

Unsupervised Deformable Ultrasound Image Registration and Its Application for Vessel Segmentation

FNU Abhimanyu¹, Andrew L. Orekhov¹, Ananya Bal¹, John Galeotti¹, Howie Choset¹

Abstract—This paper presents a deep-learning model for deformable registration of ultrasound images at online rates, which we call U-RAFT. As its name suggests, U-RAFT is based on RAFT, a convolutional neural network for estimating optical flow. U-RAFT, however, can be trained in an unsupervised manner and can generate synthetic images for training vessel segmentation models. We propose and compare the registration quality of different loss functions for training U-RAFT. We also show how our approach, together with a robot performing force-controlled scans, can be used to generate synthetic deformed images to significantly expand the size of a femoral vessel segmentation training dataset without the need for additional manual labeling. We validate our approach on both a silicone human tissue phantom as well as *in-vivo* porcine images. We show that U-RAFT generates synthetic ultrasound images with 98% and 81% structural similarity index measure (SSIM) to the real ultrasound images for the phantom and porcine datasets, respectively. We also demonstrate that synthetic deformed images from U-RAFT can be used as a data augmentation technique for vessel segmentation models to improve intersection-over-union (IoU) segmentation performance.

I. INTRODUCTION

Vascular access and subsequent placement of central venous and arterial catheters is an essential first step for delivering life-saving medical care to trauma patients, e.g. administering anesthesia, monitoring vitals, and delivering resuscitative treatments like Resuscitative Endovascular Balloon Occlusion of the Aorta (REBOA). Accessing a blood vessel, commonly done via the Seldinger technique [1], requires insertion of a needle into the center of the vessel, which is typically done by a highly skilled clinician using ultrasound to determine where to insert the needle.

The work in this paper is motivated by the potential benefits of supporting human-guided vascular access with a robot so as to enable personnel away from centers of medical excellence to gain vascular access while avoiding vessel wall damage and hematomas caused by failed needle insertion attempts. This would be especially impactful on battlefields and in mass casualty disasters where there is limited access to trained medical personnel and hospital facilities. Examples of recent work towards the goal of robot-assisted femoral vessel access under ultrasound guidance include a hand-held device [2] and our group’s system using a robot manipulator [3].

Ultrasound imaging is an important modality for vascular access because it is safe, portable, and low-cost. However, the

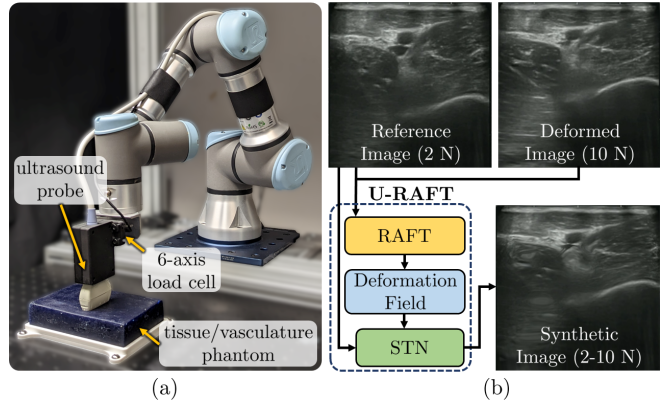


Fig. 1. (a) The robot arm we used for capturing the ultrasound images with force-controlled scanning, together with a human tissue and vasculature phantom model we used to test U-RAFT. (b) Our U-RAFT model registers a deformed image to a reference image with RAFT, creating deformation field that is then used with a spatial transformer network (STN) to generate new deformed images, shown here with example *in-vivo* porcine images.

ultrasound probe needs to press against the skin to maintain acoustic coupling contact while capturing images, causing anatomical deformations during the scan. These deformations present a challenge for training deep neural networks that segment vessels from the ultrasound images.

Convolutional neural networks (CNNs), U-Net [4], and variants of U-Net [5] are commonly used in medical image segmentation, including vessel segmentation in ultrasound [2], [6]–[8]. Training these models requires time-consuming labeling of the vessel contours in each image by personnel trained to interpret ultrasound images. Furthermore, the training set includes only a small subset of deformed vessel shapes that could occur, limiting the ability of the model to generalize to probe forces and deformed vessel shapes outside the training dataset. Although nonlinear warping augmentations could be applied to images [4], this technique is not guaranteed to generate physically realistic image augmentations.

