

EchoDFKD: Data-Free Knowledge Distillation for Cardiac Ultrasound Segmentation using Synthetic Data

Grégoire Petit^{1,2,*}, Nathan Palluau^{*}, Axel Bauer², Clemens Dlaska^{1,2}

¹Digital Cardiology Lab, Medical University of Innsbruck, A-6020 Innsbruck, Austria

²University Clinic of Internal Medicine III, Cardiology and Angiology,
Medical University of Innsbruck, A-6020 Innsbruck, Austria

g.petit360@gmail.com, nathan.palluau@gmail.com, clemens.dlaska@i-med.ac.at

Abstract

The application of machine learning to medical ultrasound videos of the heart, i.e., echocardiography, has recently gained traction with the availability of large public datasets. Traditional supervised tasks, such as ejection fraction regression, are now making way for approaches focusing more on the latent structure of data distributions, as well as generative methods.

We propose a model trained exclusively by knowledge distillation, either on real or synthetic data, involving retrieving masks suggested by a teacher model. We achieve state-of-the-art (SOTA) values on the task of identifying end-diastolic and end-systolic frames. By training the model only on synthetic data, it reaches segmentation capabilities close to the performance when trained on real data with a significantly reduced number of weights. A comparison with the 5 main existing methods shows that our method outperforms the others in most cases.

We also present a new evaluation method that does not require human annotation and instead relies on a large auxiliary model. We show that this method produces scores consistent with those obtained from human annotations. Relying on the integrated knowledge from a vast amount of records, this method overcomes certain inherent limitations of human annotator labeling.

Code: <https://github.com/GregoirePetit/EchoDFKD>

1. Introduction

In Deep Learning applied to computer vision, one critical task is segmentation [33], which is essential for accurately interpreting and analyzing visual data. However, segmentation as an annotation task is often resource-intensive and time-consuming, particularly in the medical domain. Knowledge distillation (KD) [12] typically aims to transfer

knowledge, such as the ability to segment data, from a large, complex model (often referred to as the “teacher model”) to a smaller, more efficient model (the “student model”). This process involves training the student model to replicate the outputs of the teacher model on certain tasks. By doing so, the student model inherits the teacher’s capabilities relevant to this task, requiring fewer computational resources. Beyond model compression, KD has demonstrated its versatility in various applications, such as adversarial robustness [27], ensemble fusion [45], continual learning [40], partial or missing labels [4], multi-task learning [18], and cross-modal learning [1].

It often happens that the medical data used to optimize a model is not accessible while the model itself is shared. In particular, regarding computer vision applied to echocardiography, several recent models are shared with the public while the corresponding datasets are not [6, 46]. It’s a safe bet that many laboratories are reluctant even to share their models because of the risk of reconstruction attacks. Knowledge distillation offers protection against attacks on sensitive biometric data that attempt to reconstruct training examples from models, making it a robust solution in medical applications [26]. Owners of a sensitive dataset could, therefore, share only the outputs of their models on several examples or a fusion of outputs from restricted models instead of sharing their model directly. Furthermore, KD is an efficient way of transferring knowledge from one modality to another and thus managing the situation common in medicine, where the different modalities of an example are not always the same from one dataset to another (some datasets might include, for instance, other views or patient metadata). Finally, it is effective for model compression when the student model is lightweight, reducing noise and forcing the student model to focus on the knowledge of interest while removing, for instance, labeler-specific style.

Compared with KD methods based on real data, Data-Free Knowledge Distillation (DFKD) methods enable

*equal contribution

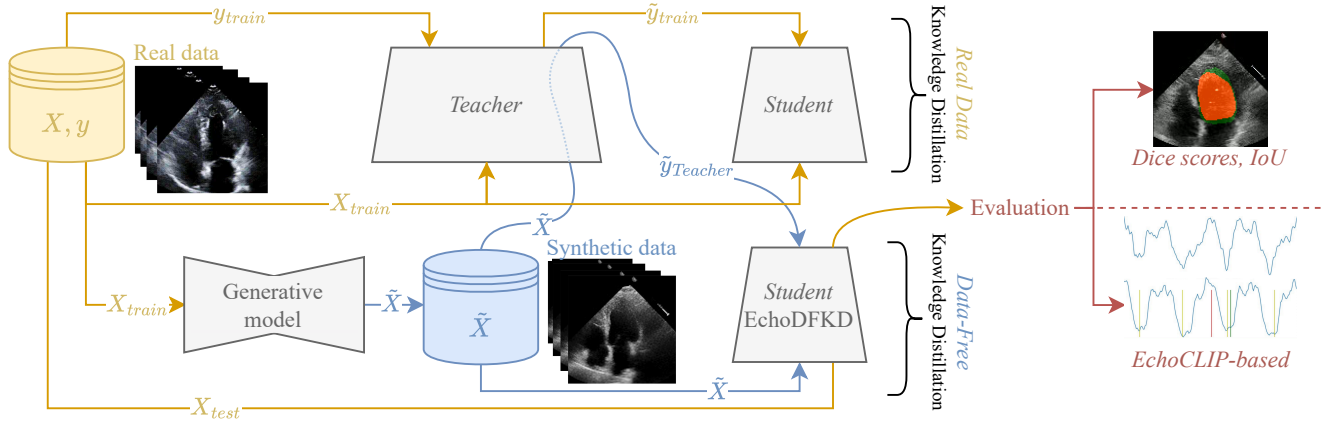


Figure 1. Overview of EchoDFKD. In KD, *real data*, e.g., from EchoNet-Dynamic [25], is often used to train a *Teacher* model and then used to train a *Student* model. EchoDFKD is using *synthetic data* from EchoNet-Synthetic dataset [30] to distill knowledge. By analyzing the mask generated by our ConvLSTM-based EchoDFKD segmentator or the similarity outputs of custom EchoCLIP [6] prompts and EchoNet-Dynamic raw images, we can predict the average Frame Distance (aFD). Additionally, our EchoDFKD segmentator described in Subsec. 3.2, is evaluated against EchoCLIP knowledge to assess its segmentation quality alongside traditional metrics (meanIoU, dice score) against human labels.

