mm), IP, bovine	5	0.17
	17G Tuohy (1.5 mm, 90 mm), IP, bovine	5	0.17
	22G BD (0.7 mm, 90 mm), IP, bovine	5	0.17
	22G BD (0.7 mm, 90 mm), OP, bovine	5	0.17
	17G SonixGPS (1.5 mm, 70 mm), OP, porcine	5	0.17
	17G Tuohy (1.5 mm, 90 mm), IP, porcine	5	0.17
	17G SonixGPS (1.5 mm, 70 mm), IP, chicken	5	0.17
	22G BD (0.7 mm, 90 mm), OP, chicken	5	0.17
Clarius C3	17G SonixGPS (1.5 mm, 70 mm), IP, bovine	15	0.24
	17G Tuohy (1.5 mm, 90 mm), IP, bovine	15	0.24
	22G BD (0.7 mm, 90 mm), OP, porcine	5	0.24
	22G BD (0.7 mm, 90 mm), OP, chicken	5	0.24

IP in-plane insertion, *OP* out-of-plane insertion

$US(x, y, t)$, where t represents the time index and (x, y) are the spatial indexes. We want to broadly categorize the pixels in each frame as either a foreground (needle tip) or background (tissue). To achieve this, we utilize a dynamic background subtraction model which we first introduced in [21]. To enhance the needle tip in frame $US(x, y, t)$, we consider $US(x, y, t-1)$, as a background, and perform the operation:

$$US_E(x, y) = US(x, y, y) \land US(x, y, t-1)^c. \quad (1)$$

Here, $\land$ represents the bitwise AND logical operation, and (1) calculates the conjunction of pixels in the current frame and the logical complement of the preceding frame. Therefore, the output, $US_E(x, y)$, contains only pixels that are in the present frame and not in the previous frame. (1) is remarkably efficient at extracting the needle tip since it considers any spatiotemporal intensity variation between consecutive frames. To further enhance the needle tip, the output of (1) is passed through a median filter with a 7×7 kernel. Figure 2 illustrates a typical output of this enhancement approach on four consecutive frames (yielding three consecutive enhanced frames). The process of needle tip enhancement is almost cost-free and takes 0.0016 s on a 512×512 frame. Certainly, there could be other motion
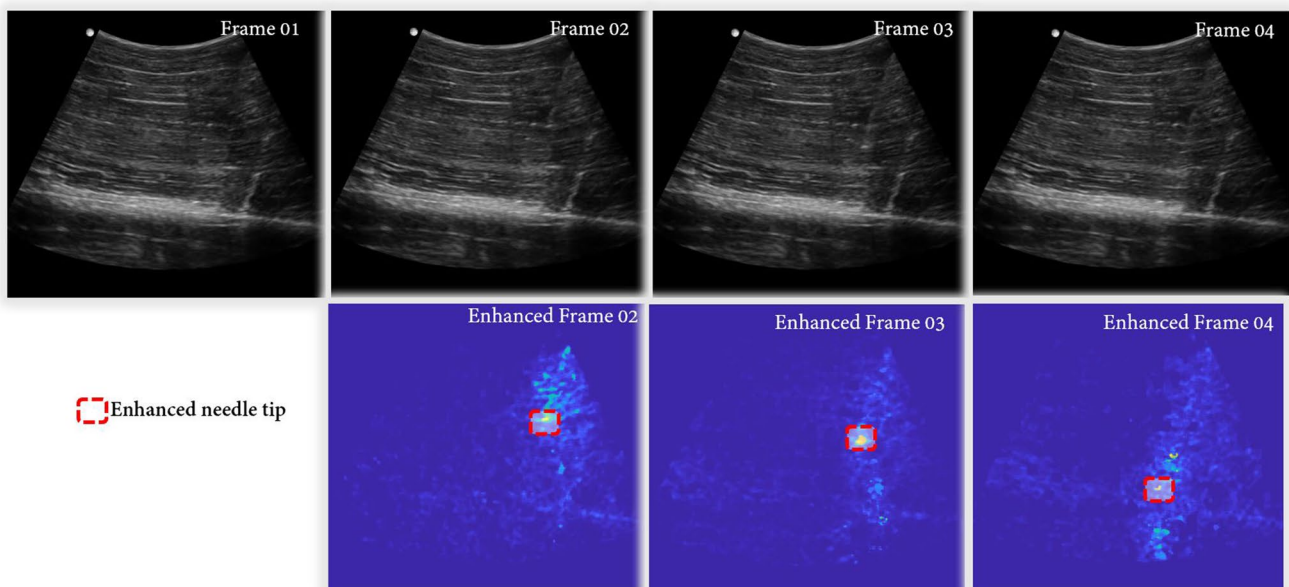

