

Extraction of volumetric indices from echocardiography: which deep learning solution for clinical use?

Hang Jung Ling, Nathan Painchaud, Pierre-Yves Courand, Pierre-Marc Jodoin, Damien Garcia, Olivier Bernard

► To cite this version:

Hang Jung Ling, Nathan Painchaud, Pierre-Yves Courand, Pierre-Marc Jodoin, Damien Garcia, et al.. Extraction of volumetric indices from echocardiography: which deep learning solution for clinical use?. Functional Imaging and Modeling of the Heart, Jun 2023, Lyon, France. hal-04087268v1

HAL Id: hal-04087268 https://hal.science/hal-04087268v1

Submitted on 3 May 2023 (v1), last revised 8 May 2023 (v2)

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

Extraction of volumetric indices from echocardiography: which deep learning solution for clinical use?

 $\begin{array}{c} {\rm Hang\ Jung\ Ling^{1}[0000-0003-0475-9121]},\ Nathan \\ {\rm Painchaud^{1,2}[0000-0001-8269-5704]},\ Pierre-Yves \\ {\rm Courand^{1,3,4}[0000-0003-3199-7977]},\ Pierre-Marc\ Jodoin^{2}[0000-0002-6038-5753], \\ {\rm Damien\ Garcia^{1}[0000-0002-8552-1475]},\ and\ Olivier\ Bernard^{1}[0000-0003-0752-9946]}\end{array}$

¹ Univ Lyon, INSA-Lyon, Université Claude Bernard Lyon 1, UJM-Saint Etienne, CNRS, Inserm, CREATIS UMR 5220, U1294, F-69621, Lyon, France

² Dept. of Computer Science, University of Sherbrooke, Sherbrooke, QC, Canada

³ Cardiology Dept., Hôpital Croix-Rousse, Hospices Civils de Lyon, Lyon, France

⁴ Cardiology Dept., Hôpital Lyon Sud, Hospices Civils de Lyon, Lyon, France

Abstract. Deep learning-based methods have spearheaded the automatic analysis of echocardiographic images, taking advantage of the publication of multiple open access datasets annotated by experts (CAMUS being one of the largest public databases). However, these models are still considered unreliable by clinicians due to unresolved issues concerning i) the temporal consistency of their predictions, and ii) their ability to generalize across datasets. In this context, we propose a comprehensive comparison between the current best performing methods in medical/echocardiographic image segmentation, with a particular focus on temporal consistency and cross-dataset aspects. We introduce a new private dataset, named CARDINAL, of apical two-chamber and apical four-chamber sequences, with reference segmentation over the full cardiac cycle. We show that the proposed 3D nnU-Net outperforms alternative 2D and recurrent segmentation methods. We also report that the best models trained on CARDINAL, when tested on CAMUS without any fine-tuning, still manage to perform competitively with respect to prior methods. Overall, the experimental results suggest that with sufficient training data, 3D nnU-Net could become the first automated tool to finally meet the standards of an everyday clinical device.

Keywords: Ultrasound \cdot cardiac segmentation \cdot temporal segmentation \cdot deep learning \cdot CNN.

1 Introduction

Echocardiographic imaging has undergone major advances in recent years thanks to artificial intelligence, especially the deep learning (DL) paradigm. In particular, the automated extraction of clinical indices from the segmentation of cardiac structures has been the subject of intense research leading to major breakthroughs. A key component of these advances has been the publication of open access annotated datasets, including CETUS (45 patients, 3D images annotated at End-Diastole - ED and End-Systole - ES) [1], CAMUS (500 patients, 2D images annotated at ED and ES in apical two-chamber - A2C - and apical four-chamber - A4C - views) [5], EchoNet-Dynamic (10,036 patients, 2D sub-sampled images annotated at ED and ES in A4C view) [7], HMC-QU (109 patients, 2D sequences annotated in A4C view) [2] and TED (98 patients from the CAMUS dataset, 2D sequences annotated in A4C view) [8].

