Deep Learning for Automated Echocardiogram Analysis

ABSTRACT

Heart disease is the #1 global killer. Left ventricular ejection fraction (EF) analysis via echocardiography is crucial for heart failure detection. However, traditional EF calculation is a time-consuming manual process with a ~30% error rate. My project aimed to develop an interpretable deep learning workflow to automatically analyze echocardiogram videos and calculate EF for quick assessment of cardiac function. Using Transfer learning, five PyTorch deep learning models were trained, validated, and tested using 10,030 echocardiogram videos with ~1 million image frames in the EchoNet Dynamic Dataset. My workflow mimics that of a cardiac sonographer so that results can be easily interpreted by doctors. The error rates for the models ranged from 10-16%, comparable to an expert cardiac sonographer (13.9% error rate), and significantly outperformed qualitative analysis by physicians (~30% error rate). Results can differentiate heart failure with reduced EF from normal cardiac function. MobileNet was identified as the best deep learning model for web applications and portable devices. I further deployed MobileNet-based workflow onto a mobile app, Raspberry Pi, laptop, and AWS, allowing physicians to upload echocardiogram videos and obtain EF results within seconds. In collaboration with a handheld ultrasound company, my model has been deployed with EchoNous' handheld ultrasound device for clinical testing. Automated echocardiogram analysis can dramatically speed up image analysis, reduce the burden on cardiologists, and democratize cardiac care by enabling non-experts to quickly and accurately assess cardiac functions at the point of care, in rural areas, or in developing countries, where cardiology expertise is limited.

Introduction

Heart disease is a leading cause of adult death/disability and the #1 global killer. According to the Centers for Disease Control and Prevention (CDC), one person dies every 36 seconds in the United States due to cardiovascular disease. Heart disease alone costs our healthcare more than $363 billion a year. Every year, about 805,000 people in the United States have a heart attack [1]. About 1 in 5 heart attacks is silent – the damage is done, but the person is not aware of it [1]. As a safe and non-invasive method, echocardiography is often the first and one of the most important tools for the diagnosis of heart diseases. It uses ultrasound reflections of cardiac structures to generate images of the heart and vessels to provide real-time imaging and diagnosis of heart problems, such as damaged cardiac tissue, stiffening of the heart muscle, heart chamber enlargement, blood clots, fluid around the heart, and damaged or poorly functioning heart valves [2].

Traditional echocardiogram imaging interpretation relies on skilled sonographers/cardiologists examining echocardiograms for pathologies. During a single routine echocardiogram, approximately 10–50 videos (around 3,000 images) are acquired to assess heart anatomy and function [3]. In clinical practice, human experts have limited time to analyze this large amount of imaging data with numerous other data such as laboratory results, vital signs and additional imaging studies (radiography, magnetic resonance imaging, nuclear imaging and computed tomography, etc.) [3]. As a result, echocardiographic assessment inaccuracy rate can be as high as 30% [4,5]. Manual and subjective echocardiogram analysis is a major barrier to the accurate diagnosis of heart disease [6]. This is especially problematic in the emergency room (ER) setting, where doctors have only 90 seconds to assess the cardiac function and make critical triage decisions for patients. Accurate and rapid assessment of echocardiograms is crucial for improving patient outcomes. Therefore, it is important to harness artificial intelligence (AI) technologies to address the challenges in echocardiogram image analysis.

AI is revolutionizing our lives. Deep learning and computer vision techniques particularly have had a profound impact on health care and disease diagnosis. Deep learning is a subset of AI machine learning. It uses neural networks to mimic the human brain's neural processes and can learn features from examples and detect deviations from the data patterns. In recent years, deep learning has advanced quickly in medical image analysis [7]. However, its application in echocardiogram analysis is lagging, mainly due to the lack of availability of large datasets, since deep learning algorithms require massive amounts of labeled data to achieve human-level classification performance [7].

Automated echocardiogram analysis with deep learning could dramatically speed up image analysis and significantly lower the cost of echocardiography. More importantly, it can democratize cardiac care by simplifying and providing more consistent and automated data interpretation, enabling non-experts to quickly assess cardiac functions, while reducing the burden on cardiologists.

In addition, the advent of handheld ultrasounds has allowed for mobile acquisition of echocardiograms. Rather using bulky traditional ultrasound machines, which can cost upwards of $50,000, handheld ultrasounds offer echocardiogram acquisition in a variety of previously untapped settings for a much lower cost. A mobile app can work in concert with handheld ultrasounds to provide rapid cardiac screening at the point of care, emergency rooms, rural areas, in developing countries, and homes.

With the prevalence of heart disease, simplified and automated echocardiogram analysis can provide readily accessible cardiac assessment, especially in rural areas or developing countries, and save lives by democratizing cardiac care.

Left ventricular ejection fraction (EF) analysis using an echocardiogram is the most powerful predictor of heart failure. Conventionally, EF measurement requires a laborious review of echocardiogram images and manual tracing of the endocardial border [8]. This time-consuming manual process has a high inaccuracy rate and large interobserver variability, leading to misdiagnosis of heart disease [6].

Two public datasets of echocardiograms became available very recently. One is Cardiac Acquisitions for Multi-structure Ultrasound Segmentation (CAMUS), which only contains still images [9]. The other is EchoNet, released in 2020, which contains 10,030 echocardiographic videos with expert tracings and measurements [10]. A few recent publications demonstrated the potential of deep learning for echocardiogram image analysis, but they require custom-built, sophisticated end-to-end models that are black-box solution and difficult to interpret [3,9,11]. This project aims to build deep learning pipelines using free online resources to automate echocardiogram analysis with a transparent and interpretable architecture.

The overall objective is to build a deep learning pipeline to automatically interpret echocardiogram videos and calculate EF to enable a quick and accurate assessment of cardiac functions. Specific engineering goals include:

1) Build transparent and interpretable deep learning models with improved accuracy and faster analysis time
2) Mimic the workflow of a cardiac sonographer (i.e., use Simpson's method to estimate the left ventricular volume), so that the deep learning pipeline is easier to understand and interpret by the doctors
3) Identify the best deep learning model for web applications and portable devices
4) Deploy this model using a mobile app, Amazon Web Services (AWS), standalone PC, and Raspberry Pi for remote access in concert with handheld ultrasounds

Methods

Database Preprocessing

The EchoNet dataset was downloaded from Stanford Artificial Intelligence in Medicine and Imaging (AIMI) Center Shared Datasets Portal [9]. As shown in Figure 1, over 10,000 videos were split into three subset development groups: training, validation, and testing subsets.