Fig. 2 The needle tip enhancement process applied to four consecutive ultrasound frames, from data collected with in-plane insertion of a 17G needle in porcine tissue. Row 1: Original B-mode ultrasound frames, $US(x, y, t)$. Row 2: Needle tip enhanced images, $US_E(x, y)$. Notice that without the enhancement step, the needle tip is not easy to visualize with the naked eye

artifacts picked up by (1), and therefore, the learning-based approach described next is important, to accurately localize the needle tip from $US_E(x, y)$.

Needle tip localization

In "Needle tip feature extraction from ultrasound frame sequences" section, we have derived enhanced tip images $US_E(x, y)$ from consecutive frames in a B-mode ultrasound sequence. It is expected that the tip feature in $US_E(x, y)$ will exhibit a high intensity. Nevertheless, there are often motion artifacts or high-intensity artifacts arising from anatomy, which could be equally significant in the enhanced image. For this reason, we cannot rely on the highest intensity in $US_E(x, y)$ to be the needle tip. To accurately localize the needle tip, we feed a plurality of fused images, comprising a combination of the tip enhanced images and the corresponding B-mode images, to a CNN–LSTM network, which associates spatial needle tip features in each enhanced frame, with the temporal information across the frame sequence. Next, we describe this deep learning framework, emphasizing the aspects that are unique to our approach.

CNN architecture

We introduce a new deep neural network for needle tip localization, whose architecture, shown in Fig. 3 and Table 2, combines convolutional and recurrent layers. The convolutional layers extract abstract representations of the input image data in feature maps. The recurrent layers implemented as LSTM layers pass previous hidden states to the next step of the sequence. The overall network holds information on previously seen data and uses it to make decisions.

The input to the network consists of a sequence of five fused images, with each image consisting of the enhanced tip

Table 2 Architecture of the CNN–LSTM network

N	Layer name	Output dimensions	Kernel	Filter number
0	Input	$5 \times 512 \times 512 \times 1$	–	–
1	Conv 1	$5 \times 512 \times 512 \times 16$	3	16
2	Maxpool	$5 \times 256 \times 256 \times 16$	2	–
3	Conv 2	$5 \times 256 \times 256 \times 16$	3	16
4	Maxpool	$5 \times 128 \times 128 \times 16$	2	–
5	Conv 3	$5 \times 128 \times 128 \times 32$	3	32
6	Maxpool	$5 \times 64 \times 64 \times 32$	2	–
7	Conv 4	$5 \times 64 \times 64 \times 32$	3	32
8	Maxpool	$5 \times 32 \times 32 \times 32$	2	–
9	LSTM 1	$5 \times 32 \times 32 \times 16$	3	16
10	Maxpool	$5 \times 16 \times 16 \times 16$	2	–
11	LSTM 2	$5 \times 16 \times 16 \times 8$	3	8
12	Maxpool	$5 \times 8 \times 8 \times 8$	2	–
13	LSTM 3	$8 \times 8 \times 4$	3	4
14	Maxpool	$4 \times 4 \times 4$	2	–
15	FC 1	20	–	–
16	FC 2	2	–	–

