

Epicardium Prompt-guided Real-time Cardiac Ultrasound Frame-to-volume Registration

Long Lei^{*1}, Jun Zhou^{*2}, Jialun Pei^{1(✉)}, Baoliang Zhao³, Yueming Jin⁴,
Yuen-Chun Jeremy Teoh¹, Jing Qin², and Pheng-Ann Heng¹

¹ The Chinese University of Hong Kong, Hong Kong, China

² The Hong Kong Polytechnic University, Hong Kong, China

³ Shenzhen Institute of Advanced Technology, CAS, Shenzhen, China

⁴ National University of Singapore, Singapore, Singapore

jialunpei@cuhk.edu.hk

Abstract. Real-time fusion of intraoperative 2D ultrasound images and the preoperative 3D ultrasound volume based on the frame-to-volume registration can provide a comprehensive guidance view for cardiac interventional surgery. However, cardiac ultrasound images are characterized by a low signal-to-noise ratio and small differences between adjacent frames, coupled with significant dimension variations between 2D frames and 3D volumes to be registered, resulting in real-time and accurate cardiac ultrasound frame-to-volume registration being a very challenging task. This paper introduces a lightweight end-to-end **Cardiac Ultrasound frame-to-volume Registration** network, termed **CU-Reg**. Specifically, the proposed model leverages epicardium prompt-guided anatomical clues to reinforce the interaction of 2D sparse and 3D dense features, followed by a voxel-wise local-global aggregation of enhanced features, thereby boosting the cross-dimensional matching effectiveness of low-quality ultrasound modalities. We further embed an inter-frame discriminative regularization term within the hybrid supervised learning to increase the distinction between adjacent slices in the same ultrasound volume to ensure registration stability. Experimental results on the reprocessed CAMUS dataset demonstrate that our CU-Reg surpasses existing methods in terms of registration accuracy and efficiency, meeting the guidance requirements of clinical cardiac interventional surgery. Our code is available at <https://github.com/LLEIHIT/CU-Reg>.

Keywords: Cardiac interventional surgery · Frame-to-volume registration · Ultrasound image.

1 Introduction

Cardiac interventional surgery has been widely used in the treatment of structural heart diseases, such as congenital heart disease and valvular heart disease

^{*} Equal contribution.




Fig. 1. Schematic of cardiac ultrasound frame-to-volume registration.

[2]. Compared to DSA (Digital Subtraction Angiography) and CT, 2D ultrasound imaging has the advantages of low equipment requirements, easy operation, real-time imaging, and no radiation exposure, so ultrasound-guided cardiac interventional surgery has become a new trend [1,15]. However, 2D ultrasound imaging can only display one section of the heart at a time. Doctors need to determine the position of the section in the heart structure reconstructed in their mind, and further fusion the real-time ultrasound images with the virtual cardiac anatomy to guide the surgical instruments [7], which requires extremely high levels of doctor experience. Currently, 3D ultrasound imaging is also becoming increasingly popular to obtain the complete anatomical structure of the heart [3,18]. To provide a complete guidance view for cardiac interventions, it is necessary to explore frame-to-volume registration that fuse intraoperative 2D ultrasound images and preoperative 3D ultrasound volumes in real time, which shortens the learning curve of ultrasound-guided cardiac interventions.

The ultrasound frame-to-volume registration aims to seek a transformation that optimally aligns the resampled slice from the given volume by the transformation with the 2D input image [6,5], as shown in Fig. 1. Existing registration methods are divided into mathematical methods and deep learning-based methods. Mathematically, the registration task is usually modeled as an optimization problem [17,14]. Although iteration-based methods can yield reasonable accuracy, they cannot meet the real-time requirements of cardiac surgical guidance due to the slow registration speed. Currently, various deep learning-based methods are widely applied to the image registration task, such as directly learning target transformations [19,4], keypoint descriptors [16], and image similarity metrics [8]. In the field of frame-to-volume registration, Hou *et al.* [10] utilized a CNN-based model to predict rigid transformation of arbitrary 2D image slices from 3D volumes, but only attained an average alignment error of 7 mm on simulated MRI brain data. Yeung *et al.* [20] also employed a CNN to predict the position of 2D ultrasound fetal brain scans in 3D atlas space. However, the method only takes a set of images rather than image-volume pairs as input,