knowledge to be distilled on a potentially infinite number of artificial examples, making it possible to achieve a near-perfect imitation of the teacher model, at least in the domain well represented by the generative model [5, 19, 44]. In addition, it is possible to focus on a particular section of the possible examples space, e.g., by conditioning on a particular part of a patient’s health data (such as age, resting heart rate, etc.), thus generating a very large number of examples under these conditions to ensure efficient distillation of one or more teachers to a specialized student. Moreover, DFKD methods are not necessarily incompatible with KD methods based on real data since synthetic data can be used in addition to real data, in pre-training, in composite batches, or even in the form of data augmentation when examples are created from real examples [2, 11].

In the present work, we introduce EchoDFKD, the first DFKD model for echocardiography. We use this paradigm by relying on the EchoNet-Synthetic dataset [30]. This dataset was generated following the example of EchoNet-Dynamic (see Figure 1 and Section 3), whose test set we use to demonstrate the suitability of our approaches for subsequent tasks on real data.

We first provide the performance of our models using traditional methods based on human labels to provide a fair comparison with models achieving state-of-the-art (SOTA) performance and show that the very lightweight convolutional long short-term memory (ConvLSTM) architecture [38, 39], as some previous results using this architecture (but with more weights) had suggested [17], are relevant for apical-4-chamber echocardiograms. Furthermore, we establish a line on the Pareto (score, weight) frontier (see Figure 6), which demonstrates that, although the need to process frames one by one requires relatively large num-

bers of floating point operations (FLOPs) [21], it is possible to achieve competitive performance levels utilizing a significantly lower number of FLOPs.

To take our model-centric approach to its full extent, we propose removing humans from the process altogether by proposing evaluation methods by other models. This demonstrates that starting from initial models that have been developed involving human knowledge, it is then possible to successfully tackle machine learning tasks by letting these models interact with each other. In the context of unannotated medical data availability, this offers the advantage of enabling rigorous evaluation of masks provided by another agent on this data and fair comparison with other mask proposals. Indeed, while model-based evaluation is highly error-prone, these errors are less likely to be systematic or style-specific.

2. Related Work

2.1. Data-Free Knowledge Distillation

Data-Free Knowledge Distillation is a technique in which a student model is trained to replicate the behavior of a teacher model without access to the original training data. An alternative twisted approach involves aggregating outputs to train a student model and sharing only the student model, as seen in PATE [26], which discusses medical data. Another method creates virtual images from the teacher model, using ‘metadata’ (e.g., means and standard deviations of activations from each layer) recorded from the original training dataset. However, these metadata are often not provided for well-trained CNNs [19].

A standard approach utilized in DFKD involves using a GAN (including multiple variants) and distilling a teacher

model. This approach updates the generator based on the produced distribution at each step, extracts the penultimate layer from the teacher, and then evaluates it [5]. Similarly, KegNet [44] focuses on extracting trained deep neural network knowledge and generating artificial data points to replace missing training data in knowledge distillation.

2.2. Large Vision Models evaluation

Vision Transformers are known to be effective mask auto-labelers [15], which suggests that using pseudo-labels to evaluate models is not unreasonable. Additionally, some approaches aim to find a latent ground truth among multiple annotators, acknowledging that disagreement among annotators is not necessarily problematic [10]. Furthermore, it is also common to use Large Language Models (LLMs) for labeling tasks [41].

In our work, we examine the quality of Large Vision Models (LVMs) labeling. Since our segmentations do not originate from LVMs, we can use them to evaluate our segmentations bias-freely.

2.3. Segmentation

The field of ultrasound image segmentation using machine learning methods is currently very active. Efforts in fully automated echocardiogram interpretation are highlighted in [46], which, although trained on a private data set, laid the foundations for future research. The introduction of the CAMUS dataset, as detailed in [16], provided a well-labeled but smaller dataset for segmentation tasks.

A significant milestone was the release of the EchoNet-Dynamic dataset [25], which has become a benchmark for many studies [3, 20, 29, 36]. The segmentation and prediction of ejection fraction (EF) using this dataset has been explored extensively [14].

Recent advancements have seen the application of transformer models and other novel architectures. To predict end-systolic (ES) and end-diastolic (ED) frames, [32] employed a BERT model by treating the sequence like a series of words. Similarly, [9, 23] focused on both segmentation and EF prediction, the later achieving SOTA results on EF predictions.

Methods using pre-training [35] seem to be particularly effective. Notably, [21], from the same laboratory as EchoCoTr, achieves the best segmentation results to date.

Innovations in lightweight and explainable models have also emerged. For EF estimation, [22] used graph neural networks with a focus on ED and ES frames, while [24] achieved commendable segmentation and EF prediction using a lightweight Mobile U-Net [33].

3. Proposed Method

The overview of the proposed method, illustrated in Figure 1 is detailed in the following sections.

3.1. Datasets

The release of EchoNet-Dynamic dataset [25] by the Stanford University Medical Center, which contains 10,024 sparsely annotated¹ using a classic apical-4-chamber view image tracing method, has stimulated deep learning-oriented research on the interpretation of echocardiography videos. Annotations would consist of approximating the segmentation mask of the left ventricle using segments. Generally, one segment is along the principal axis, and then twenty segments cross this axis. Some examples have been annotated multiple times. The dataset contains examples with varied sampling rates and different image qualities. Some clips are corrupted (see supplementary material).