In this paper, we propose an approach to use deformable registration for augmenting images for training vessel segmentation models with small training datasets. Deformable registration is the problem of how to register pairs of images, one referred to as the *reference image* and the other as the *deformed image*, where the two images are of the same anatomy, but exhibit different deformations. By registering images captured at different forces, we will show that we can generate synthetic images at intermediate forces, ensuring that the augmented images are physically realistic. A similar

¹Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA (abhiman2, aorekhov, abal, choset, jgaleotti@andrew.cmu.edu)

idea was used in [9], [10] for brain MRI images, and in [11] with lung MRI images.

Previous approaches to deformable registration include hand-crafted, iterative nonlinear optimization methods with a variety of cost function definitions and parameterizations of deformations [12]–[14]. These methods, however, typically have non-convex cost functions and are slow due to the large number of optimization parameters. To overcome these difficulties, deep-learning methods have also been presented [15]–[17], but work on deep-learning methods for ultrasound-to-ultrasound registration is limited. Compared to other higher-quality imaging modalities studied in most deformable registration work (e.g. CT-to-MRI or CT-to-ultrasound), ultrasound-to-ultrasound deformable registration is uniquely challenging due to noise, speckle, shadows, and mirror image artifacts [12]. Despite this, we will show that this problem is amenable to a deep-learning approach using our proposed model.

In this paper, we present a deep-learning model called U-RAFT (Unsupervised Recurrent All-pairs Field Transforms) for ultrasound-to-ultrasound deformable image registration and synthetic ultrasound image generation. As shown in Fig. 1b, U-RAFT uses RAFT [18], a CNN for optical flow estimation, to register images and create a deformation field (DF). It then uses a Spatial Transformer Network (STN) [19] to generate new synthetic images. This approach allows for unsupervised training of U-RAFT, which we use to apply RAFT to ultrasound images for the first time, as well as generate realistic deformations for expanding vessel segmentation training datasets.

Compared to prior work, our work is unique in that it tackles ultrasound-to-ultrasound deformable registration in images of vasculature, we are able to register images at a rate suitable for online use (~ 33 Hz), and our training is unsupervised. We note that the utility of our deformable registration approach is not limited to vessel segmentation, since deformable registration in ultrasound images is broadly important for longitudinal studies/diagnosis, population studies, and intra-operative registration to anatomy [12], [17].

In Section II, we present the network architecture of U-RAFT and describe three loss functions we considered for training this network in an unsupervised manner. In Section III, we present experimental results using U-RAFT on a benchtop silicone phantom model as well as *in-vivo* porcine images of femoral arteries and veins. We compare the registration quality among the three loss functions we defined, and we demonstrate how the synthetic images that U-RAFT generates can be used as a data augmentation technique to improve the performance of a CNN for vessel segmentation. Section IV presents our conclusions and discussion on future directions.

II. APPROACH

This section discusses the network architecture used for predicting the deformation field (DF) and the loss function used to train this network in an unsupervised manner. Furthermore, we discuss the use of the DF to generate new

synthetic ultrasound images and their use to improve vessel segmentation.

A. U-RAFT network architecture

Let I_r and I_d be the reference and deformed ultrasound images, respectively, collected at forces $F_r \in \mathbb{R}$ and $F_d \in \mathbb{R}$. We denote a DF as $u_{dr} = g_\theta(I_d, I_r)$, where g_θ is the function we seek to model with our network and the subscript θ denotes the network parameters used. Here, we use the state-of-art RAFT network [18] to model $g_\theta(I_d, I_r)$. We chose RAFT over other CNN-based networks like FlowNet [20], FlowNet2 [21], and PWC-Net [22] because of its superior performance on the Sintel [23] and KITTI [24] datasets.

RAFT has been shown to outperform other optical flow methods in the RGB domain [20]–[22], but no prior work has shown the application of RAFT on medical ultrasound images, which are inherently noisier than RGB images [12]. RAFT is also a supervised method that needs a ground truth displacement field for training. Acquiring ground truth for ultrasound images is time-consuming and labor-intensive, so we seek to make the training unsupervised. We do this by passing the output of RAFT through a Spatial Transformer Network (STN) [19] to generate a reconstructed deformed image $I'_d = \text{STN}(u_{dr}, I_r)$. This enables us to incorporate the similarity of I'_d and I_d in our training loss function, which, as we will show below, provides improved registration performance. We refer to the RAFT architecture together with STN as U-RAFT.

B. Loss functions for unsupervised training

We now define three different choices of loss functions to train the U-RAFT network in an unsupervised manner. We discuss the formulation and the advantages/disadvantages of each and perform a quantitative comparison between them in Sec. III-B.

