



HAL
open science

Trackerless Volume Reconstruction from Intraoperative Ultrasound Images

Sidaty El Hadramy, Juan Verde, Karl-Philippe Beaudet, Nicolas Padoy,
Stéphane Cotin

► **To cite this version:**

Sidaty El Hadramy, Juan Verde, Karl-Philippe Beaudet, Nicolas Padoy, Stéphane Cotin. Trackerless Volume Reconstruction from Intraoperative Ultrasound Images. MICCAI 2023, Oct 2023, Vancouver, France. hal-04108248v1

HAL Id: hal-04108248

<https://inria.hal.science/hal-04108248v1>

Submitted on 26 May 2023 (v1), last revised 20 Jul 2023 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Trackerless Volume Reconstruction from Intraoperative Ultrasound Images

Sidaty El hadramy^{1,2,3}, Juan Verde^{2,3}, Nicolas Padoy^{2,3}, and Stéphane Cotin^{1,3*}

¹ INRIA, Strasbourg, France

² IHU Strasbourg, Strasbourg, France

³ University of Strasbourg, Strasbourg, France.

Abstract. This paper proposes a method for trackerless ultrasound volume reconstruction in the context of minimally invasive surgery. It is based on a Siamese architecture, including a recurrent neural network that leverages the ultrasound image features and the optical flow to estimate the relative position of frames. Our method does not use any additional sensor and was evaluated on *ex-vivo* porcine data. It achieves translation and orientation errors of 0.449 ± 0.189 mm and 1.3 ± 1.5 degrees respectively for the relative pose estimation. In addition, despite the predominant non-linearity motion in our context, our method achieves a good reconstruction with final and average drift rates of 23.11% and 28.71% respectively. To the best of our knowledge, this is the first work to address volume reconstruction in the context of intravascular ultrasound. Code will be publicly released.

Keywords: Intraoperative Ultrasound · Liver Surgery · Volume Reconstruction · Recurrent Neural Networks

1 Introduction

Liver cancer is the most prevalent indication for liver surgery, and although there have been notable advancements in oncologic therapies, surgery remains as the only curative approach overall [20]. Liver laparoscopic resection has demonstrated fewer complications compared to open surgery [21], however, its adoption has been hindered by several reasons, such as the risk of unintentional vessel damage, as well as oncologic concerns such as tumor detection and margin assessment. Hence, the identification of intrahepatic landmarks, such as vessels, and target lesions is crucial for successful and safe surgery, and intraoperative ultrasound (IOUS) is the preferred technique to accomplish this task. Despite the increasing use of IOUS in surgery, its integration into laparoscopic workflows (i.e., laparoscopic intraoperative ultrasound) remains challenging due to combined problems. Performing IOUS during laparoscopic (LIOUS) liver surgery poses significant challenges, as laparoscopy has poor ergonomics and narrow fields of view, and on the other hand, IOUS demands skills to manipulate the probe and analyze images. At the end, and **despite its real-time capabilities**,

* Corresponding author. E-mail: stephane.cotin@inria.fr

IOUS images are **intermittent and asynchronous** to the surgery, requiring multiple iterations and repetitive steps (probe-in - instruments-out, probe-out - instruments-in). Therefore, any method enabling a continuous and synchronous US assessment throughout the surgery, with minimal iterations required would significantly improve the surgical workflow, as well as its efficiency and safety.

To overcome these limitations, the use of intravascular ultrasound (IVUS) images has been proposed, enabling **continuous and synchronous inside-out imaging** during liver surgery [19]. With an intravascular approach, an overall view and full-thickness view of the liver can quickly and easily be obtained through mostly rotational movements of the catheter, while this is constrained to the lumen of the *inferior vena cava*, and with no interaction with the tissue (contactless, a.k.a. standoff technique) as illustrated in Figure 1.




Fig. 1: **left**: IVUS catheter positioned in the lumen of the *inferior vena cava* in the posterior surface of the organ, and an example of the lateral firing and longitudinal beam-forming images; **middle**: anterior view of the liver and the rotational movements of the catheter providing full-thickness images; **right**: inferior view showing the rotational US acquisitions

However, to benefit from such a technology in a computer-guided solution, the different US images would need to be tracked and possibly integrated into a volume for further processing. External US probes are often equipped with an electromagnetic tracking system to track its position and orientation in real-time. This information is then used to register the 3D ultrasound image with the patient’s anatomy. The use of such an electromagnetic tracking system in laparoscopic surgery is more limited due to size reduction. The tracking system may add additional complexity and cost to the surgical setup, and the tracking accuracy may be affected by metallic devices in the surgical field [22].