We also use the EchoNet-Synthetic dataset [30], which provides the same data type but is generated by the recent XSCM generative model [31]. This external, synthetic dataset enables us to train data-free. Because of its generative nature, the video clips are not that close to the original EchoNet-Dynamic dataset on which it was trained. However, the displacement fields across the samples are good enough to represent the heart motion’s basics accurately. Furthermore, as the EF conditions the generative network, the synthetic distribution ultimately becomes close enough for our KD setup.

3.2. Architecture

The model architecture we use extends the classic U-Net [33] design by integrating ConvLSTM [38] layers, forming a hybrid structure particularly well suited for spatiotemporal processing. The U-Net component handles spatial feature extraction through its characteristic encoder-decoder structure, utilizing convolutional layers for down-sampling and up-sampling, thereby preserving fine-grained details essential for segmentation. The ConvLSTM layers are embedded within the U-Net to capture temporal dependencies across sequences of echocardiographic images. We systematically investigated configurations ranging from a single block of one ConvLSTM layer to four blocks of four ConvLSTM layers, aiming to determine the optimal balance between model complexity and performance. We denote by (Bb,l) the EchoDFKD model where the downsampling part comprises b blocks of l ConvLSTM layers. By maintaining a hidden state that evolves over time, these layers enable the model to learn temporal patterns and transitions between consecutive frames, which is critical for accurate segmentation in sequences. This hybrid architecture allows the model to effectively map input sequences of images to output sequences of segmentation masks, enhancing its ability to per-

¹Six additional videos can be found in the original publicly available dataset, but they are not annotated. Examination of the code from the original paper and the codes shared in subsequent research works on the dataset shows that they are always filtered at the beginning of the pipeline, meaning that, in practice, only 10,024 clips are processed.

form tasks where both spatial and temporal information are crucial, such as in the segmentation of cardiac cycles.

3.3. Lightweight Models for Specific Tasks

Using the lightest possible model capable of performing the specific assigned task is advantageous for typical healthy cases. Heavier models represent the function $f : \text{record} \rightarrow \text{verdict}$ in a high-dimensional space, which can lead to two main issues: (1) Overfitting to specific annotation styles: This can result in the model adapting to systematic human errors or institution-specific recording peculiarities, as seen in current models whose performance drops significantly across different datasets. Many medical datasets, including EchoNet-Dynamic, compile several composite sources with recognizable formats (due to cropping methods, sampling rates, etc.). Each source may be associated with particular annotators. A heavy model tends to learn links between input specificities and annotation specificities, which can be mitigated during compression. (2) Focusing excessively on outliers or corrupted examples: Heavier models may disproportionately emphasize particular, often corrupted examples that may not strictly pertain to the defined task.

It is important to note that the models we use here never exceed 4 million parameters.

3.4. Large Vision Models and Specialized Models Evaluation

The emergence of large models like CLIP [28] and some specialized versions of it like EchoCLIP [6], capable of performing multiple tasks, offers the potential for generating pseudo-labels that complement human labels. These large models can be viewed as valuable agents in producing opinions. Their outputs, obtained at a lower cost than labor-intensive human segmentation labels, can be used to curate training datasets or refine evaluation sets by distinguishing between examples that fit a defined framework and those that do not.

Large models can also assess the understanding of specialized smaller models regarding the signals they are supposed to master. For example, they can verify that these smaller models can retrieve task-related metrics. This approach helps avoid issues related to noise and biases inherent in human labels, particularly those requiring high precision and susceptible to annotator fatigue. Additionally, imperfect pseudo-labels can integrate basic knowledge from private learning and multimodal practitioner analysis, such as annotations made with access to multiple perspectives.

3.5. Practical Considerations in Assessing and Comparing Models

A model should be evaluated based on the specific task it is designed to perform. It is unrealistic to expect a model to

handle all out-of-distribution cases effectively. In medical computer vision applications, practitioners generally prefer a model that recognizes its limitations and alerts the user when the input signal is unusual rather than a model trained to provide mediocre results in poorly captured cases.

It is important to remember that clips used for machine learning are often short extracts from longer signals. A user can extend the recording if a model indicates it cannot interpret the signal in real-time. This justifies the transition to human examination for patients with anomalies that cause deviation from the usual distribution. Thus, these examples can be considered positives in the context of pre-screening with the model.

This approach ensures the model remains reliable and useful in real-time applications, distinguishing between typical cases and those requiring further human intervention.

3.6. Data-Free Knowledge Distillation Setting

Synthetic Data Generation The student model is trained on synthetically generated data crafted to resemble the distribution of the original training data used for the teacher model.

Knowledge Transfer The teacher model provides supervision by generating pseudo-labels or masks for the synthetic data. The student model then learns to replicate these outputs, effectively distilling the teacher's knowledge without direct access to real data.

Robustness Assessment Additional evaluations may include robustness tests to ensure the student model performs well across diverse and unseen data subsets and highlights its generalization capabilities.

3.7. Training

We conducted a limited exploration of hyperparameters without delving too deeply. Among the hyperparameters, we included choices between different parameter initialization methods, various depths (number of channels) in intermediate layers, different losses, different batch sizes, different sequence lengths, and various learning rate management strategies. For models with multiple blocks, we also allowed ConvLSTM to create residual connections between the outputs of different blocks. Notably, hyperparameter optimization selected a residual connection for the last block in our best-performing model.

As the teacher model, we chose the model shared when EchoNet-Dynamic was released, based on a DeepLabv3 architecture [25].

The models converged in approximately 400 epochs. We selected the model from the epoch that yielded the best validation loss. The validation set was used for hyperparameter optimization and for choosing the thresholds to convert floating-point masks to boolean masks.