The first loss function we consider, denoted as \mathcal{L}_{us} , consists of two parts \mathcal{L}_{ssim} , a multi-scale structural similarity (SSIM) loss that penalizes the differences in appearance between I_d and I'_d , and \mathcal{L}_{smooth} , which penalizes abrupt changes in the neighboring pixels of I'_d (generated from u_{dr}):

$$\mathcal{L}_{us}(I_d, I'_d, u_{dr}) = \beta \mathcal{L}_{ssim}(I_d, I'_d) + (1 - \beta) \mathcal{L}_{smooth}(u_{dr}) \quad (1)$$

where $\beta \in \mathbb{R}$ is a parameter to adjust the relative weight of \mathcal{L}_{ssim} and \mathcal{L}_{smooth} . \mathcal{L}_{ssim} and \mathcal{L}_{smooth} are given by:

$$\begin{aligned} \mathcal{L}_{ssim}(I_d, I'_d) &= 1 - \text{SSIM}(I_d, I'_d) \\ \mathcal{L}_{smooth}(u_{dr}) &= \text{mean}_{xy} \left(\frac{\nabla^2 u_{dr}(x, y)}{\nabla x^2} + \frac{\nabla^2 u_{dr}(x, y)}{\nabla y^2} \right) \end{aligned} \quad (2)$$

where x, y are the pixel location of a 2D-deformation field, and mean_{xy} denotes the mean over all pixels.

The second loss function we consider is a cyclic version of (1), denoted as $\mathcal{L}_{us-cyclic}$. In this loss function, we register the reference image I_r to the reconstructed deformed image I'_d to generate a new reconstructed reference image $I'_r = \text{STN}(g_\theta(I_r, I'_d), I'_d)$, as shown in Fig. 2. We then add to the

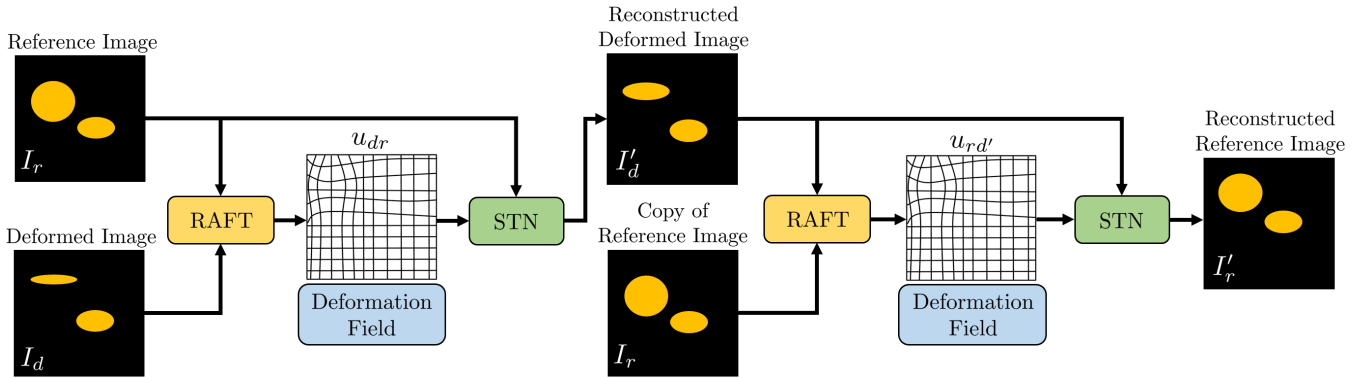


Fig. 2. Pipeline used to generate image reconstructions for training U-RAFT in an unsupervised manner. Our proposed cyclic loss function in (3) improves registration quality by comparing the reconstructed reference image to the original reference image.

loss function in (1) an additional term that calculates \mathcal{L}_{us} for I_r , I'_r , and $u_{rd'}$.

$$\mathcal{L}_{us-cyclic} = \mathcal{L}_{us}(I_d, I'_d, u_{dr}) + \mathcal{L}_{us}(I_r, I'_r, u_{rd'}) \quad (3)$$

Finally, the third loss function we consider is designed to improve flow prediction in the vicinity of important anatomical features like veins, arteries, etc. We denote this feature-aware, cyclic, multi-scale SSIM loss function as $\mathcal{L}_{fa-cyclic-us}$. We use the scale-invariant feature transform (SIFT) algorithm [25], as implemented in OpenCV, to extract keypoints in an ultrasound image and construct a binary feature map around those keypoints. We then multiply each image by its binary feature map to create \tilde{I}_d , \tilde{I}'_d , \tilde{I}_r , \tilde{I}'_r for the deformed, reconstructed deformed, reference, and reconstructed reference images, respectively. We then calculate the loss using the cyclic loss in (3) but with the images with features extracted:

$$\mathcal{L}_{fa-cyclic-us} = \mathcal{L}_{us}(\tilde{I}_d, \tilde{I}'_d, u_{dr}) + \mathcal{L}_{us}(\tilde{I}_r, \tilde{I}'_r, u_{rd'}) \quad (4)$$

