

DETECTION OF CHRONIC HEART FAILURE USING ML & DL

SHAIK SHAJUMA*1, K. PURNA CHANDRA RAO*2, MINDALA CHANDRIKA*3

* 1,3 B. Tech Students, *2 Associate Professor
Dept. of Computer Science and Engineering,
RISE Krishna Sai Prakasam Group of Institutions

ABSTRACT:

Chronic heart failure (CHF) affects over 26 million of people worldwide, and its incidence is increasing by 2% annually. Despite the significant burden that CHF poses and despite the ubiquity of sensors in our lives, methods for automatically detecting CHF are surprisingly scarce, even in the research community. We present a method for CHF detection based on heart sounds. The method combines classic Machine-Learning (ML) and end-to-end Deep Learning (DL). The classic ML learns from expert features, and the DL learns from a spectro-temporal representation of the signal. The method was evaluated on recordings from 947 subjects from six publicly available datasets and one CHF dataset that was collected for this study. Using the same evaluation method as a recent PhysoNet challenge, the proposed method achieved a score of 89.3, which is 9.1 higher than the challenge's baseline method. The method's aggregated accuracy is 92.9% (error of 7.1%); while the experimental results are not directly comparable, this error rate is relatively close to the percentage of recordings labeled as “unknown” by experts (9.7%). Finally, we identified 15 expert features that are useful for building ML models to differentiate between CHF phases (i.e., in the decompensated phase during hospitalization and in the recompensated phase) with an accuracy of 93.2%. The proposed method shows promising results both for the distinction of recordings between healthy subjects and patients and for the detection of different CHF phases. This may lead to the easier identification of new CHF patients and the development of home-based CHF monitors for avoiding hospitalizations.

1. INTRODUCTION:

Chronic heart failure (CHF) is a chronic, progressive condition underscored by the heart's inability to supply enough perfusion to target tissues and organs at the physiological filling pressures to meet their metabolic demands. CHF has reached epidemic proportions in the population, as its incidence is increasing by 2% annually. In the developed world, CHF affects 1-2% of the total population and 10% of people older than 65 years. Currently, the diagnosis and

treatment of CHF uses approximately 2% of the annual healthcare budget. In absolute terms, the USA spent approximately 35 billion USD to treat CHF in 2018 alone, and the costs are expected to double in the next 10 years . Despite the progress in medical- and device-based treatment approaches in the last decades, the overall prognosis of CHF is still dismal, as 5-year survival rate of this population is only approximately 50%. In the typical clinical course of CHF, we observe alternating episodes of compensated phases, when the patient feels well and does not display symptoms and signs of fluid overload, and decompensated phases, when symptoms and signs of systemic fluid overload (such as breathlessness, orthopnea, peripheral edema, liver congestion, pulmonary edema) can easily be observed. During the latter episodes, patients often require hospital admission to receive treatment with intravenous medications (diuretics, inotropes) to achieve a successful negative fluid balance and return to the compensation state. Early detection of HF worsening would allow a treating physician to adjust the patient's medical management on an outpatient basis in a timely manner and thus avoid the need for a hospital admission. Currently, an experienced physician can detect the worsening of HF by examining the patient and by characteristic changes in the patient's heart failure biomarkers, which are determined from the patient's blood. Unfortunately, clinical worsening of a CHF patient likely means that we are already dealing with a fully developed CHF episode that will most likely require a hospital admission. Additionally, in some patients, characteristic changes in heart sounds can accompany heart failure worsening and can be heard using phonocardiography. An example of a phonocardiogram (PCG) recording of a healthy subject is presented. In healthy subjects, 2 heart sounds are typically heard (called S1 and S2). S1 is caused by the closure of the mitral valve and ventricular wall in the early systole, S2 is caused by the closure of the aortic and pulmonary valves at the beginning of the diastole. Here, the interval between S1 and S2 is called systole, i.e., the contraction phase of the cardiac cycle, and the interval between S2 and S1 is called diastole, i.e., the relaxation phase of the cardiac cycle. Additional heart sounds (such as S3 and S4) can be heard in certain cardiac conditions and are never regarded as normal. In the case of CHF (in the course of decompensation), we can often hear a third sound (S3) that typically appears 0.1-0.2 s after the second sound, i.e., S2. Recently, it has been demonstrated that some physiological parameters, such as the occurrence of additional heart sounds or increased blood pressure in the pulmonary circulation, already start to appear several weeks before the CHF