Several approaches have been proposed to address this limitation by proposing a trackerless ultrasound volume reconstruction. Physics-based methods have exploited speckle correlation models between different adjacent frames [6–8] to estimate their relative position. With the recent advances in deep learning, recent works have proposed to learn a higher order nonlinear mapping between adjacent frames and their relative spatial transformation. *Prevost et al* [9] first demonstrated the effectiveness of a convolution neural network to learn the relative motion between a pair of US images. *Xie et al* [10] proposed a pyramid warping

layer that exploits the optical flow features in addition to the ultrasound features in order to reconstruct the volume. To enable a smooth 3D reconstruction, a case-wise correlation loss based on 3D CNN and Pearson correlation coefficient was proposed in [10, 12]. *Qi et al* [13] leverages past and future frames to estimate the relative transformation between each pair of the sequence; they used the consistency loss proposed in [14]. Despite the success of these approaches, they still suffer significant cumulative drift errors and mainly focus on linear probe motions. Recent work [15, 16] proposed to exploit the acceleration and orientation of an inertial measurement unit (IMU) to improve the reconstruction performance and reduce the drift error. Motivated by the weakness of the state-of-the-art methods when it comes to large non-linear probe motions, and the difficulty of integrating IMU sensors in the case of minimally invasive procedures, we introduce a new method for pose estimation and volume reconstruction in the context of minimally invasive trackerless ultrasound imaging. We use a Siamese architecture based on a Sequence to Vector(Seq2Vec) neural network that leverages image and optical flow features to learn relative transformation between a pair of images.

Our method improves upon previous solutions in terms of robustness and accuracy, particularly in the presence of rotational motion. Such motion is predominant in the context highlighted above and is the source of additional non-linearity in the pose estimation problem. To the best of our knowledge, this is the first work that provides a clinically sound and efficient 3D US volume reconstruction during minimally invasive procedures. The paper is organized as follows: section 2 details the method and its novelty, section 3 presents our current results on *ex-vivo* porcine data, and finally, we conclude in section 4 and discuss future work.

2 Method

In this work, we make the assumption that the organ of interest does not undergo deformation during the volume acquisition. This assumption is realistic due to the small size of the probe. Let $I_0, I_1 \dots I_N$ be a sequence of N frames. Our aim is to find the relative spatial transformation between each pair of frames I_i and I_j with $0 \leq i \leq j \leq N$. This transformation is denoted $T_{(i,j)}$ and is a six degrees of freedom vector representing three translations and three Euler angles. To achieve this goal, we propose a Siamese architecture that leverages the optical flow in the sequences in addition to the frames of interest in order to provide a mapping with the relative transformation in the space. The overview of our method is presented in the Fig 2.

We consider a window of $2k + 3$ frames from the complete sequence of length N , where $0 \leq k \leq \lfloor \frac{N-3}{2} \rfloor$ is a hyper-parameter that denotes the number of frames between two frames of interest. Our method predicts two relative transformations between the pairs of frames (I_1, I_{k+2}) and (I_{k+2}, I_{2k+3}) . The input window is divided into two equal sequences of length $k + 2$ sharing a common frame. Both deduced sequences are used to compute a sparse optical flow allowing to track the trajectory of M points. Then, Gaussian heatmaps are used to describe the

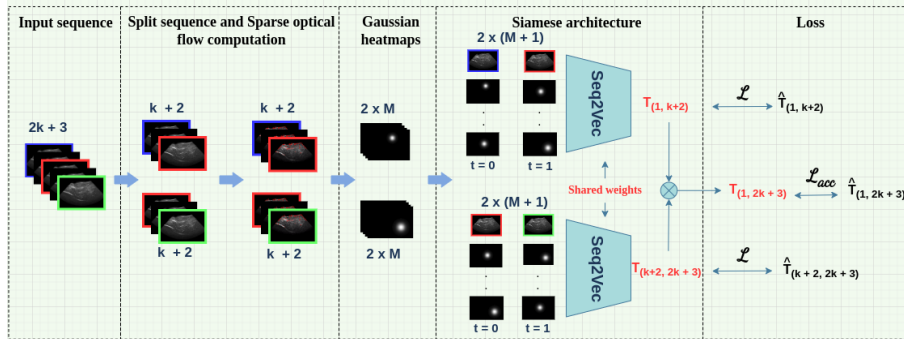


Fig. 2: Overview of the proposed method. The input sequence is split into two equal sequences with a common frame. Both are used to compute a sparse optical flow. Gaussian heatmaps tracking M points are then combined with the first and last frame of each sequence to form the network’s input. We use a Siamese architecture based on Sequence to Vector (Seq2Vec) network. The learning is done by minimising the mean square error between the output and ground truth transformations.

