

Ultrasound Video Transformers for Cardiac Ejection Fraction Estimation

Hadrien Reynaud¹, Athanasios Vlontzos¹, Benjamin Hou¹, Arian Beqiri²,
Paul Leeson^{2,4}, and Bernhard Kainz^{1,3}

¹Department of Computing, Imperial College London, London, UK

²Ultromics Ltd, Oxford, UK ³Friedrich–Alexander University Erlangen–Nürnberg, DE

⁴John Radcliffe Hospital, Cardiovascular Clinical Research Facility, Oxford, UK
hadrien.reynaud19@imperial.ac.uk

Abstract. Cardiac ultrasound imaging is used to diagnose various heart diseases. Common analysis pipelines involve manual processing of the video frames by expert clinicians. This suffers from intra- and inter-observer variability. We propose a novel approach to ultrasound video analysis using a transformer architecture based on a Residual Auto-Encoder Network and a BERT model adapted for token classification. This enables videos of any length to be processed. We apply our model to the task of End-Systolic (ES) and End-Diastolic (ED) frame detection and the automated computation of the left ventricular ejection fraction. We achieve an average frame distance of 3.36 frames for the ES and 7.17 frames for the ED on videos of arbitrary length. Our end-to-end learnable approach can estimate the ejection fraction with a MAE of 5.95 and R^2 of 0.52 in 0.15s per video, showing that segmentation is not the only way to predict ejection fraction. Code and models are available at <https://github.com/HReynaud/UVT>.

Keywords: Transformers · Cardiac · Ultrasound

1 Introduction

Measurement of Left Ventricular Ejection Fraction (LVEF) is a commonly used tool in clinical practice to aid diagnosis of patients with heart disease and to assess options for life-prolonging therapies. The LVEF is the ratio between the stroke volume, which is the difference between End-Diastolic (ED) and End-Systolic (ES) volumes, and the ED volume of the left ventricle.

In primary and secondary care, 2D Ultrasound (US) video acquisition of the standard apical four-chamber view is used to approximate LVEF from manually delineated luminal areas in the left ventricle in one chosen ED and one ES frame. The biplane method of disks, which requires both 2-chamber and 4-chamber views, is the currently recommended two-dimensional method to assess LVEF [6], with its limitations known in literature [20]. The laborious nature of the data processing and substantial inter- and intra-operator variability makes this approach infeasible for high throughput studies, *e.g.*, population health screening

applications. Current clinical practice already neglects the recommendation to repeat this process on at least five heartbeats [15] and commonly only a single measurement is acquired to mitigate clinicians’ workload. Our objective is to estimate LVEF accurately from US video sequences of arbitrary length, containing arbitrarily many heart-beats, and to localize relevant ES and ED frames.

This problem has been recognized in the medical image analysis community. Initially, techniques to automatically segment the left ventricle have been proposed [21,22] to support the manual LVEF estimation process. Recently, robust step-by-step processing pipelines have been proposed to identify relevant frames, segment them, calculate the LVEF and predict the risk for cardiac disease [17].

To the best of our knowledge, all existing techniques for automatically processing US video sequences follow the paradigm of discrete frame processing with limited temporal support. US videos, however, can be of arbitrary length and the cardiac cycle varies in length. Frame-by-frame [2,1] processing neglects the information encoded in the change over time, or requires heuristic frame sampling methods to form a stack of selected frames to enable spatio-temporal support in deep convolutional networks [17].

In this paper, we postulate that processing US videos should be considered more similar to processing language. Thus, we seek a model that can interpret sequences across their entire temporal length, being able to reason through comparison between heartbeats and to solve tasks on the entirety of the acquired data without introducing too many difficult to generalize heuristics. Following this idea, we propose a new US video ES/ED recognition and LVEF prediction network. We use an extended transformer architecture to regress simultaneously the LVEF and the indices of the ES and ED frames.