C. Experimental setup and training details

To train U-RAFT, we have collected ultrasound images from two different subjects: a human tissue/vasculature gel phantom model (CAE Blue Phantom), which we refer to as the *blue-gel dataset* and two different live pigs, which we refer to as the *live-pig dataset*. The IACUC-approved *in-vivo* porcine experiments were done in a controlled lab setting under the supervision of clinicians. The data was collected using a robotic ultrasound system which includes a UR3e manipulator (Universal Robots) with a Fukuda Denshi portable point-of-care ultrasound scanner (POCUS) using a 5-12 MHz 2D linear transducer and a six-axis force/torque sensor (ATI) mounted on the end effector, as shown in Fig. 1. The datasets were collected either in a “scanning mode”, where the robot scanned between two pre-defined points on the surface of the subject with a hybrid force motion controller similar to the controller described in [26], or in “palpation mode”, where the robot was commanded with a sinusoidal force profile at a single point on the skin surface. The minimum and maximum force used for both the modes were 2 N and 10 N, respectively.

Using the loss functions introduced in Section II-B, we train the U-RAFT network separately for the blue-gel and live-pig datasets. For all these datasets, the RAFT weights were initialized with pre-trained KITTI weights [24] and were trained for 150 additional epochs. The implementation is highly parallelized and performs full-batch gradient descent using the Stochastic Gradient Descent [27] optimizer in the Pytorch Autograd library [28], with a batch size of 12 with a learning rate of 0.0001.

D. Synthetic data generation

Now that we can use U-RAFT to predict the DF between deformed and reference ultrasound images, we can generate synthetic ultrasound images at additional probe force values. Suppose ultrasound images at two different force values F_r and F_d are I_{F_r} and I_{F_d} . We use U-RAFT to find the DFs $u_{F_r, F_d} = g_\theta(I_{F_r}, I_{F_r})$ and $u_{F_d, F_r} = g_\theta(I_{F_d}, I_{F_r})$. Then we use linear interpolation to find the intermediate DF $u_{F_{new}, F_r} = \alpha u_{F_r, F_r} + (1 - \alpha) u_{F_d, F_r}$ for an intermediate force F_{new} , where $\alpha = (F_{new} - F_{min}) / (F_{max} - F_{min})$, with F_{min} and F_{max} being the minimum and maximum forces between which we interpolate. We then pass u_{F_{new}, F_r} through *STN* to generate $I_{F_{new}} = \text{STN}(u_{F_{new}, F_1}, I_{F_1})$. We then use synthetic images to augment the vessel segmentation dataset. We will show in Section III-B that this data augmentation technique helps a U-Net vessel segmentation model generalize to different forces.

III. ANALYSIS AND RESULTS

A. Deformable registration results with U-RAFT

In this section, we evaluate U-RAFT’s performance on registering ultrasound images from the blue-gel and live-pig datasets using the three loss functions described in Section II-B. We use the image similarity between the original deformed image I_d and the reconstructed deformed image I'_d to measure the efficacy of our method. Figure 3 shows an example of reference and deformed images from the blue-gel and the live-pig datasets along with the reconstructed reference images from the U-RAFT model. We use two metrics to compare the different loss functions: 1) SSIM




Fig. 3. Example deformable registration result from the blue-gel dataset, showing (a) the reference image I_r , the deformed image I_d , and the reconstructed deformed image I'_d , and (b) the deformation field calculated between I_d and I_r using U-RAFT, with a displacement vector plotted on a 4×4 pixel grid. The vessel walls for the blue-gel images are manually annotated for visibility.




Fig. 4. Live-pig results showing reference, deformed, and reconstructed deformed images for (a) small deformation, (b) medium deformation, (c) large deformation, (d) vessel collapse, and (e) an atypical case of a vessel decollapsing. The similarity between the deformed and the reconstructed deformed images shows the efficacy of U-RAFT. Registration performance drops for the case of a vessel decollapsing, but even in this atypical scenario our approach fails gracefully.

[29] and 2) a feature-aware-SSIM (F-SSIM), which is SSIM applied to the images after using SIFT to extract features as described in Section II-B.

