Deep Cascade Networks for Single 2D US Slice to 3D CT/MRI Image Registration

Wei Wei

University of Magdeburg

Xu Haishan

Zhejiang University

Marko Rak

University of Magdeburg

Christian Hansen (✉ hansen@isg.cs.uni-magdeburg.de)

University of Magdeburg

Research Article

Keywords: Ultrasound (US) devices, Computed Tomography/Magnetic Resonance Imaging (CT/MRI), InitNet-LoopNet-MultiChannel (ILM).

Posted Date: December 10th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-1091982/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Deep cascade networks for single 2D US slice to 3D CT/MRI image registration

Wei Wei^{1,+}, Haishan Xu^{2,+}, Marko Rak¹, and Christian Hansen^{1,*}

¹University of Magdeburg, Faculty of Computer Science & Research Campus STIMULATE, Magdeburg, 39106, Germany

²Zhejiang University, Sir Run Run Shaw Hospital School of Medicine, Hangzhou, 10016, VR China

*hansen@isg.cs.uni-magdeburg.de

+these authors contributed equally to this work

ABSTRACT

Background and Objective: Ultrasound (US) devices are often used in percutaneous interventions. Due to their low image quality, the US image slices are aligned with preoperative Computed Tomography/Magnetic Resonance Imaging (CT/MRI) images to enable better visibilities of anatomies during the intervention. This work aims at improving the deep learning one shot registration by using less loops through deep learning networks.

Methods: We propose two cascade networks which aim at improving registration accuracy by less loops. The InitNet-Regression-LoopNet (IRL) network applies the plane regression method to detect the orientation of the predicted plane derived from the previous loop, then corrects input CT/MRI volume orientation and improves the prediction iteratively. The InitNet-LoopNet-MultiChannel (ILM) comprises two cascade networks, where an InitNet is trained with low resolution images to perform coarse registration. Then, a LoopNet wraps the high resolution images and result of the previous loop into a three channel input and trained to improve prediction accuracy in every loop.

Results: We benchmark the two cascade networks on 1035 clinical images from 52 patients, yielding an improved registration accuracy with LoopNet. The IRL achieved an average angle error of 13.3° and an average distance error of 4.5 millimeter. It outperforms the ILM network with angle error 17.4° and distance error 4.9 millimeter and the InitNet with angle error 18.6° and distance error 4.9 millimeter. Our results show the efficiency of the proposed registration networks, which have the potential to improve the robustness and accuracy of intraoperative patient registration.

Please note: Abbreviations should be introduced at the first mention in the main text – no abbreviations lists. Suggested structure of main text (not enforced) is provided below.

Introduction

Liver tumor ablation is most often used technique when treating metastase, especially for secondary tumors following resection procedures. Additionally, metastases are frequently spread in both left and right liver lobes. For cases which can not be resected, tumor ablation is the most efficient treatment. The liver tumor ablation procedure can be guided by CT and US devices, where the former takes CT scans to visualize the tumor and ablation needle in thick slabs. The radiologist can then estimate the needle trajectories according to the image data and correct needle insertion. The CT scans are repeated many times until the needle tip is placed into the tumor centroid. This CT-guided procedure requires a relatively long time and exposes the patient to a large amount of radiation. On the other hand, the US is most frequently used devices in liver tumor ablation. Compared to the CT, US devices are more readily available, are real-time capable, and are radiation-free. However, US provides low image quality and two-dimensional images, which is a challenge for an unexperienced clinician addressing the tumor. To resolve this issue, surgical navigation systems can be used to align pre-operative, high quality images with US images and guide the clinician in performing needle insertion during the ablation procedure.

In the last twenty years, minimal invasive navigation systems have been proposed for guiding ablation procedures. Some early works^{1,2} extended US transducer with tracking markers, reconstructing US slices to a 3D volume and aligning it to 3D volume data. This approach requires breath-holding during the US sweep, which brings an additional step in the ablation workflow. Once the patient has recovered their breath, the registration becomes invalid due to the liver shift secondary to the respiration. In recent works, Spinczyk et al.³ and Pohlman et al.⁴ developed similar needle navigation systems, but for general purposes. As discussed in these works, the most important task in their navigation system is multi-modal image registration. Towards multi-modal image registration, a conventional strategy with a similarity measurement and optimizer has been discussed frequently. The similarity measurement has the goal of calculating differences between cropped images at a

position and the target image. The optimizer aims at minimizing the difference. Wein et al.⁵ proposed a LC2 metric, which presents the gray value transformations from US images to CT slices. His method showed promising results, but required a training session to calculate the gray value transformation. The training must be completed once for every US parameter setting. Koernig et.al.⁶ have modified the similarity measurement to handle the non-rigid registration problem. Their optimization strategy simulated the behaviour of a clinician, which solves the registration problem iteratively. However, due to a large number of loops, the optimization can take up to several minutes before converging.

