TU Delft

Delft University of Technology

Document Version
Final published version

Citation (APA)
Zhao, Y., Zhang, Y., Simonetti, O., Han, Y., & Tao, Q. (2024). Lost in Tracking: Uncertainty-Guided Cardiac Cine MRI Segmentation at Right Ventricle Base. In M. G. Linguraru, Q. Dou, A. Feragen, S. Giannarou, B. Glocker, K. Lekadir, & J. A. Schnabel (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2024: 27th International Conference, Marrakesh, Morocco, October 6–10, 2024, Proceedings, Part IX* (pp. 415-424). (Lecture Notes in Computer Science; Vol. 15009 LNCS). Springer. https://doi.org/10.1007/978-3-031-72114-4_40

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright
In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse
Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy
Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

https://www.openaccess.nl/en/you-share-we-take-care

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Lost in Tracking: Uncertainty-Guided Cardiac Cine MRI Segmentation at Right Ventricle Base

Yidong Zhao[1], Yi Zhang[1], Orlando Simonetti[2], Yuchi Han[2], and Qian Tao[1(✉)]

[1] Department of Imaging Physics, Delft University of Technology, Lorentzweg 1, 2628 CJ Delft, The Netherlands
q.tao@tudelft.nl

[2] Cardiovascular Division, The Ohio State University Wexner Medical Center, Columbus, OH, USA

Abstract. Accurate biventricular segmentation of cardiac magnetic resonance (CMR) cine images is essential for the clinical evaluation of heart function. However, compared to left ventricle (LV), right ventricle (RV) segmentation is still more challenging and less reproducible. Degenerate performance frequently occurs at the RV base, where the in-plane anatomical structures are complex (with atria, valve, and aorta) and vary due to the strong interplanar motion. In this work, we propose to address the currently unsolved issues in CMR segmentation, specifically at the RV base, with two strategies: first, we complemented the public resource by reannotating the RV base in the ACDC dataset, with refined delineation of the right ventricle outflow tract (RVOT), under the guidance of an expert cardiologist. Second, we proposed a novel dual encoder U-Net architecture that leverages temporal *incoherence* to inform the segmentation when interplanar motions occur. The inter-planar motion is characterized by *loss-of-tracking*, via Bayesian uncertainty of a motion-tracking model. Our experiments showed that our method significantly improved RV base segmentation taking into account temporal incoherence. Furthermore, we investigated the reproducibility of deep learning-based segmentation and showed that the combination of consistent annotation and loss of tracking could enhance the reproducibility of RV segmentation, potentially facilitating a large number of clinical studies focusing on RV.

Keywords: Cardiac MRI · segmentation · right ventricle · uncertainty

1 Introduction

Automatic segmentation of heart chambers is crucial for quantitatively assessing heart functions from cardiac magnetic resonance (CMR). Besides the left ventricle (LV), there is a growing clinical interest in accurate assessment of the right ventricle (RV), given its significance in heart and lung diseases [10,11,16,22]. In recent years, several challenges have been dedicated to evaluating automatic

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2024
M. G. Linguraru et al. (Eds.): MICCAI 2024, LNCS 15009, pp. 415–424, 2024.
https://doi.org/10.1007/978-3-031-72114-4_40

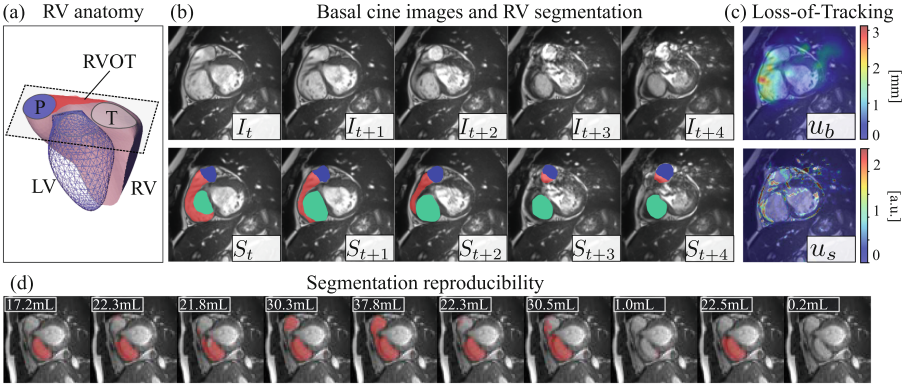


Fig. 1. (a) Anatomy of LV and RV. The basal imaging plane covers the right ventricle outflow tract (RVOT), pulmonary valve (P), and tricuspid valve (T) [20]. (b) A short-axis basal slice contains atria (in green), P (in blue), and RVOT (in red), with complex and varying layouts. (c) Motion tracking has high uncertainty here (u_b and u_s, defined in Sect. 2.1), indicating loss-of-tracking. (d) RV segmentation by 10 Bayesian ensembles exhibits high uncertainty, resulting in a poorly reproducible volume estimation ranging from 0.2 to 37.8 mL. (Color figure online)