The results are summarized in Table I. In both SSIM




Fig. 5. (a) Zoomed-in view of an example live-pig reconstructed deformed image I'_d using \mathcal{L}_{us} . (b) Zoomed-in view of the reconstructed deformed image I'_d using $\mathcal{L}_{fa-cyclic-us}$. The use of the feature-aware cyclic loss function helps remove the optical distortions observed in (a) for large deformations.

and F-SSIM, the cyclic loss function $\mathcal{L}_{cyclic-us}$ outperforms \mathcal{L}_{us} , and the feature-aware cyclic loss function $\mathcal{L}_{fa-cyclic-us}$ outperforms the cyclic loss function $\mathcal{L}_{cyclic-us}$. The cyclic function outperforms the multi-scale structural similarity as it adds a better regularization of the flow prediction. We have also observed qualitatively, as shown in the example in Fig. 5, that the combination of feature extraction and cyclic loss leads to improved registration, particularly for larger deformations.

TABLE I

COMPARISON OF REGISTRATION ERROR FOR DIFFERENT LOSS FUNCTIONS. $\mathcal{L}_{fa-cyclic-us}$ OUTPERFORMS THE OTHER TWO LOSS FUNCTIONS IN TERMS OF BOTH SSIM AND F-SSIM.

Loss Function	SSIM (blue-gel)	F-SSIM (blue-gel)	SSIM (live-pig)	F-SSIM (live-pig)
\mathcal{L}_{us}	0.905	0.966	0.870	0.918
$\mathcal{L}_{cyclic-us}$	0.907	0.967	0.883	0.927
$\mathcal{L}_{fa-cyclic-us}$	0.909	0.969	0.886	0.931

B. Synthetic data generation for training deformed vessel segmentation models

In this set of experiments, we evaluate the effect of using synthetic images to improve the results of vessel segmentation under tissue deformations. In the first experiment, we evaluated the realism of the synthetic ultrasound images created using the method described in Section II-C on the blue-gel and live-pig datasets. Using the scanning mode

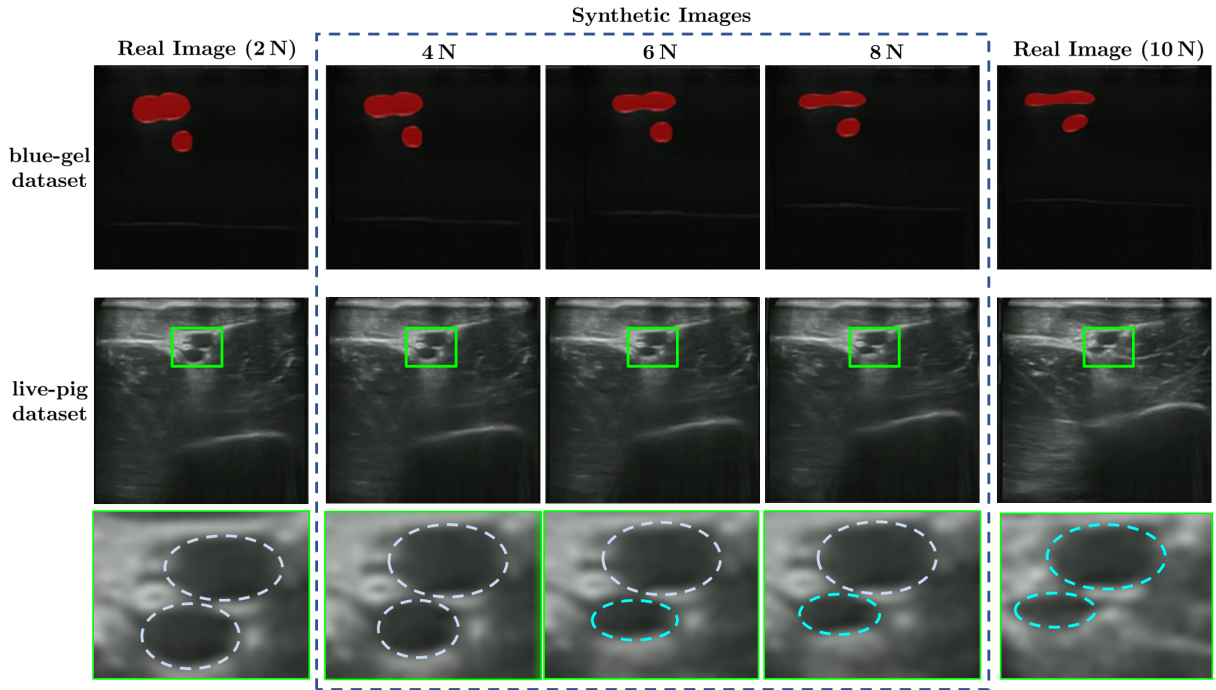


Fig. 6. Real and synthetic images from the blue-gel dataset and live-pig dataset are shown on the first and second row, respectively. The third row shows the zoomed-in view of vessels from the second row. From left to right, the images show a gradual compression of the vessel shape as the force is increased.




