

Motion-aware Needle Segmentation in Ultrasound Images

Raghavv Goel* Cecilia Morales Manpreet Singh Artur Dubrawski Jonh Galeotti Howie Choset
Carnegie Mellon University

Abstract

Segmenting a moving needle in ultrasound images is challenging due to the presence of artifacts, noise, and needle occlusion. This task becomes even more demanding in scenarios where data availability is limited. In this paper, we present a novel approach for needle segmentation for 2D ultrasound that combines classical Kalman Filter (KF) techniques with data-driven learning, incorporating both needle features and needle motion. Our method offers three key contributions. First, we propose a compatible framework that seamlessly integrates into commonly used encoder-decoder style architectures. Second, we demonstrate superior performance compared to recent state-of-the-art needle segmentation models using our novel convolutional neural network (CNN) based KF-inspired block, achieving a 15% reduction in pixel-wise needle tip error and an 8% reduction in length error. Third, to our knowledge we are the first to implement a learnable filter to incorporate non-linear needle motion for improving needle segmentation.

1. Introduction

Needle insertion is commonly used in clinical diagnosis and treatment, providing a minimally invasive method for accessing a variety of medical conditions[1]. Providing precise needle location during percutaneous ultrasound-guided needle insertion procedures can provide valuable insights to the clinician and improve the procedures' success [7]. However, segmenting needles in ultrasound images is a challenging task. Ultrasound images are prone to having artifacts, noise, and shadows, which lead to a low signal-to-noise ratio [15]. Moreover, the thin structure of the needle can be easily occluded by nerves or become partially visible due to bending out of the ultrasound image plane. Some of the aforementioned challenges are shown in Fig. 1.

Common deep learning based medical image segmentation methods have an encoder and a decoder. The encoder transforms the input image, capturing the image features, while the decoder further transforms and upsamples these features, outputting the corresponding prediction




Figure 1. Sample images from our ultrasound dataset. The top row shows the raw images and the bottom row shows the same images with the ground truth needle position overlaid in red. The images show (a) unclear initial half of needle, (b) multiple artifacts that look like the needle, (c) needle occluded in the middle, and (d) needle-tip has poor visibility

mask. This architecture can be seen in U-Net [21], U-Net with Attention gates [19], TransUnet [2], Excitation Network [12], among others. While these methods work well for segmenting static and complex anatomical structures, their efficacy in segmenting moving needles within ultrasound images has been found lacking as no needle motion information is used, ie. excitation networks are used to generate static needle masks. We believe that needle motion serves as a robust learning signal for needle segmentation.

Recent deep-learning methods for needle segmentation have also exploited relative change between consecutive image frames by using a stack of the current image together with prior frames as input to the encoder [16], pixel-wise difference of consecutive image frames as an additional input to the encoder [17], or using two different encoders for consecutive images leading to a single decoder [3, 4], or adding separate upsampling blocks for needle-tip and needle-shaft [5]. However, these simple operations for capturing needle movement perform poorly when the background is non-stationary. Some methods [18, 24] track only the needle-tip, however, these methods are useful only in cases when needle-tip movement and brightness are more distinct than the rest of the image, as is the case when the needle is inserted orthogonal to the ultrasound image view, which is not consistent with our data.

Algorithms for needle tracking involving KF [6, 25] make impractical assumptions by naively modeling needle dynamics as linear with needle moving at a constant velocity. In general, during skin penetration, the needle follows a complex motion with non-linear dynamics as it interacts

*Correspondence to raghavvg@alumni.cmu.edu.

with tissues, organs, and arteries which exert forces at the needle tip and along the shaft which can result in needle bending. Hence, needle motion can't be derived analytically.

We therefore propose a KF-inspired block (Fig. 2) which learns the needle motion features as a dynamics model in the feature space of the input images and, then combines these motion features with the visual features for improved needle segmentation. Our block stands out in that the feature space of the input images has features corresponding to needle-like structures; however, artifacts are present, therefore it combines the needle motion by acting as a filter to remove undesirable artifacts.