patient develops a clinically evident decompensation episode. This is also an important therapeutic window where outpatient-based treatment interventions can reverse CHF deterioration and return the patient to the compensated state without the need for a hospital admission. In recent years, many studies have proposed Machine Learning (ML) approaches for the automatic detection of different heart conditions using PCG signals recorded with a digital stethoscope. Nevertheless, methods that explicitly focus on CHF detection are quite scarce. The typical ML pipeline for the detection of different heart conditions is as follows: segmentation of the signals by detecting the “typical” heart sounds (i.e., S1 and S2), denoising of the signals, extracting individual frequency-domain and time-domain features, and learning a feature-based ML model (e.g., using ML algorithms, such as Random Forest or Support Vector Machine - SVM) that is capable of classifying healthy vs. unhealthy sounds. Most of the features currently used are based on medical and audio/signal analysis knowledge. However, a PCG recording that sounds unhealthy to one expert may sound healthy to another one; therefore, doctors never diagnose a CHF patient using only heart sounds, but rather use a holistic view of the patient instead (i.e., extensive medical history, blood pressure, laboratory tests, etc.). This uncertainty is one reason why 9.7% of the recordings in the recent PhysioNet cardiology challenge were actually labeled as “unknown” by experts, while the rest of the recordings were labeled as healthy or unhealthy. The recent advancements in Deep Learning (DL) suggest that end-to-end learning (i.e., ML models that learn directly from the raw data and no features are needed) can outperform the classic, feature-based ML. For example, DL has achieved breakthrough performance in tasks such as pattern recognition problems, image processing, natural language processing, speech and audio processing, and sensor data processing. For CHF detection, a successful combination of classic ML and end-to-end DL can outperform each single approach. The classic ML approach learns from a large body of expert-defined features, and the DL approach learns both from a time-domain (the raw PCG signal) representation of the signal and a temporal domain representation (the spectrogram) of the signal. This approach was successful in our previous study of human activity recognition from smartphone sensor data. In addition to distinguishing the CHF patients and healthy individuals, we focus on detecting the CHF state (compensated vs. decompensated) based on the analysis of heart sound recordings. Our work builds upon the initial studies, where we demonstrated that it is possible to distinguish between

healthy individuals and patients in a decompensated CHF episode using a stack of machinelearning classifiers and expert features, showing promising results on a limited dataset . We expand upon this approach using a considerably larger patient dataset, including six additional PhysioNet datasets, and an improved ML method that uses end-to-end DL. Furthermore, we investigate the differences in the heart sounds during the transition between the decompensated and recompensated states of CHF, with the aim of developing personalized monitoring models. Early detection of the worsening of CHF has the potential to reduce hospitalizations due to the worsening of the condition, which both improves the quality of life of patients and decreases the financial and logistic burden on the patient and the health system

1.1 Objective of the project:

Chronic heart failure (CHF) affects over 26 million of people worldwide, and its incidence is increasing by 2% annually. Despite the significant burden that CHF poses and despite the ubiquity of sensors in our lives, methods for automatically detecting CHF are surprisingly scarce, even in the research community. We present a method for CHF detection based on heart sounds. The method combines classic Machine-Learning (ML) and end-to-end Deep Learning (DL). The classic ML learns from expert features, and the DL learns from a spectro-temporal representation of the signal. The method was evaluated on recordings from 947 subjects from six publicly available datasets and one CHF dataset that was collected for this study. Using the same evaluation method as a recent PhysioNet challenge, the proposed method achieved a score of 89.3, which is 9.1 higher than the challenge's baseline method. The method's aggregated accuracy is 92.9% (error of 7.1%); while the experimental results are not directly comparable, this error rate is relatively close to the percentage of recordings labeled as "unknown" by experts (9.7%). Finally, we identified 15 expert features that are useful for building ML models to differentiate between CHF phases (i.e., in the decompensated phase during hospitalization and in the recompensated phase) with an accuracy of 93.2%. The proposed method shows promising results both for the distinction of recordings between healthy subjects and patients and for the detection of different CHF phases. This may lead to the easier identification of new CHF patients and the development of home-based CHF monitors for avoiding hospitalizations.