Fig. 7. (a) Ultrasound image at 10N. (b) Ultrasound image with predicted segmentation mask using a U-Net model trained using real data collected at 2N. (c) Ultrasound image with predicted segmentation mask using U-Net model train using real and synthetic data. The inclusion of synthetic data improves segmentation accuracy. We also show a zoomed-in view of the live-pig dataset where we are able to segment an almost collapsed vein with the augmentation of the synthetic dataset.

of the robot, we collected ~ 775 images on the blue-gel phantom, and ~ 70 images on a live pig at 2N and 10N each. We then created synthetic images at 3 intermediate force values $F = 4N$, $F = 6N$, and $F = 8N$, denoted as $D_{4,\text{syn}}$, $D_{6,\text{syn}}$, and $D_{8,\text{syn}}$, respectively. Example synthetic images are shown in Fig. 6. We use SSIM to measure the similarity between $D_{4,\text{syn}}$, $D_{6,\text{syn}}$, $D_{8,\text{syn}}$ and $D_{4,\text{real}}$, $D_{6,\text{real}}$ and $D_{8,\text{real}}$, respectively. To show that our method does not introduce excessive additional noise into the generated images, we

also evaluate the peak-signal-to-noise ratio (PSNR) [29] to compare the noise-to-signal ratio between D_{syn} and D_{real} , expressed in decibels (dB). Note that the maximum possible PSNR value for 8-bit images is ~ 48 dB. As shown in the results in Table II, we see high SSIM scores as well as high PSNR for both datasets indicating high realism of the synthetic images. We observe slightly lower SSIM and PSNR scores for the live-pig dataset compared to the blue-gel dataset because of noisy *in-vivo* ultrasound images and movement of the live animal during data collection.

TABLE II
SSIM AND PSNR BETWEEN REAL IMAGES AND SYNTHETIC IMAGES
AT INTERPOLATED FORCE VALUES

	blue-gel		live-pig	
	SSIM	PSNR	SSIM	PSNR
F=4N	0.98	42.7	0.81	25.6
F=6N	0.97	40.8	0.81	25.6
F=8N	0.97	38.8	0.78	23.7

We now show how the synthetic images generated by U-RAFT can be used as a data augmentation technique to improve a vessel segmentation model with respect to intersection-over-union (IoU). We use the ~ 775 blue-gel images from Section III-B as our training dataset. We collected an additional ~ 175 images each for testing and validation datasets. Regarding the live-pig dataset, we have obtained data from two porcine models. Our training dataset consists of the same ~ 70 images per force value as in Section III-B, with an additional ~ 260 validation and ~ 300 testing images collected using the palpation mode of the robot. We denote

our training datasets as $D_{2,\text{real}}^{\text{train}}$, $D_{4,\text{real}}^{\text{train}}$, $D_{6,\text{real}}^{\text{train}}$, $D_{8,\text{real}}^{\text{train}}$, $D_{10,\text{real}}^{\text{train}}$. The live-pig dataset has ten times fewer images than the blue-gel dataset because of high cost and effort involved in conducting *in-vivo* porcine experiments. The vessel labels for these images are manually annotated by trained annotators with at least 6 months of experience in labeling vessels in ultrasound images.

For this experiment, we train individual U-Net models on $D_{2,\text{real}}^{\text{train}}$, $D_{10,\text{real}}^{\text{train}}$ and a combination of $D_{2,\text{real}}^{\text{train}} + D_{10,\text{real}}^{\text{train}}$ dataset. Further, we augment the $D_{2,\text{real}}^{\text{train}} + D_{10,\text{real}}^{\text{train}}$ dataset with multiple-force synthetic data ($D_{4,\text{syn}}$, $D_{6,\text{syn}}$, and $D_{8,\text{syn}}$) and use it to train a U-Net model. We also compare our augmentation technique to the random elastic augmentation mentioned in [4] to underscore the significance of augmenting using multiple-force synthetic data. We show the improvement in segmentation on both the blue-gel and live-pig dataset.

TABLE III

COMPARING PERFORMANCE ON THE TEST DATASETS OF A U-NET MODEL WITH AND WITHOUT DATA AUGMENTION USING THE SYNTHETIC DATASET. THE RESULTS FROM THE BEST PERFORMING DATASET ARE HIGHLIGHTED.