Fig. 2. (a) Overview of the proposed CU-Reg, where VLGA is the voxel-wise local-global aggregation; (b) The proposed prompt-guided gated cross-dimensional attention.

which results in poor generalization ability of the model among individuals. For the ultrasound frame-to-volume registration, Guo *et al.* [6] introduced an end-to-end registration network to align a 2D TRUS frame with a 3D TRUS volume. However, this method extracts features from ultrasound images only using 2D and 3D convolutions and directly concatenates them, which can be further enhanced by epicardium mask prompts to provide sufficient critical anatomical cues and adequate cross-dimensional feature interactions for the registration of ultrasound samples with low signal-to-noise ratios.

In this paper, we aim to accomplish a real-time and accurate cardiac ultrasound frame-to-volume registration to provide a complete guidance view for cardiac interventional surgery under the beating heart. To address the feature extraction difficulties caused by low signal-to-noise ratios and low tissue contrast in ultrasound images, we introduce epicardium mask prompts to provide sufficient critical anatomical information. Specifically, a bi-directional prompt-guided gated cross-dimensional attention (PGCA) operation is introduced to produce abundant structure features and perform efficient interaction between 2D frame and 3D volume features. Further, we propose a voxel-wise local-global aggregation (VLGA) module to efficiently integrate dense local-global features across dimensions. To avoid the large registration errors caused by small differences between adjacent frame images, we embed an inter-frame discriminative regularization term within our hybrid supervised learning to increase the distinction between adjacent slices in the same ultrasound volume to ensure registration stability. Additionally, we build a simulated cardiac frame-to-volume registration dataset through post-processing the CAMUS dataset [13]. The experimental results demonstrate that our model achieves superior performance compared to the state-of-the-art methods, *e.g.*, a runtime of over 35 FPS and a DistErr of 3.91 mm, which can meet the 5 mm accuracy requirement for many cardiac catheterizations [12]. We hope our model can be applied for real-time and accurate guidance in cardiac interventions.

2 Method

2.1 Overview of the proposed CU-Reg

Fig. 2(a) illustrates the proposed lightweight end-to-end cardiac frame-to-volume registration network, called CU-Reg. We consider real-time 2D ultrasound frame images and 3D ultrasound volumes as fixed and moving images respectively [6], and take them as inputs to our framework. CU-Reg outputs six parameters $\{t_x, t_y, t_z, r_x, r_y, r_z\}$ that uniquely determine the spatial transformation of the 2D image coordinate system relative to the 3D volumetric coordinate system. Therein, the first three parameters determine the relative displacement between the origin of the coordinate system, and the last three parameters determine the rotation transformation matrix.

Given the moving image $U_v \in \mathbb{R}^{D \times H \times W}$ and fixed images $U_f^i \in \mathbb{R}^{H \times W}$, $i = \{0, 1, 2, 3\}$, where U_f^0 denotes the current anchor frame to be estimated and $U_f^{1,2,3}$ denotes adjacent frames within the same volume. We initially utilize two independent encoding branches to extract 2D slice features $\mathcal{F}_s$ and 3D volume features $\mathcal{F}_v$. For the 2D frame branch, we first use two *Conv1d* layers to normalize the channel dimension by increasing the number of channels from 4 to 64 followed by reducing it to 3, then feed frames into the CNN-based encoder [21] for multi-level features. To encode the 3D volume, we employ three 3D convolutional blocks with different kernel sizes to extract coarse-to-fine multi-scale features. Subsequently, we introduced the epicardium prompt-guided cross-dimensional attention operation that leverages the epicardium mask prompt with the bi-directional gated cross-dimensional attention block to spotlight critical anatomical features, providing informative alignment cues for ultrasound images. The enhanced features are processed by our voxel-wise local-global aggregation module to boost the fine-grained fusion of cross-dimensional representations. Finally, transformation parameters are estimated via the pose predictor. Additionally, our model embeds an inter-frame discriminative regularization term to highlight the discrimination between adjacent slices within the same ultrasound volume, yielding a hybrid-supervised training strategy to ensure registration stability.