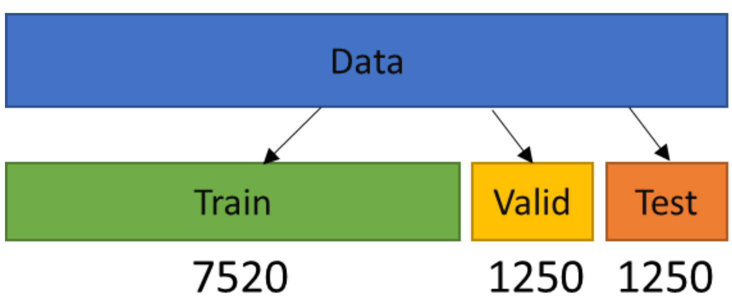

Figure 1. Echocardiogram data was split into training, validation, and testing subsets.

Semantic Segmentation of Left Ventricle

Manual segmentation of the cardiac boundary requires zooming in on left ventricle, manually tracing the endocardial border, and calculating the end-diastolic volume (EDV) and end-systolic volume (ESV) using Simpson's method [8].

Image semantic segmentation by deep learning involves identifying the location of interest within an image, so that subsequence classification will only focus on relevant pixels. This technique has been widely used for tasks such as facial segmentation and autonomous driving [12]. A common architecture used for semantic segmentation is convolutional neural networks (CNNs) because CNNs can identify local patterns such as edges, shape in an image and effectively use this information to segment an image.

Accurate segmentation is essential to evaluate cardiac structure from echocardiogram. Transfer learning and fine-tuning were applied to train five PyTorch deep learning models [13] (MobileNet, FCN50, FCN101, DeepLab 50, and DeepLab101) with the EchoNet dataset using the manual tracings by cardiac sonographer as the ground truth. These five segmentation models were pretrained on the COCO-Stuff dataset [14], which is a general-purpose semantic segmentation dataset containing over 1 million objects. The output layers of the models were then restructured and optimized to fit the EchoNet dataset using transfer learning. Figure 2 illustrates the workflow for semantic segmentation of left ventricle from the echocardiograms.

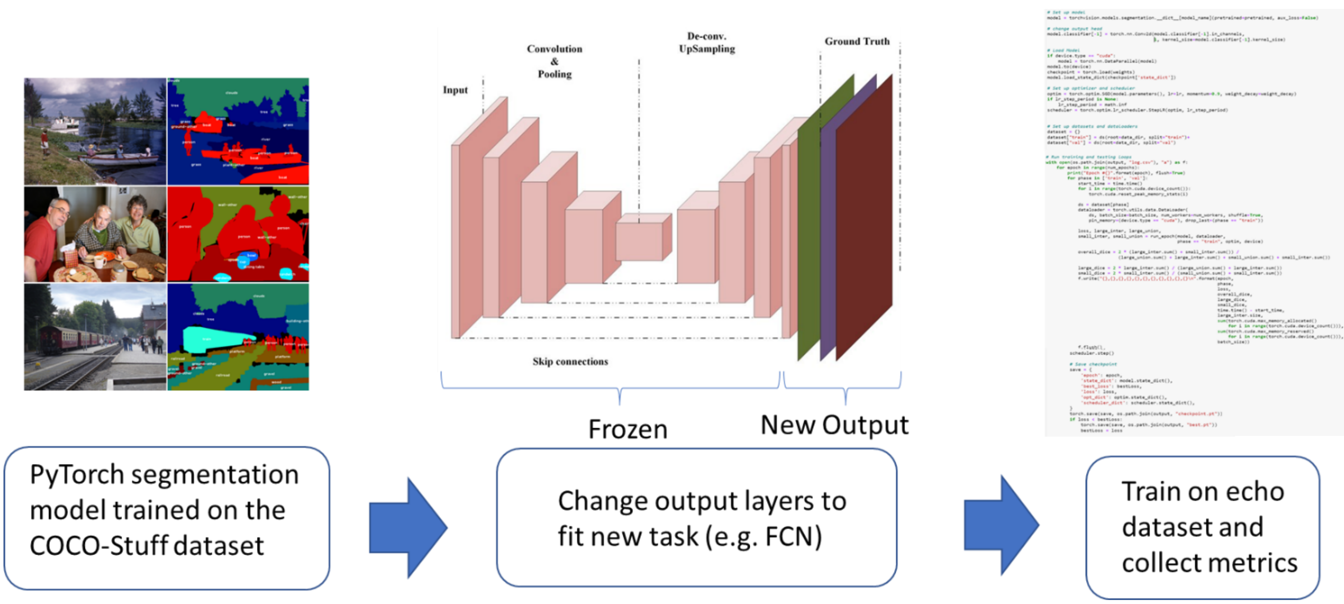

Figure 2. Semantic segmentation of left ventricle.

Five segmentation models from three main architectures were used for semantic segmentation. MobileNet is a model designed for edge devices like smart phones. It is faster and lighter than most deep learning models. Fully Convolutional Networks (FCN) combine fine and coarse layers to balance location and feature information. DeepLab is an encoder and decoder architecture. The encoder reduces features to capture higher-level semantic information, while the decoder recovers the spatial information [13].

The model performance was compared using the dice similarity coefficient (DSC), which compares the ground truth and predicted segmentation maps for each frame [15]. The performance comparison of different segmentation models is shown in Table 1. All five models have comparable DSCs, but the amount of training time and memory requirements differ significantly. MobileNet requires the least resources, only taking 2.5 hours of training time and requiring 0.8 GB of memory. With MobileNet, it also only takes 5 seconds to analyze each echocardiogram video.

Due to its superior memory and inference efficiency, MobileNet was optimized further by tuning various hyperparameter. In addition, the sampling method was changed from the default shuffle to a WeightedClassBased Shuffler. This shuffler was configured to ensure that there was a balanced sampling distribution between abnormal and normal EFs, as in the original dataset there was a bias towards normal EFs. Before the new sampler was implemented, the model performed much better on datapoints with a normal ground truth EF than it did with those with an abnormal ground truth EF. However, with the new sampler, the overall error rate decreased and the error rate and DSC on abnormal EFs also improved. In addition, the model was able to reduce its number of false positives.

Table 1. Performance comparison among five different segmentation models.