Recently, deep learning-based methods have been employed to solve the registration problem. A review on the deep learning-based medical image registration methods has been made by Haskins et al.⁷. For instance, in his early research⁸, convolutional neural networks (CNNs) was employed to learn the image similarities. Just like conventional similarity measures, these CNNs are used inside classical iterative optimization frameworks. Fan et al.⁹ and Yan et al.¹⁰ applied generative adversarial networks (GANs) to perform non-rigid image registration. Deep reinforcement learning was also introduced to solve this registration task, which involves an agent to find the best path to the target position through a reward system¹¹. However, none of these strategies were used for any challenging slice-to-volume registration task.

Regarding slice-to-volume registration,^{12,13} employed a regression network to predict registration parameters. Specifically, they used CNNs to extract image features and then fed these into a fully connected network to estimate image poses. Later, the regression networks were adapted by¹⁴⁻¹⁶, who were the first to present a working solution for slice-to-volume registration problems. They achieved promising results on monomodal fetal MR brain images, which can align motion corrupted MR slices and reduce artifacts in 3D reconstruction. In recent study, Ernst et al.¹⁷ proposed a segmentation network to predict the target plane for solving the registration problem. Inspired by his work, our previous work^{18,19} proposed U-Net based network yielding a promising registration result. However, to enable robust registration, input volumes are down-sampled such that they fit into video memory during the training. Moreover, we found it difficult to solve registration by a one-shot network. In this work, we extend the U-Net to a loop-based architect, which is intermediate between one-shot and iterative approaches. This is because we use only a few registration refinement loops, keeping calculations tractable during intervention.

The remainder is organized as follows: In Section 2, we analyse the challenges in slice-to-volume registration and describe our method. Section 3 presents our experimental study and the results. Section 4 provides a discussion on the proposed method, followed by conclusion in Section 5.

Registration Network

In this work, we propose two cascade networks to improve registration accuracy. As proposed in our previous work^{18,19}, the slice-to-volume registration can be achieved by applying a segmentation network to detect the true slice pose in the 3D volume. To this end, the registration is performed in a one-shot manner. Despite the high robustness reported in our previous works, the registration accuracy is limited by the downsampling of the input data. To circumvent the problem, we recognized that registration is often regarded as an optimization problem which is usually solved iteratively. This inspired us to cascade our previous network by a small number of refinement loops. First, a rough initialization is estimated by the so-called InitNet and then the registration is refined by the so-called LoopNet with high input resolution in only a local region of the image.

Figure 1. The IRL network comprises an InitNet on the left and a LoopNet on the right side. The input image is downsampled to a low resolution image. The prediction result is put into a plane regression module to detect plane parameters. A resample module takes the plane parameters to correct the input image and send the image data to the LoopNet. The LoopNet prediction is fed back to the regression module for the next loop.

Figure 2. The ILM network applies InitNet as coarse registration, then concatenates the prediction plane with a high resolution input image to send to LoopNet. The prediction of LoopNet is fed back again for the next loop.

InitNet

The InitNet has goal of detecting coarse pose of US plane in CT/MRI image. Unlike other state-of-the-art deep registration methods¹³, we do not use a regression network to predict the slice pose parameters. Rather, we used a 3D U-Net segmentation network as a baseline to solve the registration problem. The network is implemented according to our previous works^{18,19}. The Input images comprise of the 3D CT/MRI vessel tree and the replicated US vessel tree. As a prediction result, we expect segmented vessels laying on the US transducer plane in a 3D output volume. The DICE coefficient between prediction and the ground truth is defined as loss function for the training. Unlike other segmentation networks which apply image patches for training, registration needs to consider global features and space relations between vessel structures. Therefore, the network takes down-sampled CT/MRI and replicated US images as input. On the other hand, the InitNet should handle large deviations in translation and rotations as initial pose. Consequently, the network is configured with a large depth and a wide feature base. After prediction, the output volume presents segmented voxels on the US plane in a 3D volume. The plane parameters can be derived by plane regression algorithms. However, as discussed in our previous works, because of the limited output resolution, the plane regression can result in large angle errors. Additionally, false positives in the predicted volume can have a negative impact on accuracy, especially with output resolution.