2. LITERATURE SURVEY:

“Chronic heart failure detection from heart sounds using a stack of machine-learning classifiers,”

Chronic heart failure represents a global pandemic, currently affecting over 26 million of patients worldwide. It is a major contributor in the death rate of patients with cardiovascular diseases and results in more than 1 million hospitalizations annually in Europe and North America. Methods for chronic heart failure detection can be utilized to act preventive, improve early diagnosis and avoid hospitalizations or even life-threatening situations, thus highly enhance the quality of patient’s life. In this paper, we present a machine-learning method for chronic heart failure detection from heart sounds. The method consists of: filtering, segmentation, feature extraction and machine learning. The method was tested with a leave-one-subject-out evaluation technique on data from 122 subjects, gathered in the study. The method achieved 96% accuracy, outperforming a majority classifier for 15 percentage points. More specifically, it detects (recalls) 87% of the chronic heart failure subjects with a precision of 87%. The study confirmed that advanced machine learning applied on real-life sounds recorded with an unobtrusive digital stethoscope can be used for chronic heart failure detection.

“Classification of normal/abnormal heart sound recordings: the PhysioNet/Computing in Cardiology Challenge 2016,”

In the past few decades heart sound signals (i.e., phono-cardiograms or PCGs) have been widely studied. Automated heart sound segmentation and classification techniques have the potential to screen for pathologies in a variety of clinical applications. However, comparative analyses of algorithms in the literature have been hindered by the lack of a large and open database of heart sound recordings. The PhysioNet/Computing in Cardiology (CinC) Challenge 2016 addresses this issue by assembling the largest public heart sound database, aggregated from eight sources obtained by seven independent research groups around the world. The database includes 4,430 recordings taken from 1,072 subjects, totalling 233,512 heart sounds collected from both healthy subjects and patients with a variety of conditions such as heart valve disease and coronary artery disease. These recordings were collected using heterogeneous equipment in both clinical and nonclinical (such as in-home visits). The length of recording varied from several seconds to several minutes. Additional data provided include subject demographics (age and gender),

recording information (number per patient, body location, and length of recording), synchronously recorded signals (such as ECG), sampling frequency and sensor type used. Participants were asked to classify recordings as normal, abnormal, or not possible to evaluate (noisy/uncertain). The overall score for an entry was based on a weighted sensitivity and specificity score with respect to manual expert annotations. A brief description of a baseline classification method is provided, including a description of open source code, which has been provided in association with the Challenge. The open source code provided a score of 0.71 (Se=0.65 Sp=0.76). During the official phase of the competition, a total of 48 teams submitted 348 open source entries, with a highest score of 0.86 (Se=0.94 Sp=0.78).

"Speed up deep neural network based pedestrian detection by sharing features across multi-scale models,"

Deep neural networks (DNNs) have now demonstrated state-of-the-art detection performance on pedestrian datasets. However, because of their high computational complexity, detection efficiency is still a frustrating problem even with the help of Graphics Processing Units (GPUs). To improve detection efficiency, this paper proposes to share features across a group of DNNs that correspond to pedestrian models of different sizes. By sharing features, the computational burden for extracting features from an image pyramid can be significantly reduced. Simultaneously, we can detect pedestrians of several different scales on one single layer of an image pyramid. Furthermore, the improvement of detection efficiency is achieved with negligible loss of detection accuracy. Experimental results demonstrate the robustness and efficiency of the proposed algorithm.

"ImageNet classification with deep convolutional neural networks,"

We trained a large, deep convolutional neural network to classify the 1.3 million high-resolution images in the LSVRC-2010 ImageNet training set into the 1000 different classes. On the test data, we achieved top-1 and top-5 error rates of 39.7% and 18.9% which is considerably better than the previous state-of-the-art results. The neural network, which has 60 million parameters and 500,000 neurons, consists of five convolutional layers, some of which are followed by max-pooling layers, and two globally connected layers with a final 1000-way softmax. To make

training faster, we used non-saturating neurons and a very efficient GPU implementation of convolutional nets. To reduce overfitting in the globally connected layers we employed a new regularization method that proved to be very effective.