Our **contribution** is two-fold. (a) We evaluate a new paradigm for US video analysis based on sequence-to-sequence transformer models on a large database of more than 10,000 scans [17] for the clinically relevant task of LVEF estimation. (b) We introduce a modified transformer architecture that is able to process image sequences of variable length.

Related Works: Early automatic detection algorithms embed videos on a manifold and perform low-dimensional clustering [7]. Other methods use convolutional neural networks (CNN) to locate and monitor the center of the left ventricle (LV) of the heart to extract the necessary information to categorize frames as ES or ED [26]. In recent work, [11] introduces a CNN followed by an RNN that utilizes Doppler US videos to detect ED.

For the estimation of both the ES/ED frame indices and the LVEF, [14] assess three algorithms performing LVEF estimation in a multi-domain imaging scenario. Deep convolutional networks have been extensively used in various steps. In [21], the LV is segmented from standard plane US images with the help of a U-Net [19]. ES, ED frames are heuristically identified. In [5] the authors extract spatio-temporal features directly from the input video, classifying whether the frames belong to a systole or a diastole, identifying ES and ED as the switching points between the two states. The authors of [22] leverage deep learning techniques throughout their proposed method and directly learn to identify the

ES and ED frames which they subsequently feed into a U-Net to segment the LV followed by a metrics-based estimation of the LVEF. Segmentation has been the most explored method for the analysis of cardiac functions with neural networks [18,24]. [17] propose an end-to-end method that leverages spatio-temporal convolutions to estimate the LVEF. The extracted features in conjunction with semantic segmentation of the LV enable a beat-by-beat assessment of the input. Direct estimation of the Systolic Volume (SV) has been explored on cine MRI [13] and cine US [4]. Trained to leverage temporal information, their network directly estimates the SV, and the ES/ED frames from single-beat cardiac videos. However, these methods require a fixed number of input frames containing a single cardiac cycle and use LSTMs for temporal feature extraction, which are well known for forgetting initial elements when the sequences become longer [23]. Our approach is more closely related to the latter three approaches, which we will be comparing against. We explore the strength of transformers applied to video data like [16,12,8] have done recently, and build our own architecture.

2 Method

Just as a piece of text in natural language can have variable length and is required in its entirety to be understood, a US video can be of arbitrary duration and is needed in full for accurate interpretation and reasoning. To this end, we propose to use a new transformer model to interpret US videos. Our end-to-end trainable method comprises three distinct modules: (a) An encoder tasked with dimensionality reduction, (b) a Bidirectional Encoder Representations from Transformers (BERT)-based [3] module providing spatio-temporal reasoning capabilities, and (c) two regressors, one labeling ES and ED frames and the other estimating the LVEF. An overview of this model is shown in Fig. 1.

Dimensionality Reduction: Using full size image frames directly as inputs to a transformer is infeasible, as it would require an excessive amount of computational power. To make the problem computationally tractable and to allow BERT to understand the spatial structure of the images, we make use of the encoding part of a ResNetAE [10] to distil the US frames into a smaller dimensional embedding.

The network uses multi-scale residual blocks [9] in both the encoder and the decoder in order to incorporate information across dimensions. We optimize the hyper-parameters associated with the AE architecture, such as the depth of the encoder, the size of the latent space, by first performing a reconstruction task on the US dataset. The encoder setup of the optimal architecture is then used as the encoding module of our method. Each frame from the clip is distilled by the encoder into a $1024D$ vector. The resulting embeddings are stacked together to produce the initial embedding of the clip, characterized by a shape $Batch \times N_{frames} \times 1024$. We use these latent embeddings as the input of the transformer and the combined architecture is trained as a whole end-to-end. Cascading the encoder with the transformer model presents an important benefit as the weights of the encoder are learned to optimize the output from the transformer. This