Segmentation Model	Overall DSC	Systolic DSC	Diastolic DSC	Training Time (hours)	Memory (GB)
Improved MobileNet	0.925	0.919	0.931	3.1	1.1
MobileNet	0.912	0.892	0.925	2.5	0.8
FCN50	0.926	0.911	0.936	5.8	2.7
FCN101	0.926	0.912	0.936	11.0	4.7
DeepLab50	0.905	0.887	0.919	10.4	3.0
DeepLab101	0.928	0.914	0.937	15.0	4.7

Detect End-Diastole and End-Systole Video Frames

The heart has two phases: diastole and systole [16]. The end-diastolic and end-systolic volumes are needed to calculate the ejection fraction (EF). Since echocardiograms are video-based, the video frames that correspond to the end-diastolic and end-systolic need to be detected.

The pixel size of the left ventricle segmentation should be roughly proportional to its actual size. As shown in Figure 3, the pixel size was plotted against time; the apex and valley of each cycle enabled identification of the end-diastole and end-systole, respectively. Multiple end-diastole and end-systole video frames were identified in each video file and included in the subsequent calculations.

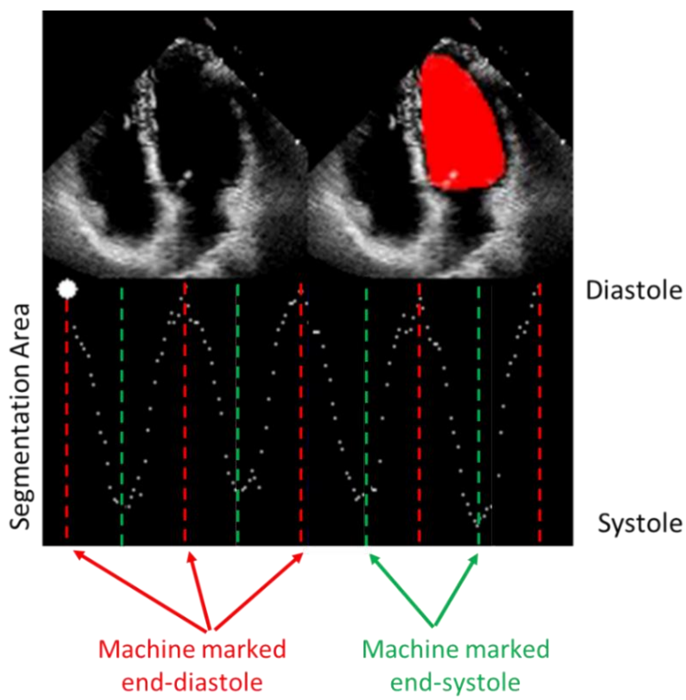

Figure 3. Detecting the end-diastole and end-systole video frames by plotting left ventricle segmentation pixel size against time and then using the peak finding algorithm.

Estimate the End-Diastolic Volume and End-Systolic Volume using Simpson's Method

As shown in Figure 4, once the video frames that correspond to the heart at the end-systole and end-diastole states were identified, the contours of the left ventricles in the relevant frames were drawn using OpenCV2 [17]. Afterwards, the end-diastolic volume (EDV) and end-systolic volume (ESV) were estimated using Simpson's method, which is a way to estimate the 3D volume of an object from just a 2D image. It is a commonly used method in the left ventricle volume calculation [18]. Simpson's method involves the following steps:

1) Measure the length of LV
2) Divide the LV cavity into 20 disk cylinders of equal height
3) The volume of each cylinder is calculated based on the diameter and the height
4) The LV volume is the summed volume of all 20 cylinders

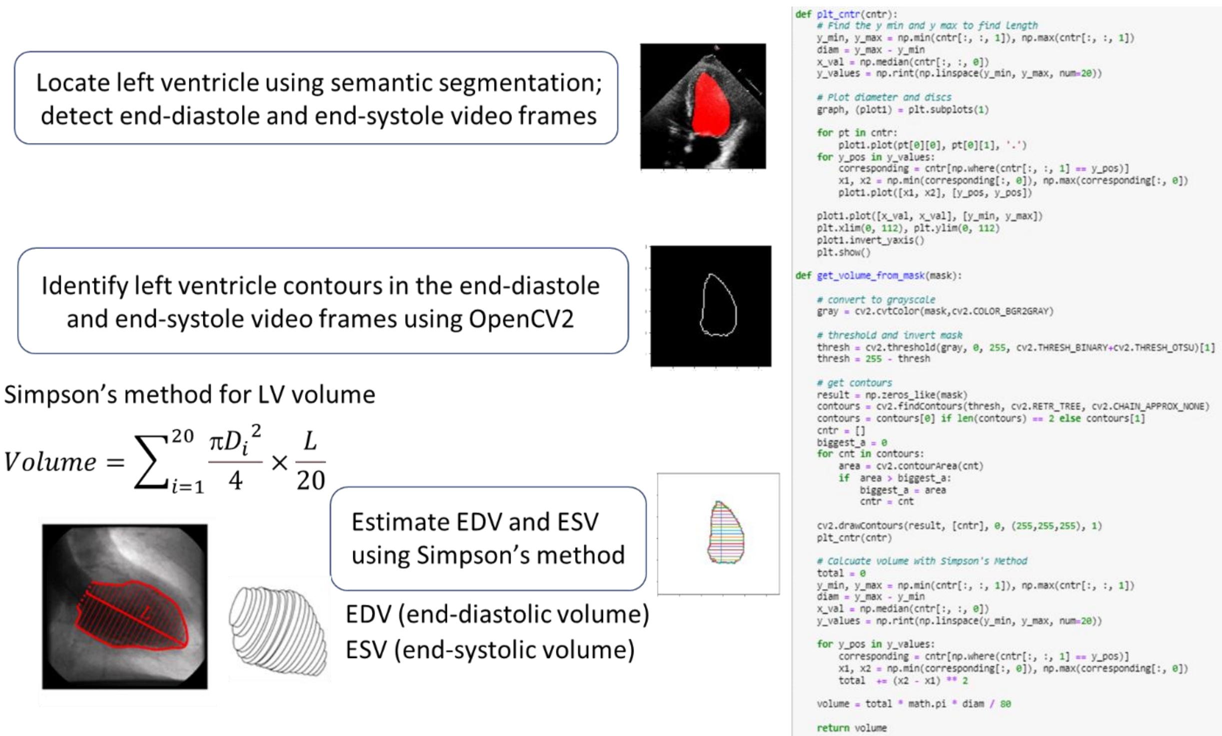

Figure 4. Estimating EDV and ESV using Simpson's method.