2.2 Epicardium Prompt-guided Cross-dimensional Interaction

Due to the low contrast and signal-to-noise ratio of cardiac ultrasound slices, relying solely on the encoder is insufficient to provide critical anatomical information for intraoperative and preoperative registration. In this regard, we exploit epicardium masks as prompts to pinpoint tissue landmarks for better alignment. Specifically, the features $\mathcal{F}_s$ extracted from the 2D image encoder are passed through a *Softmax* layer to epicardium mask prompts. The predicted epicardium mask is supervised by the ready-made ground truth during training. After passing the epicardium mask prompt through two 1×1 convolutions for matching the dimensions of 3D volumetric features, we introduce a prompt-guided gated cross-dimensional attention (PGCA) to improve the interaction among 2D slice features, 3D volume features, and epicardium prompt features.

Inspired by gated attention [11], PGCA dynamically regulates the feature dependencies between features of different dimensions, thereby enabling more efficient cross-dimensional interactions for capturing local-global features. Here, we embed bi-directional PGCA operations, and the three inputs of each PGCA are the cross-dimensional input $\mathbf{C} \in \mathbb{R}^{d \times DHW}$, the primary input $\mathbf{P} \in \mathbb{R}^{d \times DHW}$, and the epicardium prompt $\mathbf{E} \in \mathbb{R}^{d \times DHW}$, where $\mathbf{P}$ represents the current branch features ($\mathcal{F}_s$ or $\mathcal{F}_v$), and $\mathbf{C}$ means other corresponding branching features ($\mathcal{F}_v$ or $\mathcal{F}_s$). In addition, d and DHW denote the feature channel dimension and the size of each feature map, respectively. As described in Fig. 2(b), we first perform a linear projection of $\mathbf{P} \in \mathbb{R}^{d \times DHW}$ and $\mathbf{C} \in \mathbb{R}^{d \times DHW}$ with the *SiLU* function to produce the queries Q , keys K , values V and gated vectors G :

$$Q = W_q \cdot \mathbf{C}, K = W_k \cdot \mathbf{P}, V = \phi(W_v \cdot \mathbf{P}), G = \phi(W_g \cdot \mathbf{P}), \quad (1)$$

where $W_q, W_k, W_v, W_g \in \mathbb{R}^{d \times d}$ denote the projection matrix, ϕ is the *SiLU* function. Then, we obtain the enhanced 2D slice feature $z_s \in \mathbb{R}^{d \times DHW}$ and 3D volume feature $z_v \in \mathbb{R}^{d \times DHW}$ via the prompt-guided cross-attention, which can be formulated as follows:

$$z_{i \in \{s, v\}} = \mathbf{P} + f(G \cdot \theta(Q \cdot K^T / \sqrt{d_k}) \cdot V) + f(\mathbf{E}), \quad (2)$$

where $f(\cdot)$ denotes the convolution operations, $\theta(\cdot)$ is the standard *Softmax* function, $1/\sqrt{d_k}$ is a scaling factor and d_k is the number of channels.