These datasets allowed effective and fair comparisons of methods, whether they are generic image segmentation models [5,6] or were specifically designed to process echocardiographic images [11,9]. Thus, the performance of current state-of-the-art (SOTA) methods on the CAMUS dataset confirmed the dominance of the DL-based methods, which finally achieved inter- and intra-observer variability for most of the geometric (Dice score, Hausdorff distance, mean absolute distance - MAD) and clinical metrics (ejection fraction - EF, volumes at ED/ES).

Although these results are extremely promising and represent a crucial step towards the automation of echocardiographic image analysis, they are not sufficient to justify confidence in fully automated methods in a clinical context. Indeed, two crucial challenges on the path to the practical application of these algorithms remain understudied in the field: i) the frame-by-frame temporal consistency of the predictions, and ii) the generalization of the methods across datasets. Based on this observation, we propose the following contributions:

- 1. We study the performance of two generic architectures based on common temporal data processing techniques on 2D echocardiography sequences, and compare them to current SOTAs in the same field;
- 2. We present a new private dataset called CARDINAL (240 patients, 2D sequences annotated in A4C and A2C views), and report the performance impact of training our methods exclusively on CARDINAL and testing on CAMUS.

2 Benchmarked Methods

CAMUS is currently the only dataset where an evaluation platform has been established to effectively compare the performance of segmentation methods⁵. We therefore relied on this dataset to select the methods we retained in this study.

2.1 2D DL Methods

The currently best performing method on the CAMUS dataset exploits the nnU-Net formalism [6]. This model is based on the U-Net architecture and implements several successful DL tricks, such as a patch-wise approach to preserve image resolution, data augmentation during both training and inference to enforce

⁵ https://www.creatis.insa-lyon.fr/Challenge/camus/results.html

generalization and automatic hyperparameter search of the U-Net architecture to increase accuracy. Note that the 2D version uses only one U-Net model [4].

Recently, Sfakianakis *et al.* [9] developed a DL solution called GUDU based on three key aspects. First, they proposed to use data augmentations tailored to ultrasound acquisition, i.e. variation of the contrast between the myocardial tissue and the left ventricular (LV) cavity, random rotation from the origin of the sectorial shape to mimic different probe positioning, and perspective transformations to simulate probe twisting. Inspired by ensemble models, the authors also trained 5 U-Nets with different architectures and averaged their outputs during inference in order to compute the final prediction. Finally, a new loss function was proposed that takes into account the relative position of the cardiac structures with respect to each other. The authors demonstrated the usefulness and complementarity of each contribution in an ablation study.

2.2 2D+t DL Methods

Despite the success of 2D methods in producing accurate segmentations for individual echocardiographic frames, they often fail to maintain temporal consistency between frames [8]. Prior to the publication of HMC-QU and TED, there were no publicly available datasets to train and compare methods that incorporate the temporal dimension in 2D+time echocardiography. This explains why so few papers have focused on this topic.

Due to the lack of 2D+time annotated datasets, Wei *et al.* [11] proposed a method to leverage the limited ED/ES annotated frames and propagate them to unannotated frames. This was achieved by training a 3D U-Net designed to predict both the deformation fields between each pair of consecutive frames and the segmentation masks at each frame of the sequence. The deformation fields are used to propagate the ED/ES reference annotations forward and backward in time through the sequence. The corresponding propagated masks are then used as targets for self-supervised segmentation of the entire sequence. This encourages the model to learn consistent temporal dynamics to find the best match between the predicted segmentation masks and the propagated annotations.

More recently, Smistad *et al.* [10] and Hu *et al.* [3] added convolutional longterm memory blocks to each layer of the encoder of a 2D U-Net (this type of model is hereafter referred to as U-Net LSTM). Thus, instead of processing a single frame, these methods take a series of frames as input and store the extracted features over time to produce the final segmentation of the entire sequence. Results show that such a strategy tends to reduce segmentation shifts from one frame to another.

