DEEP LEARNING-ENABLED ECHOCARDIOGRAPHIC ASSESSMENT OF BIVENTRICULAR EJECTION FRACTIONS

Ph.D. thesis

Ádám Szijártó

Semmelweis University Doctoral School
Cardiovascular Medicine and Research Division

Supervisors: Attila Kovács, MD, DSc,
Márton Tokodi, MD, PhD

Official reviewers: Bence Ágg, MD, PhD
Márton Áron Goda PhD

Head of the Complex Examination Committee:
Alán Alpár, MD, DSc
Members of the Final Examination Committee:
László Cervenák, MD, PhD
Róbert Langer, MD, PhD

Budapest, 2026

1. Introduction

Convolutional neural networks (CNNs) have revolutionised image processing by extracting spatially invariant features through localised filters, enabling effective classification, regression, and segmentation for many years. However, as data volumes and model sizes increased, the focus of CNNs on local structures limited their ability to capture global dependencies.

Transformers addressed this issue using the self-attention mechanism, which constructs a relationship matrix between all input elements, making them capable of learning global structures of the input. Thanks to this, transformers are highly scalable, with models reaching hundreds of billions in parameter size.

However, these architectures' requirement for enormous datasets necessitated the use of pre-training techniques. Self-supervised learning (SSL) is a pre-training technique that enables models to learn meaningful representations from unlabeled data. It is achieved by defining ML tasks, for which the input and output can be generated from the raw data in an automated fashion. Masked autoencoders are state-of-the-art SSL techniques and have performed exceptionally well in various computer vision tasks.

Echocardiography plays a pivotal role in modern cardiology, providing real-time visualization of cardiac structure and function. A key functional parameter measured using echocardiography is ejection fraction (EF). To this day, two-dimensional (2D) echocardiography (2DE) remains the most commonly used imaging modality for assessing left and right ventricular systolic function. While three-dimensional echocardiography (3DE) offers clear incremental value over 2DE by directly measuring ventricular volumes, it is still

underutilized due to the lack of dedicated training, time constraints, and the complexity of post-processing.

Due to the importance of EF in clinical cardiology and the challenges related to manual annotation and interpretation of echocardiographic images and videos, AI-enabled automated estimation of EF has been the focus of many studies in recent years.

2. Objectives

1. Adapting VideoMAE for analyzing echocardiographic videos

VideoMAE is a state-of-the-art SSL method designed specifically for video transformers. Nevertheless, this method was developed for the general domain of natural videos. Accordingly, we hypothesized that echocardiography-specific modifications could markedly improve the method's performance, and we proposed a novel region-of-interest (ROI)-aware masking strategy that accounts for the characteristic sector-shaped ROI during mask sampling.

2. Developing an end-to-end DL pipeline to predict LVEF and RVEF from A4C echocardiographic videos

We proposed QUEST-EF (QUantification of Echocardiographic STudies – Ejection Fraction), a dual-task DL pipeline trained with SSL incorporating our novel ROI-aware masking strategy for automated quantification of 3DE-derived LVEF and RVEF from A4C echocardiographic videos. Besides internal testing, we sought to evaluate QUEST-EF across a diverse spectrum of acquired and congenital cardiac diseases and various geographic regions.

3. Methods

3.1. Preprocessing of echocardiographic videos

The preprocessing phase comprised three distinct steps: (1) algorithmically cleaning and cropping the frames, (2) confirming that the video is an A4C and determining the orientation of the A4C, and (3) splitting each video into cardiac cycles.

3.1.1 Cleaning and cropping of frames

Echocardiographic videos were exported in a DICOM file format. All the frames were converted to gray-scale images, then cropped using the bounding box of the sector-shaped ROI. Next, a binary mask was created by tracking frame-to-frame changes in each pixel's intensity value.

3.1.2 Confirming view and determining orientation

To confirm the view and determine orientation, two classifiers of the same architecture were trained: Model-V for view classification and Model-O for orientation classification. Both models operated at the frame level, and frames used as input were resized to 256x256 pixels.