Figure 2. Segmentation examples from an EchoNet-Dynamic video [(a), (b)] and an EchoNet-Synthetic video [(c)]. In (a), EchoDFKD predictions are shown in the green (G) channel, while DeepLabv3 predictions are displayed in the red (R) channel. Image (b) is the original, unaltered EchoNet-Dynamic frame. Image (c) represents a segmentation mask of DeepLabv3 on the EchoNet-Synthetic dataset used to train EchoDFKD.

3.8. Direct Evaluation of Segmentation: Assessing mask quality

Creating a segmentation label is more complex than producing labels for tasks like classification or regression. This complexity arises because a segmentation mask requires assigning a label to every pixel in an image, effectively working in a much higher-dimensional space. Unlike classification or regression, where a single decision is made for the entire image, segmentation demands precise labeling of intricate details across the whole picture. Due to this complexity, even expert human annotators often find it challenging to produce perfectly accurate segmentation masks, which can lead to inconsistencies and a lack of satisfaction with the labels they create.

For instance, if we refer to the value of the left ventricular ejection fraction (LVEF), which is often derived from multiple segmentation masks, it is interesting to note that the best performance in estimating the LVEF on EchoNet-Dynamic, currently achieved by EchoCoTr [23] (the coefficient of determination R^2 is 0.82, meaning that the squared correlation coefficient between outputs and targets is at least 0.82), surpasses the intra-annotator variability (which squared correlation coefficient is $r^2 = 0.81$) reported in [16]. Although these values are not truly comparable, as the annotator named O1 from CAMUS does not necessarily have the same consistency as that of DeepLabv3, and EchoNet-Dynamic is different from that of CAMUS, which also includes 2-chamber views, this comparison still raises methodological questions. It should be noted that the intra-annotator correlation coefficient cannot in any way constitute a theoretical ceiling for the performance of a model; the model outputs could be extremely close to the mean annotation of the annotator while being less noisy, the proof can be found in the supplementary materials. However, cases of exceedances warrant attention.

In EchoNet-Dynamic, the human labels consist of 21 segments, which cannot be strictly classified as segmentation but as a polygonal annotation. It is easy to understand

that a label produced this way cannot be perfectly satisfactory, as some regions are not labeled at the pixel level. When EchoNet-Dynamic was released, the authors showed the segmentations produced by their model and the segmentations produced by humans, and some of the segmentations produced by the machine were preferred to that of the human. This approach can be likened to second-order labeling, like that used to train a reward model.

Inspired by this observation, we chose to use the large EchoCLIP [6] model to assess the quality of the masks produced by the models. Large vision models trained by contrastive learning have already demonstrated their ability to perform segmenting tasks [8, 42, 43]. Without going that far, we are simply trying to use such a model to judge the quality of semantic segmentation.

We rely on EchoCLIP’s ability to judge the quality of a segmentation mask based on a few well-chosen prompts. To do this, we adopt a strategy that evaluates, on the one hand, that the mask does not significantly exceed the boundaries of the left ventricle and, on the other hand, that it adequately fills the entire area of the left ventricle.

To evaluate whether the mask does not overflow, we check that EchoCLIP can correctly distinguish the walls around the left ventricle. For this purpose, we blacken the area of the image under the mask and show the frames to EchoCLIP, using the prompt “WALL” to verify that it correctly distinguishes the walls of the left ventricle. To eval-

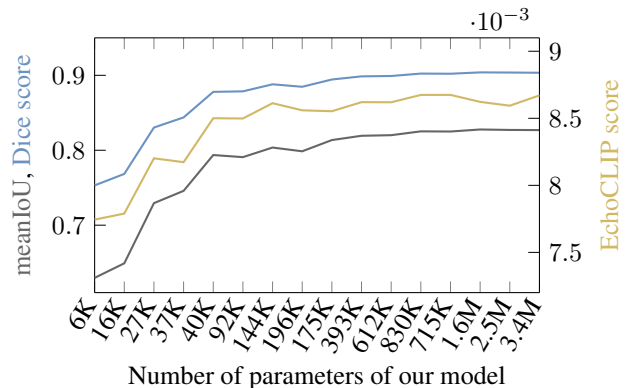


Figure 3. Illustration of the relationship between the number of model parameters and three key performance metrics: mean Intersection over Union (meanIoU), Dice score, and our custom EchoCLIP score. The meanIoU and Dice score, displayed on the left y-axis, show how segmentation accuracy against human annotators improves with increased model complexity. Our EchoCLIP score, shown on the right y-axis, reflects the segmentation quality without needing any annotator. In the EchoCLIP segmentation quality assessment, the segmentation quality is determined by the difference between the prompts “LEFT VENTRICLE” and “NOTHING” applied on raw masks that have been expanded by a few pixels.

uate whether the mask is not too small, we take the op-

		B1	B2	B3	B4
meanIoU	11	66.92%	77.31%	82.47%	83.70%
	12	66.23%	81.08%	83.30%	83.96%
	13	73.38%	81.19%	83.53%	83.89%
	14	74.47%	81.59%	83.62%	83.72%
Dice score	11	78.92%	86.50%	90.13%	90.93%
	12	78.21%	89.17%	90.68%	91.11%
	13	83.68%	89.23%	90.83%	91.07%
	14	84.48%	89.58%	90.89%	90.96%
DeepLabv3 [25]		Dice score		92.26%	
simLVSeg [21]		Dice score		93.32%	
MU-UNET [24]		Dice score		90.50%	
nnU-net ² [13]		Dice score		92.86%	

Table 1. Traditional performance metrics across EchoDFKD configurations (B*b*,*l*)³, against human annotators. DeepLabv3 [25], simLVSeg [21], MU-UNET [24] and nnU-net [13] are trained on real data, whereas EchoDFKD is trained via knowledge distillation on synthetic data.

posite approach; we enlarge the mask by a few (5) pixels at the borders, then blacken the enlarged mask. We check that EchoCLIP recognizes the resulting image by differentiating the responses to the prompts “*LEFT VENTRICLE*” and “*NOTHING*”. Figure 3 represents the evolution of EchoCLIP answers depending on the model size. We observe a plateau beyond which the model no longer improves, confirmed by the evolution of the dice scores and meanIoU, detailed in Table 1, measured against the frames initially labeled in EchoNet-Dynamic.