“Inception-v4, inception-ResNet and the impact of residual connections on learning,”

Very deep convolutional networks have been central to the largest advances in image recognition performance in recent years. One example is the Inception architecture that has been shown to achieve very good performance at relatively low computational cost. Recently, the introduction of residual connections in conjunction with a more traditional architecture has yielded state-of-the-art performance in the 2015 ILSVRC challenge; its performance was similar to the latest generation Inception-v3 network. This raises the question: Are there any benefits to combining Inception architectures with residual connections? Here we give clear empirical evidence that training with residual connections accelerates the training of Inception networks significantly. There is also some evidence of residual Inception networks outperforming similarly expensive Inception networks without residual connections by a thin margin. We also present several new streamlined architectures for both residual and non-residual Inception networks. These variations improve the single-frame recognition performance on the ILSVRC 2012 classification task significantly. We further demonstrate how proper activation scaling stabilizes the training of very wide residual Inception networks. With an ensemble of three residual and one Inception-v4 networks, we achieve 3.08% top-5 error on the test set of the ImageNet classification (CLS) challenge.

“Recent trends in deep learning based natural language processing,”

Deep learning methods employ multiple processing layers to learn hierarchical representations of data, and have produced state-of-the-art results in many domains. Recently, a variety of model designs and methods have blossomed in the context of natural language processing (NLP). In this paper, we review significant deep learning related models and methods that have been employed for numerous NLP tasks and provide a walk-through of their evolution. We also summarize, compare and contrast the various models and put forward a detailed understanding of the past, present and future of deep learning in NLP.

“A neural probabilistic language model,”

A goal of statistical language modeling is to learn the joint probability function of sequences of words in a language. This is intrinsically difficult because of the curse of dimensionality: a word sequence on which the model will be tested is likely to be different from all the word sequences seen during training. Traditional but very successful approaches based on n-grams obtain generalization by concatenating very short overlapping sequences seen in the training set. We propose to fight the curse of dimensionality by learning a distributed representation for words which allows each training sentence to inform the model about an exponential number of semantically neighboring sentences. The model learns simultaneously (1) a distributed representation for each word along with (2) the probability function for word sequences, expressed in terms of these representations. Generalization is obtained because a sequence of words that has never been seen before gets high probability if it is made of words that are similar (in the sense of having a nearby representation) to words forming an already seen sentence. Training such large models (with millions of parameters) within a reasonable time is itself a significant challenge. We report on experiments using neural networks for the probability function, showing on two text corpora that the proposed approach significantly improves on state-of-the-art n-gram models, and that the proposed approach allows to take advantage of longer contexts.

“Recognition of echolalic autistic child vocalisations utilising convolutional recurrent neural networks,”

Autism spectrum conditions (ASC) are a set of neurodevelopmental conditions partly characterised by difficulties with communication. Individuals with ASC can show a variety of atypical speech behaviours, including echolalia or the ‘echoing’ of another’s speech. We herein introduce a new dataset of 15 Serbian ASC children in a human-robot interaction scenario, annotated for the presence of echolalia amongst other ASC vocal behaviours. From this, we propose a four-class classification problem and investigate the suitability of applying a 2D convolutional neural network augmented with a recurrent neural network with bidirectional long short-term memory cells to solve the proposed task of echolalia recognition. In this approach, log Mel-spectrograms are first generated from the audio recordings and then fed as input into the

convolutional layers to extract high-level spectral features. The subsequent recurrent layers are applied to learn the long-term temporal context from the obtained features. Finally, we use a feed forward neural network with softmax activation to classify the dataset. To evaluate the performance of our deep learning approach, we use leave-onesubject-out cross-validation. Key results presented indicate the suitability of our approach by achieving a classification accuracy of 83.5 % unweighted average recall.

“Bag-of-Deep-Features: Noise-robust deep feature representations for audio analysis,”

In the era of deep learning, research into the classification of various components of the acoustic environment, especially in-the-wild recordings, is gaining in popularity. This is due in part to the increasing computational capacities and the expanding amount of real-world data available on social multimedia. However, the noisy nature of this data can add an additional complexity to the already complex deep learning systems. Herein, we tackle this issue by quantising deep feature representations of various in-the-wild audio data sets. The aim of this paper is twofold: 1) to assess the feasibility of the proposed feature quantisation task, and 2) to compare the efficacy of various feature spaces extracted from different fully connected deep neural networks to classify six real-world audio corpora. For the classification, we extract two feature sets: i) DEEP SPECTRUM features which are derived from forwarding the visual representations of the audio instances, in particular mel-spectrograms through very deep task-independent pre-trained Convolutional Neural Networks (CNNs), and ii) Bag-of-Deep-Features (BODF) which is the quantisation of the DEEP SPECTRUM features. Using BODF, we show the suitability of quantising the deep representations for noisy in-the-wild audio data. Finally, we analyse the effect of early and late fusion of the CNN features and models on the classification results.