3.1.2.1 View classification – Model-V

Model-V aimed to predict the view of 2DE videos. It was designed to classify each frame of the given video into three categories: (1) standard A4C, (2) RV-focused A4C, or (3) other view.

The model was trained, validated, and tested on an internal dataset of 13,864 videos from 1,031 transthoracic echocardiographic examinations performed at the Heart and Vascular Center of Semmelweis University (Budapest, Hungary). The model was also tested externally on 5,650 videos from 150 healthy adults enrolled in the World Alliance of

Societies of Echocardiography (WASE) study and 300 participants of the Budakalász Epidemiology Study.

3.1.2.2 Orientation classification – Model-O

Model-O was trained to classify whether the A4C videos were acquired in Stanford (LV on the right and RV on the left side) or Mayo orientation (LV on the left and RV on the right side).

A dual-center dataset comprising 5,513 A4C videos from 1,418 transthoracic echocardiographic studies was used to train, validate, and test the model internally. The model was also tested externally in 1,200 A4C videos from the same 450 individuals used for the external validation of Model-V.

3.1.3 Splitting videos into cardiac cycles

Given that echocardiographic videos may substantially differ in length and contain different numbers of cardiac cycles, each video was split into segments containing frames from exactly one cardiac cycle. For this task, a DL model – Model-CC – was trained using Ultrasound Video Transformers (UVT), a previously published architecture.

Model-CC was trained, validated, and tested on an internal dataset containing 3,108 A4C videos from 991 transthoracic echocardiographic examinations performed between November 2013 and March 2021 at the Heart and Vascular Center of Semmelweis University (Budapest, Hungary).

Model-CC's predictions were analyzed using signal processing techniques to identify the peak values corresponding to ED and ES frames. ED frames were identified as local maxima in the model's prediction signal.

3.2 Self-supervised pre-training

To achieve high performance in predicting LVEF and RVEF, we first performed self-supervised pre-training using VideoMAE on a large unlabeled echocardiographic dataset before proceeding to the supervised training phase in a smaller labeled dataset. Although such pre-training techniques have achieved outstanding results in various computer vision tasks, the conventional masking strategy of the VideoMAE pipeline may perform sub-optimally in the specific domain of echocardiographic videos. Motivated by this, we proposed a novel ROI-aware masking method that considers the unique characteristics of this domain.

3.2.1 Previous masking strategies

In VideoMAE, the input videos $\mathrm{I} \in \mathbb{R}^{\mathrm{T}\times\mathrm{C}\times\mathrm{H}\times\mathrm{W}}$ are transformed into a token sequence using cube embedding $S = \Phi(I)$. Then, a binary masking map $\mathrm{M_e}$ is generated for the sequence using a custom tube masking strategy to mask certain parts. The unmasked tokens $\boldsymbol{S^u} = \{S_i\}_{i\in\neg M_e}$ are used to generate a latent representation of the input $\mathrm{Z} = \Phi_{\mathrm{enc}}(\mathrm{S^u})$. This encoded representation is then combined with the learnable masking tokens A, and the decoder is tasked with reconstructing the masked tokens in the pixel space as $\hat{I} = \Phi_{dec}(Z, A)$.

3.2.2 Challenges in analyzing echocardiographic videos

In a typical medical ultrasound video, only a fraction of pixels in each frame represent the actual image, while the remaining pixels display textual information and technical markers that, although useful for clinicians during the examination, are considered noise for DL algorithms.

3.2.3 Concept and technical details of ROI-aware masking

To overcome the challenge detailed above, we proposed a novel ROI-aware masking strategy that considers the peculiar shape and location of the ROI in the echocardiographic videos.

A binary segmentation mask $I_s \in R^{HxW}$ could be created around the ROI.

As the ROI is static across all frames of the given video, we could stack the mask along the temporal dimension to map the same binary mask to each input frame.

Cube embedding was then applied to this stacked mask, resulting in a token masking map $\mathrm{M_s}$ that masked the embedded patches outside the segmentation mask.