image + the corresponding B-mode image. Using the fusion strategy instead of using only the enhanced tip image input is important because in case the needle tip does not move within the five-frame consecutive sequence, tip information is still available in the input, derived from the original B-mode frame (if there is no tip motion, $US_E(x, y)$ is ideally all zeros, and does not contain any tip information). The frame number of 5 has been empirically determined based on optimizing computational efficiency of the network, while considering typical ultrasound frame rate and needle insertion speed for the data. Each input image is resized to 512×512. The input data feeds a series of four convolutional layers, which apply the respectively defined convolutional operations to each temporal slice in the input. The
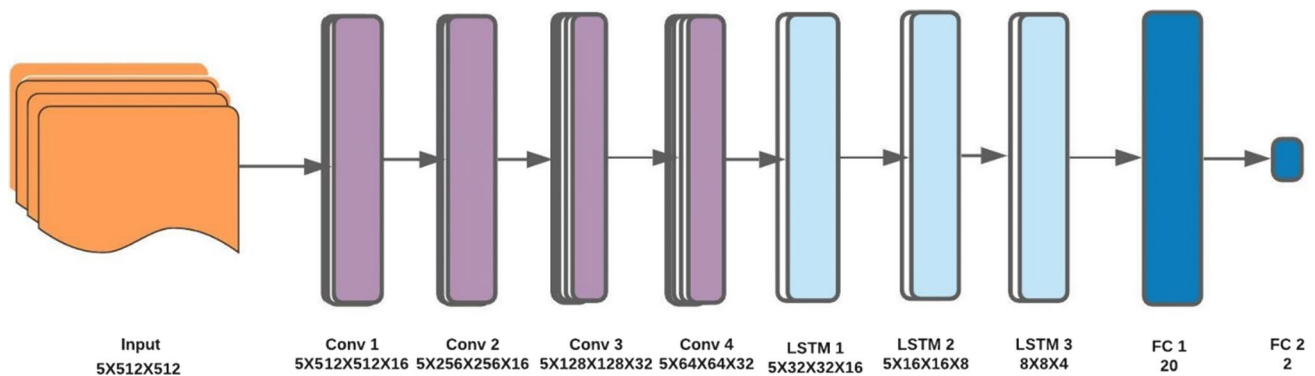

Fig. 3 Architecture of the deep CNN–LSTM network for needle tip localization. L-R: Input data from five fused images (enhanced tip image + corresponding B-mode image) are processed by four time-distributed convolutional layers. These are followed by convolutional LSTM layers which model temporal dynamics associated with needle tip motion from the prior extracted activation maps, and lastly, two fully connected layers, whose final output is the tip location (x, y)

size of feature maps varies in different convolutional layers varies as shown in Table 1. All convolutional layers employ rectified linear units (ReLUs) activations, whose nonlinear function is defined as $\sigma(x) = \max(0, x)$. Each convolutional layer is followed by a max pooling layer, which also applies the max pooling operation to each temporal slice in the input at that stage. The convolutional max pool sequence is then followed by three convolutional LSTM layers, whose output size mirrors the input temporal sequence. Like the convolutional layers, the convolutional LSTM layers are interspersed with max pool layers. The last LSTM layer is followed by another max pooling layer, and two fully connected layers of size 20 and 2, respectively, since the desired model output is the tip position (x, y).

Training details

For data collected with the SonixGPS system, we derive ground-truth needle tip locations (x, y) in each frame using the inbuilt electromagnetic tracking system, cross-checked by an expert interventional radiologist with more than 25 years of experience. For data collected with the Clarius C3 system (which does not have a tracking solution), the ground-truth tip locations are determined via manual labeling by the expert radiologist. The labels are rescaled to be in the range [0, 1]. Following our desired output, we train our network as a regression CNN–LSTM, using Adam optimizer and mean squared error (MSE) loss. We trained and evaluated the model using Tensorflow in Google Colab, powered by the 12 GB Tesla K80 GPU parallel computing platform. The needle tip enhancement method was implemented in MATLAB 2019b on a 3.6 GHz Intel (R) Core™ i7 16 GB CPU Windows PC.

Experimental results and discussion

We evaluated the performance of the proposed method by comparing the automatically localized tip with the ground truth obtained from the electromagnetic tracking system for data collected with the SonixGPS needle. For data collected with needles without tracking capability (Tuohy and BD needles), the ground truth was determined by an expert sonographer. Tip localization accuracy was determined from the Euclidean distance between the corresponding measurements [20, 21].

Qualitative results

In Fig. 4, we illustrate the needle tip localization results for three consecutive frames, for in-plane needle insertion. These results show that the needle tip is accurately localized, even when it is not easily discernible with

the naked eye. The proposed method performs well in the presence of high-intensity artifacts in the rest of the ultrasound image. Meanwhile, the proposed method is not sensitive to the type and size of the needle used in the experiments.

Model comparison

We use two metrics to compare the performance of the proposed method with existing state-of-the-art methods and variants of the current approach: tip localization error and total processing time. This comparison is shown in Table 3.