CMR segmentation, including the ACDC challenge [4], M&M [5], and M&M-v2 [16] which particularly focuses on RV. Results reveal that the nnUNet families [1,12] have overall superior performance in biventricular segmentation, but RV has degenerated performance compared to that of LV [4,5,16,21]. Degeneration is especially pronounced at the RV base, due to the irregular RV shape, large variability, and complex anatomical context [27].

Segmentation of basal slices is intrinsically challenging, because clinical cine MRI is 2D+t, with each plane imaged in a separate breath-hold, unable to capture the complex 3D+t spatiotemporal motion at the base. Ventricles, atria, and valves all have inter-planar movements [16,26], as shown in Fig. 1(b). This complicates basal segmentation, resulting in quantitative errors of RV assessment [4,16].

More specifically, RV segmentation error stems from the region of the right ventricle outflow tract (RVOT). RVOT is a pathway where the blood exits RV and enters the pulmonary artery [11], spanning from the right side of the tricuspid valve to the pulmonary valve [8] (Fig. 1a). RVOT needs to be included for accurate RV quantification, but it is often overlooked in the annotations of public CMR datasets [4,5,16,26]. Common protocols delineate RV when the full cavity is covered [4,5,16] while RVOT is labeled as RV or background depending on cases or observers. This inconsistency in the annotation can affect the confidence of the neural network [13]. Segmentation models, even when trained on the same dataset [9,13,15], will have high uncertainty on basal slices, resulting in low *reproducibility* of volume quantification. A typical example is shown in Fig. 1 (d), where different segmentation models make varying RV predictions,

indicating low reproducibility of RV volume estimation [9,28,29]. This undermines the reliability of the assessment of RV function.

Traditionally, temporal coherence is leveraged to improve the segmentation performance of CMR, because segmentation tends to be continuous in time and space [2,7,17,18,24,25]. Nilsson *et al.* proposed a spatial-temporal Gated Recurrent Unit (ST-GRU) [17] to promote coherence of segmentation maps. In CMR segmentation, joint motion estimation and segmentation proved to be mutually beneficial [18]. Yan *et al.* proposed a flow-based feature fusion framework [24,25] to integrate temporal coherence. Similarly, Wu *et al.* explicitly encoded the flow as an additional feature for segmentation [7,23]. Bai *et al.* leveraged recurrent neural network (RNN) for cine segmentation with registration-based pseudo-labels [2].

However, the same principle does not apply to the RV base, because of the strong in-plane anatomy change (Fig. 1b). Intuitively, estimating the motion between temporal frames at RV base is ill-posed, i.e. a well-trained motion tracking model will fail to track due to the inter-planar motion, a phenomenon we hereafter call *loss-of-tracking*. Hence, instead of leveraging the temporal coherence, we propose the opposite: we make use of the temporal *incoherence*, which can be identified by motion tracking uncertainty. This uncertainty highlights the inter-planar motion of different structures (Fig. 1c) and is highly informative.

We propose a novel loss-of-tracking-based method to tackle the currently unsolved RV base segmentation in CMR analysis, with the following contributions:

- For a more accurate RV definition, we complemented current public resources by providing refined RV base annotations for the ACDC dataset [4], under the guidance of an expert cardiologist. This complemented community resource can be used to train and evaluate RV segmentation algorithms.
- We propose a Bayesian motion tracking framework for CMR cine, to estimate the tracking uncertainty (loss-of-tracking) which can identify the interplaner cardiac motion in an unsupervised manner.
- We integrate this tracking uncertainty into a Dual-Encoder UNet architecture to enhance segmentation performance in the challenging regions of RV base.
- In addition, we demonstrated that the low reproducibility of deep learning segmentation can stem in part from the annotation inconsistency. Our work improved the RV segmentation reproducibility with refined RV annotation and loss-of-tracking.