Using this new map, we defined the encoder ($\widehat{\mathrm{M_e}}$) and decoder masks ($\widehat{\mathrm{M_d}}$) as follows:

$$\widehat{\mathrm{M_e}} = \mathrm{M_e} \cup \neg \mathrm{M_s}$$

$$\widehat{M_d} = M_e \cap M_s$$

where M_e was the tube mask defined for the original VideoMAE. The pre-training pipeline with the ROI-aware masking strategy is demonstrated in Figure 1.

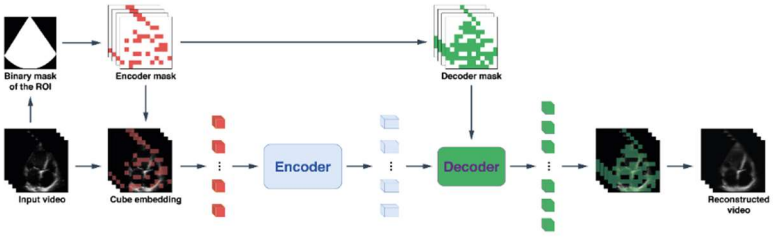

Figure 1 Schematic illustration of self-supervised pre-training with the novel ROI-aware masking strategy.

3.2.4 Evaluating the performance of ROI-aware masking

We evaluated the proposed ROI-aware masking strategy in comparison with simple tube masking and no pre-training, both quantitatively (by evaluating performance in downstream

tasks) and qualitatively (by visually inspecting the reconstructed frames).

3.2.4.1 Pre-training using the novel ROI-aware masking strategy

We pre-trained a VideoMAE using a large dataset of unlabeled echocardiographic videos acquired at the Heart and Vascular Center of Semmelweis University between 2006 and 2023. From this dataset, 29,876 A4C videos were selected with the help of Model-V, corresponding to 15,661 transthoracic echocardiographic studies and encompassing 57,691 cardiac cycles in total.

3.2.4.2 Evaluating the impact of ROI-aware masking on model performance in downstream tasks

In the subsequent supervised training phase, we replaced the decoder of the pre-trained model with either a regressor or a classification head and fine-tuned separate models on the smaller publicly available RVENet dataset for two downstream tasks: (1) predicting LVEF, and (2) assigning a primary diagnosis (healthy, athlete, heart failure with reduced LVEF, aortic valve disease, or mitral valve disease) to each video.

The model's performance was evaluated using mean absolute error (MAE), root mean squared error (RMSE), and the coefficient of determination (R^2) for the regression task (i.e., predicting LVEF) and accuracy for the classification task (i.e., predicting the primary diagnosis).

3.3 Supervised training for predicting ejection fractions

After performing pre-training, we proceeded to the supervised training phase by discarding the decoder of the pre-trained model and attaching a feed-forward regressor network to the encoder. Separate models were trained to predict LVEF and RVEF. Accordingly, the final QUEST-EF pipeline consisted of

two encoders, each initialized with the pre-trained weights, and each connected to a regressor head.

3.3.1 Datasets used for training and evaluation

The LVEF prediction model of QUEST-EF was trained on the publicly available EchoNet-Dynamic dataset comprising 10,030 A4C videos with 2DE-derived labels and a dual-center 3DE dataset comprising 5,341 A4C videos with 3DE-derived labels, whereas the RVEF prediction model was trained only on the latter dataset.

Beyond testing QUEST-EF internally on 15% of the dual-center dataset, its performance was also evaluated in a labeled external test set, which included (1) 238 A4C videos of 238 patients with mixed cardiac diseases from an Italian center of whom 187 had available data regarding heart failure hospitalizations and all-cause mortality during follow-up, (2) 177 A4C videos of 90 adults with congenital heart disease from a British center, (3) 183 A4C videos (with LVEF labels only) of 183 patients with mixed cardiac diseases from a German center, (4) 20 A4C videos (with RVEF labels only) of 20 patients with severe tricuspid regurgitation from another German center, and (5) 4,695 A4C videos of 901 healthy adults enrolled in the WASE study. Last, the associations between the predictions and 10-year all-cause mortality were also investigated in a Hungarian, low-risk, community-based cohort (1,166 unlabeled A4C videos of 1,166 individuals)(Figure 2).