motion of the M points in an image-like format(see section 2.2). Finally, a Siamese architecture based on two shared weights Sequence to Vector (Seq2Vec) takes as input the Gaussian heatmaps in addition to the first and last frames and predicts the relative transformations. In the following we detail our pipeline.

2.1 Sparse Optical Flow

Given a sequence of frames I_i and I_{i+k+1} , we aim at finding the trajectory of a set of points throughout the sequence. We choose the M most prominent points from the first frame using the feature selection algorithm proposed in [3]. Points are then tracked throughout each pair of adjacent frames in the sequence by solving the equation 1 which is known as the Optical flow equation. We use the pyramidal implementation of Lucas-Kanade method proposed in [4] to solve the equation. Thus, yielding a trajectory matrix $A \in \mathbb{R}^{M \times (i+k+2) \times 2}$ that contains the position of each point throughout the sequence. The figure 3 illustrates an example where we track two points in a sequence of frames.

$$I_i(x, y, t) = I_i(x + dx, y + dy, t + dt) \quad (1)$$

2.2 Gaussian Heatmaps

After obtaining the trajectory of M points in the sequence $\{I_i | 1 \leq i \leq k+2\}$ we only keep the first and last position of each point, which corresponds to the positions in our frames of interest. We use Gaussian heatmaps $\mathcal{H} \in \mathbb{R}^{H \times W}$ with the same dimension as the ultrasound frames to encode these points, they are more suitable as input for the convolutional networks. For a point with a position (x_0, y_0) , the corresponding heatmap is defined in the equation 2.



(a) First frame in the sequence

(b) Last frame in the sequence

Fig. 3: Sparse Optical tracking of two points in a sequence, red points represent the chosen points to track, while the blue lines describe the trajectory of the points throughout the sequence.

$$\mathcal{H}(x, y) = \frac{1}{\sigma^2 \sqrt{2\pi}} e^{-\frac{(x-x_0)^2 + (y-y_0)^2}{2\sigma^2}} \quad (2)$$

Thus, each of our M points are converted to a pair of heatmaps that represent the position in the first and last frames of the ultrasound sequence. These pairs concatenated with the ultrasound first and last frames form the recurrent neural network sequential input of size $(M + 1, H, W, 2)$, where $M + 1$ is the number of channels (M heatmaps and one ultrasound frame), H and W are the height and width of the frames and finally 2 represents the dimension over time.

2.3 Network Architecture

The Siamese architecture is based on a sequence to vector network. Our network maps a sequence of two images having $M + 1$ channel each to a six degrees of freedom vector (three translations and three rotation angles). The architecture of Seq2Vec is illustrated in the Fig 4. It contains five times the same block composed of two Convolutional LSTMs (ConvLSTM) [5] followed by a Batch Normalisation. Their output is then flattened and mapped to a six degrees of freedom vector through linear layers; ReLU is the chosen activation function for the first linear layer. We use an architecture similar to the one proposed in [5] for the ConvLSTM layers. Seq2Vec networks share the same weights.

2.4 Loss function

In the training phase, given a sequence of $2k+3$ frames in addition to their ground truth transformations $\hat{T}_{(1,k+2)}$, $\hat{T}_{(k+2,2k+3)}$ and $\hat{T}_{(1,2k+3)}$, the Seq2Vec's weights are optimized by minimising the loss function given in the equation 3. The loss contains two terms. The first represents the mean square error (MSE) between the estimated transformations $(T_{(1,k+2)}, T_{(k+2,2k+3)})$ at each corner point of the frames and their respective ground truth. The second term represents the accumulation loss that aims at reducing the error of the volume reconstruction, the effectiveness of the accumulation loss have been proven in the literature [13]. It is written as the MSE between the estimated $T_{(1,2k+3)} = T_{(k+2,2k+3)} \times T_{(1,k+2)}$ at the corner points of the frames and the ground truth $\hat{T}_{(1,2k+3)}$.