Training Dataset	blue-gel	live-pig		
	IoU	IoU All Vessels	IoU Arteries	IoU Veins
$D_{2,\text{real}}^{\text{train}}$	0.77	0.53	0.58	0.48
$D_{10,\text{real}}^{\text{train}}$	0.73	0.49	0.54	0.45
$D_{2,10,\text{real}}^{\text{train}}$	0.86	0.56	0.59	0.51
$D_{2,10,\text{real}}^{\text{train}} + D_{4,6,8,\text{syn}}^{\text{train}}$	0.89	0.62	0.64	0.58
$D_{2,10,\text{real}}^{\text{train}} + D_{\text{rand-syn}}^{\text{train}}$	0.86	0.56	0.60	0.52

As shown in Table III, augmenting the real dataset with multi-force synthetic data outperforms the training done using only real images. We also see it outperform the model trained with random elastic deformations, highlighting the need for realistic force-based augmentation using deformable registration. We also compare the affect of force-based augmentation on veins and arteries separately in Table III, as arteries and veins experience different levels of deformation, i.e. veins collapse easier than arteries. We see that there is a larger improvement in IoU for vessels that experience larger deformations, demonstrating that data augmentation with deformable registration is of particular importance when deformations are large. An example of this is shown in Fig. 7, where a model trained using only on 2N images fails to segment a nearly collapsed vein, but the model trained with our synthetic images is able to segment the vein.

IV. CONCLUSIONS AND FUTURE WORK

In conclusion, this paper proposes an unsupervised approach for training an ultrasound-to-ultrasound deformable registration model. We propose and compare three different loss functions and show that a loss function based on feature-aware cyclic loss performs best. We also demonstrate how our approach, combined with a force-controlled robot, can be used to generate synthetic deformed images to expand the

size of a femoral vessel segmentation training dataset and improve vessel segmentation performance. We also validate our approach on both a benchtop human tissue/vessel phantom and *in-vivo* porcine images, highlighting the practical application of our deformable registration model in real-world medical imaging tasks. Overall, this paper presents an innovative solution to address the challenges associated with ultrasound imaging, particularly in the areas of image registration and segmentation which are critical for accurate vascular access.

As a part of future work, we plan to extend the U-RAFT algorithm with a recurrent model to learn a physics model in order to improve the deformation prediction especially for predicting scenarios like vessel decollapsing. To improve the live-pig IoU score, we will also like to further expand the training dataset by using the spatial augmentation mentioned in [8], along with force-based augmentation.

ACKNOWLEDGMENT

We would like to thank Nico Zevallos, Dr. Michael R. Pinsky, Dr. Hernando Gomez, Lisa Gordon, Antonio Gumucio, Ted Lagattuta, and Andrew Schoenling, for their help in dataset collection and labelling. We would also like to thank Yizhu Gu for help with the design and fabrication of the system components.

REFERENCES

- [1] S. Ivar Seldinger, "Catheter replacement of the needle in percutaneous arteriography: a new technique," *Acta radiologica*, vol. 49, no. sup434, pp. 47–52, 2008.
- [2] L. Brattain, T. Pierce, L. Gjestebj, M. Johnson, N. DeLosa, J. Werblin, J. Gupta, A. Ozturk, X. Wang, Q. Li, B. Telfer, and A. Samir, "AI-enabled, ultrasound-guided handheld robotic device for femoral vascular access," *Biosensors*, vol. 11, no. 12, p. 522, Dec. 2021.
- [3] N. Zevallos, E. Harber, K. Patel, Y. Gu, K. Sladick, F. Guyette, L. Weiss, M. R. Pinsky, H. Gomez, J. Galeotti, *et al.*, "Toward robotically automated femoral vascular access," in *2021 International Symposium on Medical Robotics (ISMR)*. IEEE, 2021, pp. 1–7.
- [4] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*. Springer, 2015, pp. 234–241.
- [5] N. Siddique, S. Paheding, C. P. Elkin, and V. Devabhaktuni, "U-Net and its variants for medical image segmentation: a review of theory and applications," *IEEE Access*, vol. 9, pp. 82 031–82 057, 2021.
- [6] E. Chen, H. Choset, and J. Galeotti, "Uncertainty-based adaptive data augmentation for ultrasound imaging anatomical variations," in *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*. IEEE, 2021, pp. 438–442.
- [7] N. Marahrens, B. Scaglioni, D. Jones, R. Prasad, C. S. Biyani, and P. Valdastrri, "Towards autonomous robotic minimally invasive ultrasound scanning and vessel reconstruction on non-planar surfaces," *Frontiers in Robotics and AI*, p. 178, 2022.
- [8] C. Morales, J. Yao, T. Rane, R. Edman, H. Choset, and A. Dubrawski, "Reslicing ultrasound images for data augmentation and vessel reconstruction," in *2023 International Conference on Robotics and Automation (ICRA)*. IEEE, 2023 (Accepted).
- [9] E. Krivov, M. Pisov, and M. Belyaev, "MRI augmentation via elastic registration for brain lesions segmentation," in *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries: Third International Workshop, BrainLes 2017, Held in Conjunction with MICCAI 2017, Quebec City, QC, Canada, September 14, 2017, Revised Selected Papers 3*. Springer, 2018, pp. 369–380.