3.9. Evaluation on a Downstream Task

Because judging a model based on the segmentation task is difficult, as human labels are imperfect in their own opinion, it is natural to judge the knowledge acquired by the model based on its ability to perform a task whose completion seems possible for an agent who knows how to perform the main task, such as reconstructing features that can be deduced from the segmentation. In particular, the ability to identify the end-diastolic (ED) and end-systolic (ES) frames constitutes a robust test, which has already been employed on other datasets [7], and seems to have recently attracted attention for the EchoNet-Dynamic dataset [22, 32]. These frames have a concrete utility as they allow the calculation of LVEF when segmented.

We noticed that methods reported in the literature for determining aFD tend to be inconsistent. It can be calculated using sequences of different lengths and different strategies. For example, in [22], sub-clips of 64 frames are used, and results are only provided for examples where different phases are identified. They discard from the evalu-

²reported from [21]

³described in Subsection 3.2




Figure 4. Relationship between sampling rate and mean aFD (average Frame Distance) error for EchoDFKD: This plot shows the mean aFD error of the EchoNet-Dynamic dataset as a function of the sampling rate used for EchoDFKD.

ation examples where their model does not find one of the two phases. In [32] sub-clips of size 128 are extracted using a method that is difficult to assess in terms of statistical artifacts and also remove some examples that the model is unable to process (between 3 and 6 in the case of the scores reported here). We, therefore, propose a simple, direct method for deducing these frames from any model that can produce, in zero-shot, a sequence of values (one per frame) correlated with the alternation of the two phases and hope that it can constitute a standard. In addition, we advise not to remove any more examples during evaluation, to facilitate comparison between models. We report our performance both on the entire test set (we take these values to compare our approach with other works) and, for information purposes only, on a subset that excludes examples with corrupted labels.

Methodology (1) We take the entire video clip, (2) We examine the masks produced by the model and calculate the mask area in each frame. We calculate the median of the resulting sequence. (3) We identify the contiguous blocks of values below (or above) the median (this roughly corresponds to systolic or diastolic phases), choose the contiguous block closest to the reference value (to position ourselves on the same beat), and select the smallest (or highest) value within this block from the series of areas.

Note that, as might be expected, the errors in the aFD computing, shown in Figure 4, appear to be proportional to the sampling rate. Consequently, comparing aFD scores from one dataset to another is delicate, as this requires a certain homogeneity in the sampling rate distribution.

Alternative way of evaluating performances Comparing the frame labeled as a systole by the human annotator to the corresponding frame automatically labeled as a systole based on the mask widths from the original EchoNet-Dynamic’s DeepLabv3, we observe that the human-identified frame often appears before the one identified by DeepLabv3, on average, 2.3 frames earlier.

Our interest in finding an unbiased ground truth motivates our approach to using EchoCLIP to discriminate be-

tween ED and ES frames and other frames. To do this, we attempted to find prompts that varied simultaneously with the size of the DeepLabv3 masks and could capture this quasi-cyclical movement of the image.

We tried different combinations of prompts. Differentiating the result of the prompt “THE MITRAL VALVE IS CLOSED” from the prompt “THE MITRAL VALVE IS OPEN” and then integrating this signal yields a convincing plot that is consistent with that of other models and with human labels on many examples, as shown on Figure 5.

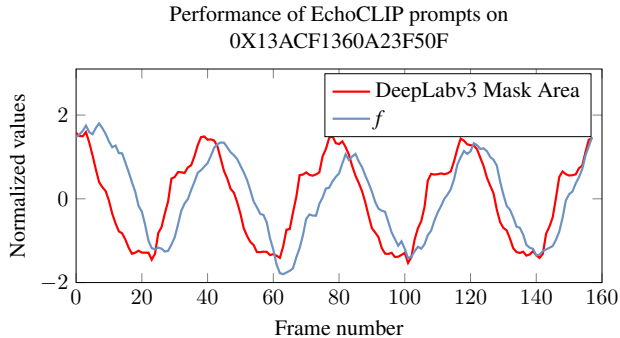


Figure 5. Comparison of DeepLabv3 masks area and prompts similarities. f is the cumulative sum of the difference between the prompts “THE MITRAL VALVE IS CLOSED” and “THE MITRAL VALVE IS OPEN” with the linear trend removed.

Due to the novelty of this approach and the lack of previous scores to compare against, we report the results in the supplementary material. Note that we find the same trend when comparing to human labels.

3.10. Results

	#	aFD _{ED}	aFD _{ES}
DeepLabv3* [25]	40M	8.63	3.66
UVT _R [32]	347M	7.88	2.86
UVT _M [32]	347M	7.17	3.35
EchoGNN [22]	1.7M	3.68	4.15
DRNNEcho [7]	(>10M) ⁴	3.7	4.1
EchoDFKD	0.72M	2.72	2.83
EchoDFKD*	0.72M	<u>2.77</u>	<u>2.83</u>

Table 2. Number of parameters and aFD obtained using various segmentators. All results reported in the table are computed on a subset⁵ of the EchoNet-Dynamic dataset. We also provide results of the entire test set (*). **Best results - in bold**, second best - underlined.

Comparison to real-data trained methods Table 2 shows that EchoDFKD outperforms almost all compared

⁴The authors claimed to use ResNets with 126 layers and some LSTM on top, so this number of parameters is very approximated.