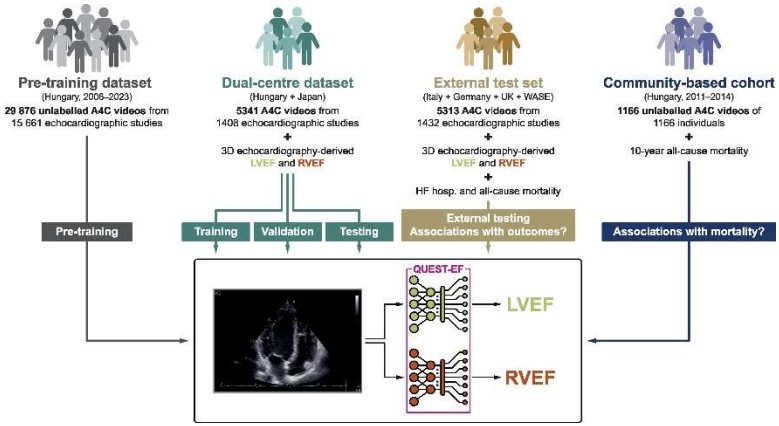

Figure 2 Training and testing datasets used for the development and testing of QUEST-EF.

3.3 Explainability

Besides achieving prime performance in predicting LVEF and RVEF, we also sought to explore how QUEST-EF makes its final predictions using the tools of explainability.

Due to its transformer architecture, the performance of QUEST-EF could be assessed on subsets of cube-embedded patches corresponding to distinct cardiac structures in the video. Based on the change in performance compared with that obtained using the entire video, the relative importance of each structure could be determined: the smaller the drop in performance, the more important the structure. The following formula was used to calculate an importance score for each structure:

$$Importance\ score = 1 - \frac{MAE_S - MAE_O}{MAE_O}$$

where MAES is the MAE when only cube-embedded patches corresponding to the given cardiac structures were used, and

MAEO is the MAE when all cube-embedded patches (i.e., the entire video) were used. In our experiments, we utilized a previously published segmentation model to automatically identify five cardiac structures – LV blood pool, LV myocardium (LV_{MYO}), left atrium (LA), RV, and right atrium (RA) – in the A4C videos

4. Results

4.1 Performance of the models used in preprocessing

4.1.1 View classification – Model-V

In the three-class classification task, Model-V achieved balanced accuracies of 0.931 (95% CI: 0.915-0.947) and 0.835 (95% CI: 0.816-0.853) during internal and external testing (Table 2), whereas in binary classification, it disciminated A4C from non-A4C with balanced accuracies of 0.991 (95% CI: 0.987-0.994) and 0.895 (95% CI: 0.884-0.907) in the internal and external test sets, respectively.

4.1.2 Orientation classification – Model-O

In classifying the orientation of the echocardiographic videos, Model-O achieved balanced accuracies of 0.995 (95% CI: 0.990-1.000) and 0.960 (95% CI: 0.948-0.971) in the internal and external test sets

4.1.3 Splitting videos into cardiac cycles – Model-CC

Model-CC identified ED frames with MAEs of 4.041 (95% CI: 3.878-4.210) and 3.117 (95% CI: 3.044-3.193) frames and ES frames with MAEs of 2.799 (95% CI: 2.708-2.894) frames and 4.827 (95% CI: 4.738-4.917) frames in the internal and external test sets, respectively. More importantly, the model was unable to identify 7.3 and 4.2% of the ED frames and 1.2 and 3.6% of the ES frames in the internal and external test sets, respectively (Table 5). The actual heart rate could be estimated with MAEs of 4.106 (95% CI: 3.586-4.698) and 1.911 (95% CI:

1.665-2.187) beats per minute based on the predicted ED frames and with MAEs of 2.460 (95% CI: 2.190-2.759) and 1.757 (95% CI: 1.517-2.030) beats per minute based on the predicted ES frames in the internal and external test sets, respectively.