Calculate LVEF for Cardiac Assessment

The left ventricle ejection fraction (EF) was calculated using the average end-diastolic volume (EDV) and end-systolic volume (ESV) based on the following equation [19].

$$EF(\%) = \frac{EDV - ESV}{EDV} \times 100\%$$

EF = Ejection fraction, EDV = End-diastolic volume, ESV = End-systolic volume

As shown in Table 2, LVEF-based cardiac function assessment is based on the guidelines provided by American College of Cardiology (ACC) [19]. For example, a high EF of 71% means the heart has a normal function, while a low EF of 44% means moderate dysfunction (Figure 5).

Table 2. Cardiac function classification based on EF as per the American College of Cardiology (Ref. 19).

Cardiac Function	EF
Hyperdynamic	>70%
Normal Function	50-70%
Mild Dysfunction	40-49%
Moderate Dysfunction	30-39%
Severe Dysfunction	< 30%

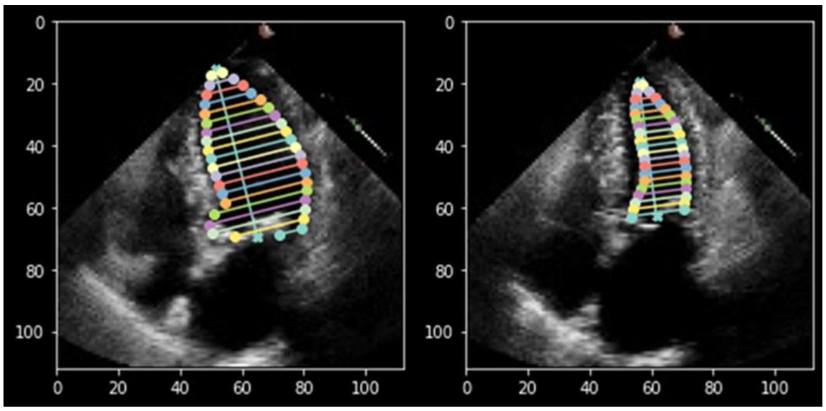

$$EF = \frac{69.28 - 19.89}{69.28} \times 100\% = 71\%$$

High Function (EF > 70%)

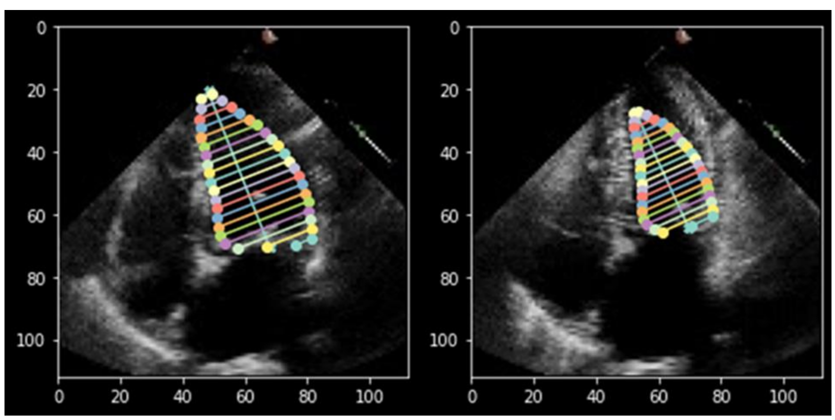

$$EF = \frac{60.72 - 33.70}{60.72} \times 100\% = 44\%$$

Low Function (EF 40-55%)

Figure 5. Using EDV and ESV to calculate EF for cardiac assessment.

Results and Discussion

Model Comparison

The machine predicted EF values were plotted against the ground truth from human experts to determine the correlations using linear regression analysis. The equations defined by linear regression analysis for each segmentation model are shown in Figures 6. In addition, the absolute differences between the machine predictions and the human expert calculations for EF were calculated for each data point. Mean absolute error reflects the average of absolute differences/errors from ground truth the entire dataset as a measurement of the magnitude of errors. Results are shown in Table 3.

The error rates for five original deep learning models ranged from 14-16%, comparable to an expert cardiac sonographer (13.9% error rate) [20], and significantly outperformed qualitative analysis by physicians (~30% error rate) [4]. FCN101 is the best model with a 14% error rate but requires more computing power and a longer (~15 second) analysis time. Compared to the other models, MobileNet has a similar error rate (16%) and DSC (0.91), but significantly faster run time (5 seconds per video) and lower memory requirement (0.8 GB), and thus is the best model for deployment to laptops and phones. After further tuning, Mobile Net was able to exceed FCN101's performance while still having a significantly faster run time and lower memory requirement compared to the other models.

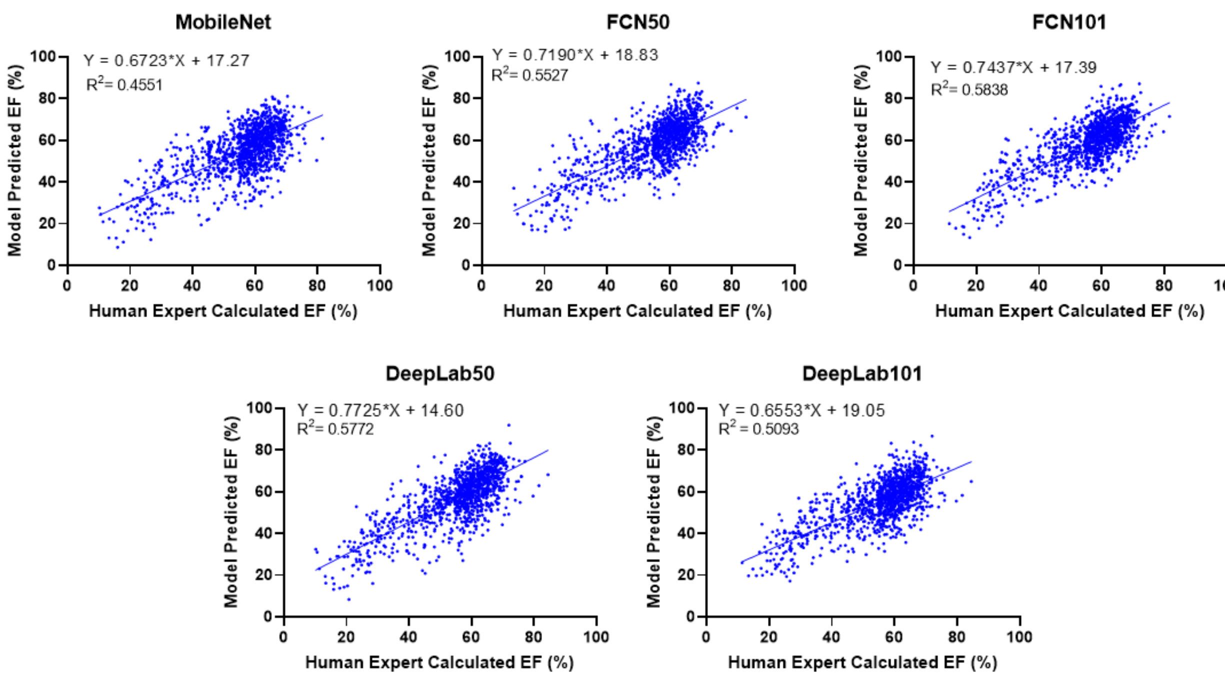

Figure 6. Comparison between model predictions and human calculations.