2 Method

2.1 Loss-of-Tracking Detection via Registration Uncertainty

We formulate the motion estimation problem as the registration between two temporal frames I_t and $I_{t+\delta t}$ in a cine MRI. A VoxelMorph model [3] V is trained to predict the deformation field $\phi_t = V(I_t, I_{t+\delta t})$. Since the temporal

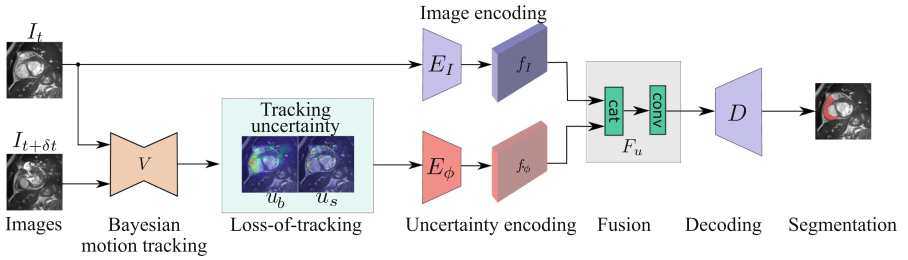


Fig. 2. The Dual-Encoder UNet architecture for CMR segmentation: the upper path encodes the original image, while the lower path encodes the "loss-of-tracking" from I_t to $I_{t+\delta t}$, identified by a Bayesian motion-tracking model. The decoder D predicts the segmentation map from the fused image feature f_I and loss-of-tracking feature f_ϕ.

cine images have similar contrast, we use the simple mean square error (MSE) u_s as the image similarity metric, and regularize the smoothness of the estimated field ϕ_t via a gradient-based regularizer R. The loss function for motion tracking model training is:

$$\begin{aligned}\mathcal{L}_{reg}(\phi_t; I_t, I_{t+\delta t}) &= u_s(\phi_t \circ I_t, I_{t+\delta t}) + \lambda R(\phi_t) \\ &= \|\phi_t \circ I_t - I_{t+\delta t}\|_2^2 + \lambda \|\nabla \phi_t\|_2^2,\end{aligned} \tag{1}$$

where $\circ$ denotes the warping operation. At test time, the registration uncertainty can be inferred in two ways. First, we can evaluate the similarity between the warped image and the target image with u_s:

$$u_s(\phi_t, I_t, I_{t+\delta t}) = \|\phi_t \circ I_t - I_{t+\delta t}\|_2^2, \tag{2}$$

where an elevated level of u_s indicates regions of registration failure. Second, the model uncertainty of the trained registration network V can be estimated from a Bayesian perspective [9,29], which derives the tracking uncertainty from the posterior $p(w|\mathcal{D})$ of network weights w, with $\mathcal{D}$ being the training dataset. The Bayesian uncertainty of motion u_b is estimated via:

$$u_b(\phi_t, I_t, I_{t+\delta t}) \approx \mathsf{std}\left\{V_{w_i}(I_t, I_{t+\delta t})\right\}_{i=1}^{M}, \tag{3}$$

where std is the standard deviation operator, $\{w_i \sim p(w|\mathcal{D})\}_{i=1}^{M}$ is a set of M weights drawn from the posterior distribution, and V_{w_i} denotes the trained model with weight w_i. We draw the posterior samples $\{w_i \sim p(w|\mathcal{D})\}_{i=1}^{M}$ via the Hamiltonian Monte Carlo (HMC) method [6,28,29].

2.2 Uncertainty-Guided Segmentation

Network Architecture. We propose a Dual-Encoder UNet architecture that takes both the image I_t and its motion uncertainty u_ϕ, which highlights the basal areas with interplanar motion, as input. Specifically, we use an encoder E_I

for image encoding and an additional encoder E_ϕ for loss-of-tracking encoding. The image encoder E_I encodes the current temporal frame I_t and outputs the feature $f_I = E_I(I_t)$. The loss-of-tracking encoder E_ϕ learns a representation $f_\phi = E_\phi(u_b, u_s)$ of the estimated motion uncertainty. The feature fusion is then performed by a learnable convolutional layer F_u, and the aggregated feature f_a is expressed by

$$f_a(I_t, u_b, u_s) = F_u\left(\mathsf{cat}\left(f_I, f_\phi\right)\right). \tag{4}$$

Subsequently, f_a is fed into the decoder D for the final segmentation prediction. Skip connections between the encoder and decoder are preserved as in the original U-Net [19]. The overall segmentation model is hence expressed as $\mathcal{S} = D \circ F_u \circ \mathsf{cat} \circ (E_I, E_\phi)$. We show the proposed network architecture in Fig. 2.