By their very nature, echocardiographic sequences exhibit regular properties along the time axis. Therefore, it seems logical to consider 2D ultrasound sequences as complete volumes containing coherent 3D shapes and to extract 3D features using 3D convolutional layers to promote temporal consistency. Thus, in this paper, we propose to train a 3D nnU-Net to segment the complete cardiac sequences in a single run. We hypothesize that this model will inherently learn temporal consistency while maintaining a high level of segmentation accuracy.

3 Experimental Setup

3.1 CARDINAL Dataset

Acquisition Protocol: The proposed dataset consists of clinical examinations of 240 patients, acquired at the University Hospital of Lyon (Croix-Rousse Lyon Sud, France) under the regulation of the local ethics committee of the hospital. The complete dataset was acquired with GE ultrasound scanners. For each patient, 2D A4C and A2C view sequences were exported from the EchoPAC analysis software. Each exported sequence corresponds to a set of B-mode images expressed in polar coordinates. The same interpolation procedure as used for the CAMUS dataset was applied to express all sequences in Cartesian coordinates with a single grid resolution of 0.31 mm². Each sequence in the CARDINAL dataset corresponds to a complete cardiac cycle defined as the interval between peaks of maximal LV cavity surface area.

Reference Annotations: To tackle the total number of frames to be annotated, an experienced observer first delineated the different contours using semiautomatic tools to ensure temporal consistency of the segmented shapes. Each corresponding output was then checked/corrected by two other experienced observers. To identify the ED/ES frames in the sequence, the ED frames correspond (by definition) to the beginning and end of each sequence, and the ES frame corresponds to the frame where the LV cavity surface is smallest.

3.2 Implemented DL Methods

For a fair comparison, we implemented the 2D nnU-Net, U-Net LSTM, and 3D nnU-Net described in section 2 using the same Python library called ASCENT ⁶. These models shared the following training hyperparameters: batch size of 2, SGD optimizer with a learning rate of 0.01 coupled with a polynomial decay scheduler, and 1000 training epochs. The 2D and 3D nnU-Net used a patch-wise approach to avoid resizing the images, thus preserving the native image resolution. To train the U-Net LSTM, the input images were resized to 256×256 and 24 consecutive frames were randomly selected and fed to the model to produce the corresponding segmentations. For inference, the sliding window approach with a Gaussian importance map was used. The prediction was given by the average of the *softmax* probabilities of all windows. To improve segmentation accuracy, the final prediction was obtained by averaging the predictions of the original and mirrored images along different axes. More implementation details for each model can be found in table 1.

4 Results

We evaluate the methods described in section 2 using three types of measures to get a complete picture of their performance in terms of segmentation accuracy

⁶ https://github.com/creatis-myriad/ASCENT

Table 1: Details of the implementation of the three methods evaluated in this study. Lowest resolution: Size of the lowest resolution of feature maps in pixels. Optimization scheme: Optimizer + initial learning rate + learning rate scheduler used. Training duration (hours): number of hours required to train each model for 1000 epochs. The configurations shared between models are only shown once in their respective rows.

Configurations	3D nnU-Net	2D nnU-Net	U-Net LSTM			
Patch size (pixels) Batch size	$\begin{array}{cccccccccccccccccccccccccccccccccccc$					
Nb. feature maps Lowest resolution	$10 \times 8 \times 6$	$8 \times 8 \times 24$				
Downsampling scheme Upsampling scheme Normalization scheme	Stride pooling Deconvolution Instance normalization					
Optimization scheme Loss function	SGD + 0.01 + polynomial decay Cross entropy + Dice					
Number of parameters Training duration (hours)	41.3 M 22.8	30.4 M 8	49 M 69.5			

(table 2), extraction of clinical indices (table 3) and temporal consistency (table 4). In each of these tables, we group the methods according to the datasets on which they were trained and tested (CARDINAL is abbreviated as CL and CAMUS is abbreviated as CS) to make it easier to observe the change in performance when generalizing to a new dataset.