⁵The authors removed complicated samples to compute their aFD prediction.

methods in setups where the real data is accessible to train on directly. However, the scores for each experiment are close. Additionally, it can be observed that training on synthetic data requires 3 to 4 times fewer weights. This suggests that the intrinsic complexity of the synthetic dataset concerning the tasks we are performing is lower, in similar proportions. However, this hypothesis should be taken cautiously, and quantifying each dataset’s dimensionality would require more rigorous exploration. Without delving into such deep considerations, it can be noted that when the model trained on synthetic data is evaluated on real data, most of the total error is concentrated on a small portion of the test examples (see supplementary materials for details).




Figure 6. Log-log display of the performances of different EchoDFKD models as a function of model size, as well as several models from the literature. The performances of the EchoDFKD models reach a ceiling as they approach their teacher’s performance.

The deep learning community has established that the evolution of model performance scores as a function of their size is generally well-described by exponential scaling laws, at least in a regime preceding the saturation imposed by the dataset size or other factors [34, 37]. We therefore examined, on the one hand, $-\log(\text{aFD}_{ED} + \text{aFD}_{ES})$ as a function of $\log(\text{Model_size})$, and, on the other hand, $\log(1 - \text{meanIoU})$ as a function of $\log(\text{Model_size})$. These analyses were conducted both for evaluation against human annotations and using EchoCLIP rewards (see plots in the supplementary material). We observed clear linear trends up to the largest models, where performance begins to stagnate as referred in Figure 6.

FLOPs In evaluating our model’s efficiency, we consider the number of FLOPs required for inference. EchoNet-Dynamic’s DeepLabv3 operates at 7.84 GFLOPs, reflecting its computational intensity and resource demands. In contrast, our models demonstrate significantly lower computational requirements, with GFLOPs values of 0.25 and 1.56, respectively, for the results reported in Table 2.

4. Limitations and future work

Our experiments show that it is possible to obtain results close to the SOTA in the segmentation stage and to outper-

form the SOTA in the frame identification stage with tiny models. Training models on synthetic data rather than real data also yields competitive results. The fact that we managed to surpass the SOTA with so few parameters may be due to our model producing smoothed outputs precisely because of its compression and its recurrent nature. Applying low-pass filters to a signal can sometimes improve peak detection, and we have probably reproduced a similar effect.

It is worth noting that we use ConvLSTMs in a mode where there is only one pass, so they have no hindsight and must make decisions in real-time, frame by frame, as soon as they arrive. This mainly implies poor performance in the very first frames (see supplementary materials), as the model likely needs to grasp the situation. It also means that it is more difficult for the model to identify the ES and ED frames since it cannot know if the next frame will show a stagnation or a change in the observed motion.

In this work, we only used 10,000 artificial examples of EchoNet-Synthetic, the generative model being recent and the generation of new examples time-consuming. In addition, we only used one teacher model rather than aggregating outputs from various models. It would be interesting to extend the data-free experiment into a large-scale one, using a cohort of teachers that could include models trained on other datasets and distillation datasets much more significant than 10,000 examples.

Our results suggest that it does not take much complexity to model the variables of interest from standard ultrasound scans. Moreover, the trend observed in model performance when varying the number of weights suggests that models may quickly reach a glass ceiling imposed by a certain vagueness around input standardization conventions and the limits of human labels. Last but not least, we must not lose sight of the fact that the tasks used as benchmarks for the many papers published on EchoNet-Dynamic are part of an approach that ultimately consists of evaluating LVEF. Ultrasound scans likely contain other interesting information from a medical point of view, which would be of more interest than scraping a few decimal points on a conventional score, and this is particularly true of examples where the video clearly shows patients that have other significant dysfunctions, as it was pointed out by [25].

This has several critical implications. Firstly, it implies that, given the simplicity of the inputs and targets, choosing one architecture over another is unlikely to impact the results obtained significantly. Multiplying work by trying out a whole panoply of different architectures, which are rarely original but adapted from models already well established in computer vision, is therefore likely to have limited scientific interest in this dataset and these tasks in the future.

If researchers want to continue improving model scores on these tasks, seeking richer or more precise targets will likely be necessary. In this paper, we proposed an approach

involving a large model as an evaluator. The large model does not have its own style, as it has seen data processed by various practitioners, and the fact that it has seen data in multiple modalities simultaneously may give it the ability to access, from ultrasound scans, information that is difficult for humans to perceive with this single modality. We generally seek information from other modalities; what happens along the third dimension is usually analyzed with a complementary 2-chamber view. However, the EchoCLIP verdicts remain pretty noisy, implying that many examples would be needed to discriminate between models reliably. EchoCLIP-like models may improve, and the approaches we explored with synthetic data for the training phase could be extended to the evaluation phase. A more straightforward and traditional way to obtain rich targets would be to collect multiple human labels for each evaluation example and combine these labels. It should also be noted that less downsampled data would also provide more precise labels. We have seen this with the choice of ES and ED frames, which makes less sense for videos with a low sampling rate. Still, more generally, downsampling videos to 112×112 pixels compresses much information on the ability to capture contours.

5. Conclusion

In this work, we have demonstrated the potential of using DFKD in video segmentation applied to cardiac ultrasound. First, we highlighted the challenges of using real data to train our model and how, with synthetic data, KD can overcome data access issues and effectively train a much lighter network. Second, we proposed a robust method to compute the aFD scores accurately. This method only requires an estimate of the inner left ventricle mask and thus is very easily transferrable. Finally, we have demonstrated a method for evaluating a model's ability to segment medical images using a large model, which gives results consistent with an evaluation of the fully labeled test set. Unlike the latter approach, ours does not require human labels and challenges the need to produce segmentation labels on test sets. A possibility to generalize this approach to training is using reinforcement learning to take advantage of feedback on mask quality.