2.3 Voxel-wise Dense Local-Global Aggregation

After obtaining the sufficient interaction between ultrasound frames and volume features, it is essential to fuse cross-dimensional features for a cohesive synthesis of critical structural details. To accommodate ultrasound registration with multiple noises, we introduce a voxel-wise local-global aggregation module (VLGA) to efficiently associate local dense cues with global geometric information. Given the enhanced 2D slice features $z_s \in \mathbb{R}^{d \times DHW}$ and 3D volume features $z_v \in \mathbb{R}^{d \times DHW}$ derived from bi-directional PGCA operations, we map the volume feature of each voxel and its spatial corresponding slice feature to the same size through 3D convolution and 2D convolution operations to generate voxel-wise pairs of features. Subsequently, these feature pairs are concatenated and fed into an MLP to obtain a global feature vector $\mathcal{F}_{glo}$. Lastly, $\mathcal{F}_{glo}$ is concatenated with the paired features, facilitating the acquisition of local-global context insights. Our VLGA module can be summarized as follows:

$$Z = \mathcal{C}[z_s; z_v; \text{Max}(\text{MLP}(\mathcal{C}[z_s; z_v]))], \quad (3)$$

where $\mathcal{C}[\cdot]$ is the concatenation operation and $\text{Max}(\cdot)$ denotes max-pooling.

2.4 Hybrid Supervised Learning

In the training phase, we employ a hybrid loss function to supervise our model. Unlike existing methods that jointly regress the pose parameters, we propose to

predict the pose parameters separately by decoupling the rotation and translation branches so as to avoid discontinuity in the rotational space from disturbing the prediction of translation parameters. Here, we use two MLP layers to regress the rotation and translation parameters along with a smoothed L1 loss to supervise the pose parameters. The translation loss $\mathcal{L}_{trans}$ and rotation loss $\mathcal{L}_{rot}$ are used to make the network converge quickly, and they can be formulated as

$$\mathcal{L}_{trans} = \frac{1}{N} \sum_{i=1}^N \text{smooth}_{L1}(\mathbf{T} - \mathbf{T}^*), \quad \mathcal{L}_{rot} = \frac{1}{N} \sum_{i=1}^N \text{smooth}_{L1}(\mathbf{R} - \mathbf{R}^*), \quad (4)$$

where N is the total number of samples. $\mathbf{T}$ and $\mathbf{R}$ means the predicted poses parameters, $\mathbf{T}^*$ and $\mathbf{R}^*$ are the ground truth. For prompt learning, we utilize the MSE loss to directly supervise the predicted epicardium prompt $\mathbf{E}$:

$$\mathcal{L}_{prompt} = \frac{1}{N} \sum_{i=1}^N \|\mathbf{E} - \mathbf{E}^*\|_2, \quad (5)$$

where $\mathbf{E}^*$ is the ground truth of mask prompts. Moreover, to enlarge the discrimination between neighboring slices in the same volume, we embed an inter-frame discriminative regularization term $\mathcal{L}_{reg}$ to ensure the stability of the registration:

$$\mathcal{L}_{reg} = \frac{1}{N} \sum_{i=1}^N \text{smooth}_{L1}(D_f - D_f^*), \quad (6)$$

where $D_f \in \mathbb{R}^{1 \times 3}$ is the estimated inter-frame distance, which is predicted from the aggregated feature Z by a separated MLP network, and D_f^* is the ground truth. The inter-frame distance is defined as the Euclidean distance between the translation vectors of the two adjacent frames. In addition, we leverage the MS-SSIM loss $\mathcal{L}_{sim}$ for self-supervised training by constraining the similarity between the resampled and input frames to make the training process more stable and avoid overfitting. Overall, the hybrid loss function of CU-Reg is computed as

$$\mathcal{L} = \lambda_1 \mathcal{L}_{trans} + \lambda_2 \mathcal{L}_{rot} + \lambda_3 \mathcal{L}_{prompt} + \lambda_4 \mathcal{L}_{reg} + \lambda_5 \mathcal{L}_{sim} \quad (7)$$

where $\lambda_n, n = 1, 2, \dots, 5$ are the hyper-parameters. In our implementation, we set $\lambda_1 = \lambda_2 = 1.0$, $\lambda_3 = \lambda_4 = 0.1$, and $\lambda_5 = 0.5$, respectively.