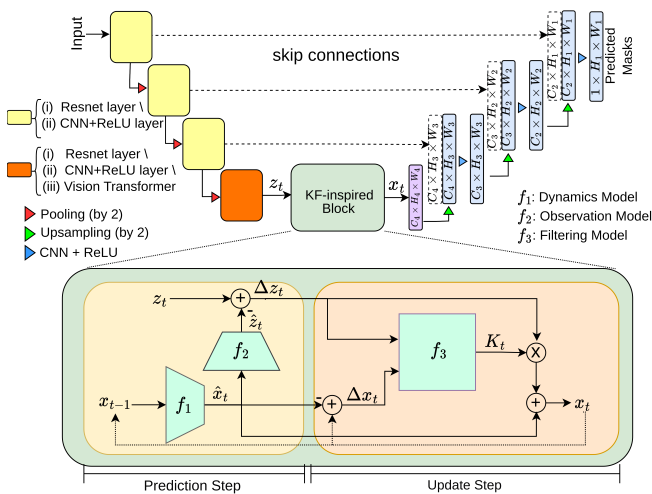


Figure 2. Overall architecture diagram showing the variety of encoders in yellow and orange. The yellow blocks can be either (i) residual network blocks or, (ii) convolutional neural network (CNN) and ReLU blocks. The following orange block can be either (i), (ii), or (iii) a vision transformer block. The red triangles are the max-pooling layer to downsample the features. The right side in light blue shows decoder with convolutional and ReLU layers (blue triangle) and upsampling (green triangle). Finally, our KF-inspired block (light green) is in between the encoder and decoder having two stages: prediction step and update step, with three learnable blocks represented by f_1 , f_2 , f_3 .

2. Method

The overall method for segmenting needles in ultrasound images consists of three stages: (i) an **encoder** for extracting visual features of the needle from individual frames, (ii) a **KF-inspired block** for capturing needle motion across successive image frames and combining motion with visual features, and (iii) a **decoder** with skip-connections to combine high-level features with the feature representation infused with needle motion to produce segmentation masks. These three stages are depicted in Fig. 2.

As shown in Fig. 2, the initial 3 encoder blocks (yellow) consist of either convolutional and ReLU layers, similarly to the encoder blocks in U-Net architectures [21], or residual-network (ResNet) blocks [9]. If the first three blocks are based on the encoder in U-Net, then the last block (orange) can be either a block of convolutional and ReLU

layer or a Vision transformer (ViT), resembling the encoder in TransUNet [2] where the placement of ViT block follows [2]. In contrast, for ResNet blocks, the last block is also a residual network block, and we use ResNet-18 due to computing budget. For the decoder, we follow the decoder mentioned in [2] which consists of deconvolutional, convolutional and ReLU layers. All implementation has been done in Python using PyTorch for ease of reproducibility.

2.1. KF-inspired Block

We propose a learnable KF-inspired block improving upon KalmanNet [20] with two key modifications. First, using CNN layer instead of linear layer, and second, placing the KF-inspired block in between an encoder and a decoder for the purpose of segmentation. The CNN helps to maintain and use the spatial structure of the feature space as no reshaping of the features is required which could result in loss of spatial information, as is the case for linear layer. The KF-inspired block is placed after the encoder to learn needle motion using lower level needle-features.

As seen from Fig. 2, the KF-inspired block follows similar high-level steps to the conventional KF [11]: prediction-step and update-step. For a time-step t , first, the prediction-step consists of the following: the observation (z_t), which represents the current image features, the state-estimate (\hat{x}_t), obtained by propagating output state at the previous time-step (x_{t-1}) through the dynamics model (f_1), and observation estimate (\hat{z}_t), produced by propagating \hat{x}_t through the observation model (f_2). Subsequently, the update-step consists of the following: computing dynamics error ($\Delta x_t = x_{t-1} - \hat{x}_t$) and observation error ($\Delta z_t = z_t - \hat{z}_t$), computing the Kalman gain (K_t) using history of dynamics and observation error propagated through f_3 . The Kalman gain acts as a filter combining the observation, dynamics, and computing the output state ($x_t = \hat{x}_t + K_t \odot \Delta z_t$).

Generally, the observation, represented by the current image feature, has no information of needle motion and can even consist of undesirable features due to low signal to noise ratio of ultrasound images. We hypothesize that by incorporating both the needle motion and needle features, we can improve the needle segmentation mask prediction. The KF is useful in such situations which involve noisy observations and approximate dynamics. The KF combines them both and produces an improved state through the update-step.