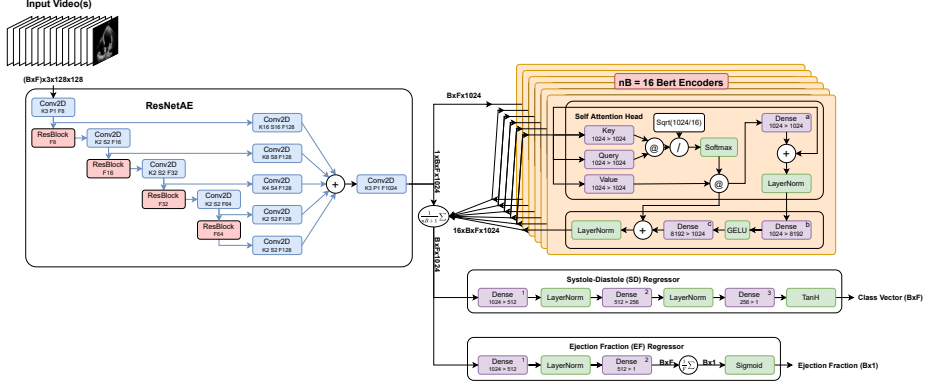


Fig. 1: Overview of the proposed architecture; Left to right: Clips are reduced in dimensions through the ResAE, spatio-temporal information is extracted through the BERT and then passed to each regression branch. The Systole-Diastole (SD) Regressor predicts the ES/ED indices while the EF Regressor predicts the LVEF. The @ operation is the dot product.

is in contrast to using a pretrained dimensionality reduction network where the encoder would be task-agnostic.

Spatio-Temporal Reasoning: In order to analyze videos of arbitrary length we use a BERT encoder [25] to which we attach a regression network to build a Named Entity Recognition (NER) model for video. This acts as a spatio-temporal information extractor. The extracted embeddings E from the ResNetAE encoding step, are used as inputs for the BERT encoder. As in [3], the k^{th} encoder can be characterized as $B_k(E) = \text{LayerNorm}(D_{k,c}(\text{GELU}(D_{k,b}(A_k(E)))) + S_k(E)$, where,

$$S_k(E) = \text{Softmax} \left(\frac{Q_k(E)K_k^T(E)}{\sqrt{\frac{nD}{nB}}} \right) V_k(E),$$

$$A_k(E) = \text{LayerNorm}(D_{k,a}(S_k(E)) + E).$$

$S_k(E)$ and $A_k(E)$ describe the Self-Attention block and Attention block respectively. The Query, Key and Value are parameterized as linear layers Q_k , K_k , V_k . While, $D_{k,\{a,b,c\}}$ are the intermediate linear (dense) layers in the BERT Encoder. We keep the key parameters nB , the number of BERT encoders, and nD the dimensionality of the embeddings, similar to [3], setting them to $nB = 16$ and $nD = 1024$ respectively. During training, dropout layers act as regularizers with a drop probability of 0.1.

Regressing the Output: Following the spatio-temporal information extraction, the resulting features from the k BERT encoders $B_k(E)$ are averaged to-

gether with the ResNetAE output to

$$M(E) = \frac{1}{nB + 1} \left(E + \sum_k^{nB} B_k(E) \right), \quad (1)$$

and passed through two regressors tasked with predicting ES, ED frame indices, $R_{SD}(M(E))$, and the LVEF, $R_{EF}(M(E))$, respectively. We define the output of $R_{SD}(M(E))$ as the output of three linear layers interleaved with layer-normalization, with a *tanh* activation at the end. LVEF is characterized by:

$$R_{EF}(M(E)) = \text{Sigmoid} \left(\frac{1}{nF} \sum_f^{nF} (\text{DEF}_2(\text{LayerNorm}(\text{DEF}_1(M(E)))))) \right), \quad (2)$$