Our work has shown possibilities in knowledge dissemination that circumvent limits imposed by the need to keep data private, as well as the possibility of extracting knowledge from heavy models on the best-understood part of the latent data space. We have also shown that unconventional evaluation methods can avoid costly and time-consuming human segmentation labels by taking advantage of a large existing model. We hope that by developing these methods, it will be possible to better share the significant but scattered efforts made in medical imaging.

References

- [1] Yusuf Aytar, Carl Vondrick, and Antonio Torralba. Soundnet: Learning sound representations from unlabeled video. *Advances in neural information processing systems*, 29, 2016. **1**
- [2] Shekoofeh Azizi, Simon Kornblith, Chitwan Saharia, Mohammad Norouzi, and David J Fleet. Synthetic data from diffusion models improves imagenet classification. *arXiv preprint arXiv:2304.08466*, 2023. **2**
- [3] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021. **3**
- [4] Guglielmo Camporese, Pasquale Coscia, Antonino Furnari, Giovanni Maria Farinella, and Lamberto Ballan. Knowledge distillation for action anticipation via label smoothing. *arXiv preprint arXiv:2004.07711*, 2020. **1**
- [5] Hanting Chen, Yunhe Wang, Chang Xu, Zhaohui Yang, Chuanjian Liu, Boxin Shi, Chunjing Xu, Chao Xu, and Qi Tian. Data-free learning of student networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3514–3522, 2019. **2, 3**
- [6] Matthew Christensen, Milos Vukadinovic, Neal Yuan, and David Ouyang. Vision–language foundation model for echocardiogram interpretation. *Nature Medicine*, pages 1–8, 2024. **1, 2, 4, 5**
- [7] Fatemeh Taheri Dezaki, Neeraj Dhungel, Amir H Abdi, Christina Luong, Teresa Tsang, John Jue, Ken Gin, Dale Hawley, Robert Rohling, and Purang Abolmaesumi. 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: Third International Workshop, DLMIA 2017, and 7th International Workshop, ML-CDS 2017, Held in Conjunction with MICCAI 2017, Québec City, QC, Canada, September 14, Proceedings 3*, pages 100–108. Springer, 2017. **6, 7**
- [8] Junsong Fan, Zhaoxiang Zhang, Chunfeng Song, and Tieniu Tan. Learning integral objects with intra-class discriminator for weakly-supervised semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4283–4292, 2020. **5**
- [9] Lhuqita Fazry, Asep Haryono, Nuzulul Khairu Nissa, Naulfal Muhammad Hirzi, Muhammad Febrian Rachmadi, Wisnu Jatmiko, et al. Hierarchical vision transformers for cardiac ejection fraction estimation. In *2022 7th International Workshop on Big Data and Information Security (IW BIS)*, pages 39–44. IEEE, 2022. **3**
- [10] Melody Guan, Varun Gulshan, Andrew Dai, and Geoffrey Hinton. Who said what: Modeling individual labelers improves classification. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018. **3**
- [11] Ruifei He, Shuyang Sun, Xin Yu, Chuhui Xue, Wenqing Zhang, Philip Torr, Song Bai, and XIAOJUAN QI. Is synthetic data from generative models ready for image recognition? In *The Eleventh International Conference on Learning Representations*, 2023. **2**
- [12] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. **1**
- [13] Fabian Isensee, Paul F Jaeger, Simon AA Kohl, Jens Petersen, and Klaus H Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods*, 18(2):203–211, 2021. **6, 7**
- [14] Mohammad Mahdi Kazemi Esfeh, Christina Luong, Delaram Behnami, Teresa Tsang, and Purang Abolmaesumi. A deep bayesian video analysis framework: towards a more robust estimation of ejection fraction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 582–590. Springer, 2020. **3**
- [15] Shiyi Lan, Xitong Yang, Zhiding Yu, Zuxuan Wu, Jose M Alvarez, and Anima Anandkumar. Vision transformers are good mask auto-labelers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23745–23755, 2023. **3**
- [16] Sarah Leclerc, Erik Smistad, Joao Pedrosa, Andreas Østvik, Frederic Cervenansky, Florian Espinosa, Torvald Espeland, Erik Andreas Rye Berg, Pierre-Marc Jodoin, Thomas Grenier, et al. Deep learning for segmentation using an open large-scale dataset in 2d echocardiography. *IEEE transactions on medical imaging*, 38(9):2198–2210, 2019. **3, 5**
- [17] Ming Li, Weiwei Zhang, Guang Yang, Chengjia Wang, Heye Zhang, Huafeng Liu, Wei Zheng, and Shuo Li. Recurrent aggregation learning for multi-view echocardiographic sequences segmentation. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part II 22*, pages 678–686. Springer, 2019. **2**
- [18] Wei-Hong Li and Hakan Bilen. Knowledge distillation for multi-task learning. *arXiv preprint arXiv:2007.06889*, 2020. **1**
- [19] Raphael Gontijo Lopes, Stefano Fenu, and Thad Starner. Data-free knowledge distillation for deep neural networks. *arXiv preprint arXiv:1710.07535*, 2017. **2**
- [20] Jun Ma, Yuting He, Feifei Li, Lin Han, Chenyu You, and Bo Wang. Segment anything in medical images. *Nature Communications*, 15(1):654, 2024. **3**
- [21] Fadillah Maani, Asim Ukaye, Nada Saadi, Numan Saeed, and Mohammad Yaqub. Simlvseg: Simplifying left ventricular segmentation in 2d+time echocardiograms with self- and weakly-supervised learning, 2024. **2, 3, 6, 7**
- [22] Masoud Mokhtari, Teresa Tsang, Purang Abolmaesumi, and Renjie Liao. Echognn: explainable ejection fraction estimation with graph neural networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 360–369. Springer, 2022. **3, 6, 7**
- [23] Rand Muhtaseb and Mohammad Yaqub. Echocotr: Estimation of the left ventricular ejection fraction from spatiotemporal echocardiography. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 370–379. Springer, 2022. **3, 5**