Table 3. Comparison of error rates between five deep learning models and human experts.

Segmentation Model	Mean Absolute Error (%)
Improved MobileNet	10.2
MobileNet	16.0
FCN50	15.9
FCN101	14.0
DeepLab50	14.2
DeepLab101	14.6
Cardiac Sonographer[20]	13.9
Physician[4]	~30

Videos with the most discordance between the model prediction and human label of ejection fraction were further examined. These outliers were mainly caused by human errors in the ground truth or poor image quality. Since a human sonographer picks only one set of end-diastole and end-systole, they sometimes incorrectly assign the end-diastole and end-systole, leading to errors in the ground truth, as shown in the left panel of Figure 7. In contrast, the machine identifies multiple end-diastoles and end-systoles and uses averaged results for more reliable EF calculation.

There are also cases where echocardiogram videos have poor image quality, which can impact image interpretation, as shown in the right panel of Figure 7. Since it only takes 5-15 seconds to process each echocardiogram video, a future improvement could be using machines to judge the image quality and guide sonographers to acquire better echocardiogram images.

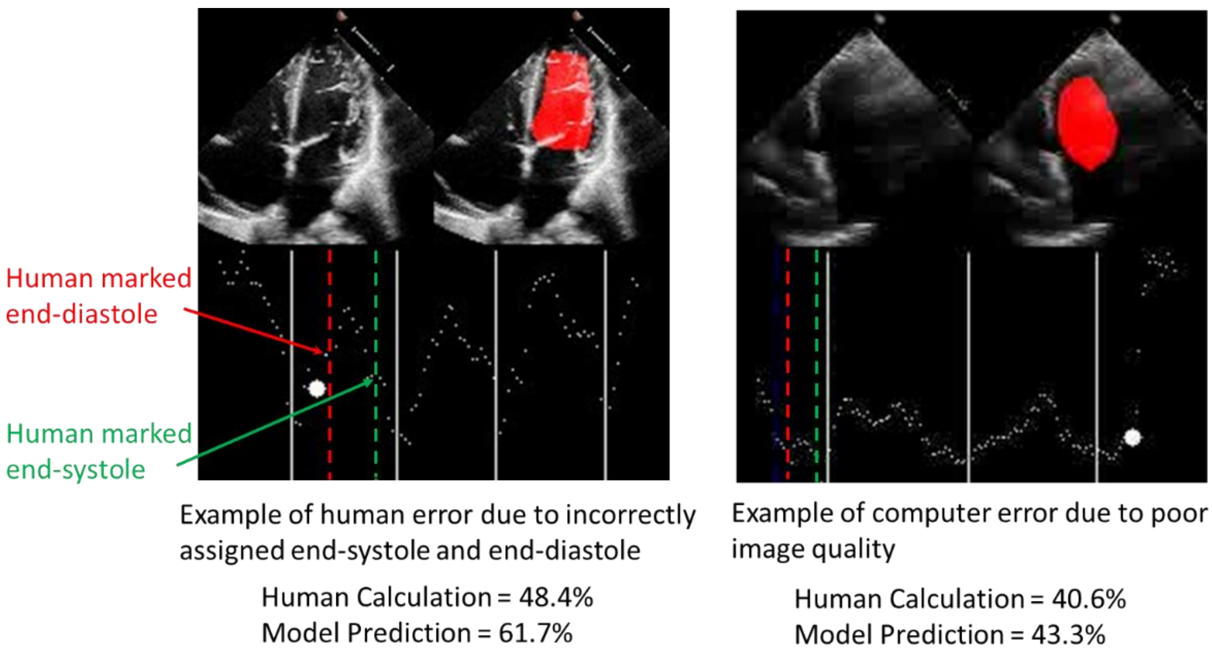

Figure 7. Discordance between model predictions and human calculations.

Model Deployment using Web-based App

To make this deep learning pipeline into a software system accessible by doctors and patients, a web-based app with MobileNet deep learning model was developed using Python and Flask, micro web framework. As shown in Figure 8, the web-based app's graphic user interface allows users to select an echocardiogram video. Once the video is uploaded, the app will display the echocardiogram video while executing the model to produce the left ventricle segmentation, the end-diastole and end-systole video frames, the averaged EDV and ESV estimated using Simpson's method, the calculated EF results, and the cardiac function assessment. The app also displays these results in a stepwise fashion which will further help its interpretability. Using an LG Gram laptop (8th Gen. Intel Core i7 processor and 16GB RAM), these calculations take only ~5 seconds per video. The graphic user interface (GUI) makes the entire workflow transparent, so that it can be readily audited at each step, which is important in a clinical setting.