2.2. Dataset

Our dataset consists of 17 ultrasound videos of the femoral area of different participants, which were captured using a Butterfly handheld ultrasound probe with patients undergoing adductor canal (AC) blocks at the Brooke Army Medical Center (BAMC). The videos were collected as part of routine clinical trials, with approval from the Institutional Re-

view Board (IRB). During these clinical procedures, doctors often have to exert jerking motions with the needle in order to visualize it properly. It is worth noting that even though a rigid needle is used during the procedures, when interacting with the tissue the needle can sometimes buckle due to resistance from the surrounding tissues. The ultrasound probe had a maximum depth of 10cm. Each video contains between 250 to 600 frames captured at a resolution of around 580×585 and a frame-rate of 20 fps. Although images were resized to 256×256 . The annotations were created by expert clinicians using the open-source Computer Vision Annotation Tool (CVAT) [22] to ensure accuracy. The annotations provide segmentation masks for various features including nerve structures and needle positions. It is important to note that the ultrasound data is obtained by a non-stationary ultrasound probe.

2.3. Comparison Models

To evaluate the efficacy of our method in an encoder-decoder framework, we show the performance of our KF-inspired block with different encoders: (i) encoder from vanilla U-Net architecture [21], (ii) ResNet-18, and (iii) encoder from TransUnet architecture [2]. Using these encoders and decoder from [2], we make a variety of different architectures by replacing/removing the KF-inspired block: (a) no block, (b) no block with attention gates ¹, (c) no block but inputting a stack of consecutive images to the encoder instead of a single image, following [16], (d) Long-Short Term Memory (LSTM) [10] unit block, (e) ConvLSTM block [23]. Note that no block present ((a), (b) and (c)) implies that $x_t = z_t$ from Fig. 2. These different combinations can recreate known architectures, for example, the encoder (i) with (a) or (b) and the given decoder gives the vanilla U-Net or U-Net with gates Attention architecture respectively. Similarly, the encoder (iii) with (a) and given decoder gives the TransUnet from [2]. We utilized LSTM and ConvLSTM blocks to demonstrate that using CNNs to capture motion directly in the feature space (utilizing adjacency information), as done in ConvLSTM and our method, outperforms capturing motion in a smaller embedding space, without using adjacency information, as done in LSTM and KalmanNet. This is particularly evident in the challenging task of needle segmentation, as shown in Table 1 below.

2.4. Training Setup

We utilized 5-fold cross-validation for both training and testing. Our dataset consisted of 17 videos, which were divided into folds based on the number of frames each video had. Each fold contained approximately an equal number of frames, with a range of 1200-1520 frames per fold. We randomized the splits at the video level, ensuring no frames

¹<https://github.com/ozan-oktay/Attention-Gated-Networks>

from the same video were present in both training and testing sets to prevent data leakage. During preprocessing, frames were subjected to variable rotation and jitter for data augmentation. In each fold, the model was trained on four sets and tested on the remaining set, repeating this process five times to ensure robust results. The final performance metrics were derived by averaging the results and are presented in Table 1. Due to the limited size of the dataset, no validation set was made for fine-tuning.

The encoder (i) from Subsection 2.3 was initialized with Kaiming initialization [8] and ResNet-18 was initialized using Imagenet weights ². We used Binary Cross-Entropy (BCE) loss³, and AdamW [14] optimizer with default hyper-parameters, with a learning rate of 0.001, and *ReduceLROnPlateau* learning-rate scheduler with a factor of 0.7 and patience of 20 for training the models end-to-end. We used a batch size of 4 and 8 for the models with ((d),(e),(f) in Subsection 2.3) and without ((a),(b),(c) in Subsection 2.3) a block in between the encoder-decoder respectively. We couldn't use a higher batch size for the case with blocks due to limited hardware resources and back-propagation through time (BPTT) requiring more memory. A batch for models with a block ((d),(e),(f) in Subsection 2.3) has a sequence of consecutive frames (7-10) where each frame is passed into the model one at a time and then after the last frame is passed, the model weights are updated through BPTT to capture the needle motion. On average, each model is trained for 10 epochs which comprises more than 6000 optimization steps in total. All the models are trained using a 24 GB Nvidia RTX 3090Ti. During inference, a single ultrasound image is propagated through the network. The frameworks having a block between the encoder-decoder maintain a hidden state which captures the dependence of previously propagated frames and is used to provide additional context to the current frame.

2.5. Metrics

We use common image segmentation metrics: Dice Score-Coefficient (DSC), Precision, and Recall along with needle-specific metrics such as needle tip error ($\Delta x, \Delta y$) and needle length error (ΔL). The additional metrics are useful for segmenting needle-like objects which occupy a small number of pixels as compared to total pixels in the image. Needle tip error is calculated based on the absolute difference between last pixel along the needle shaft of ground truth mask and predicted mask. Similarly, needle length can be determined by calculating the L_2 norm between the coordinates of the first and last pixel along the needle shaft, then we take absolute difference of length between the ground truth and predicted mask to compute the needle length er-

²Vision-Transformer initialized from github.com/Beckschen/TransUNet

³Pytorch's BCELossWithLogits is used.