On our test data of 600 frames extracted from 30 ultrasound sequences, the proposed method achieved a tip localization error of 0.52 ± 0.06 mm and an overall computation time of 0.064 s (0.0016 s for frame enhancement and 0.062 s for model inference). Here, we consider the computational cost for enhancing one frame since we used a sliding window approach with a frame overlap of four frames in our model. We trained a similar CNN–LSTM model with input of raw B-mode ultrasound frames (without the tip enhancement step), while keeping the network's architecture and training detail constant. The resulting model performed poorly, with a tip localization error of 5.92 ± 1.5 mm. This was not unexpected: Without enhancement, the needle tip feature is often not distinct. Thus, the model could not learn the associated features, and this led to poor performance.

Next, we tried the approach described in [21], where the needle tip is enhanced, using a similar approach to the one described in this paper, and the resulting image is fed to a network derived from the YOLO architecture [22] for needle tip detection. For a fair comparison, localization errors above 2 mm (24% with [21] compared to 6% of the test data with the proposed method) were not considered. This model achieves a localization error of 0.79 ± 0.15 mm which is higher than that of the proposed method. A one-tailed paired t test shows that the difference between the localization errors from the proposed method and the method in [21] is statistically significant ($p < 0.005$).

We also compared the proposed method with the approach of [20], where the needle tip is also first enhanced, and then fed to a model cascade of classifier and a location regressor. This approach achieves a localization error of 0.74 ± 0.08 mm (81% of the test data below 2 mm error), with statistically significant under-performance compared to the proposed method ($p < 0.005$). It is not hard to understand why the proposed method is superior to the methods in [20, 21]: The current approach takes as input an enhanced sequence of needle tip images and hence learns spatiotemporal information related to both structure and motion behavior of the needle tip. The previous methods, on the other hand, take in one frame at a time and do not learn any temporal
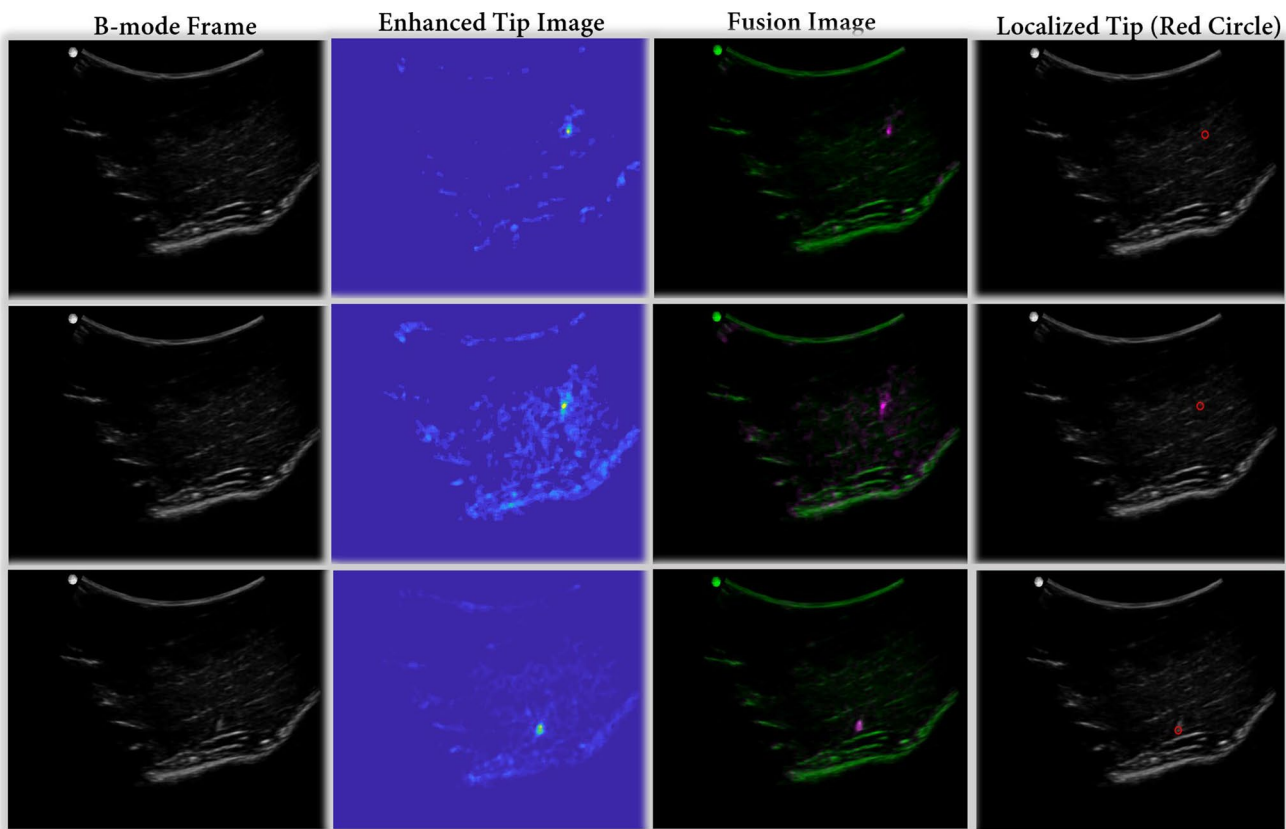
Fig. 4 Needle tip enhancement and localization results in three frames from one ultrasound sequence, obtained with in-plane insertion of a 22G needle in chicken tissue. Column 1 shows the original B-mode ultrasound frames. Note that the needle tip is difficult to observe and shaft information is unavailable. Column 2 shows the needle tip enhanced image $US_E(x, y)$ obtained using the method described in "Needle tip feature extraction from ultrasound frame sequences" section. Here, the tip appears as a characteristic high intensity in the image. Column 3 shows the fused image, derived from $US_E(x, y)$ and the corresponding B-mode image. A consecutive sequence of five fused images is input to the CNN–LSTM network. Column 4 shows the tip localization result obtained from the CNN–LSTM model described in "Needle tip localization" section