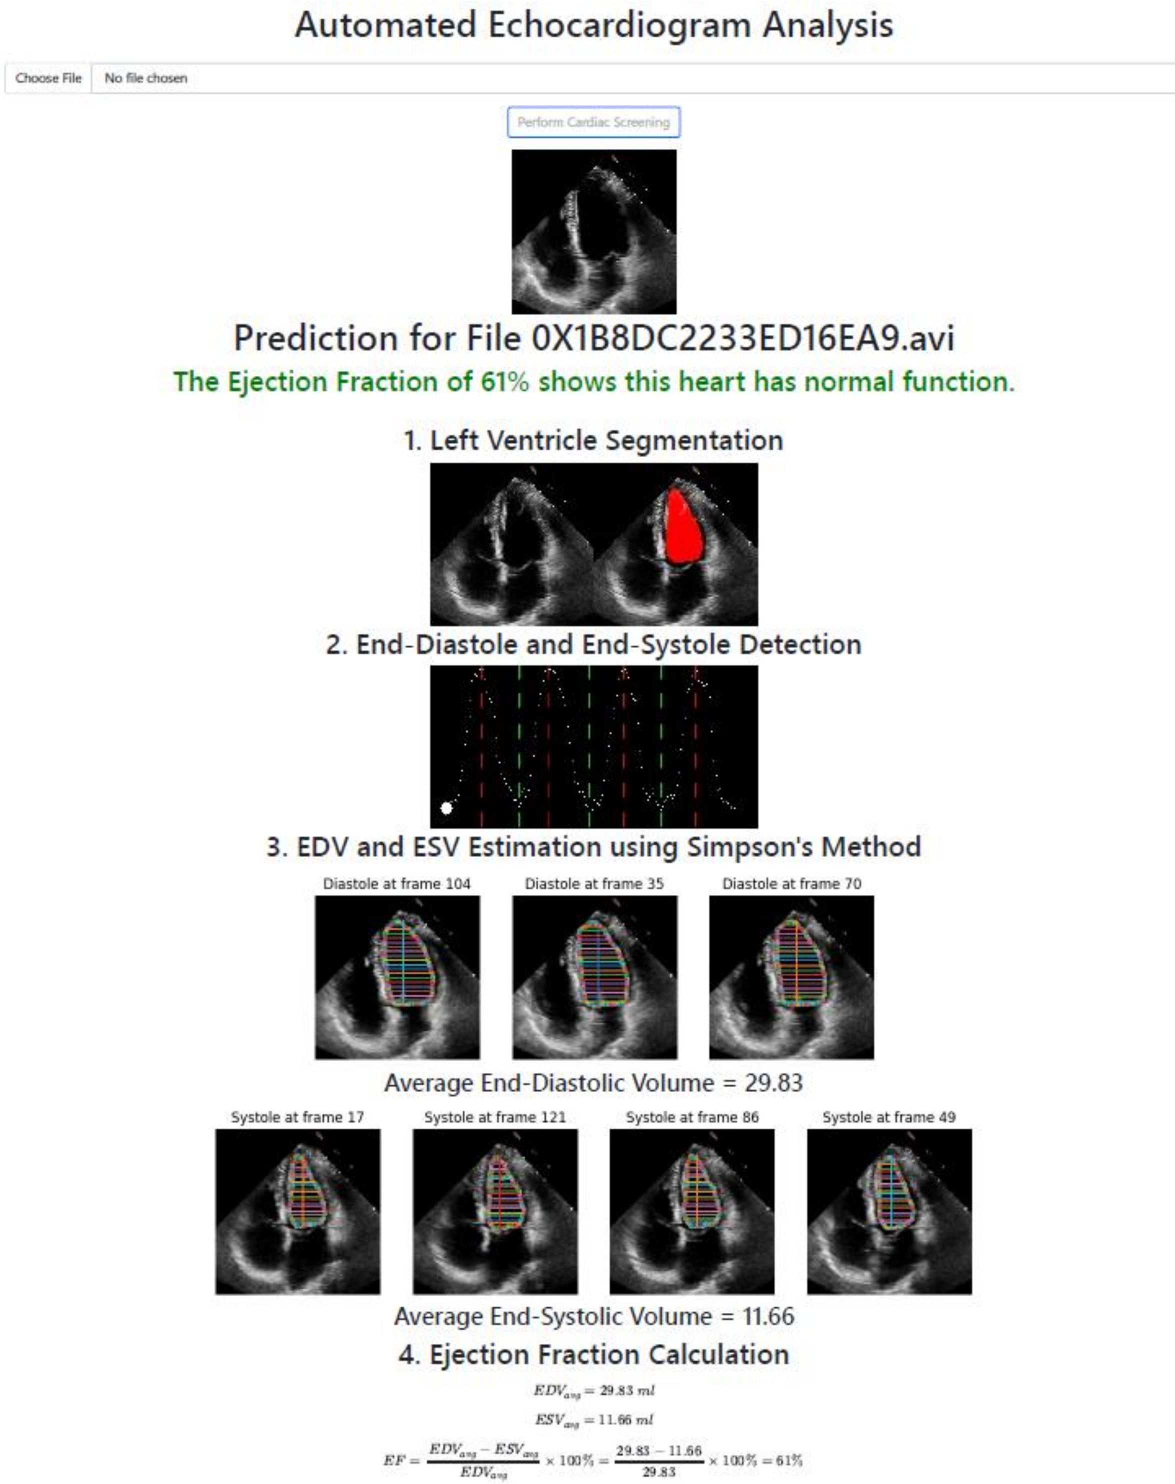

Figure 8. Automated echocardiogram analysis using a web-based app.

As shown in Figure 9, this web-based app was further deployed to a Raspberry Pi. One of the main objectives of this project is to provide access to accurate cardiac screening in rural areas or developing countries, where cardiology expertise is limited, and internet service may not be reliable. Therefore, the device affordability and usability are important. A low-cost Raspberry Pi processor ($75/EA) can meet these needs. To improve the calculation speed, MobileNet model was optimized through recompiling the model with PyTorch JIT, which is an experimental feature that can double the calculation speed. The calculation speed on Raspberry Pi 4 base model (1.5 GHz 64-bit quad core ARM Cortex-A72 processor) is about 1 min per video, understandably slower than a more powerful laptop (8th Gen. Intel Core i7 processor and 16GB RAM), but still 10x faster than manual analysis by an expert cardiac sonographer. Despite the longer processing time, this Raspberry Pi implementation demonstrated that it is possible to deploy web-based implementation onto smartphones, which have much higher computing power than Raspberry Pis. It is also feasible to integrate this deep learning model with a low computing power portable handheld ultrasound device as an integral solution for mobile echocardiogram screening. Work is on-going to deploy the model using Amazon Web Services (AWS) to further expand the model accessibility.