2.5 Implementation Details

Our model is trained by an Adam optimizer on a single RTX3090 GPU for 500 epochs with a batch size of 16. The inputs of CU-Reg are ultrasound slices with the size of $128 \times 128 \times 1 \times 1$ and corresponding volumes with the size of $128 \times 128 \times 32 \times 1$. In the 2D ultrasound frame branch, we employ a ResNet-34 [9] as the 2D encoder. In the training phase, the inputs are a current anchor frame and three adjacent frames, *i.e.*, $128 \times 128 \times 1 \times 4$. During inference, we manually adjust the input dimension by repeating the number of channels four times. The hyperparameter values are the optimal results obtained through ablation experiments.

Table 1. Comparative results and ablation analysis of our CU-Reg in terms of mean values of quantitative metrics. $\uparrow / \downarrow$ indicates the higher/lower the score, the better.

Methods	DistErr (mm) $\downarrow$	Img -NCC($\%$) $\uparrow$	Img -SSIM($\%$) $\uparrow$	Transformation parameters			Run-time (FPS) $\uparrow$
				TE(mm) $\downarrow$	RE($^\circ$) $\downarrow$	Para -NCC($\%$) $\uparrow$	
MRF-based [17]	4.02	87.14	60.06	2.51	6.49	72.83	0.1
FVR-Net [6]	5.84	65.49	47.77	4.22	7.87	56.60	36
CU-Reg	3.91	88.07	60.53	2.48	6.24	74.07	37
w/o PGCA	4.06	87.91	59.89	2.63	6.43	72.10	38
w/o VLGA	4.04	87.90	59.90	2.61	6.30	72.28	38
w/o PGCA&VLGA	5.06	82.36	53.28	3.07	6.93	67.11	39
w/o PGCA&VLGA& $\mathcal{L}_{reg}$	5.47	77.17	48.26	3.83	7.40	61.98	39
Baseline	7.99	60.39	44.39	4.81	8.54	49.72	39

3 Experiments

3.1 Dataset and Evaluation Metrics

To evaluate the cardiac frame-to-volume registration network, a simulated dataset is generated by post-processing the public CAMUS dataset [13]. CAMUS contains 2D echocardiographic sequences with two- and four-chamber views of 500 patients, along with the masks of the left ventricular epicardium, these 2D sequences are expressed as 3D volumes in Cartesian coordinates with a unique grid resolution using the same interpolation procedure. For each original 3D volume, four transformations are generated by add random deviations to the identity transformation, the deviations of translation parameters of each transformation are within the range of 10 mm, and the deviations of the rotation parameters are within the range of 20 degrees. Based on these transformations, 2D slices and corresponding masks are sampled from the original volume. The sampled slices include 128×128 pixels, and the pixel spacing is $0.62 \text{ mm} \times 0.62 \text{ mm}$. Meanwhile, a new volume is sampled from the original volume based on the identity transformation with a volume size of $128 \times 128 \times 32$ and a voxel spacing of $0.62 \text{ mm} \times 0.62 \text{ mm} \times 0.62 \text{ mm}$. In this way, we can obtain four volume-frame-mask pairs with true transformations for one original volume. All data are split at the patient level, with 3,600 volume-frame-mask pairs for training and 400 volume-frame-mask pairs for testing.

For evaluation, we adopt the distance error (DistErr) to represent the average distance of the center and four corners between the input slice and the predicted slice. Normalized cross-correlation (NCC) and structure similarity index measure (SSIM) are used as image similarity metrics. In addition, the translation error (TE) denotes the L1 distance between the true and the predicted translation vectors $[t_x, t_y, t_z]$, and the rotation error (RE) denotes the L1 distance between the true and the predicted rotation vectors $[r_x, r_y, r_z]$.