LoopNet

Due to the high down-sampling rate of the input image, the accuracy of InitNet prediction has the potential to be improved. To this end, we propose LoopNet architectures which connect to the InitNet and focus on improving the registration accuracy with few loops.

IRL-Net

The IRL-Net starts with InitNet on the left side of a cascaded network (see Fig. 1). A plane regression module is implemented to detect the plane parameters with $\vec{p}$ and $\vec{n}$ which denote a point inside a plane and a normal vector perpendicular to that plane, respectively. This plane normal vector is put into a resample module to correct the orientation of the input volume with higher resolution. The LoopNet is implemented based on a U-Net, which has less depth layers than the InitNet. Thus, the network can be trained with higher resolution images. As discussed in the previous work, the registration accuracy has a dependency on its initial pose. The average registration error is propotional to the deviation of US slice rotated away from XY plane of 3D CT/MRI image. Therefore, the output of the LoopNet is fed back to the regression module to improve the angle between CT/MRI XY plane and the slice plane. To enhance the accuracy of the LoopNet, we used high resolution data to train the LoopNet.

ILM-Net

The IRL-Net represents an architecture that can achieve registration in few loops. However, it applies the plane registration algorithm, which is a non-trainable module with additional hyper parameters. Furthermore, since the robust plane estimation algorithm is an iterative process, this can slow down the performance of whole registration. To address this, a fully convolution network ILM-Net (see Fig. 2) is proposed. The ILM-Net starts with the InitNet, then up-samples the prediction to a higher resolution. This high resolution output is stacked with CT/MRI and replicated US images to a multi-channel image. The

LoopNet is implemented with less layers in depth and takes higher resolution images as input. In this case, the InitNet result is used as region-of-interest to guide the LoopNet. The DICE coefficient is employed as loss function to train the ILM network.

Experimental Results

Data Preparation

For the experiments, we used image data from 52 patients acquired at Sir Run Run Shaw Hospital, School of Medicine, Zhejiang University. The data collection had been approved by the ethics committee of the hospital. Each patient had one 3D image volume scanned by CT from GE LightSpeed VCT system or MR from Siemens MAGNETOM Skyra system, and series of 2D US image data acquired by Telemed SmartUS system with C5-2R60HI-5 probe. The CT/MR 3D volumes contained 150 – 250 slices with 256×256 voxels. The voxel spacing was normalized to $[1.0, 1.0, 1.0]$ mm.

From 28 patients we acquired US in transcostal position and from the other 24 patients in medial position. 1035 US image data were collected in total. For each US image, field experts annotated manually the ground truth including the class label of the US slice pose and the transformation from a US slice to the 3D CT/MR volume. In details, field experts assigned the slice pose labels manually based on US contents. Subsequently, manual slice-to-volume registration was performed with a self developed software, the appropriate rotations and translations between the 2D slice and the 3D volume were determined using key and mouse functions. With an approximate ratio of 10:2:1, the image data and patients were split into training, validation, and testing dataset.

Network settings

The training images were down-scaled to $40 \times 40 \times 40$ voxels with a spacing of 4 mm and $80 \times 80 \times 80$ voxels with a spacing of 2 mm, respectively. The former low resolution images were used to train the InitNet. The images with higher resolution were used to train the LoopNet. The InitNet was configured with depth of 4, filter size of $3 \times 3 \times 3$, batch size of 20, a learn rate of 0.0003 and Adam optimizer. The LoopNet was implemented to process high resolution images and configured with depth of 3, filter size of $3 \times 3 \times 3$, batch size of 4, a learn rate of 0.0003 and Adam optimizer. To prevent overfitting, we used a large amount of synthetic vessel images resliced from 3D vessel volumes instead of limited number of real US vessel images. Data augmentation schemes were applied to vary US plane pose from ground truth position, including a random rotation of $[-20, 20]^\circ$ and random translation of $[-20, 20]$ mm.