Fig. 4: **Architecture of Seq2Vec network.** We use five blocks that contains each two ConvLSTM followed by Batch Normalisation. The output is flattened and mapped to a six degrees of freedom translation and rotation angles through linear layers. The network takes as input a sequence of two images with $M + 1$ channel each, M heatmaps and an ultrasound frame. The output corresponds to the relative transformation between the blue and red frames.

$$\mathcal{L} = \|T_{(1,k+2)} - \hat{T}_{(1,k+2)}\|_2 + \|T_{(k+2,2k+3)} - \hat{T}_{(k+2,2k+3)}\|_2 + \|T_{(1,2k+3)} - \hat{T}_{(1,2k+3)}\|_2 \quad (3)$$

3 Results and Discussion

3.1 Dataset and implementation details

To validate our method, six tracked sequences were acquired from an *ex-vivo* swine liver. A manually manipulated IVUS catheter was used (8 Fr lateral firing AcuNav™, 4–10 MHz) connected to an ultrasound system (ACUSON S3000 HELX Touch, Siemens Healthineers, Germany), both commercially available. An electromagnetic tracking system (trakSTAR™, NDI, Canada) was used along with a 6 DoF sensor (Model 130) embedded close to the tip of the catheter, and the PLUS toolkit [17] along with 3D Slicer [18] were used to record the sequences. The frame size was initially 480×640 . Frames were cropped to remove the patient and probe characteristics, then down-sampled to a size of 128×128 with an image spacing of 0.22 mm per pixel. First and end stages of the sequences were removed from the six acquired sequences, as they were considered to be largely stationary, and aiming to avoid training bias. Clips were created by sliding a window of 7 frames (corresponding to a value of $k = 2$) with a stride of 1 over each continuous sequence, yielding a data set that contains a total of 13734 clips. The tracking was provided for each frame as a 4×4 transformation matrix. We have converted each to a vector of six degrees of freedom that corresponds to three translations in *mm* and three Euler angles in *degrees*. For each clip, relative frame to frame transformations were computed for the frames number 0, 3 and 6. The distribution of the relative transformation between the frames in our clips is illustrated in the figure 5. It is clear that our data mostly contains rotations, in particular over the axis x . Heatmaps were calculated for two points

($M = 2$) and with a quality level of 0.1, a minimum distance of 7 and a block size of 7 for the optical flow algorithm (see [4] for more details). The number of heatmaps M and the frame jump k were experimentally chosen among 0, 2, 4, 6. The data was split into train, validation and test sets by a ratio of 7:1.5:1.5. Our method is implemented in *Pytorch*⁴ 1.8.2, trained and evaluated on a *GeForce RTX 3090*. We use an Adam optimizer with a learning rate of 10^{-4} . The training process converges in 40 epochs with a batch size of 16. The model with the best performance on the validation data was selected and used for the testing.

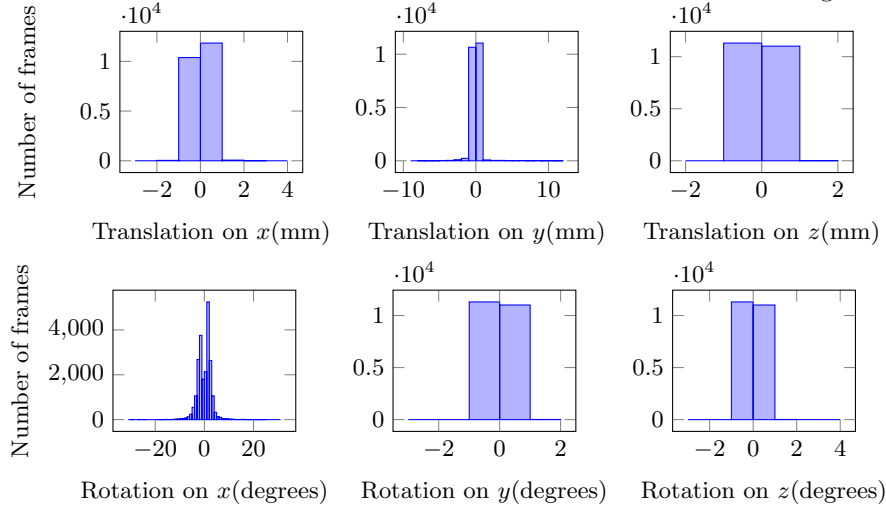


Fig. 5: The distribution of the relative rotations and translations over the data