“Deep affect recognition from R-R intervals,”

Affect recognition is an important task in ubiquitous computing, in particular in health and human-computer interaction. In the former, it contributes to the timely detection and treatment of emotional and mental disorders, and in the latter, it enables indigenous interaction and enhanced user experience. We present an inter-domain study for affect recognition on seven different datasets, recorded with six different sensors, three different sensor placements, 211 subjects and

nearly 1000 hours of labelled data. The datasets are processed and translated into a common spectro-temporal space. The data represented in the common spectro-temporal space is used to train a deep neural network (DNN) for arousal recognition that benefits from the large amounts of data even when the data are heterogeneous (i.e., different sensors and different datasets). The DNN approach outperforms the classical machine-learning approaches in six out of seven datasets

“Learning deep physiological models of affect,”

Feature extraction and feature selection are crucial phases in the process of affective modeling. Both, however, incorporate substantial limitations that hinder the development of reliable and accurate models of affect. For the purpose of modeling affect manifested through physiology, this paper builds on recent advances in machine learning with deep learning (DL) approaches. The efficiency of DL algorithms that train artificial neural network models is tested and compared against standard feature extraction and selection approaches followed in the literature. Results on a game data corpus - containing players' physiological signals (i.e., skin conductance and blood volume pulse) and subjective self-reports of affect - reveal that DL outperforms manual ad-hoc feature extraction as it yields significantly more accurate affective models. Moreover, it appears that DL meets and even outperforms affective models that are boosted by automatic feature selection, for several of the scenarios examined. As the DL method is generic and applicable to any affective modeling task, the key findings of the paper suggest that ad-hoc feature extraction and selection - to a lesser degree - could be bypassed.

3. SYSTEM ANALYSIS

SYSTEM ARCHITECTURE

3.1 Existing System

Chronic heart failure (CHF) affects over 26 million of people worldwide, and its incidence is increasing by 2% annually. Despite the significant burden that CHF poses and despite the ubiquity of sensors in our lives, methods for automatically detecting CHF are surprisingly scarce, even in the research community

Disadvantages of Existing System:

- Less Accuracy
- A soft first heart sound is present in congestive heart failure or with prolonged atrioventricular (AV) conduction

3.2 Proposed System

Chronic heart failure (CHF) is a chronic, progressive condition underscored by the heart's inability to supply enough perfusion to target tissues and organs at the physiological filling pressures to meet their metabolic demands [1]. CHF has reached epidemic proportions in the population, as its incidence is increasing by 2% annually. In the developed world, CHF affects 1-2% of the total population and 10% of people older than 65 years. Currently, the diagnosis and treatment of CHF uses approximately 2% of the annual healthcare budget

Advantages of Proposed System:

- High Accuracy.
- For emergency department patients with shortness of breath and a risk of heart failure, physicians usually grab one thing first: a stethoscope.
- It allows them to hear the S3, an abnormal third sound in the heart's rhythm strongly associated with cardiac disease and heart failure.

Modules Information:

To implement this project we have designed following modules

- 1) Upload Physionet Dataset: using this module we will upload dataset to application
- 2) Dataset Preprocessing: using this module we will extract audio recording features and systolic and diastolic features from dataset and then normalize values
- 3) Run ML Segmented Model with FE & FS: using this module we will extract and select systolic and diastolic features from dataset and then train with Random Forest Classic ML model and then apply test data to calculate its prediction accuracy

- 4) Run DL Model on Raw Features: using this module we will extract RAW features from recording and then train with deep learning model and then this model will be applied on test data to calculate its accuracy
- 5) Run Recording ML Model: using this module we will extract features from Classic ML model and deep learning model and then retrain with 3rd classifier to get its prediction accuracy
- 6) Predict CHF from Test Sound: using this module we will upload Test Heart Sound file and then classifier model will predict whether given recording file is Normal or Abnormal