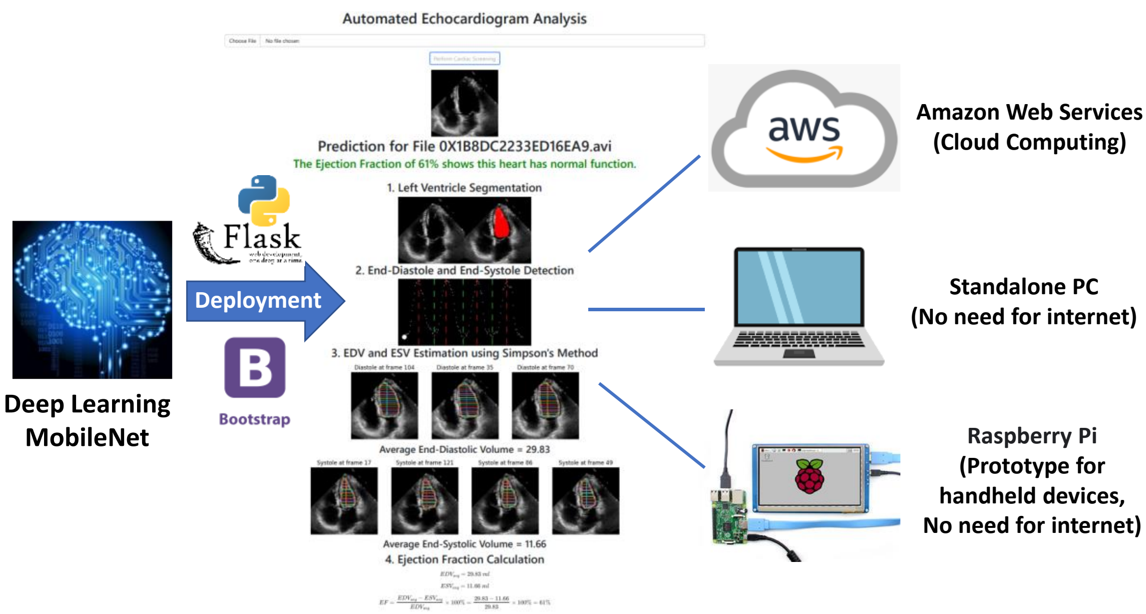

Figure 9. Deployment of MobileNet deep learning model using AWS, standalone PC, and Raspberry Pi for remote access.

Model Deployment using Mobile App

To make this deep learning pipeline even more accessible to doctors and patients, a mobile app was developed using Swift, CoreML, and the Metal Framework. This involved first converting the model to a format that could run on iOS devices and work inside of a Swift program. This conversion involved a transformation from the original PyTorch converted format to a CoreML compatible model using the coremltools library. Then, the model could be run inside of the Swift based mobile application.

In addition to the model, the other steps in the model, such as the systole and diastole frame detection, the Simpson's method-based volume calculation, and EF calculation algorithms all had to be ported from Python. This included the rewrite of functions from libraries like SciPy that had been used in the Python version of the program for peak detection and researching alternative versions for things like finding contours in the ventricular segmentation.

The mobile app's interface allows users to select an echocardiogram video. Once the video is uploaded, the app will display the echocardiogram video while executing the model to calculate EF results and obtain the cardiac function assessment. On an iPhone, these calculations take only ~5 seconds per video.

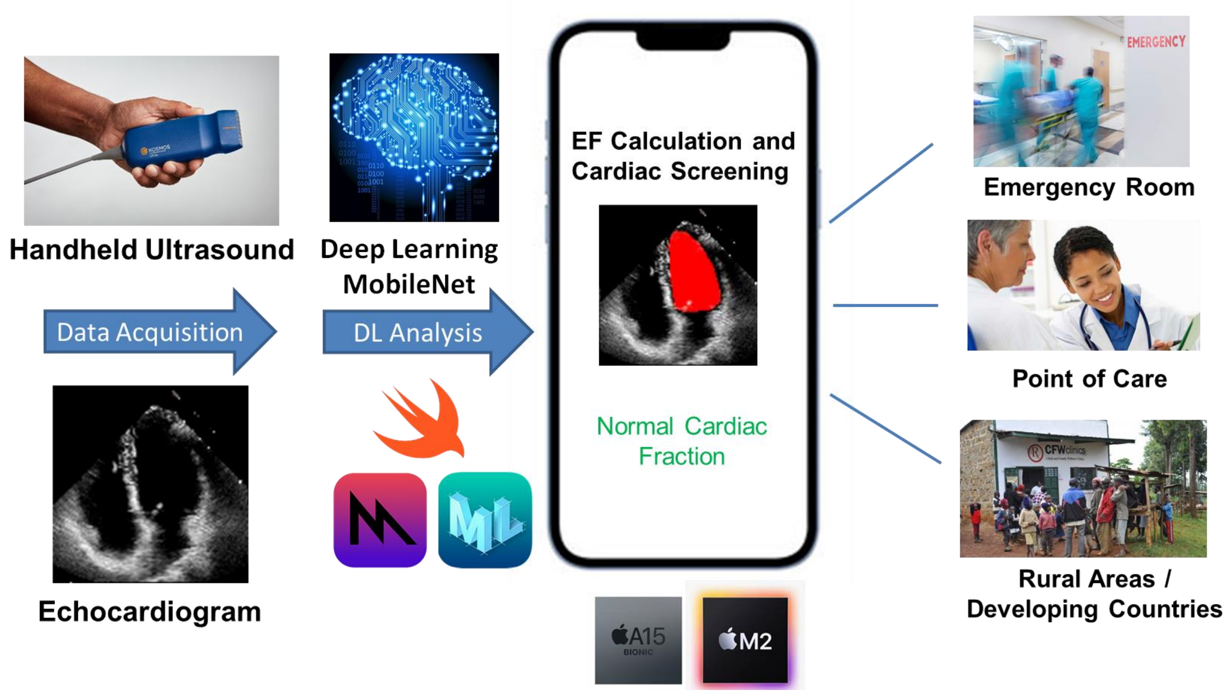

Figure 10. Workflow for MobileNet deep learning model deployed using a mobile app

The workflow for how the mobile app would integrate into healthcare workflows is shown in Figure 10. First, doctors would acquire echocardiogram data using handheld ultrasounds. A relatively new invention, handheld ultrasounds allow for the mobile acquisition of echocardiogram, while also costing a fraction of the price of conventional bulky ultrasound machines. Then, using the acquired data, the mobile app could run analysis in just seconds to deliver a cardiac function prognosis within seconds. This has a multitude of applications, particularly in emergency rooms, points of care, and rural/developing settings, where cardiac expertise is limited and echocardiograms are currently not used or are analyzed qualitatively (a very error-prone method).

Conclusion

Using deep learning, each echocardiogram video can be analyzed in 5-15 seconds with significantly improved accuracy in EF calculation compared to a qualitative analysis by physicians. Results can differentiate heart failure with reduced ejection fraction from normal cardiac function. a MobileNet-based deep learning workflow was deployed using Raspberry Pi and standalone PC for remote access, allowing physicians to upload and analyze echocardiogram videos and obtain EF calculation results within seconds without the need of internet access. Particularly, the successful model deployment using Raspberry Pi demonstrated the feasibility of integrating the deep learning model onto low computing power portable devices such as handhold ultrasound as an integral device for mobile echocardiogram screening. Work is on-going to deploy the model using AWS to further expand the model access.