3.2 Evaluation metrics and results


The test data was used to evaluate our method, it contains 2060 clips over which our method achieved a translation error of $\epsilon_{translation}$ of 0.449 ± 0.189 mm, and an orientation error of $\epsilon_{orientation}$ 1.3 ± 1.5 degrees. We have evaluated our reconstruction with a commonly used in state-of-the-art metric called final drift error, which measures the distance between the center point of the final frame according to the real relative position and the estimated one in the sequence. On this basis, each of the following metrics was reported over the reconstructions of our method. **Final drift rate (FDR)**: the final drift divided by the sequence length. **Average drift rate (ADR)**: the average cumulative drift of all frames divided by the length from the frame to the starting point of the sequence. Table 1 shows the evaluation of our method over these metrics compared to the state-of-the-art methods MoNet [15] and CNN [9]. Both state-of-the-art methods use IMU sensor data as additional input to estimate the relative transformation between two relative frames. Due to the difficulty of including an IMU sensor in our IVUS catheter, the results of both methods were reported from the MoNet paper where the models have been trained on arm scans, see [15] for more details.

⁴ <https://pytorch.org/docs/stable/index.html>


Models	FDR(%)	ADR(%)
CNN [9]	31.88(15.76)	39.71(14.88)
MoNet [15]	15.67(8.37)	25.08(9.34)
Ours	23.11(11.6)	28.71(12.97)

Table 1: The mean and standard deviation FDR and ADR of our method compared with state-of-the-art models MoNet [15] and CNN [9]

As the table 1 shows, our method is comparable with state-of-the-art methods in terms of drift errors without using any IMU and with non-linear probe motion as one may notice in our data distribution in the Fig 5. Fig 6 shows the volume reconstruction of two sequences of different sizes with our method in red against the ground truth slices. Despite the non-linearity of the probe motion, the relative pose estimation results obtained by our method remains very accurate. However, one may notice that the drift error increases with respect to the sequence length. This remains a challenge for the community even in the case of linear probe motions.



(a) Sequence of length 50 frames



(b) Sequence of length 300 frames

Fig. 6: The reconstruction of two sequences of lengths 50 and 300 respectively with our method in red compared with the ground truth sequences.

4 Conclusion

In this paper, we proposed the first method for trackerless ultrasound volume reconstruction in the context of minimally invasive surgery. Our method does not use any additional sensor data and is based on a Siamese architecture that leverages the ultrasound image features and the optical flow to estimate relative transformations. Our method was evaluated on *ex-vivo* porcine data and achieved translation and orientation errors of 0.449 ± 0.189 mm and 1.3 ± 1.5 degrees respectively with a fair drift error. In the future work, we will extend our work to further improve the volume reconstruction and use it to register a pre-operative CT image in order to provide guidance during interventions.