4. SCREENSHOTS:

Due to chronic heart failure many peoples are losing their lives worldwide and to reduce this lives lost we need to have expert physicians and sometime if such experts not available then it's difficult to save life and to overcome from such issue author of this paper is combining different algorithms such as Classic Random Forest and End-End Deep Learning model and then extracting features from both algorithms and then retraining with Random forest by taking AVERAGE Aggregate Recording features from Classic ML and end - end deep learning models. Average Aggregate Recording model giving better accuracy compare to other algorithms.

In propose paper author is using heart sound dataset from PHYSIONET website and this dataset contains PCG signals data and we are extracting systolic and diastolic features from this PCG signals and training with Classic ML algorithms and then PCG recording voice data will get trained with deep learning algorithm.

ML cannot train on RAW features so we are extracting systolic and diastolic features from PCG RAW data and training with Classic ML and then Raw features get trained with Deep learning. From both models we will extract average recordings and then retrain with 3rd classifier which will give more accuracy.

Below is the dataset screen used in this project




In above screen we have 3 files where .hea file contains class label as Normal or Abnormal and .dat file contains PCG signals and .wav file contains heart sound recording and by using all files we will train all algorithms

To run project double click on 'run.bat' file to get below screen



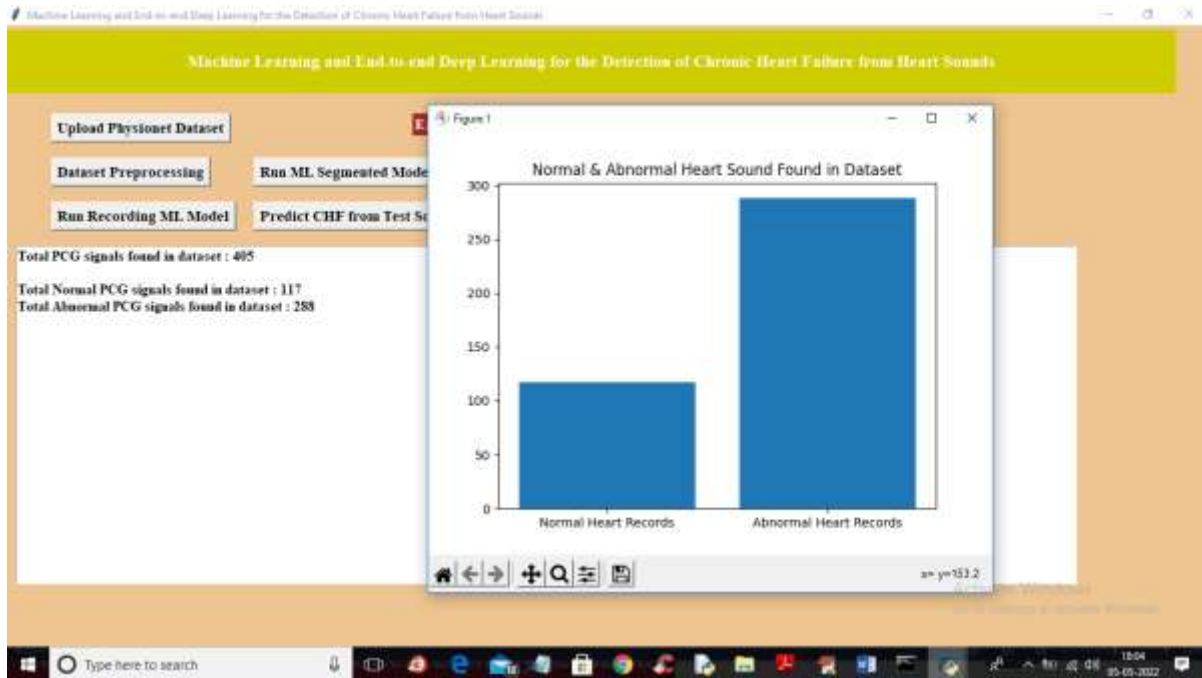
In above screen click on 'Upload Physionet Dataset' button to upload dataset



In above screen selecting and uploading 'Dataset' folder and then click on 'Select Folder' button to load dataset and to get below output