Bayesian Segmentation and Reproducibility. We used the same Bayesian HMC principle to generate a range of models [9,28,29]. This Bayesian ensemble of segmentation networks is denoted by $\{\mathcal{S}_{\theta_j}\}_{j=1}^{M}$, parameterized by weight posterior samples $\{\theta_j \sim p(\theta|\mathcal{D})\}_{j=1}^{M}$. Bayesian segmentation is performed via

$$S_t = \frac{1}{M}\sum_{j=1}^{M} \mathcal{S}_{\theta_j}(I_t), \quad \sigma_{\mathrm{v}}(I_t) \approx \mathsf{std}\left\{\mathcal{V} \circ S_{\theta_j}(I_t)\right\}_{j=1}^{M}, \tag{5}$$

where S_t is the segmentation of I_t, $\mathcal{V}$ is the volume calculation operator, and σ_v additionally quantifies the model reproducibility as the standard deviation (SD) of RV volume predicted by the Bayesian ensembles.

With ACDC standard labels, previous work has reported degenerate performance at the RV base, accompanied by high uncertainty [4,5,16,29]. In the literature, degraded performance is often attributed to network generalizability or domain shift. However, a less explored hypothesis is that the training data can also play a role: if the annotation is inconsistent, the prediction from multiple trained models is also inconsistent. To validate the hypothesis and demonstrate the benefit of the refined RV annotation, we evaluated the RV segmentation reproducibility using the standard ACDC and our complemented ACDC annotation.

3 Data and Experiments

Dataset. We evaluated our method on the publicly available ACDC dataset [4]. It consists of cine images of 150 subjects, of which 100 serve as the training set, and the remaining 50 subjects are reserved for testing. Under the guidance of an expert cardiologist, we reannotated RV, including RVOT, on basal slices using the 3D Slicer [14] for all 150 subjects. In total, we manually refined the segmentation map on 240 slices of the original datasets (135 on the training split and 105 on the test split). In the following, we denote the original ACDC dataset as **Original**, and our relabeled dataset as **New**. We will open-source the new RV annotations on GitHub.

Table 1. Segmentation accuracy of RV measured by the Dice coefficient [%] on basal, middle, apical slices, and the full volume. Improvements with statistical significance ($p < 0.05$) using the Wilcoxon signed-rank test are labeled with *.

Methods	Base		Mid		Apex		Full	
	ED	ES	ED	ES	ED	ES	ED	ES
U-Net	90.1 (±9.0)	80.7 (±22.4)	94.5 (±2.7)	90.2 (±5.1)	86.9 (±10.4)	70.9 (±26.1)	92.9 (±2.6)	88.5 (±4.7)
ST-GRU	87.3 (±9.1)	79.2 (±17.1)	91.0 (±4.2)	87.0 (±5.6)	78.1 (±15.2)	63.7 (±21.6)	89.4 (±3.4)	84.9 (±4.9)
Proposed	**91.3*** (±7.3)	**84.0*** (±16.5)	**95.0*** (±2.6)	**91.0*** (±5.0)	**88.1*** (±9.8)	**71.8*** (±26.8)	**93.6** (±2.3)	**89.5*** (±4.5)

Experimental Settings. We used nnUNet [12] as our segmentation backbone, sticking to its original loss function, optimizer, and network plan. Specifically, we trained a 2D nnUNet with all training samples, to be the baseline. To identify loss-of-tracking, we trained a motion tracking network using the VoxelMorph backbone [3], on images with a phase difference of $\delta t = 4$. Furthermore, we trained an ST-GRU network [17], as a contrastive baseline that leverages temporal coherence for refined segmentation. For our Dual-Encoder UNet, we used the same encoder architecture for the loss-of-tracking input as for the image input, following the nnUNet design. We used $M = 10$ HMC samples to build Bayesian ensembles, both for VoxelMorph (motion-tracking) and Dual-Encoder UNet (segmentation).