4.2 ROI-aware masking

4.2.1 Improvement in video reconstruction

Training with the proposed ROI-aware masking strategy enabled the model to capture the more granular structure of the heart seen in the original frame, whereas the model trained with tube masking focused on correctly identifying and reconstructing the border of the ROI and learned only the local black-to-white gradients. These findings confirmed that by allowing tokens consisting of irrelevant pixels to be fed into the encoder and reconstructed by the decoder, a considerable amount of resources is wasted on learning the representation of such tokens.

4.2.2 Impact on performance in downstream tasks

For LVEF prediction, the model without pre-training achieved an MAE of 6.212 (95% CI: 5.913-6.503) percentage points. Pre-training with the traditional tube masking strategy provided only a marginal improvement (MAE: 5.933 [95% CI: 5.659-6.234]), whereas pre-training with the ROI-aware masking strategy outperformed both, yielding an MAE of 4.373 (95% CI: 4.196-4.570). For predicting the primary diagnosis, the tube masking strategy performed worse than the model with no pre-training (accuracy: 0.422 [95% CI: 0.392-0.451] vs. 0.581 [95% CI: 0.551-0.610]), while the ROI-aware masking strategy substantially improved performance, reaching an accuracy of 0.800 (95% CI: 0.776-0.824).

4.3 Performance of QUEST-EF

4.3.1 Performance in predicting LVEF and RVEF

In predicting LVEF, QUEST-EF achieved study-level MAEs of 4.564 (95% CI: 4.105-5.087) and 4.601 (95% CI: 4.409-4.797) percentage points in the internal and external test sets, respectively. Bland-Altman analysis showed a study-level bias of 0.130 (95% CI: -0.648 to 0.908) percentage points with upper and lower LOAs of 11.739 (95% CI: 10.407-13.070) and -11.478 (95% CI: -12.810 to -10.146) percentage points in the internal, and a study-level bias of -1.471 (95% CI: -1.772 to -1.170) percentage points with upper and lower LOAs of 9.832 (95% CI: 9.317-10.347) and -12.774 (95% CI: -13.289 to -12.260) percentage points in the external test set.

In predicting RVEF, QUEST-EF achieved MAEs of 4.815 (95% CI: 4.329-5.409) and 5.417 (95% CI: 5.183-5.653) on the study level in the internal and external test sets, respectively. Bland-Altman analysis showed a study-level bias of -0.211 (95% CI: -1.047 to 0.625) percentage points with upper and lower LOAs of 12.090 (95% CI: 10.660-13.521) and -12.513 (95% CI: -13.943 to -11.082) percentage points in the internal, and a study-level bias of 0.157 (95% CI: -0.223 to 0.537) percentage points with upper and lower LOAs of 13.568 (95% CI: 12.919-14.217) and -13.254 (95% CI: -13.903 to -12.605) percentage points in the external test set.

4.3.2 Performance in predicting LV and RV dysfunction

In the internal test set, QUEST-EF identified LV and RV systolic dysfunction (i.e. LVEF <50% and RVEF <45%) with AUCs of 0.944 (95% CI: 0.911-0.969) and 0.879 (95% CI: 0.821-0.928), respectively (Table 10 and Figure 13), whereas, in the labeled external test set, it achieved AUCs of 0.940 (95% CI: 0.919-0.958) and 0.791 (95% CI: 0.736-0.841) in these tasks.