First, we trained LoopNet for IRL-Net. The network was fed with a two channel input containing 3D vessel volume and replicated vessel image. We used DICE coefficient as loss function to optimize the model parameters. Then, we trained LoopNet for ILM-Net. We fed the network with a 3 channel input comprising 3D vessel volume, replicated vessel image and the prediction from the previous loop in the third channel. More specifically, the training of loop 0 put the prediction of InitNet into the third input channel. The training at loop N loaded the trained model from loop N-1 and put the prediction of loop N-1 into the third channel to optimize parameters for loop N. In total, LoopNet are trained up to 10 loops to ensure the convergence. After training, IRL-Net and ILM-Net were tested with 85 images. Both training and testing was performed on a NVIDIA Geforce GTX 1080Ti.

Figure 3. Registration error of IRL compared with network loops. Left: Angle error in degree; Right: distance error in millimeter. The first column on the left in each chart shows the results of InitNet.

IRL-Net

In the evaluation, the LoopNet was trained with 10 loops. The evaluation metrics included a rotation error and a distance error. The rotation error was calculated by considering the angle between normal vectors of the predicted and the ground truth plane. The distance error was determined by calculating the average distance between vessel voxels in the ground truth and prediction plane. As shown in Figure 3, the angle and distance error decreases with loops. After the 4th loop, both errors converged. Therefore, the hyperparameter loop number is assigned with 4, which shows the best registration accuracy with an average angle error of 13.3° and an average distance error of 4.5 mm. This shows a significant improvement compared to the InitNet, with an average angle error of 18.6° and an average distance error of 4.9 mm.

ILM-Net

For each test image, the evaluation starts with InitNet for a coarse alignment. The prediction result is then upsampled, concatenated with input CT/MRI and replicated volume and used later as input for the LoopNet. The LoopNet prediction is launched up to 10 loops to determine the best loop in an experimental way.

The Figure 4 shows the DICE scores on loops. The left column represents the evaluation on loops with low initial angle error within 20 degrees. In this case, the DICE score begins with high average value at InitNet but does not show an improvement on loops. The middle and right column represents the evaluation with relative high initial angle error in the range of $[20^\circ, 40^\circ]$ degrees. The results show a significant DICE score improvement along the loops. More specifically, the DICE scores increase

Figure 4. DICE score of ILM prediction results. Left: Prediction with initial US plane angle error in range of $[10^\circ, 20^\circ]$; Middle: Prediction with initial US plane angle error in range of $[20^\circ, 30^\circ]$; Right: Prediction with initial US plane angle error in range of $[30^\circ, 40^\circ]$.

significantly in the first two loops. Then the DICE scores converge to a stable value after the 4th loop. Further, the prediction results are postprocessed by a plane regression method. The registration accuracy is evaluated with the angle error and the distance error between the prediction plane and the ground truth plane. As shown in Figure 5 with the initial angle error $[20^\circ, 30^\circ]$, the InitNet prediction results achieve a low registration error both in angle and distance. The error is reduced slightly on further loops. On the other hand, when a large initial angle error $[30^\circ, 40^\circ]$ is given, the registration error of InitNet is relative high. This error is reduced significantly in the first two loops and converged to a stable value at the 4th loop. The overall registration error can be summarized in Figure 6. The best results are obtained at loop 2, presenting an average angle error of 17.4° and an average distance error of 4.9 mm.

Discussion

Compared to the InitNet, both IRL and ILM networks can improve the registration accuracy by loops. As shown in the first row in Figure 7, IRL and ILM registration results are improved by loops visually. In addition, the improvement can be recognized at left portal vein branches marked with circles in the overlay images in the second row. The overlay image is generated by overlapping resliced images after plane regression over the ground truth image. The plane regression determines one point on the plane and the plane normal vector, therefore, the cropped out 2D images can have a slight deviation in in-plane translation and rotation. To enable the comparison in 2D images, the resliced images are aligned with those resliced images on ground truth position manually. The result of InitNet shows a coarse alignment with a rare overlap on the left portal vein vessel branch. This is improved in further registration loops with IRL and ILM methods. Figure 8 visualizes a sample of registration results with high, middle, and low accuracy respectively. On the left column, the overlap covers both major and minor vessels. The results with middle accuracy show overlap only on the major vessels. Results on the right column show a poor alignment