4.1 Geometric and Clinical Accuracy

Table 2 shows the segmentation accuracy computed from the CARDINAL and CAMUS datasets for the 5 algorithms described in section 2. The values in bold correspond to the best scores for each metric for a given training/test dataset setup. From the results on the CARDINAL dataset (CL/CL case), we can see that the 3D nnU-Net has the best segmentation scores for all metrics, for both ED and ES. It is also interesting to note that the two temporal consistency methods (3D nnU-Net and U-Net LSTM) produce better results than the 2D nnU-Net method. This can be explained by the fact that the reference segmentation has regular properties along the temporal axis due to the annotation process. Methods that integrate the temporal dimension into their architecture are therefore more likely to produce segmentation results that are closer to the manual references.

It is worth mentioning that methods trained and tested on the same dataset (sections CL/CL and CS/CS in Table 2) get overall better results up to 1.7x for the Hausdorff and MAD metrics. One reason for such an improvement is the larger amount of annotated training images for CARDINAL (18,793 im-

6 H. J. Ling *et al.*

ages from 190 training/validation patients, reference frames for the full cardiac cycle in A2C and A4C views) than for CAMUS (1,800 images from 450 training/validation patients, reference frames at ED and ES in A2C and A4C views).

Table 3 reports the clinical metrics for the 5 methods. As in Table 2, the methods enforcing temporal consistency gets the best results on CARDINAL, especially for the ejection fraction for which temporal consistency is essential (mean correlation score of 0.917). Furthermore, the best models trained on CARDINAL or CAMUS produce similar results for volume estimation (average correlation of 0.978), revealing a limit reached by these approaches, certainly due to the resolution of the imaging systems.

Table 2: LV segmentation accuracy of the benchmarked methods, on different subsets of frames. The columns *All*, *ED* and *ES* indicate results averaged over all frames, only ED frames, and only ES frames, respectively. Since CAMUS only provides annotation for ED/ES frames, results over *all* frames are not available when testing on it. (CL:CARDINAL, CS:CAMUS)

		Dice		Hausdorff (mm)		MAD (mm)				
Methods	Train/test	All	ED	ES	All	ED	ES	All	ED	ES
3D nnU-Net		.969	.968	.960	2.3	2.7	2.5	0.7	0.8	0.7
2D nnU-Net	CL/CL	.957	.961	.942	2.9	3.1	3.1	0.9	1.0	1.1
U-Net LSTM	,	.964	.964	.956	2.5	2.8	2.6	0.8	0.9	0.8
3D nnU-Net		-	.939	.926	-	5.2	4.6	-	1.6	1.5
2D nnU-Net	CL/CS	-	.934	.921	-	4.9	4.6	-	1.8	1.6
U-Net LSTM		-	.925	.903	-	6.0	5.8	-	2.1	2.1
2D nnU-Net		-	.952	.935	-	4.3	4.2	-	1.3	1.3
CLAS	CS/CS	-	.947	.929	-	4.6	4.6	-	1.4	1.4
GUDU		-	.946	.929	-	4.7	4.7	-	1.4	1.4

4.2 Integration of Temporal Consistency

Table 4 allows a better investigation of the temporal performance of the methods by providing additional information on the number/percentage of frames considered temporally inconsistent w.r.t. their neighboring frames. As expected, the methods incorporating temporal persistence produced fewer temporal errors. Looking at the number of sequences with at least one temporally inconsistent frame, the 3D nnU-Net clearly outperforms U-Net LSTM, with only 4 inconsistent sequences over 100 compared to 98 sequences for U-Net LSTM. This result illustrates the greater ability of features computed from 3D convolutional layers to extract relevant spatio-temporal information. The few remaining temporal errors for the 3D nnU-Net are more an indication that the metrics we

Table 3: Clinical metrics of the benchmarked methods. The ED/ES volumes were computed from both the predicted and reference masks using Simpson's biplane method. *Corr.*: Correlation between the ejection fraction (EF) derived from the predicted/reference segmentation. *MAE*: Mean Absolute Error between the EF derived from predicted/reference segmentation.