with nF the number of input frames. Thus, LVEF is estimated through a regression network which reduces the embedding dimension to 1 for each input frame. We then take the average of the predictions over the frames to output a single LVEF prediction per video. The LVEF prediction training is done with a combination of losses and regularization to address imbalance in the distribution of LVEF in the training set. We use both Mean Squared Error $\mathcal{L}_{MSE}(\hat{y}, y) = \frac{1}{nF} \sum_f^{nF} (\hat{y}_f, y_f)^2$ and Mean Average Error $\mathcal{L}_{MAE} = \frac{1}{nF} \sum_f^{nF} \|\hat{y}_f, y_f\|$ to ensure that the network will be penalized exponentially when the error is large, but that the loss will not decrease too much when reaching small errors. Thus, the network will continue to learn even if its predictions are already close to the ground truth. A regularization term $\mathcal{R}(y) = (1 - \alpha) + \left(\alpha \cdot \frac{\|y - \gamma\|}{\gamma} \right)$ helps the network to emphasize training on the LVEF objective, weighing down LVEF estimates which are away from γ , where γ is chosen close to the mean of all the LVEF on the training set. The α parameter is a scalar which adjusts the maximum amount of regularization applied to LVEF close to γ . Thus, our overall objective loss can be written as $\mathcal{L}_{EF} = (\mathcal{L}_{MSE} + \mathcal{L}_{MAE}) \cdot \mathcal{R}(y)$.

3 Experimentation

Dataset: For all of our experiments, we use the Echonet-Dynamic dataset [17] that consists of a variety of pathologies and healthy hearts. It contains 10,030 echocardiogram videos of varied length, frame rate, and image quality, all containing 4-chamber views. We pad all the frames with zeros to go from 112×112 to 128×128 pixel size frames. The videos represent at least one full heart cycle, but most of them contain three or more cycles. In each clip, however, only one cycle has the ES and ED frames labelled. For each labelled frame, the index in the video and the corresponding ground truth segmentation of the left ventricle is available. We have access to the frame rate, left ventricular volumes in milliliters for the ES and ED frames and the ejection fraction, computed from these heart volumes via $\text{LVEF} = \frac{\text{EDV} - \text{ESV}}{\text{ESV}} * 100$. Following Echonet [17], we

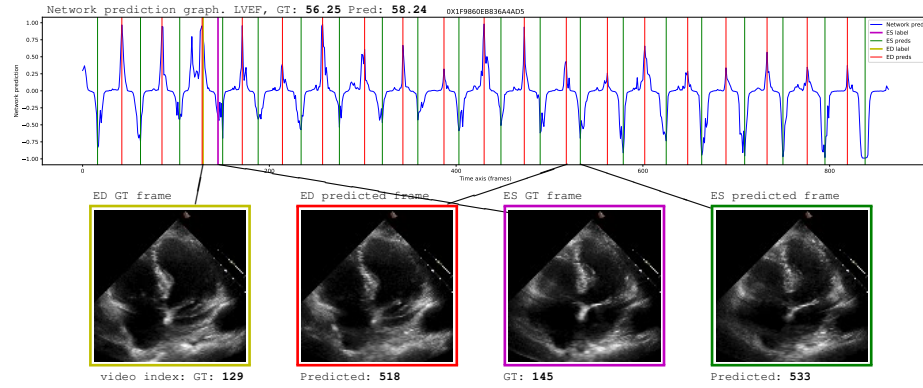


Fig. 2: Example network response (top), labelled pair and one predicted pair of US frames (bottom) for a very long video. The frames are shown for the labelled ED and ES indices and one of the multiple sets of the ES/ED frames produced by the network from a different heartbeat than those of the labelled indices.