Figure 5. Registration error of ILM prediction results. First row: Angle error with 10 loops; Second row: Distance error with 10 loops; Columns from left to right represent the errors grouped by the initial angle error in range $[20^\circ, 30^\circ]$ and $[30^\circ, 40^\circ]$

Figure 6. Overall registration error of ILM on loops. Left: angle error in degree; Right: distance error in millimeter.

between the resliced images on the ground truth position and the prediction position. The low registration accuracy is caused by less availability of vessels.

In addition, the evaluation result in Figure 3 indicates that the IRL network achieves a significant improvement of registration accuracy within the first four loops and converges at the 4th loop. However, the IRL network employs a plane regression module to calculate the plane pose within each loop. Therefore, the IRL is a deep network-based hybrid approach. On the other hand, the ILM uses a fully deep network approach to solve registration in an iterative manner. However, experimental results show that its performance is depend on the initial angle error. Starting with a high initial angle error as shown in Figure 5, the InitNet shows a high angle and distance error, the registration results are improved after each loop. More specifically, the results with $[30-40]$ initial angle error on the right column show a more significant improvement on angle error than those with $[20-30]$ initial angle error on the left column.

To compare our LoopNet methods with a state-of-the-art, we implemented SVR Deep method proposed by¹⁴, which is very closely related to our methods, by means of applying regression strategy to align 2D and 3D medical images. We evaluated the SVR Deep method with the same dataset. The results in Table 1 show that the InitNet, IRL, and ILM outperform the SVR method both in angle error and distance error. Moreover, the IRL method shows a significant improvement to the InitNet both

Figure 7. Sample of resliced images on registration loops. The first row shows the results on loop samples in 3D view, with red planes indicating the InitNet results, green planes indicating the ground truth and yellow planes indicating the loop samples. The second row compares loop samples on resliced planes. Green: ground truth, yellow: prediction, blue: overlapped area.

Method	SVR ¹⁴	InitNet	IRL	ILM
Angle Error Mean	37.0	18.6	13.3	17.4
Angle Error Stdev	15.6	8.2	6.3	7.9
Distance Error Mean	19.0	4.9	4.4	4.9
Distance Error Stdev	11.6	2.0	2.4	1.9

Table 1. Comparison of state-of-the-art SVR method with InitNet, IRL at the 4th. loop, and ILM at the second loop.

in angle error and distance error. The registration results are visualized on three samples in the Figure 9. The SVR Deep method with case 1 shows a low distance error but a high angle error. The resliced image shows vessels in yellow close to the ground truth vessel structures labeled in green. The SVR Deep result shows a poor alignment on case 2 and case 3, where the distance error is more than 20 mm. With these large distance errors, the predicted planes are positioned out of the region of the target vessels.

According to the experimental evaluation results shown in Figure 3 and Figure 6, the IRL and ILM converge at the 4th and the second loop, respectively. In the further loops, both networks show a relatively high robustness at the converged state. Therefore, we can choose the 4th loop and the second loop as early stop criteria for IRL and ILM respectively.

Conclusions

This work proposed two loop network architectures which aim at improving registration accuracy with less loops. The IRL network achieves a high registration accuracy at the 4th loop with an average angle error of 13.3° and an average distance error of 4.5 mm. It outperforms the ILM network at the second loop with angle error of 17.4° and distance error of 4.9 mm and the InitNet with angle error of 18.6° and distance error of 4.9 mm. The evaluation results show the efficiency of the proposed IRL and ILM registration networks, which have the potential to improve the robustness and accuracy of intraoperative patient

Figure 8. Sample of IRL and ILM registration results with high, middle, and low accuracy. First row compares the results in 3D views with green planes indicating the ground truth and yellow planes indicating the predictions. The second row shows the results on the resliced planes. Green: ground truth, yellow: prediction, blue: overlapped area.

registration.

Acknowledgement

This work has been funded by Zhejiang Provincial Natural Science Foundation of China under Grant No.LSY19H180012.