3.2 Comparison with State-of-the-Art methods

We compare the proposed CU-Reg with the MRF-based conventional method [17] and the deep model FVR-Net [6] on the test set of our simulated data. As shown




Fig. 3. Qualitative comparison on the registration results of different methods, including predicted slices and their difference heatmaps with the ground truth.

in Table 1, our model significantly outperforms FVR-Net for registration accuracy, *e.g.*, about 33% decrease for the DistErr and 34% improvement for the Img-NCC, which can be attributed to the interaction of cross-dimensional features by the proposed PGCA and the augmentation of structural information by the epicardium mask prompt. The visualization results in Fig. 3 also illustrate that our model can perform remarkable registration outcomes. For registration efficiency, CU-Reg is significantly faster than conventional MRF-based methods (requiring multiple optimization iterations) by over 35 FPS, further confirming the superiority of our model in enhancing registration speed.

3.3 Ablation Study

We conduct thorough ablation experiments on each key component of the proposed CU-Reg, including the epicardium prompt supervision, a prompt-guided gated cross-dimensional attention (PGCA), a voxel-wise local-global aggregation module (VLGA), and an inter-frame discriminative regularization term $\mathcal{L}_{reg}$. For our baseline, we utilize the regular 2D frame and 3D volume encoders to extract frame and volume features and directly concatenate them for feeding into the pose predictor. As shown in the last two rows of Table 1, when we embed the epicardium prompt supervision to the baseline, there is a significant improvement in the perception of anatomical features of cardiac ultrasound images by CU-Reg, *e.g.*, about 17% increase in the Img-NCC metric. With the addition of $\mathcal{L}_{reg}$ to our total loss function, the registration accuracy of our model is further improved. Moreover, the proposed PGCA and VLGA play an indispensable role in the overall model and drive our model to optimal performance when used in synergy. Additionally, the last column of Table 1 illustrates the advantage of our model in inference speed, thanks to the lightweight design of CU-Reg.

4 Conclusion

In this study, we present a novel lightweight end-to-end model, termed CU-Reg, for real-time and accurate cardiac ultrasound frame-to-volume registration. Launched from the epicardium mask prompt, we present a bi-directional prompt-guided gated cross-dimensional attention together with a voxel-wise local-global aggregation module to efficiently interact and integrate 2D sparse features and 3D dense features to obtain sufficient registration information. Further, we also introduce inter-frame discriminative regularization to increase the discrimination of similar frames by our model. The experimental results demonstrate that the proposed CU-Reg outperforms the current state-of-the-art methods in both precision and efficiency. Significantly, our model provides indispensable real-time guidance view for cardiac interventional surgery. Furthermore, it can serve as a bridge for ultrasound-CT/MRI registration and showcase the potential for immediate application in cross-modal ultrasound-CT/MRI registration fields.

Acknowledgments. This work was supported in part by the Research Grants Council of the Hong Kong Special Administrative Region under Grant T45-401/22-N, in part by the Hong Kong Innovation and Technology Fund under Grants GHP/080/20SZ and GHP/050/20SZ, in part by the National Natural Science Foundation of China under Grants U23A20391 and 62273328, and in part by the Regional Joint Fund of Guangdong under Grants 2021B1515130003 and 2021B1515120011.

Disclosure of Interests. The authors have no competing interests to declare.