split the dataset in training, validation and testing subsets; dedicating 75% to our training subset and equally splitting validation and testing with 12.5% each. **Video sampling process:** For training, we need videos of fixed size length to enable batch training and we need to know where the ES and ED frames are. In our data, only single ES and ED frames are labelled for each video, leaving many of the other true ES and ED frames unlabeled. We choose to create 128 frames long sequences, based on the distribution of distances between consecutive ES and ED frames in the training set. Any sequence longer than 128 frames is sub-sampled by a factor of 2. As training the transformer with unlabeled ES and ED frames would considerably harm performance, we tried two approaches to mitigate the lack of labels: (1) Guided Random Sampling: We sample the labeled frames and all frames in-between. We then add 10% to 70% of the distance between the two labelled frames before and after the sampled frames. To match the desired 128 length clip, where appropriate we pad it with black frames. The resulting clip is masked such that the attention heads ignore the empty frames. (2) Mirroring: We augment our clips by mirroring the transition frames between the two labelled frames, and placing them after the last labelled frame. Given the sequence $S = [f_{ES}, f_{T_1}, \dots, f_{T_N}, f_{ED}]$, where f_{T_N} stands for the N th transition frame, we augment the sequence such that it reaches a length superior to 128 frames by mirroring the frames around f_{ED} creating the new sequence $S' = [f_{ES}, f_{T_1}, \dots, f_{T_N}, f_{ED}, f_{T_N}, \dots, f_{T_1}, f_{ES} \dots]$. Once the sequences exceed 128 indices, we randomly crop them to 128 frames. Doing so ensures that the first frame is not always labelled. This augmentation follows common practices when using a transformer model in order to provide seamless computational implementation. In addition, it ensures that all the ES and ED frames that our transformer sees, are labelled as such, while having no empty frames and retaining spatio-temporal coherence.

| Video sampling | Single heartbeat | | | Full video | | |
|----------------|----------------------------------|--------------------|--|--------------------|---------------------|------------|
| | <i>ED and ES index detection</i> | | | | | |
| | ES ↓ | ED ↓ | | ES ↓ | ED ↓ | Rejected ↓ |
| US [4] | 4.1 | 3.7 | | / | / | / |
| MRI [13] | 0.44 (0.46) | 0.38 (0.39) | | / | / | / |
| R./Cla. (ours) | 1.84 (2.73) | 2.34 (3.15) | | 2.86 (6.43) | 7.88 (11.03) | 3 |
| R./Reg. (ours) | 2.14 (2.46) | 3.05 (3.74) | | 3.66 (7.96) | 9.60 (30.71) | 4 |
| M./Cla. (ours) | 0.09 (1.25) | 0.14 (1.49) | | 5.49 (12.94) | 9.21 (14.15) | 31 |
| M./Reg. (ours) | 0.08 (1.53) | 0.15 (1.89) | | 3.35 (6.79) | 7.17 (12.92) | 6 |

| | <i>LVEF prediction</i> | | | | | |
|-----------------|------------------------|--------|------------------|-------------|--------|------------------|
| | MAE ↓ | RMSE ↓ | R ² ↑ | MAE ↓ | RMSE ↓ | R ² ↑ |
| EchoNet(1) [17] | 4.22 | 5.56 | 0.79 | 4.05 | 5.32 | 0.81 |
| EchoNet(2) [17] | 7.35 | 9.53 | 0.40 | / | / | / |
| R3D [17] | 7.63 | 9.75 | 0.37 | / | / | / |
| MC3 [17] | 6.59 | 9.39 | 0.42 | / | / | / |
| R. (ours) | 5.54 (5.17) | 7.57 | 0.61 | 6.77 (5.47) | 8.70 | 0.48 |
| M. (ours) | 5.32 (4.90) | 7.23 | 0.64 | 5.95 (5.90) | 8.38 | 0.52 |

Table 1: *Top:* Results for ES and ED detection on the test set (1024 videos) stated in **average Frame Distance (aFD) (standard deviation)** over all frame distances. Video sampling methods are equivalent for training and testing in each experiment. The rejected column indicates the number of videos for which the network did not find clear index positions, *i.e.*, not enough zero crossings in the output sequence. Our clinical expert confirmed not ideal image quality for these. US [4] and MRI [13] both train and test their approaches on single beat videos of fixed length from private datasets. *Bottom:* Results for LVEF prediction, compared to ground truth, labeled LVEF. Echonet (1 & 2), R3D and MC3 all come from [17] and use combinations of neural networks and heuristics. "R." and "M." refer to the random and mirror video sampling methods; "/" means not available, *e.g.*, from literature. Echonet(1) is restrictive through the use of segmentations and a fixed number of frames per beat. It aggregates per-beat predictions into a final estimate for entire sequences. Processing time is 1.6s per beat and a minimum number of frames, *e.g.*, 32, have to be available for each beat during retrospective analysis. Echonet(2) uses segmentation and all the available frames over a single, selected beat, with no further processing. Our average processing time is 0.15s for entire videos and our method runs in real-time during examination. ↑/↓ mean higher/lower values are better.