Figure 3. Qualitative performance with encoder from vanilla U-Net architecture and decoder from [2] which results in (a) U-Net, (b) U-Net with Attention gates, (c) stacked input, (d) LSTM block, (e) ConvLSTM block, (f) Ours, and (g) ground truth label.

ror. The reported errors are averaged over the all the frames in the test data.

3. Results

Table 1. Comparison of dice score-coefficient (DSC), precision, recall, needle tip error (Δx , Δy) and needle length error (ΔL) with different encoders averaged over 5 folds. The overall best are **bolded** while the best in each encoder are in **blue**. Note: An upward-pointing arrow (\uparrow) indicates better performance at higher values, while a downward-pointing arrow (\downarrow) indicates better performance at lower values. Naming convention for encoders is V = encoder from vanilla U-Net architecture, R = ResNet-18, and H = encoder from TransUnet architecture. In addition to the encoders, there are variations in the model’s architecture (a) no block, (b) no block with attention gates, (c) no block with inputting a stack of consecutive 5 images, (d) Long-Short term Memory unit block, (e) ConvLSTM block

Method	DSC(D) \uparrow	Precision(P) \uparrow	Recall(R) \uparrow	Δx \downarrow	Δy \downarrow	ΔL \downarrow
V	0.32 \pm 0.15	0.37 \pm 0.19	0.33 \pm 0.14	54 \pm 18	91 \pm 29	110 \pm 43
V+(b)	0.32 \pm 0.14	0.38 \pm 0.17	0.34 \pm 0.15	48 \pm 20	74 \pm 35	76 \pm 38
V+(c)	0.34 \pm 0.13	0.42 \pm 0.18	0.33 \pm 0.11	48 \pm 18	83 \pm 35	110 \pm 50
V+(d)	0.22 \pm 0.10	0.34 \pm 0.11	0.28 \pm 0.10	59 \pm 9	80 \pm 20	87 \pm 28
V+(e)	0.39 \pm 0.17	0.44 \pm 0.16	0.43 \pm 0.19	43 \pm 7	67 \pm 17	80 \pm 27
V+Ours	0.39 \pm 0.16	0.48 \pm 0.19	0.39 \pm 0.15	35 \pm 12	55 \pm 26	72 \pm 32
R	0.27 \pm 0.14	0.39 \pm 0.22	0.27 \pm 0.10	53 \pm 17	73 \pm 21	74 \pm 14
R+(b)	0.29 \pm 0.13	0.41 \pm 0.18	0.27 \pm 0.09	57 \pm 8	79 \pm 20	87 \pm 18
R+(c)	0.28 \pm 0.15	0.41 \pm 0.20	0.26 \pm 0.12	43 \pm 12	65 \pm 29	73 \pm 23
R+(d)	0.24 \pm 0.13	0.40 \pm 0.14	0.30 \pm 0.16	64 \pm 23	72 \pm 21	89 \pm 23
R+(e)	0.34 \pm 0.15	0.44 \pm 0.18	0.34 \pm 0.18	49 \pm 17	73 \pm 28	78 \pm 28
R+Ours	0.36 \pm 0.16	0.45 \pm 0.20	0.36 \pm 0.19	40 \pm 17	65 \pm 32	71 \pm 27
H	0.29 \pm 0.13	0.41 \pm 0.21	0.27 \pm 0.11	40 \pm 12	63 \pm 28	72 \pm 19
H+(b)	0.29 \pm 0.13	0.39 \pm 0.17	0.29 \pm 0.10	48 \pm 19	65 \pm 32	69 \pm 28
H+(c)	0.30 \pm 0.15	0.41 \pm 0.20	0.30 \pm 0.14	37 \pm 14	64 \pm 30	67 \pm 16
H+(d)	0.27 \pm 0.12	0.40 \pm 0.15	0.38 \pm 0.14	60 \pm 23	72 \pm 19	91 \pm 17
H+(e)	0.36 \pm 0.16	0.47 \pm 0.20	0.34 \pm 0.17	38 \pm 9	56 \pm 26	60 \pm 32
H+Ours	0.36 \pm 0.16	0.47 \pm 0.20	0.35 \pm 0.18	34 \pm 8	53 \pm 24	66 \pm 23

We employed both qualitative and quantitative methods to evaluate the effectiveness of our approach in needle tracking. The performance of the three different encoders and the six different choices with and without a block in between the encoder and decoder is compared in Table 1. For the sake of brevity in the results table, we refer to the encoders as **V** for the encoder from vanilla U-Net architecture, **R** for ResNet-18, and **H** for encoder from TransUnet architecture.