Table 3 Comparing performance of the proposed method with state-of-the-art methods and alternative implementations

Method	Error (mm)	Processing time (s)	Error below 2 mm (%)
Proposed method	0.52 ± 0.06	0.064	94
Proposed method with raw input	5.92 ± 1.5	0.062	4
Method in [21]	0.79 ± 0.15	0.092	76
Method in [20]	0.74 ± 0.08	0.021	81

information. This makes them prone to artifacts that may look like the needle tip, especially when they are outside the needle trajectory.

Conclusions

In this paper, we have presented a novel approach for localization of the needle tip in 2D ultrasound, focusing on operating scenarios when the needle shaft is not visible and/or the needle tip does not have a characteristic high intensity. The proposed method achieves a better tip localization accuracy on challenging datasets, when compared to state-of-the-art methods in [20] (30% improvement) and [21] (34% improvement). Although the method in [20] achieves better computational efficiency, the current approach has a competitive processing time of 0.064 s (15 fps). Moreover, this processing time can be improved with code optimization. The proposed approach has been tested on data collected with two imaging systems: one cart-based and the other hand-held. The data were collected using three different needle types and three different tissue types. From the results, it is observed that the proposed method performs well for both in-plane and out-of-plane

needle insertions and does not require the needle tip to appear as a high intensity or a continuous feature in the ultrasound image. Furthermore, the localization accuracy does not vary significantly with needle type and size.

We believe our proposed work would augment clinicians and thus improve the target detection rate, procedure time, and success rate. Nevertheless, we would like to outline some of the aspects of our work that require further attention: (1) We have only evaluated the proposed method on ex vivo data. Although we gave attention to introducing clinically relevant and challenging situations during data collection, in vivo evaluation is still required to fully assess the clinical utility of the proposed method. (2) Although we collected scans from two different ultrasound machines, we did not investigate the domain invariance of our method. (3) The needles that are used in this study are 17G and 22G needles which have minimal bending during insertion into the ex vivo tissue. For very small needles, tissue motion introduced by the needle insertion could be exceedingly small resulting in no clear enhanced needle image. As part of our future work, we will investigate the success of our method on small bending needles. (4) The proposed work is based on the information extracted from B-mode ultrasound data. Currently, some medical device companies provide access to radio frequency (RF) ultrasound data. In the future, one potential pivot, which could provide improved localization, would be the fusion of information extracted from RF data with B-mode ultrasound data. (5) The proposed method utilizes supervised learning, which requires the labeling of needle tip location. A clinical evaluation would require the collection and labeling of large-scale data, which is a time-consuming procedure. As part of our future work, we will also investigate using a semi-supervised learning method to limit the annotation efforts.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Ethical approval This article does not contain any studies with human participants or animals performed by any of the authors.

Informed consent This article does not contain patient data.