		${ m EF}$		Volume ED		Volume ES	
Methods	Train/test	Corr.	MAE (%)	Corr.	MAE (ml)	Corr.	MAE (ml)
3D nnU-Net		.913	2.9	.978	3.3	.974	2.7
2D nnU-Net	CL/CL	.850	3.8	.967	4.4	.957	3.2
U-Net LSTM		.922	2.7	.973	3.4	.969	2.8
3D nnU-Net		.869	5.3	.974	9.6	.976	4.9
2D nnU-Net	CL/CS	.810	7.0	.970	12.8	.959	6.2
U-Net LSTM		.822	11.1	.879	15.9	.903	8.2
2D nnU-Net		.857	4.7	.977	5.9	.987	4.0
CLAS	CS/CS	.926	4.0	.958	7.7	.979	4.4
GUDU	·	.897	4.0	.977	6.7	.981	4.6

used are (overly) strict on the temporal smoothness. Indeed, the 3D nnU-Net temporal "inconsistencies" appear invisible to the expert eye. As a qualitative evaluation, fig. 1 illustrates in detail the temporal consistency of each of our own method on one patient from the CARDINAL test set. To complement this, we also provide in the supplementary material examples of temporally consistent and inconsistent segmentation results obtained by the 3D nnU-Net method for the CARDINAL and CAMUS datasets.

4.3 Generalization Across Datasets

The ability to generalize across datasets is crucially important to gauge the capacity of a method to properly analyze data affected by a distributional shift. To this end, the models trained on CARDINAL were also evaluated on the CA-MUS test set without any fine-tuning. The results are reported in the "CL/CS" sections of tables 2 to 4. Among the methods evaluated, 3D nnU-Net is the undisputed best. It even produces competitive geometric and clinical scores compared with SOTA methods trained directly on CAMUS. Thanks to the integration of temporal consistency, the 3D nnU-Net trained on CARDINAL also produces one of the best correlation scores for the EF calculated on the CAMUS dataset, even when compared to SOTA methods trained directly on CAMUS. In view of these results, and considering that the annotation process between the two databases was not identical and was carried out by different experts (which inevitably introduces a bias during the learning phase), the generalization capacity of the 3D nnU-Net model seems remarkable. Table 4: Temporal consistency of the benchmarked methods, as defined in [8]. Nb of seq. w/err.: number of sequences (out of the 100 testing sequences) where at least one frame is temporally inconsistent. % of frames w/err.: percentage of frames that are inconsistent in the sequences with at least one temporally inconsistent frame. *Err. to thresh. ratio*: average ratio between the measure used to identify temporal inconsistencies and the threshold for temporal inconsistencies. A lower value indicates "smoother" temporal segmentations.

Methods	Train/test	Nb of seq. w/ err.	% of frames w/ err.	Err. to thresh. ratio
3D nnU-Net	CL/CL	4	4	.045
2D nnU-Net		100	30	.210
U-Net LSTM		98	13	.110
3D nnU-Net	CL/CS	28	12	.095
2D nnU-Net		85	21	.162
U-Net LSTM		83	16	.114

5 Conclusion

We evaluated the ability of different methods to accurately segment echocardiographic images, with a focus on temporal consistency and cross-dataset generalization. To this end, we introduced a new private database called CARDINAL, with annotations from an expert on the full cardiac cycle for each sequence. The results show that 3D nnU-Net and U-Net LSTM produce the best geometric and clinical scores on the CARDINAL dataset due to the integration of temporal persistence. Regarding the temporal consistency metrics, 3D nnU-Net performed significantly better than U-Net LSTM with only four sequences (instead of 98) out of 100 having at least one image that was temporally inconsistent. As far as cross-dataset generalization is concerned, 3D nnU-Net is also the best performing method. When trained on CARDINAL and tested on CAMUS, it achieved comparable geometric and clinical scores to the best methods both trained and tested on CAMUS. All these results clearly show that 3D nnU-Net is a serious candidate to become the first automated tool to meet the requirements of routine clinical examinations.