Our results show that our method consistently achieves the best or second-best results across all encoders, showing that our proposed block is generalizable, trainable, and compatible with different encoder-decoder networks, and can give visible performance improvement when no KF-inspired block is used. Specifically, our method outperforms all other methods in terms of precision and needle tip error, while also achieving the highest dice score together with ConvLSTM. Knowing the precise location of the needle tip in needle-guided procedures is critical to en-

surging their success. The needle tip error metric plays a vital role in determining whether the needle has accurately reached the intended vessel. In this regard, the use of the KF-inspired block has proven to be highly effective, as it yields a smaller needle tip error compared to the diameter of the femoral vessels. This increased level of accuracy can greatly improve the success rates of procedures, leading to better patient outcomes and overall safety.

The encoder from vanilla U-Net architecture shows overall best results on DSC, Precision, and Recall scores with ConvLSTM and KF-inspired blocks. However, for the ConvLSTM and KF-inspired blocks, the encoder from TransUnet architecture performs slightly better overall in terms of needle length and tip error. In general, training the vanilla encoder from scratch has proven to be advantageous, likely due to the specific characteristics of our dataset. The use of pre-trained ImageNet weights as a warm-start for the ResNet and Vision-transformer has not been as effective, potentially due to (1) low signal-to-noise ratio in the ultrasound images, and (2) our ultrasound dataset differs significantly from the high-resolution colored images used for pre-training.

Qualitative results on a sample of ultrasound images from the test set are shown in Fig. 3, the proposed network’s needle mask resembles the ground truth the most and correctly estimates the location of the needle. Furthermore, from Fig.4 we observe compelling evidence of the effectiveness and progressive refinement inherent in our proposed method’s prediction especially in the presence of occlusion showing similarity to a KF, due to its filtering structure.




Figure 4. Qualitative assessment of the proposed method on test data. (a) Sequential Ultrasound Images: The figure showcases a series of ultrasound images capturing a needle occluded within nervous tissue. This occlusion is positioned at the center of the images. (b) Needle Mask Prediction Overlay: Overlaid in red upon the ultrasound images is the needle mask prediction generated by our proposed method. Notably, the predictive accuracy of the mask consistently improves over time, underscoring the efficacy of our approach. This observed behavior parallels that of a Kalman Filter (KF), where predictions tend to refine as time progresses.

4. Conclusion

This research introduces a novel KF-inspired block that enhances the accuracy of needle segmentation mask prediction by integrating both needle features and motion. This adaptable and learnable block is compatible with different encoders. Our experiments demonstrate that our approach achieves superior results in terms of DSC, Precision, and needle tip error, ranking second in performance of other metrics. Furthermore, our method outperforms or competes

favorably with others within each encoder type. Finally, our approach opens up avenues for further research, such as applying our method to different video domains, finding the optimal placement of the KF-inspired block, comparing the affect of different encoder blocks such as Swin transformer [13], combining other needle segmentation methods with our KF-inspired block, incorporating ideas from learnable optical flow.