In above screen dataset loaded and now click on 'Dataset Preprocessing' button to read all dataset file and then extract features from it



In above screen we can see dataset contains 405 heart sound files from 405 different person and 117 are the Normal sound and 288 are abnormal and in graph x-axis represents normal or abnormal and y-axis represents number of persons for normal or abnormal. Now close above graph and then click on 'Run ML Segmented Model with FE & FS' button to train Classic ML segmented model on above dataset and get below output

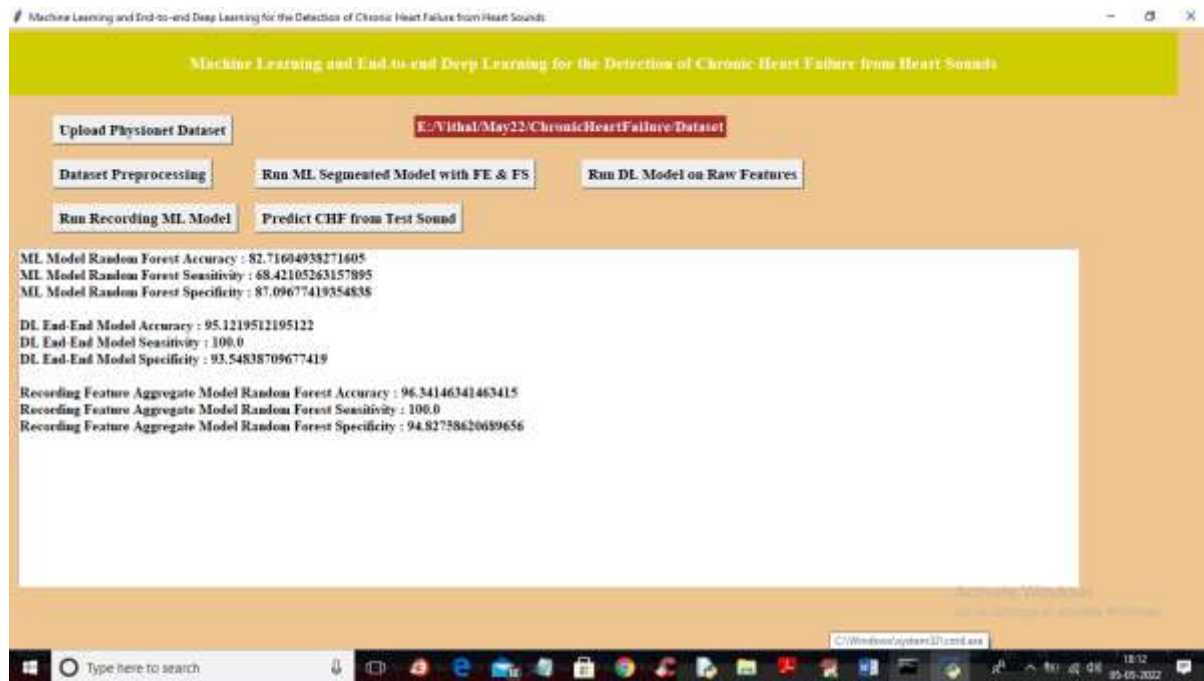


In above screen with Classic ML we got 90% accuracy and now click on 'Run DL Model on Raw Features' to get below output