The labels for these frames are defined depending on the method we use to predict the frame type. When regressing, we set the ES to -1 and the ED to 1. We adapt the heuristics used in [4,13] to smooth the transition between the ES and ED volumes. When using a classification objective, we define three classes: transition (0), ED (1), ES (2), with no form of smoothing. In the mirror-sampling

| Exp | BERTs | Em. | Linear | Seq. len. | ES ED | MAE RMSE R ² | Inference | Params. |
|-----------|-------|------|--------|-----------|------------|-------------------------|-----------|---------|
| Ours | 16 | 1024 | 8192 | 128 | 3.35 7.17 | 5.95 8.38 0.52 | 0.15s | 346.8M |
| Reduced 1 | 4 | 256 | 1024 | 64 | 4.28 10.12 | 6.06 8.37 0.51 | 0.13s | 6.8M |
| Reduced 2 | 1 | 128 | 512 | 64 | 8.18 35.42 | 7.34 9.96 0.30 | 0.13s | 2.7M |
| No EF | 16 | 1024 | 8192 | 128 | 3.71 7.24 | / | 0.15s | 346.3M |
| No SD | 16 | 1024 | 8192 | 128 | / | 5.95 8.38 0.52 | 0.15s | 346.2M |

Table 2: Ablation study. ES|ED columns show the aFD and MAE|RMSE|R² shows LVEF scores. (Em. = Embeddings; Seq. len. = Sequence length)

method we cannot use the same heuristics to smooth the volume transition when regressing. Instead, we apply an x^3 function scaled between the ES and ED to soften the loss around the class peaks (-1, 1).

ES & ED Frame Detection: We train our end-to-end model to predict the ES, ED indices and LVEF using the two video sampling methods and two architecture variations for the SD branch. The architecture variation consists of outputting a $3 \times nF$ matrix instead of a $1 \times nF$ vector for the ES/ED prediction and replacing the activation function by *Softmax*. Instead of approximating the volume, like [4,13], of the left ventricle chamber, our architecture classifies each frame in either *transition*, ES or ED. For the regression (Reg.) model, we use the MSE loss. For the classification (Cla.) model, we use the weighted Cross-Entropy Loss, with weights (1, 5, 5), corresponding to our classes [transition, ES, ED]. The results presented in the top half of Tab. 1 are obtained while training simultaneously the SD and EF branches of the network. As common in literature [4,13] we state average Frame Distance (aFD) as $aFD = \frac{1}{N} \sum_{n=1}^N | \hat{i}_n - i_n |$ and $std = sqrt\left(\frac{1}{N} \sum_{n=1}^N | \hat{i}_n - \bar{i} |^2\right)$, where i_n is the ES or ED label index for the n^{th} test video and \hat{i}_n is the corresponding predicted index.

LVEF Prediction: The prediction of the LVEF is done in parallel with the ES and ED detection. We apply a single pre-processing step which is the scaling of the value from 0-100 to 0-1. The results are presented in the 0-100 range. For the regularization term $\mathcal{R}(y)$ we empirically chose $\alpha = 0.7$ and set $\gamma = 0.65$ to match the over-represented LVEF values from the LVEF distribution in the training set. Results for LVEF are summarized in the bottom half of Tab. 1 and compared to results from literature.

Ablation Study: We tested the capacity of our M. Reg. model. From Tab. 2 we observe that using only four BERT encoders achieves results similar to the original architecture. Removing the EF or SD branch has little to no impact on the other branch. Most of the computing time is used in the ResNetAE encoder while most of the parameters are in the BERT encoders.