Conflict of interest statement

No potential conflict of interest was reported by the authors

References

1. Penney G. P., et al. Overview of an ultrasound to CT or MR registration system for use in thermal ablation of liver metastases. In *Proc Med Imag Underst Anal*, vol. 1, 6568 (2001).
2. Fusaglia, M., Tinguely, P., Banz, V., Weber, S. & Lu, H. A novel ultrasound-based registration for image-guided laparoscopic liver ablation. *Surg. innovation* **23**, 397–406 (2016).
3. Spinczyk, D. Towards the clinical integration of an image-guided navigation system for percutaneous liver tumor ablation using freehand 2D ultrasound images. *Comput. Aided Surg.* **20**, 61–72 (2015).

Figure 9. Visual comparison of SVR, IRL and ILM methods on three cases. First row shows the results in 3D view. Green planes indicates the ground truth. Yellow planes indicate the predicted planes. The second row compares results of methods in the reduced planes. Green: ground truth, yellow: prediction, blue: overlapped area.

- Pohlman, R. M. *et al.* Two-dimensional ultrasound-computed tomography image registration for monitoring percutaneous hepatic intervention. *Med. physics* **46**, 2600–2609 (2019).
- Wein, W., Brunke, S., Khamene, A., Callstrom, M. R. & Navab, N. Automatic CT-ultrasound registration for diagnostic imaging and image-guided intervention. *Med. image analysis* **12**, 577–585 (2008).
- Koenig, L., Ruehaak, J., Derksen, A. & Lellmann, J. A matrix-free approach to parallel and memory-efficient deformable image registration. *SIAM J. on Sci. Comput.* **40**, B858–B888 (2018).
- Haskins, G., Kruger, U. & Yan, P. Deep learning in medical image registration: a survey. *Mach. Vis. Appl.* **31**, 8 (2020).
- Haskins, G. *et al.* Learning deep similarity metric for 3D MR–TRUS image registration. *Int. journal computer assisted radiology surgery* **14**, 417–425 (2019).
- Fan, J., Cao, X., Xue, Z., Yap, P.-T. & Shen, D. Adversarial similarity network for evaluating image alignment in deep learning based registration. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 739–746 (Springer, 2018).

10. Yan, P., Xu, S., Rastinehad, A. R. & Wood, B. J. Adversarial image registration with application for MR and TRUS image fusion. In *International Workshop on Machine Learning in Medical Imaging*, 197–204 (Springer, 2018).
11. Ma, K. *et al.* Multimodal image registration with deep context reinforcement learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 240–248 (Springer, 2017).
12. Sloan, J. M., Goatman, K. A. & Siebert, J. P. Learning rigid image registration - utilizing convolutional neural networks for medical image registration. In *Proceedings of the 11th International Joint Conference on Biomedical Engineering Systems and Technologies*, 89–99 (Springer, 2018).
13. Hou, B. *et al.* 3-D reconstruction in canonical co-ordinate space from arbitrarily oriented 2-D images. *IEEE transactions on medical imaging* **37**, 1737–1750 (2018).
14. Salehi, S. S. M., Khan, S., Erdogmus, D. & Gholipour, A. Real-time deep registration with geodesic loss. *arXiv preprint arXiv:1803.05982* (2018).
15. Hou, B. *et al.* Predicting slice-to-volume transformation in presence of arbitrary subject motion. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 296–304 (Springer, 2017).
16. Almogadwy, B., McLeod, K. & Burger, A. A Deep Learning Approach for Slice to Volume Biomedical Image Integration. In *Proceedings of the 2019 11th International Conference on Bioinformatics and Biomedical Technology*, 62–68 (2019).
17. Ernst, P., Hille, G., Hansen, C., Tönnies, K. & Rak, M. A CNN-Based framework for statistical assessment of spinal shape and curvature in whole-body MRI images of large populations. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 3–11 (Springer, 2019).
18. Wei, W., Rak, M., Alpers, J. & Hansen, C. Towards fully automatic 2D US to 3D CT/MR Registration: A novel segmentation-based Strategy. In *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, 433–437 (IEEE, 2020).
19. Wei, W., Haishan, X., Alpers, J., Rak, M. & Hansen, C. A deep learning approach for 2d ultrasound and 3d ct/mr image registration in liver tumor ablation. *Comput. Methods Programs Biomed.* **206**, 106117 (2021).