4 Results

4.1 Segmentation Accuracy

We evaluated the segmentation performance of the three methods on RV, using the new label as ground truth. We divided each short-axis volume into basal, middle, and apical slices and evaluated segmentation accuracy on end-systolic (ES) and end-diastolic (ED) volumes separately. The accuracy measured by Dice coefficients is listed in Table 1 on the three regions and the full volume. The table shows that the vanilla nnU-Net already forms a strong baseline for RV segmentation. Using the coherence-promoting ST-GRU leads to reduced segmentation accuracy. With the loss-of-tracking encoding, our proposed method outperforms the Vanilla U-Net, especially on the basal slices with an improvement of 1.2% and 3.3% at ED and ES volumes, respectively. In comparison, the improvement on the middle and apical slices is marginal in comparison with that on basal slices.

In Fig. 3, we show some qualitative results of the loss-of-tracking detection and the predicted segmentation maps. In case (a), the right ventricle preserves its shape from I_t to $I_{t+\delta t}$, and the detected loss-of-tracking u_s and u_b stays on a relatively low level. On such images, all methods can correctly predict the RV segmentation map. However, the ST-GRU prediction still has a small deviation from the ground truth on ventricular borders (red arrow). We conjecture that ST-GRU suffers from imperfect motion tracking here. In case (b), we show a basal slice with strong interplanar motion on which the RV and valves can hardly be

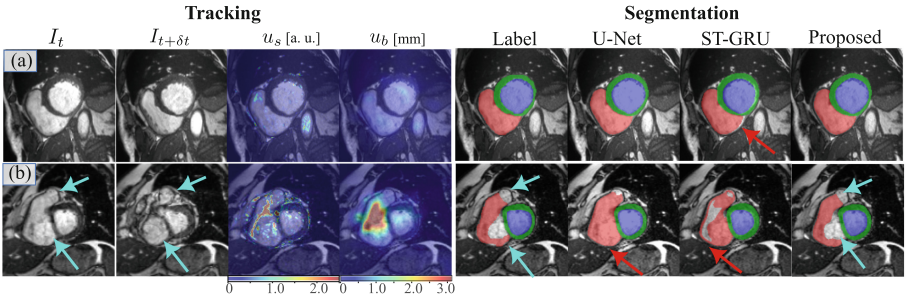


Fig. 3. Qualitative results of tracking uncertainty and segmentation. The left panel shows the tracking uncertainty u_b and u_s between I_t and $I_{t+\delta t}$. The right panel shows the segmentation labels and predictions. In case (b), the atrium and valve (cyan arrows) coexist with RV, and should not be included (c.f. the anatomy in Fig. 1). (Color figure online)

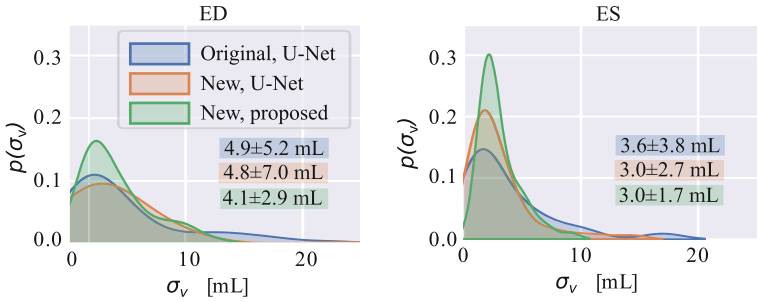


Fig. 4. Distribution of segmentation reproducibility as measured by volume standard deviation σ_v. Statistics (mean ± std) are given in corresponding colors.

distinguished from the single image I_t. The detected loss-of-tracking u_b highlights the area that cannot be tracked from I_t to $I_{t+\delta t}$, mainly in RVOT. The MSE pattern u_s approximately delineates the separation between the valves and the RV. In this slice, the U-Net has difficulty in predicting RV segmentation in a single image I_t, but the proposed method can successfully predict the RV border with loss-of-tracking taken into account. In this case, taking segmentation consistency for granted like ST-GRU can harm the segmentation accuracy.

4.2 Segmentation Reproducibility

In this section, we compare the reproducibility measured by the standard deviation of RV volume from Bayesian ensembles of segmentation networks. To validate the role of annotation, we repeated the experiments on both the Original and New ACDC annotations. Figure 4 shows the distributions of σ_v of the RV base in the ED and ES phases, respectively. The statistics are reported for the testing datasets.