References

- [1] G. A. Chapman, D. Johnson, and A. R. Bodenham. Visualisation of needle position using ultrasonography. *Anaesthesia*, 61(2):148–158, 2006.
- [2] Jieneng Chen, Yongyi Lu, Qihang Yu, Xiangde Luo, Ehsan Adeli, Yan Wang, Le Lu, Alan L Yuille, and Yuyin Zhou. Transunet: Transformers make strong encoders for medical image segmentation. *arXiv preprint arXiv:2102.04306*, 2021.
- [3] Shihang Chen, Yanping Lin, Zhaojun Li, Fang Wang, and Qixin Cao. Automatic and accurate needle detection in 2d ultrasound during robot-assisted needle insertion process. *International Journal of Computer Assisted Radiology and Surgery*, pages 1–9, 2022.
- [4] Matteo Dunnhofer, Maria Antico, Fumio Sasazawa, Yu Takeda, Saskia Camps, Niki Martinel, Christian Micheloni, Gustavo Carneiro, and Davide Fontanarosa. Siam-u-net: encoder-decoder siamese network for knee cartilage tracking in ultrasound images. *Medical Image Analysis*, 60:101631, 2020.
- [5] Jun Gao, Paul Liu, Guang-Di Liu, and Le Zhang. Robust needle localization and enhancement algorithm for ultrasound by deep learning and beam steering methods. *Journal of Computer Science and Technology*, 36:334–346, 2021.
- [6] André A Galdes and Thiago S Rocha. A neural network approach for flexible needle tracking in ultrasound images using kalman filter. In *5th IEEE RAS/EMBS International Conference on Biomedical Robotics and Biomechanics*, pages 70–75. IEEE, 2014.
- [7] Andrew A Gumbs, Vincent Grasso, Nicolas Bourdel, Roland Croner, Gaya Spolverato, Isabella Frigerio, Alfredo Illanes, Mohammad Abu Hilal, Adrian Park, and Eyad Elyan. The advances in computer vision that are enabling more autonomous actions in surgery: A systematic review of the literature. *Sensors*, 22(13):4918, 2022.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [10] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [11] R. E. Kalman. A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, 82(1):35–45, 1960.
- [12] Jia Yi Lee, Mobarakol Islam, Jing Ru Woh, TS Mohamed Washeem, Lee Ying Clara Ngoh, Weng Kin Wong, and Hongliang Ren. Ultrasound needle segmentation and trajectory prediction using excitation network. *International Journal of Computer Assisted Radiology and Surgery*, 15(3):437–443, 2020.
- [13] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021.
- [14] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [15] O.V. Michailovich and A. Tannenbaum. Despeckling of medical ultrasound images. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 53(1):64–78, 2006.
- [16] Soumik Mukhopadhyay, Praful Mathur, Aditya Bharadwaj, Yuri Son, Jun-Sung Park, Srinivas Rao Kudavelly, Sangha Song, and Hokyung Kang. Deep learning based needle tracking in prostate fusion biopsy. In *Medical Imaging 2021: Image-Guided Procedures, Robotic Interventions, and Modeling*, pages 605–613. SPIE, 2021.
- [17] Cosmas Mwikirize, John L Noshier, and Ilker Hacihaliloglu. Learning needle tip localization from digital subtraction in 2d ultrasound. *International journal of computer assisted radiology and surgery*, 14(6):1017–1026, 2019.
- [18] Cosmas Mwikirize, John L Noshier, and Ilker Hacihaliloglu. Single shot needle tip localization in 2d ultrasound. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part V*, pages 637–645. Springer, 2019.
- [19] Ozan Oktay, Jo Schlemper, Loic Le Folgoc, Matthew Lee, Mattias Heinrich, Kazunari Misawa, Kensaku Mori, Steven McDonagh, Nils Y Hammerla, Bernhard Kainz, et al. Attention u-net: Learning where to look for the pancreas. *arXiv preprint arXiv:1804.03999*, 2018.
- [20] Guy Revach, Nir Shlezinger, Xiaoyong Ni, Adria Lopez Escoriza, Ruud JG Van Sloun, and Yonina C Eldar. Kalmannet: Neural network aided kalman filtering for partially known dynamics. *IEEE Transactions on Signal Processing*, 70:1532–1547, 2022.
- [21] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [22] Boris Sekachev, Nikita Manovich, Maxim Zhiltsov, Andrey Zhavoronkov, Dmitry Kalinin, Ben Hoff, Tosmanov, and Tritin Truong. opencv/cvat: v1.1.0, 2020.
- [23] Xingjian Shi, Zhoung Chen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, and Wang-chun Woo. Convolutional lstm network: a machine learning approach for precipitation nowcasting. In *Proceedings of the 28th International Conference*

on *Neural Information Processing Systems - Volume 1*, page 802–810, Cambridge, MA, USA, 2015. MIT Press.

- [24] Joshua Shubert and Muyinatu A. Lediju Bell. Photoacoustic based visual servoing of needle tips to improve biopsy on obese patients. In *2017 IEEE International Ultrasonics Symposium (IUS)*, pages 1–4, 2017.
- [25] Wanquan Yan, Qingpeng Ding, Jianghua Chen, Yunhui Liu, and Shing Shin Cheng. Needle tip tracking in 2d ultrasound based on improved compressive tracking and adaptive kalman filter. *IEEE Robotics and Automation Letters*, 6(2): 3224–3231, 2021.