4.3.3 Associations with outcomes

Among the 187 patients with available outcome data [28 (15.0%) died or were hospitalized due to heart failure during the

median follow-up duration of 1.1 (interquartile range: 0.5-1.6) years], the QUEST-EF-predicted EF values were associated with the composite endpoint of heart failure hospitalization or all-cause death [LVEF – adjusted HR (aHR): 0.945 (95% CI: 0.913-0.979), p=0.002; RVEF – aHR: 0.927 (95% CI: 0.877-0.979), p=0.006], independent of age and sex. In the community-based cohort [10-year all-cause mortality rate: 131/1,166 (11.2%)], the predictions were also associated with 10-year all-cause death [LVEF – aHR: 0.947 (95% CI: 0.924-0.970), p<0.001; RVEF – aHR, 0.877 (95% CI: 0.845-0.909), p<0.001], independent of the Framingham Risk Score and LV filling pressures estimated by E/e′ ratio.

4.3.4 Explainability

Based on our custom segmentation-based method, the LV myocardium and blood pool were the most important segments for predicting LVEF, whereas the remaining three segments (LA, RV, and RA) contributed marginally. For RVEF prediction, the RA was the most important segment, followed by the LV myocardium, RV, LV blood pool, and LA in decreasing order of importance; however, the differences in importance scores across segments were less pronounced than in the LVEF prediction task.

5. Conclusions

1. Based on our first study, in which we proposed a novel ROI-aware masking strategy that accounts for the characteristic sector-shaped ROI in echocardiographic videos during SSL, we drew the following conclusions:

    1.1 Pre-training with ROI-aware masking yielded superior performance compared with conventional tube masking and no pre-training in self-supervised image reconstruction and two downstream tasks, namely

estimating LVEF and predicting primary diagnosis from an A4C echocardiographic video.

1.2 The advantage of the ROI-aware masking strategy was even more pronounced as the amount of labeled training data decreased, underscoring its effectiveness in data-limited settings.

2. In our second study, we developed and validated QUEST-EF, a dual-task DL pipeline trained with SSL incorporating our novel ROI-aware masking strategy for automated quantification of 3DE-derived LVEF and RVEF from A4C echocardiographic videos. The key findings of this study can be summarized as follows:

   2.1 QUEST-EF was designed as a fully automated, segmentation-free end-to-end DL pipeline featuring a multi-stage preprocessing module that algorithmically cleans and crops frames and integrates three validated DL models for view and orientation verification and splitting videos into cardiac cycles.

   2.2 QUEST-EF exhibited robust performance in predicting LVEF and RVEF and accurately detected LV and RV dysfunction during both internal and external validation across diverse patient cohorts.

   2.3 The prognostic value of QUEST-EF was also confirmed, as lower predicted LVEF and RVEF values were independently associated with a significantly higher risk of adverse outcomes in two separate cohorts.

2.4 Explainability analysis revealed that the LV myocardium and blood pool were the most important segments for predicting LVEF. In contrast, the RA contributed the most to RVEF prediction, and less pronounced differences were observed in importance scores across cardiac structures than in the LVEF task.

6. Bibliography of the candidate's publications

Publications related to the thesis:

1. **Ádám Szijártó**, Béla Merkely, Attila Kovács, Márton Tokodi on behalf of, the QUEST-EF Investigators
   Deep Learning-Enabled Echocardiographic Assessment of Biventricular Ejection Fractions: The Dual-Task QUEST-EF Model
   European Heart Journal-Cardiovascular Imaging 26 : 8 pp. 1402-1405. , 4 p. (2025) DOI: 10.1093/ehjci/jeaf147
   IF: **6.6**
2. **Ádám Szijártó**, Bálint Magyar, Thomas Á. Szeier, Máté Tolvaj, Alexandra Fábián, Bálint K. Lakatos, Zsuzsanna Ladányi, Zsolt Bagyura, Béla Merkely, Attila Kovács, Márton Tokodi
   Masked Autoencoders for Medical Ultrasound Videos Using ROI-Aware Masking
   Lecture Notes in Computer Science 15186 pp. 167-176. , 10 p. (2025) 65 : 2 pp. 5-16. , 12 p. (2020) DOI: 10.1007/978-3-031-73647-6_16

Publications not related to the thesis:

1. **Ádám Szijártó**; Péter Lehotay-Kéry; Attila Kiss
   EXPERIMENTAL STUDY OF SOME PROPERTIES OF KNOWLEDGE DISTILLATION.
   Studia Universitatis Babes-Bolyai, Informatica 65 : 2 pp. 5-16. , 12 p. (2020) DOI: 10.24193/subbi.2020.2.01
2. Máté Tolvaj, Alexandra Fábián, Márton Tokodi, Bálint Lakatos, Alexandra Assabiny, Zsuzsanna Ladányi, Kai Shiida, Andrea Ferencz, Walter Schwertner, Boglárka Veres, Annamária Kosztin, **Ádám Szijártó**, Balázs Sax, Béla Merkely, Attila Kovács
   There is more than just longitudinal strain: Prognostic significance of biventricular circumferential mechanics
   Frontiers in Cardiovascular Medicine 10 Paper: 1082725 , 9 p. (2023) DOI: 10.3389/fcvm.2023.1082725
   IF: **2.8**
3. Márton Tokodi, Bálint Magyar, Ádám Szijártó, Bálint Károly Lakatos, Attila Kovács
   Reply: Can Deep Learning Improve 2D Echocardiographic RV Assessment? First Important Steps

JACC-Cardiovascular Imaging 16 : 12 pp. 1636-1636. , 1 p. (2023) DOI: 10.1016/j.jcmg.2023.10.015

4. **Ádám Szijártó**, Alexandra Fábián, Bálint Károly Lakatos, Máté Tolvaj, Béla Merkely, Attila Kovács, and Márton Tokod
   A machine learning framework for performing binary classification on tabular biomedical data
   Imaging 15 : 1 pp. 1-6. , 6 p. (2023), DOI: 10.1556/1647.2023.00109
   IF: **0.7**
5. **Ádám Szijártó,** Ellák Somfai, András Lőrincz
   Design of a Machine Learning System to Predict the Thickness of a Melanoma Lesion in a Non-Invasive Way from Dermoscopic Images
   Healthcare Informatics Research 29 : 2 pp. 112-119. , 8 p. (2023), DOI: 10.4258/hir.2023.29.2.112
   IF: **2.3**
6. Richard Walter Schwertner, Márton Tokodi, Boglárka Veres, Anett Behon, Eperke Dóra Merkel, Richárd Masszi, Luca Kuthi, **Ádám Szijártó**, Attila Kovács, István Osztheimer, Endre Zima, László Gellér, Máté Vámos, László Sághy, Béla Merkely, Annamária Kosztin, Dávid Becker
   Phenogrouping and risk stratification of patients undergoing cardiac resynchronization therapy upgrade using topological data analysis
   Scientific Reports 13 : 1 Paper: 20594 , 13 p. (2023) DOI: 10.1038/s41598-023-47092-x
   IF: **3.8**
7. Boglárka Veres, Walter Richard Schwertner, Márton Tokodi, **Ádám Szijártó**, Attila Kovács, Eperke Dóra Merkel, Anett Behon, Luca Kuthi, Richárd Masszi, László Gellér, Endre Zima, Levente Molnár, István Osztheimer, Dávid Becker, Annamária Kosztin, Béla Merkely
   Topological data analysis to identify cardiac resynchronization therapy patients exhibiting benefit from an implantable cardioverter-defibrillator
   Clinical Research in Cardiology 113 : 10 pp. 1430-1442. , 13 p. (2024) DOI: 10.1007/s00392-023-02281-6
   IF: **3.7**
8. **Ádám Szijártó**, Alina Nicoara, Mihai Podgoreanu, Márton Tokodi, Alexandra Fábián, Béla Merkely, András Sárkány, Zoltán Tősér, Sergio Caravita, Claudia Baratto, Michele Tomaselli,

Denisa Muraru, Luigi Paolo Badano, Bálint Lakatos, Attila Kovács
Artificial intelligence-enabled reconstruction of the right ventricular pressure curve using the peak pressure value: a proof-of-concept study
European Heart Journal - Imaging Methods and Practice 2 : 4 Paper: qyae099 , 7 p. (2024) DOI: 10.1093/ehjimp/qyae099