Discussion: Manual measurement of LVEF exhibits an inter-observer variability of 7.6% – 13.9% [17]. From Tab. 1, we can speculate that transformer models are en par with human experts for LVEF estimation and ED/ES identification.

We observe that our model predicts on average a 3% higher LVEF than measured by the manual ground truth observers, especially on long videos. We

hypothesize that the model learned to correlate spatio-temporal image features to ventricular blood volume, therefore sometimes selecting different heartbeats than labelled in the ground truth, *e.g.*, Fig. 2. While guidelines state that operators should select the cycle with the largest change between the ED area and the ES area, in practice, operators will usually just pick a “good” heart cycle, which might also be the case for our ground truth. This hypothesis requires clinical studies and cross-modal patient studies in future work.

A limitation of our study is that aFD is not an ideal metric, since over-prediction would lead to inflated scores. Our model predicts distinct ES/ED frames as shown in Fig. 2. This can also be shown by calculating the average heart rate from peak predictions. When accounting for frame rate differences this results in a reasonable 53 beats per minutes.

Implementation: PyTorch 1.7.1+cu110 with two Nvidia Titan RTX GPUs.

4 Conclusion

We have discussed probably the first transformer architecture that is able to analyze US videos of arbitrary length. Our model outperforms all existing techniques when evaluated on adult cardiac 4-chamber view sequences, where the task is to accurately find ES and ED frame indices. We also outperform heuristic-free methods on LVEF prediction. In the future we would expect transformers to play a more prominent role for temporally resolved medical imaging data.

Acknowledgements: This work was supported by the UKRI Centre for Doctoral Training in Artificial Intelligence for Healthcare (EP/S023283/1) and Ultrasonics Ltd.

References

1. Baumgartner, C.F., Kamnitsas, K., Matthew, J., Fletcher, T.P., Smith, S., Koch, L.M., Kainz, B., Rueckert, D.: Sononet: real-time detection and localisation of fetal standard scan planes in freehand ultrasound. *IEEE Trans Med Imaging* **36**(11), 2204–2215 (2017)
2. Carneiro, G., Nascimento, J.C., Freitas, A.: The segmentation of the left ventricle of the heart from ultrasound data using deep learning architectures and derivative-based search methods. *IEEE Trans Image Process* **21**(3), 968–982 (2011)
3. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv:1810.04805* (2018)
4. Dezaki, F.T., Dhungel, N., Abdi, A.H., Luong, C., Tsang, T., Jue, J., Gin, K., Hawley, D., Rohling, R., Abolmaesumi, P.: Deep residual recurrent neural networks for characterisation of cardiac cycle phase from echocardiograms. In: *Deep learning in medical image analysis and multimodal learning for clinical decision support*, pp. 100–108. Springer (2017)
5. Fiorito, A.M., Østvik, A., Smistad, E., Leclerc, S., Bernard, O., Lovstakken, L.: Detection of Cardiac Events in Echocardiography Using 3D Convolutional Recurrent Neural Networks. In: *2018 IEEE IUS*. pp. 1–4 (2018)