- [10] J. Nalepa, G. Mrukwa, S. Piechaczek, P. R. Lorenzo, M. Marcinkiewicz, B. Bobek-Billewicz, P. Wawrzyniak, P. Ulrych, J. Szymanek, M. Cwiek, W. Dudzik, M. Kawulok, and M. P. Hayball, "Data augmentation via image registration," in *2019 IEEE International Conference on Image Processing (ICIP)*, Sept. 2019, pp. 4250–4254.
- [11] N. J. Tustison, B. B. Avants, Z. Lin, X. Feng, N. Cullen, J. F. Mata, L. Flors, J. C. Gee, T. A. Altes, J. P. Mugler, and K. Qing, "Convolutional neural networks with template-based data augmentation for functional lung image quantification," *Academic radiology*, vol. 26, no. 3, pp. 412–423, Mar. 2019.
- [12] C. Che, T. S. Mathai, and J. Galeotti, "Ultrasound registration: A review," *Methods*, vol. 115, pp. 128–143, 2017.
- [13] M. Wang and L. Pengcheng, "A review of deformation models in medical image registration," *Journal of Medical and Biological Engineering*, vol. 39, Apr. 2018.
- [14] A. Sotiras, C. Davatzikos, and N. Paragios, "Deformable medical image registration: a survey," *IEEE Transactions on Medical Imaging*, vol. 32, no. 7, pp. 1153–1190, July 2013, conference Name: IEEE Transactions on Medical Imaging.
- [15] H. R. Boveiri, R. Khayami, R. Javidan, and A. Mehdizadeh, "Medical image registration using deep neural networks: a comprehensive review," *Computers & Electrical Engineering*, vol. 87, p. 106767, Oct. 2020.
- [16] Y. Fu, Y. Lei, T. Wang, W. J. Curran, T. Liu, and X. Yang, "Deep learning in medical image registration: a review," *Physics in Medicine & Biology*, vol. 65, no. 20, p. 20TR01, Oct. 2020.
- [17] J. Zou, B. Gao, Y. Song, and J. Qin, "A review of deep learning-based deformable medical image registration," *Frontiers in Oncology*, vol. 12, p. 1047215, Dec. 2022.
- [18] Z. Teed and J. Deng, "RAFT: Recurrent all-pairs field transforms for optical flow," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*. Springer, 2020, pp. 402–419.
- [19] M. Jaderberg, K. Simonyan, A. Zisserman, and k. kavukcuoglu, "Spatial transformer networks," in *Advances in Neural Information Processing Systems*, C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, Eds., vol. 28. Curran Associates, Inc., 2015.
- [20] A. Dosovitskiy, P. Fischer, E. Ilg, P. Hausser, C. Hazirbas, V. Golkov, P. van der Smagt, D. Cremers, and T. Brox, "FlowNet: Learning optical flow with convolutional networks," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [21] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox, "FlowNet 2.0: Evolution of optical flow estimation with deep networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [22] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz, "PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [23] N. Mayer, E. Ilg, P. Hausser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox, "A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [24] M. Menze and A. Geiger, "Object scene flow for autonomous vehicles," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [25] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, vol. 60, pp. 91–110, 2004.
- [26] R. Goel, F. Abhimanyu, K. Patel, J. Galeotti, and H. Choset, "Autonomous ultrasound scanning using bayesian optimization and hybrid force control," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 8396–8402.
- [27] L. Bottou, "Stochastic gradient descent tricks," *Neural Networks: Tricks of the Trade: Second Edition*, pp. 421–436, 2012.
- [28] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in PyTorch," *NIPS 2017 Autodiff Workshop*, 2017.
- [29] U. Sara, M. Akter, and M. S. Uddin, "Image quality assessment through FSIM, SSIM, MSE and PSNR—a comparative study," *Journal of Computer and Communications*, vol. 7, no. 3, pp. 8–18, 2019.