References

1. Marhofer P, Fritsch G (2017) Safe performance of peripheral regional anaesthesia: the significance of ultrasound guidance. Anaesthesia 72(4):431–434
2. Huang J, Li J, Wang H (2018) The principles and procedures of ultrasound-guided anesthesia techniques. Cureus 10(7):e2980
3. Kim SY, Chung HW, Oh TS, Lee J (2017) Practical guidelines for ultrasound-guided core needle biopsy of soft-tissue lesions: transformation from beginner to specialist. Korean J Radiol 18(2):361–369
4. Elsharkawy H, Babazade R, Kolli S, Kalagara H, Soliman ML (2016) The Infiniti plus ultrasound needle guidance system improves needle visualization during the placement of spinal anesthesia. Korean J Anesthesiol 69(4):417–419
5. Lu H, Li J, Lu Q, Bharat S, Erkamp R, Chen B, Drysdale J, Vignon F, Jain A (2014) A new sensor technology for 2D ultrasound-guided needle tracking. MICCAI 17(Pt. 2):389–396
6. Xia W, West S, Finlay M, Mari J, Ourselin S, David A, Desjardins A (2017) Looking beyond the imaging plane: 3D needle tracking with a linear array ultrasound probe. Sci Rep 7(1):3674
7. Arif M, Moelker A, van Walsum T (2018) Needle tip visibility in 3D ultrasound images. Cardiovasc Intervent Radiol 41(1):145–152
8. Umbarje K, Tang R, Randhawa R, Sawka A, Vaghadia H (2013) Out-of-plane brachial plexus block with a novel SonixGPS (TM) needle tracking system. Anaesthesia 68:433–434
9. Gadsden J, Latmore M, Levine DM (2015) Evaluation of the eZono 4000 with eZGuide for ultrasound-guided procedures. Expert Rev Med Dev 12:251–261
10. McVicar J, Niazi AU, Murgatroyd H, Chin KJ, Chan VW (2015) Novice performance of ultrasound-guided needling skills: effect of a needle guidance system. Reg Anesth Pain Med 40:150–153
11. Najafi M, Abolmaesumi P, Rohling R (2015) Single-camera closed-form real-time needle tracking for ultrasound-guided needle insertion. Ultrasound Med Biol 41:2663–2676
12. Xia W, Mari JM, West SJ, Ginsberg Y, David AL, Ourselin S, Desjardins AE (2015) In-plane ultrasonic needle tracking using a fiber-optic hydrophone. Med Phys 42:5983–5991
13. Stolka PJ, Foroughi P, Rendina M, Weiss CR, Hager GD, Boctor EM (2014) Needle guidance using handheld stereo vision and projection for ultrasound-based interventions. MICCAI 17(Pt.2):684–691
14. Neubash Z, Shoham M (2010) Ultrasound-guided robot for flexible needle steering. IEEE TBME 57(4):799–805
15. Welleweerd MK, Siepel FJ, Groenhuis V, Veltman J, Stramigioli S (2020) Design of an end effector for robot-assisted ultrasound-guided breast biopsies. Int J Comput Assist Radiol Surg 15:681–690
16. Hatt CR, Ng G, Parthasarathy V (2015) Enhanced needle localization in ultrasound using beam steering and learning-based segmentation. Comput Med Imaging Graph 41:46–54
17. Beigi P, Rohling R, Salcudean S, Ng G (2017) CASPER: computer-aided segmentation of imperceptible motion—a learning-based tracking of an invisible needle in ultrasound. Int J Comput Assist Radiol Surg 12(11):1857–1866
18. Pourtaherian A, Ghazvinian Zanjani F, Zinger S, Mihajlovic N, Ng G, Korsten H, With P (2018) Robust and semantic needle detection in 3D ultrasound using orthogonal-plane convolutional neural networks. Int J Comput Assist Radiol Surg 13(9):1321–1333
19. Mwikirize C, Nosher JL, Hacihaliloglu I (2018) Convolutional neural networks for real-time needle detection and localization in 2D ultrasound. Int J Comput Assist Radiol Surg 13(5):647–657
20. Mwikirize C, Nosher JL, Hacihaliloglu I (2019) Single shot needle localization in 2D ultrasound. In: MICCAI 2019. MICCAI LNCS, 11768, pp 637–645

21. Mwikirize C, Nosher JL, Hacihaliloglu I (2019) Learning needle tip localization from digital subtraction in 2D ultrasound. Int J Comput Assist Radiol Surg 14(6):1017–1026

22. Redmon J, Divvala S, Girshick R, Farhadi A (2016) You only look once: unified, real-time object detection. In: 2016 IEEE conference on computer vision and pattern recognition (CVPR), pp 779–788