References

1. De Gottardi A, Keller P-F, Hadengue A, Giostra E, Spahr L. Transjugular intravascular ultrasound for the evaluation of hepatic vasculature and parenchyma in patients with chronic liver disease. *BMC Res Notes*. 2012;5:77.
2. Takeshi Urade, Juan Verde, Alain García Vázquez, Konstance Gunzert, Patrick Pessaux, et al.. Fluoro- less intravascular ultrasound image-guided liver navigation in porcine models. *BMC Gastroenterology*, 2021, 21, pp.24. 10.1186/s12876-021-01600-3.
3. Jianbo Shi and Carlo Tomasi. Good features to track. In *Computer Vision and Pattern Recognition, 1994. Proceedings CVPR'94., 1994 IEEE Computer Society Conference on*, pages 593–600. IEEE, 1994.
4. Jean-Yves Bouguet. Pyramidal implementation of the affine lucas kanade feature tracker description of the algorithm. Intel Corporation, 5, 2001.
5. Shi, Xingjian, et al. "Convolutional LSTM network: A machine learning approach for precipitation nowcasting." *Advances in neural information processing systems* 28 (2015).
6. Mercier, L., LangØ, T., Lindseth, F., Collins, D.L.: A review of calibration techniques for freehand 3-d ultrasound systems. *Ultrasound in Medicine Biology* 31,449–471 (2005).
7. Mohamed, F., Siang, C.V.: A survey on 3d ultrasound reconstruction techniques. *Artificial Intelligence—Applications in Medicine and Biology* (2019).
8. Mozaffari, M.H., Lee, W.S.: Freehand 3-d ultrasound imaging: a systematic review. *Ultrasound in medicine biology* 43(10), 2099–2124 (2017).
9. Prevost, R., Salehi, M., Sprung, J., Ladikos, A., Bauer, R., Wein, W.: Deep learning for sensorless 3d freehand ultrasound imaging. In: *International conference on medical image computing and computer-assisted intervention*. pp. 628–636. Springer (2017).
10. Xie Y, Liao H, Zhang D, Zhou L, Chen F. Image-Based 3D Ultrasound Reconstruction with Optical Flow via Pyramid Warping Network. *Annu Int Conf IEEE Eng Med Biol Soc*. 2021.
11. Guo, H., Xu, S., Wood, B., Yan, P.: Sensorless freehand 3d ultrasound reconstruction via deep contextual learning. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 463–472. Springer (2020).
12. H. Guo, H. Chao, S. Xu, B. J. Wood, J. Wang and P. Yan, "Ultrasound Volume Reconstruction From Freehand Scans Without Tracking," in *IEEE Transactions on Biomedical Engineering*, vol. 70, no. 3, pp. 970-979, March 2023.
13. Qi Li, Ziyi Shen, Qian Li, Dean C Barratt, Thomas Dowrick, Matthew J Clarkson, Tom Vercauteren and Yipeng Hu. Trackerless freehand ultrasound with sequence modelling and auxiliary transformation over past and future frames. arXiv:2211.04867v2. Nov 2022.
14. Kanta Miura, Koichi Ito, Takafumi Aoki, Jun Ohmiya, and Satoshi Kondo "Probe localization from ultrasound image sequences using deep learning for volume reconstruction", *Proc. SPIE 11792, International Forum on Medical Imaging in Asia 2021, 117920O* (20 April 2021).
15. Luo, M., Yang, X., Wang, H., Du, L., Ni, D. (2022). Deep Motion Network for Freehand 3D Ultrasound Reconstruction. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*. MICCAI 2022. Lecture Notes in Computer Science, vol 13434. Springer, Cham.

16. G. Ning, H. Liang, L. Zhou, X. Zhang and H. Liao, "Spatial Position Estimation Method for 3D Ultrasound Reconstruction Based on Hybrid Transformers," 2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI), Kolkata, India, 2022.
17. Lasso A, Heffter T, Rankin A, Pinter C, Ungi T, Fichtinger G. PLUS: Open-Source Toolkit for Ultrasound-Guided Intervention Systems. *IEEE Trans Biomed Eng.* 2014 Oct;61(10):2527–37.
18. Fedorov A., Beichel R., Kalpathy-Cramer J., Finet J., Fillion-Robin J-C., Pujol S., Bauer C., Jennings D., Fennessy F.M., Sonka M., Buatti J., Aylward S.R., Miller J.V., Pieper S., Kikinis R. 3D Slicer as an Image Computing Platform for the Quantitative Imaging Network. *Magnetic Resonance Imaging.* 2012 Nov;30(9):1323-41. PMID: 22770690. PMCID: PMC3466397
19. Urade T, Verde JM, García Vázquez A, et al. Fluoroless intravascular ultrasound image-guided liver navigation in porcine models. *BMC Gastroenterol.* 2021;21(1):24.
20. Aghayan DL, Kazaryan AM, Dagenborg VJ, et al. Long-Term Oncologic Outcomes After Laparoscopic Versus Open Resection for Colorectal Liver Metastases: A Randomized Trial. *Ann Intern Med.* 2021;174(2):175-182.
21. Fretland ÅA, Dagenborg VJ, Bjørnelv GMW, et al. Laparoscopic Versus Open Resection for Colorectal Liver Metastases: The OSLO-COMET Randomized Controlled Trial. *Ann Surg.* 2018;267(2):199-207.
22. Franz AM, Haidegger T, Birkfellner W, Cleary K, Peters TM, Maier-Hein L. Electromagnetic Tracking in Medicine—A Review of Technology, Validation, and Applications. *IEEE Trans Med Imaging.* 2014 Aug;33(8):1702–25.