9. Bálint K. Lakatos, Zvonimir Rako, **Ádám Szijártó**, Bruno R. Brito da Rocha, Manuel J. Richter, Alexandra Fábián, Henning Gall, Hossein A. Ghofrani, Nils Kremer, Werner Seeger, Daniel Zedler, Selin Yildiz, Athiththan Yogeswaran, Béla Merkely, Khodr Tello, Attila Kovács
Right ventricular pressure-strain relationship-derived myocardial work reflects contractility: Validation with invasive pressure-volume analysis
The Journal of Heart and Lung Transplantation 43 : 7 pp. 1183-1187. , 5 p. (2024) DOI: 10.1016/j.healun.2024.03.007
IF: **6**
10. FJolla Zhubi Bakija, Máté Tolvaj, **Ádám Szijártó**, Márton Tokodi, Andrea Ferencz, Bálint Károly Lakatos, Zsuzsanna Ladányi, Loretta Kiss, Zsolt Szelid, Pál Soós, Béla Merkely, Zsolt Bagyura, Attila Kovács & Alexandra Fábián
Long-term prognostic value of myocardial work analysis across obesity stages: insights from a community-based study
International Journal of Obesity , 10 p. (2025) DOI: 10.1038/s41366-025-01863-w
IF: **3.8**
11. Máté Tolvaj, Fjolla Zhubi Bakija, Alexandra Fábián, Andrea Ferencz, Bálint Lakatos, Zsuzsanna Ladányi, Ádám Szijártó, Borbála Edvi, Loretta Kiss, Zsolt Szelid, Pál Soós, Béla Merkely, Zsolt Bagyura, Márton Tokodi, Attila Kovács
Integrating Left Atrial Reservoir Strain Into the First-Line Assessment of Diastolic Function: Prognostic Implications in a Community-Based Cohort With Normal Left Ventricular Systolic Function
Journal of the American Society of Echocardiography 38 : 7 pp. 570-582. , 13 p. (2025) DOI: 10.1016/j.echo.2025.03.012
IF: **6**
12. Márton Tokodi, **Ádám Szijártó**
From Promise to Practice: Reducing Research Waste in Deep Learning Model Development for Cardiovascular Imaging

JACC-Cardiovascular Imaging 18 : 7 pp. 765-767. , 3 p. (2025) DOI: 10.1016/j.jcmg.2025.05.003

13. Timea Teszak, Timea Barcziova, Csaba Bödör, Lajos Hegyi, Luca Levay, Beata Nagy, Attila Fintha, **Adam Szijarto**, Attila Kovacs, Bela Merkely, Balazs Sax
Donor-Derived Cell-Free DNA Versus Left Ventricular Longitudinal Strain and Strain-Derived Myocardial Work Indices for Identification of Heart Transplant Injury
Biomedicines 13 : 4 Paper: 841 , 12 p. (2025) DOI: 10.3390/biomedicines13040841
IF: **3.9**
14. **Ádám Szijártó**, Alexandra Fábián, Bálint Károly Lakatos, Máté Tolvaj, Béla Merkely, Attila Kovács, and Márton Tokodi
Addendum to: A machine learning framework for performing binary classification on tabular biomedical data
Imaging 17 : 1 pp. 83-83. , 1 p. (2025) DOI: 10.1556/1647.2023.11109
15. Marius Keller, Alexandra Fábián, Andrea Bandini, **Ádám Szijártó**, Zoltán Tősér, Béla Merkely, Tim Heller, Marcia-Marleen Dürr, Peter Rosenberger, Attila Kovács & Harry Magunia
Impact of the right ventricular mechanical pattern assessed by three-dimensional echocardiography on adverse outcomes following cardiac surgery
Scientific Reports 15 : 1 Paper: 5623 , 13 p. (2025) DOI: 10.1038/s41598-025-89122-w
IF: **3.9**

∑IF: 43.5