This automated echocardiogram analysis can dramatically speed up image analysis, reduce the burden on cardiologists, eliminate inter-observer variability, and democratize cardiac care by enabling non-experts to quickly and accurately assess cardiac functions at the point of care, including rural areas, or in developing countries, where cardiology expertise is limited. It could also provide near real-time cardiac analysis data to emergency room doctors for patient triage, or guide sonographers to acquire better echocardiogram images. The model has been deployed to an EchoNet handheld ultrasound device for ongoing clinical validation and testing.

The current method still has limitations: it only uses the Apical 4 Chamber view (A4C) in the EF calculation due to the limitations of the EchoNet dataset. The analysis will be expanded to the Apical 2 Chamber view (A2C) to further improve the EF calculation accuracy.

References

[1] CDC Heart Disease Facts: CDC Heart Disease Statistics and Maps, from https://www.cdc.gov/heartdisease/facts.htm

[2] Cheitlin, M. D., Alpert, J. S., Armstrong, W. F., Aurigemma, G. P., Beller, G. A., Bierman, F. Z., Davidson, T. W., Davis, J. L., Douglas, P. S., & Gillam, L. D. (1997). ACC/AHA Guidelines for the Clinical Application of Echocardiography. A report of the American College of Cardiology/American Heart Association Task Force on Practice Guidelines (Committee on Clinical Application of Echocardiography). Developed in collaboration with the American Society of Echocardiography. Circulation, 95(6), 1686-1744. doi:10.1161/01.cir.95.6.1686

[3] Ulloa Cerna, A. E., Jing, L., Good, C. W., vanMaanen, D. P., Raghunath, S., Suever, J. D., Nevius, C. D., Wehner, G. J., Hartzel, D. N., Leader, J. B., Alsaid, A., Patel, A. A., Kirchner, H. L., Pfeifer, J. M., Carry, B. J., Pattichis, M. S., Haggerty, C. M., & Fornwalt, B. K. (2021). Deep-learning-assisted analysis of echocardiographic videos improves predictions of all-cause mortality. Nat Biomed Eng, 5(6), 546-554. doi:10.1038/s41551-020-00667-9

[4] Jan, M. (2010) Impact of physician training on interpretation of echocardiograms and health care costs. in ASE 2010 Abstract P2-40 2010

[5] St. Luke's review finds almost 30% echocardiograms are misread, from https://archive.jsonline.com/news/health/96945709.html/

[6] Lin, S., Xie, M., Lv, Q., Wang, J., He, L., Wang, B., Li, Y., Xu, L., & Yang, Y. (2020). Misdiagnosis of anomalous origin of the left coronary artery from the pulmonary artery by echocardiography: Single-center experience from China. Echocardiography, 37(1), 104-113. doi:10.1111/echo.14578

[7] Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A., Ciompi, F., Ghafoorian, M., van der Laak, J., van Ginneken, B., & Sánchez, C. I. (2017). A survey on deep learning in medical image analysis. Med Image Anal, 42, 60-88. doi:10.1016/j.media.2017.07.005

[8] Lang, R. M., MD, Badano, L. P., MD, & Mor-Avi, V., PhD. (2015). Recommendations for Cardiac Chamber Quantification by Echocardiography in Adults. JASE, 28(1), 1-53. Retrieved March 1, 2017, from http://asecho.org/wordpress/wp-content/uploads/2015/01/ChamberQuantification2015.pdf

[9] Leclerc, S., Smistad, E., Pedrosa, J., Østvik, A., Cervenansky, F., Espinosa, F., Bernard, O. (2019). Deep Learning for Segmentation Using an Open Large-Scale Dataset in 2D Echocardiography. IEEE Transactions on Medical Imaging, 38(9), 2198-2210. doi:10.1109/TMI.2019.2900516

[10] Ouyang, D., He, B., Ghorbani, A., Lungren, M. P., Ashley, E. A., Liang, D. H., & Zou, J. Y. (2019). EchoNet-Dynamic: a Large New Cardiac Motion Video Data Resource for Medical Machine Learning, from https://stanfordaimi.azurewebsites.net/datasets/834e1cd1-92f7-4268-9daa-d359198b310a

[11] Ouyang, D., He, B., Ghorbani, A., Yuan, N., Ebinger, J., Langlotz, C. P., Zou, J. Y. (2020). Video-based AI for beat-to-beat assessment of cardiac function. Nature, 580(7802), 252-256. doi:10.1038/s41586-020-2145-8

[12] Minaee, S., Boykov, Y. Y., Porikli, F., Plaza, A. J., Kehtarnavaz, N., & Terzopoulos, D. (2021). Image Segmentation Using Deep Learning: A Survey. IEEE Transactions on Pattern Analysis and Machine Intelligence, 1-1. doi:10.1109/TPAMI.2021.3059968

[13] Pytorch Hub for Researchers, from https://pytorch.org/hub/research-models

[14] The COCO-Stuff dataset, from https://github.com/nightrome/cocostuff

[15] Zou, K. H., Warfield, S. K., Bharatha, A., Tempany, C. M., Kaus, M. R., Haker, S. J., Wells, W. M., 3rd, Jolesz, F. A., & Kikinis, R. (2004). Statistical validation of image segmentation quality based on a spatial overlap index. Acad Radiol, 11(2), 178-189. doi:10.1016/s1076-6332(03)00671-8

[16] The Cardiac Cycle, from https://www.thoughtco.com/phases-of-the-cardiac-cycle-anatomy-373240

[17] OpenCV, from https://opencv.org/

[18] Cardiovacular Research & Training Center Cardiac Imaging Research Lab, from https://depts.washington.edu/cvrtc/simplvgm.html

[19] Kosaraju, A., Goyal, A., Grigorova, Y., & Makaryus, A. N. (2022). Left Ventricular Ejection Fraction. StatPearls Publishing

[20] Malm, S., Frigstad, S., Sagberg, E., Larsson, H., & Skjaerpe, T. (2004). Accurate and reproducible measurement of left ventricular volume and ejection fraction by contrast echocardiography: a comparison with magnetic resonance imaging. J Am Coll Cardiol, 44(5), 1030-1035. doi:10.1016/j.jacc.2004.05.068