6. Folland, E., Parisi, A., Moynihan, P., Jones, D.R., Feldman, C.L., Tow, D.: Assessment of left ventricular ejection fraction and volumes by real-time, two-dimensional echocardiography. a comparison of cineangiographic and radionuclide techniques. *Circulation* **60**(4), 760–766 (1979)
7. Gifani, P., Behnam, H., Shalbaf, A., Sani, Z.A.: Automatic detection of end-diastole and end-systole from echocardiography images using manifold learning. *Physiological Measurement* **31**(9), 1091–1103 (jul 2010)
8. Girdhar, R., Carreira, J., Doersch, C., Zisserman, A.: Video action transformer network. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 244–253 (2019)
9. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: *European conference on computer vision*. pp. 630–645. Springer (2016)
10. Hou, B.: ResNetAE (last accessed: 22/06/2021) (2019), <https://github.com/farrell1236/ResNetAE>
11. Jahren, T.S., Steen, E.N., Aase, S.A., Solberg, A.H.S.: Estimation of end-diastole in cardiac spectral doppler using deep learning. *IEEE T ULTRASON FERR* **67**(12), 2605–2614 (2020)
12. Kalfaoglu, M.E., Kalkan, S., Alatan, A.A.: Late temporal modeling in 3d cnn architectures with bert for action recognition. In: Bartoli, A., Fusiello, A. (eds.) *Computer Vision – ECCV 2020 Workshops*. pp. 731–747. Springer International Publishing, Cham (2020)
13. Kong, B., Zhan, Y., Shin, M., Denny, T., Zhang, S.: Recognizing end-diastole and end-systole frames via deep temporal regression network. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2016*. pp. 264–272. Springer International Publishing, Cham (2016)
14. Kupinski, M.A., Hoppin, J.W., Krasnow, J., Dahlberg, S., Leppo, J.A., King, M.A., Clarkson, E., Barrett, H.H.: Comparing cardiac ejection fraction estimation algorithms without a gold standard. *Academic radiology* **13**(3), 329–337 (Mar 2006)
15. Lang, R.M., Badano, L.P., Mor-Avi, V., Afilalo, J., Armstrong, A., Ernande, L., Flachskampf, F.A., Foster, E., Goldstein, S.A., Kuznetsova, T., et al.: Recommendations for cardiac chamber quantification by echocardiography in adults: an update from the american society of echocardiography and the european association of cardiovascular imaging. *Eur. Heart J. Cardiovasc. Imaging* **16**(3), 233–271 (2015)
16. Måløy, H.: Echobert: A transformer-based approach for behavior detection in echograms. *IEEE Access* **8**, 218372–218385 (2020)
17. 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 for beat-to-beat assessment of cardiac function. *Nature* **580**, 252–256 (4 2020)
18. Qin, C., Bai, W., Schlemper, J., Petersen, S.E., Piechnik, S.K., Neubauer, S., Rueckert, D.: 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.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*. pp. 472–480. Springer International Publishing, Cham (2018)
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.) *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. pp. 234–241. Springer International Publishing, Cham (2015)

20. Russo, C., Hahn, R.T., Jin, Z., Homma, S., Sacco, R.L., Di Tullio, M.R.: Comparison of echocardiographic single-plane versus biplane method in the assessment of left atrial volume and validation by real time three-dimensional echocardiography. *Journal of the American Society of Echocardiography* **23**(9), 954–960 (2010)
21. Smistad, E., Østvik, A., Salte, I.M., Leclerc, S., Bernard, O., Lovstakken, L.: Fully Automatic Real-Time Ejection Fraction and MAPSE Measurements in 2D Echocardiography Using Deep Neural Networks. In: 2018 IEEE IUS. pp. 1–4 (2018)
22. Smistad, E., Østvik, A., Salte, I.M., Melichova, D., Nguyen, T.M., Haugaa, K., Brunvand, H., Edvardsen, T., Leclerc, S., Bernard, O., Grenne, B., Løvstakken, L.: Real-time automatic ejection fraction and foreshortening detection using deep learning. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **67**(12), 2595–2604 (2020)
23. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: *Advances in neural information processing systems*. pp. 5998–6008 (2017)
24. 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: Martel, A.L., Abolmaesumi, P., Stoyanov, D., Mateus, D., Zuluaga, M.A., Zhou, S.K., Racoceanu, D., Joskowicz, L. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*. pp. 623–632. Springer International Publishing, Cham (2020)
25. Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., et al.: Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771* (2019)
26. Zolgharni, M., Negoita, M., Dhutia, N.M., Mielewczik, M., Manoharan, K., Sohaib, S.A., Finegold, J.A., Sacchi, S., Cole, G.D., Francis, D.P.: Automatic detection of end-diastolic and end-systolic frames in 2d echocardiography. *Echocardiography* **34**(7), 956–967 (2017)