- [24] Meghan Muldoon and Naimul Khan. Lightweight and interpretable left ventricular ejection fraction estimation using mobile u-net. In *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI)*, pages 1–5. IEEE, 2023. 3, 6, 7
- [25] David Ouyang, Bryan He, Amirata Ghorbani, Neal Yuan, Joseph Ebinger, Curtis P Langlotz, Paul A Heidenreich, Robert A Harrington, David H Liang, Euan A Ashley, et al. Video-based ai for beat-to-beat assessment of cardiac function. *Nature*, 580(7802):252–256, 2020. 2, 3, 4, 6, 7, 8
- [26] Nicolas Papernot, Martín Abadi, Ulfar Erlingsson, Ian Goodfellow, and Kunal Talwar. Semi-supervised knowledge transfer for deep learning from private training data. *arXiv preprint arXiv:1610.05755*, 2016. 1, 2
- [27] Nicolas Papernot, Patrick McDaniel, Xi Wu, Somesh Jha, and Ananthram Swami. Distillation as a defense to adversarial perturbations against deep neural networks. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 582–597, Los Alamitos, CA, USA, may 2016. IEEE Computer Society. 1
- [28] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 4
- [29] Pranav Rajpurkar, Emma Chen, Oishi Banerjee, and Eric J Topol. Ai in health and medicine. *Nature medicine*, 28(1):31–38, 2022. 3
- [30] Hadrien Reynaud, Qingjie Meng, Mischa Dombrowski, Arijit Ghosh, Thomas Day, Alberto Gomez, Paul Leeson, and Bernhard Kainz. Echonet-synthetic: Privacy-preserving video generation for safe medical data sharing. *arXiv preprint arXiv:2406.00808*, 2024. 2, 3
- [31] Hadrien Reynaud, Mengyun Qiao, Mischa Dombrowski, Thomas Day, Reza Razavi, Alberto Gomez, Paul Leeson, and Bernhard Kainz. Feature-conditioned cascaded video diffusion models for precise echocardiogram synthesis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 142–152. Springer, 2023. 3
- [32] Hadrien Reynaud, Athanasios Vlontzos, Benjamin Hou, Arjan Beqiri, Paul Leeson, and Bernhard Kainz. Ultrasound video transformers for cardiac ejection fraction estimation. In *MICCAI*. Springer, 2021. 3, 6, 7
- [33] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015. 1, 3
- [34] Jonathan S. Rosenfeld, Amir Rosenfeld, Yonatan Belinkov, and Nir Shavit. A constructive prediction of the generalization error across scales. In *International Conference on Learning Representations (ICLR) 2020*, 2020. 7
- [35] Mohamed Saeed, Rand Muhtaseb, and Mohammad Yaqub. Is contrastive learning suitable for left ventricular segmentation in echocardiographic images? *arXiv*, 2022. 3
- [36] Fahad Shamshad, Salman Khan, Syed Waqas Zamir, Muhammad Haris Khan, Munawar Hayat, Fahad Shahbaz Khan, and Huazhu Fu. Transformers in medical imaging: A survey. *Medical Image Analysis*, 88:102802, 2023. 3
- [37] Utkarsh Sharma and Jared Kaplan. Scaling laws from the data manifold dimension. *Journal of Machine Learning Research*, 23(9):1–34, 2022. 7
- [38] Xingjian Shi, Zhouong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. *Advances in neural information processing systems*, 28, 2015. 2, 3
- [39] Xingjian Shi, Zhihan Gao, Leonard Lausen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, and Wang-chun Woo. Deep learning for precipitation nowcasting: A benchmark and a new model. *Advances in neural information processing systems*, 30, 2017. 2
- [40] Filip Szatkowski, Mateusz Pyla, Marcin Przewiezlikowski, Sebastian Cygert, Bartłomiej Twardowski, and Tomasz Trzcinski. Adapt your teacher: Improving knowledge distillation for exemplar-free continual learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1977–1987, 2024. 1
- [41] Zhen Tan, Alimohammad Beigi, Song Wang, Ruocheng Guo, Amrita Bhattacharjee, Bohan Jiang, Mansooreh Karami, Jundong Li, Lu Cheng, and Huan Liu. Large language models for data annotation: A survey. *arXiv preprint arXiv:2402.13446*, 2024. 3
- [42] Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. Filip: Fine-grained interactive language-image pre-training. In *International Conference on Learning Representations*, 2021. 5
- [43] Muyang Yi, Quan Cui, Hao Wu, Cheng Yang, Osamu Yoshie, and Hongtao Lu. A simple framework for text-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7071–7080, 2023. 5
- [44] Jaemin Yoo, Minyong Cho, Taebum Kim, and U Kang. Knowledge extraction with no observable data. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. 2, 3
- [45] Shuyang You, Chang Xu, Chao Xu, and Dacheng Tao. Learning from multiple teacher networks. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1285–1294, 2017. 1
- [46] Jeffrey Zhang, Sravani Gajjala, Pulkit Agrawal, Geoffrey H Tison, Laura A Hallock, Lauren Beussink-Nelson, Mats H Lassen, Eugene Fan, Mandar A Aras, ChaRandle Jordan, et al. Fully automated echocardiogram interpretation in clinical practice: feasibility and diagnostic accuracy. *Circulation*, 138(16):1623–1635, 2018. 1, 3