References

1. Abbas, S., Peng, P.: Basic principles and physics of ultrasound. *Ultrasound for Interventional Pain Management: An Illustrated Procedural Guide* pp. 1–31 (2020)
2. Al-Ebrahim, E.K., Madani, T.A., Al-Ebrahim, K.E.: Future of cardiac surgery, introducing the interventional surgeon. *Journal of Cardiac Surgery* **37**(1), 88–92 (2022)
3. Avola, D., Cinque, L., Fagioli, A., Foresti, G., Mecca, A.: Ultrasound medical imaging techniques: a survey. *ACM CSUR* **54**(3), 1–38 (2021)
4. Bharati, S., Mondal, M., Podder, P., Prasath, V.: Deep learning for medical image registration: A comprehensive review. *arXiv preprint arXiv:2204.11341* (2022)
5. Ferrante, E., Paragios, N.: Slice-to-volume medical image registration: A survey. *Medical Image Anal.* **39**, 101–123 (2017)
6. Guo, H., Xu, X., Xu, S., Wood, B.J., Yan, P.: End-to-end ultrasound frame to volume registration. In: *MICCAI*. pp. 56–65 (2021)
7. Hacıhaliloglu, I., Chen, E.C., Mousavi, P., Abolmaesumi, P., Boctor, E., Linte, C.A.: Interventional imaging: Ultrasound. In: *Handbook of Medical Image Computing and Computer Assisted Intervention*, pp. 701–720. Elsevier (2020)
8. Haskins, G., Kruger, U., Yan, P.: Deep learning in medical image registration: a survey. *Machine Vision and Applications* **31**, 1–18 (2020)
9. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *IEEE CVPR*. pp. 770–778 (2016)

10. Hou, B., Alansary, A., McDonagh, S., Davidson, A., Rutherford, M., Hajnal, J.V., Rueckert, D., Glocker, B., Kainz, B.: Predicting slice-to-volume transformation in presence of arbitrary subject motion. In: MICCAI. pp. 296–304 (2017)
11. Hua, W., Dai, Z., Liu, H., Le, Q.: Transformer quality in linear time. In: ICML. pp. 9099–9117 (2022)
12. King, A.P., Jansen, C., Rhode, K.S., Caulfield, D., Razavi, R., Penney, G.P.: Respiratory motion correction for image-guided cardiac interventions using 3-d echocardiography. *Medical image analysis* **14**(1), 21–29 (2010)
13. Leclerc, S., Smistad, E., Pedrosa, J., Østvik, A., Cervenansky, F., Espinosa, F., Espeland, T., Berg, E.A.R., Jodoin, P.M., Grenier, T., et al.: Deep learning for segmentation using an open large-scale dataset in 2d echocardiography. *IEEE TIP* **38**(9), 2198–2210 (2019)
14. Lei, L., Zhao, B., Qi, X., Mi, R., Ye, H., Zhang, P., Wang, Q., Heng, P.A., Hu, Y.: Robotic needle insertion with 2d ultrasound–3d ct fusion guidance. *IEEE TASE* (2023)
15. Liu, Z., Li, W., Li, H., Zhang, F., Ouyang, W., Wang, S., Wang, C., Luo, Z., Wang, J., Chen, Y., et al.: Automated deep neural network-based identification, localization, and tracking of cardiac structures for ultrasound-guided interventional surgery. *Journal of Thoracic Disease* **15**(4), 2129 (2023)
16. Markova, V., Ronchetti, M., Wein, W., Zettinig, O., Prevost, R.: Global multi-modal 2d/3d registration via local descriptors learning. In: MICCAI. pp. 269–279 (2022)
17. Porchetto, R., Stramana, F., Paragios, N., Ferrante, E.: Rigid slice-to-volume medical image registration through markov random fields. In: MICCAIW. pp. 172–185 (2017)
18. Song, X., Guo, H., Xu, X., Chao, H., Xu, S., Turkbey, B., Wood, B.J., Wang, G., Yan, P.: Cross-modal attention for mri and ultrasound volume registration. In: MICCAI. pp. 66–75 (2021)
19. Xu, J., Moyer, D., Grant, P.E., Golland, P., Iglesias, J.E., Adalsteinsson, E.: Svort: iterative transformer for slice-to-volume registration in fetal brain mri. In: MICCAI. pp. 3–13 (2022)
20. Yeung, P.H., Aliasi, M., Papageorghiou, A.T., Haak, M., Xie, W., Namburete, A.I.: Learning to map 2d ultrasound images into 3d space with minimal human annotation. *Medical Image Anal.* **70**, 101998 (2021)
21. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: IEEE CVPR. pp. 2881–2890 (2017)