Acknowledgment

The authors gratefully acknowledge financial support of the MEGA doctoral school (ED 162), the NSERC Canada Graduate Scholarships-Doctoral Program, the FRQNT Doctoral Scholarships Program, and the LABEX PRIMES (ANR-11-LABX-0063) of Université de Lyon, within the program "Investissements d'Avenir" operated by the French ANR. The authors thank GENCI-IDRIS for providing access to HPC ressources (Grant 2022-[AD010313603]).

9



Fig. 1: Visualization of the temporal consistency of the segmentations on one patient from the CARDINAL test set. (left) Frames sampled between ED and ES, with segmentation masks from our own methods + reference. (right) Curves of the LV and myocardium surfaces w.r.t. frame in the sequence.

References

- Bernard, O., Heyde, B., Alessandrini, M., Barbosa, D., Camarasu-Pop, S., Cervenansky, F., Valette, S., Mirea, O., Galli, E., Geleijnse, M.L., Papachristidis, A., Bosch, J.G., d'Hooge, J.: Challenge on Endocardial Three-dimensional Ultrasound Segmentation (CETUS). In: MICCAI CETUS. pp. 1–8 (Sep 2014)
- Degerli, A., Zabihi, M., Kiranyaz, S., Hamid, T., Mazhar, R., Hamila, R., Gabbouj, M.: Early Detection of Myocardial Infarction in Low-Quality Echocardiography. IEEE Access 9, 34442–34453 (2021)
- 3. Hu, J., Smistad, E., Salte, I.M., Dalen, H., Lovstakken, L.: Exploiting temporal information in echocardiography for improved image segmentation. In: 2022 IEEE International Ultrasonics Symposium (IUS) (2022)
- Isensee, F., Jaeger, P.F., Kohl, S.A.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. Nature Methods 18(2), 203–211 (Feb 2021)
- Leclerc, S., Smistad, E., Pedrosa, J., Østvik, A., Cervenansky, F., Espinosa, F., Espeland, T., Berg, E.A.R., Jodoin, P.M., Grenier, T., Lartizien, C., D'hooge, J., Lovstakken, L., Bernard, O.: Deep Learning for Segmentation Using an Open Large-Scale Dataset in 2D Echocardiography. IEEE Transactions on Medical Imaging 38(9), 2198–2210 (Sep 2019)
- Ling, H.J., Garcia, D., Bernard, O.: Reaching intra-observer variability in 2-D echocardiographic image segmentation with a simple U-Net architecture. In: 2022 IEEE International Ultrasonics Symposium (IUS) (2022), https://hal.science/ hal-03979523
- Ouyang, D., He, B., Ghorbani, A., Yuan, N., Ebinger, J., Langlotz, C.P., Heidenreich, P.A., Harrington, R.A., Liang, D.H., Ashley, E.A., Zou, J.Y.: Video-based AI

10 H. J. Ling et al.

for beat-to-beat assessment of cardiac function. Nature **580**(7802), 252–256 (Apr 2020)

- Painchaud, N., Duchateau, N., Bernard, O., Jodoin, P.M.: Echocardiography Segmentation With Enforced Temporal Consistency. IEEE Transactions on Medical Imaging 41(10), 2867–2878 (Oct 2022)
- Sfakianakis, C., Simantiris, G., Tziritas, G.: GUDU: Geometrically-constrained Ultrasound Data augmentation in U-Net for echocardiography semantic segmentation. Biomedical Signal Processing and Control 82, 104557 (Apr 2023)
- Smistad, E., Salte, I.M., Dalen, H., Lovstakken, L.: Real-time temporal coherent left ventricle segmentation using convolutional LSTMs. In: 2021 IEEE International Ultrasonics Symposium (IUS) (2021)
- Wei, H., Cao, H., Cao, Y., Zhou, Y., Xue, W., Ni, D., Li, S.: Temporal-Consistent Segmentation of Echocardiography with Co-learning from Appearance and Shape. In: MICCAI. pp. 623–632 (2020)