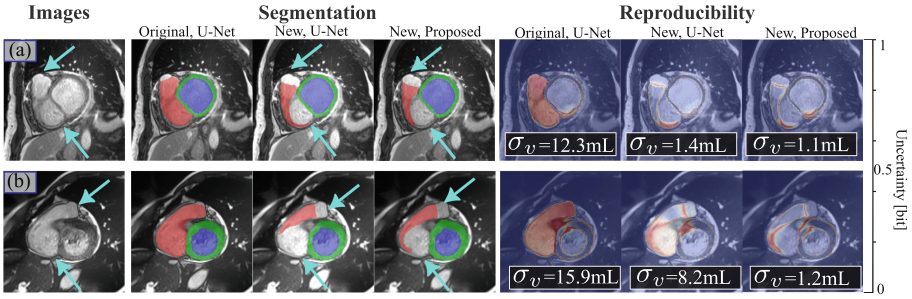


Fig. 5. Examples of RVOT segmentation and reproducibility. High uncertainty indicates strong disagreement among different ensemble models. (a) Models trained with the original labels are uncertain on basal slices with both the valve and atria (cyan arrows) inplane. (b) The reproducibility is largely improved by the new annotations, and further reduced by the proposed method. (Color figure online)

All methods have a volume variance that peaked at a relatively low value (<5 ml), especially at ES. However, we observe that the U-Nets trained with the original annotations have a longer tail than networks trained with the new annotations. The proposed method exhibits the highest reproducibility with a sharp peak. Figure 5(a) is a basal slice that covers the partial atrium and valve (cyan arrows) on the image plane. In this slice, we observe that the networks trained with the original annotations are highly uncertain, resulting in a volume SD of 12.3 ml. In contrast, the networks trained with the new annotations successfully delineate the RVOT and have a reduced SD of 1.1–1.4 mL. In Fig. 5(b), we show a case in which the vanilla U-Nets can have low reproducibility because the RVOT and atrium are not distinguishable, with $\sigma_v = 8.2$ mL. In comparison, our proposed method reduces σ_v to 1.2 mL. The results suggest that consistent annotation and loss-of-tracking can greatly improve reproducibility when segmenting difficult regions like the RV base.

5 Conclusion

Accurate biventricular segmentation of CMR cine images is important for the clinical evaluation of heart function. In this work, we set out to tackle the current challenges of segmenting RV base, for more accurate and reproducible RV assessment. We proposed a novel dual encoder U-Net architecture that leverages temporal incoherence, called *loss-of-tracking*, to identify the interplanar motion at the base that previously deteriorated segmentation. Our experiments showed that *loss-of-tracking* improved the segmentation of the RV base taking into account temporal incoherence. In addition, we complemented the public resource with refined RV base annotation including RVOT. Our work showed that the joint contribution of data and algorithm can lead to improved accuracy and reproducibility for the currently difficult regions of the RV base, potentially leading to more reliable RV assessment for future clinical studies.

Acknowledgments. The authors gratefully acknowledge the TU Delft AI Initiative for financial support.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Arega, T.W., Legrand, F., Bricq, S., Meriaudeau, F.: Using MRI-specific data augmentation to enhance the segmentation of right ventricle in multi-disease, multi-center and multi-view cardiac MRI. In: Puyol Antón, E., et al. (eds.) STACOM 2021. LNCS, vol. 13131, pp. 250–258. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-93722-5_27
2. Bai, W.: Recurrent neural networks for aortic image sequence segmentation with sparse annotations. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018, Part IV. LNCS, vol. 11073, pp. 586–594. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00937-3_67
3. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: VoxelMorph: a learning framework for deformable medical image registration. IEEE Trans. Med. Imaging **38**(8), 1788–1800 (2019)
4. Bernard, O., et al.: Deep learning techniques for automatic MRI cardiac multi-structures segmentation and diagnosis: is the problem solved? IEEE Trans. Med. Imaging **37**(11), 2514–2525 (2018)
5. Campello, V.M., et al.: Multi-centre, multi-vendor and multi-disease cardiac segmentation: the M&Ms challenge. IEEE Trans. Med. Imaging **40**(12), 3543–3554 (2021)
6. Chen, T., Fox, E., Guestrin, C.: Stochastic gradient Hamiltonian Monte Carlo. In: International Conference on Machine Learning, pp. 1683–1691. PMLR (2014)
7. Dong, S., et al.: DeU-net: deformable U-net for 3D cardiac MRI video segmentation. In: Martel, A.L., et al. (eds.) MICCAI 2020, Part IV. LNCS, vol. 12264, pp. 98–107. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59719-1_10
8. Farré, J., Anderson, R.H., Cabrera, J.A., Sánchez-Quintana, D., Rubio, J.M., Benezet-Mazuecos, J.: Cardiac anatomy for catheter mapping and ablation of arrhythmias. Catheter Ablation Cardiac Arrhythmias 74–102 (2011)
9. Gal, Y., Ghahramani, Z.: Dropout as a Bayesian approximation: representing model uncertainty in deep learning [eb/ol]. arXiv preprint arxiv:1506.02142 (2015)
10. Han, Y., et al.: Ranolazine improves right ventricular function in patients with precapillary pulmonary hypertension: results from a double-blind, randomized, placebo-controlled trial. J. Cardiac Fail. **27**(2), 253–257 (2021)
11. Ho, S., Nihoyannopoulos, P.: Anatomy, echocardiography, and normal right ventricular dimensions. Heart **92**(suppl 1), i2–i13 (2006)
12. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. Nat. Methods **18**(2), 203–211 (2021)
13. Kendall, A., Gal, Y.: What uncertainties do we need in Bayesian deep learning for computer vision? Adv. Neural Inf. Process. Syst. **30** (2017)
14. Kikinis, R., Pieper, S.D., Vosburgh, K.G.: 3D slicer: a platform for subject-specific image analysis, visualization, and clinical support. In: Jolesz, F.A. (ed.) Intraoperative Imaging and Image-Guided Therapy, pp. 277–289. Springer, New York (2014). https://doi.org/10.1007/978-1-4614-7657-3_19

15. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. arxiv e-prints, arXiv preprint arXiv:1612.01474, vol. 5 (2016)
16. Martín-Isla, C., et al.: Deep learning segmentation of the right ventricle in cardiac MRI: the M&Ms challenge. IEEE J. Biomed. Health Inform. (2023)
17. Nilsson, D., Sminchisescu, C.: Semantic video segmentation by gated recurrent flow propagation. arXiv preprint arXiv:1612.08871 (2016)
18. Qin, C., et al.: Joint learning of motion estimation and segmentation for cardiac MR image sequences. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018, Part II. LNCS, vol. 11071, pp. 472–480. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00934-2_53
19. Ronneberger, O., Fischer, P., Brox, T.: U-net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
20. Sheehan, F., Redington, A.: The right ventricle: anatomy, physiology and clinical imaging. Heart **94**(11), 1510–1515 (2008)
21. Tao, Q., et al.: Deep learning-based method for fully automatic quantification of left ventricle function from cine MR images: a multivendor, multicenter study. Radiology **290**(1), 81–88 (2019)
22. Wang, L., et al.: Diagnostic and prognostic value of right ventricular eccentricity index in pulmonary artery hypertension. Pulm. Circul. **10**(2), 2045894019899778 (2020)
23. Wu, P., et al.: Cardiac MR image sequence segmentation with temporal motion encoding. In: Bartoli, A., Fusiello, A. (eds.) ECCV 2020, Part I. LNCS, vol. 12535, pp. 298–309. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-66415-2_19
24. Yan, W., Wang, Y., van der Geest, R.J., Tao, Q.: Cine MRI analysis by deep learning of optical flow: adding the temporal dimension. Comput. Biol. Med. **111**, 103356 (2019)
25. Yan, W., Wang, Y., Li, Z., van der Geest, R.J., Tao, Q.: Left ventricle segmentation via optical-flow-net from short-axis cine MRI: preserving the temporal coherence of cardiac motion. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018, Part IV. LNCS, vol. 11073, pp. 613–621. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00937-3_70
26. Yilmaz, P., Wallecan, K., Kristanto, W., Aben, J.P., Moelker, A.: Evaluation of a semi-automatic right ventricle segmentation method on short-axis MR images. J. Digit. Imaging **31**, 670–679 (2018)
27. Zhao, Y., Simonetti, O., Han, Y., Tao, Q.: Artificial intelligence failure in cardiac magnetic resonance image segmentation: An empirical study. J. Cardiovasc. Magn. Reson. **26** (2024)
28. Zhao, Y., et al.: Bayesian uncertainty estimation by Hamiltonian Monte Carlo: applications to cardiac MRI segmentation (2024)
29. Zhao, Y., Yang, C., Schweidtmann, A., Tao, Q.: Efficient Bayesian uncertainty estimation for nnU-net. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) MICCAI 2022. LNCS, vol. 13438, pp. 535–544. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-16452-1_51