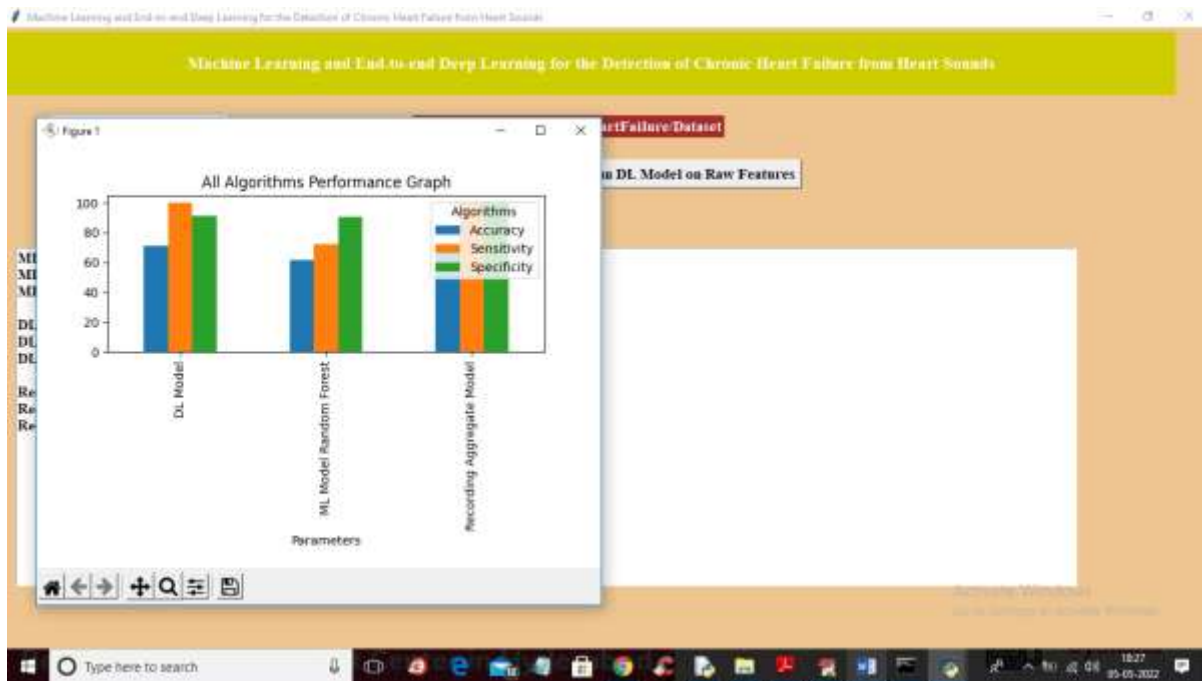


In above screen with DL model we got 93% accuracy and in graph x-axis represents epoch or iterations and y-axis represents accuracy or loss values and green line represents accuracy and


blue line represents LOSS and we can see with each increasing epoch accuracy got increase and loss got decrease and now close above graph and then click on ‘Run Recording Model’ button to get below output



In above screen with recording model we got 96% accuracy and we can see all algorithms performance graph in below screen




In above graph x-axis represents algorithm names and y-axis represents accuracy, sensitivity and specificity and in all algorithms Recording model has got high accuracy. Now close above graph and then click on 'Predict CHF from Test Sound' button to upload test sound file and get predicted output as Normal or Abnormal



In above screen selecting and uploading '1.wav' file and then click on 'Open' button to get below output



In above screen uploaded heart sound file predicted as ABNORMAL and similarly you can upload other files and test



For 2.wav' file below is the output



5. CONCLUSION:

In this paper, we presented a novel method for CHF detection from PCG audio recordings. The method combines classic ML and end-to-end DL. The classical ML learns from a large body of expert-defined features and the DL learns both from the time-domain (i.e., the raw PCG signal) representation of the signal and the spectral representation of the signal. We evaluated the method on our own dataset for CHF detection and additionally on six publicly available PhysioNet datasets used for the recent PhysioNet Cardiology Challenge. The challenge datasets allowed us to extensively evaluate the performance of the method on similar domains. The evaluation results on all the datasets showed that, compared to the challenge baseline methods, our method achieves the best performance (see the PhysioNet experiments section). The facts that most of these datasets are labeled for different types of heart-related conditions and that the PCG audio is recorded from a different body position in most of the datasets (e.g., aortic area, pulmonic area, tricuspid area, and mitral area) strongly indicate that the proposed method is quite robust and that it is useful for detecting different types of heart-sound classification problems and not just for CHF detection, as long as domain-specific labeled data are provided. Finally, we

extended the study beyond the typical healthy vs. patient classification and explored personalized models for detecting different CHF phases, i.e., the recompensated phase (i.e., when the patient feels well) and the decompensated phase (i.e., when the patient needs medical attention). We identified 15 features that have different distributions depending on the phase. By using just two of these features, we were able to build a simple and transparent decision tree classifier (see Fig. 3) that is capable of distinguishing between the recompensated and the decompensated phases with an accuracy of 93.2%, calculated using a LOSO evaluation. While we are aware that there is a risk of overfitting in these final experiments, especially since the dataset contains only 44 samples, we believe that these results are very encouraging and represent a solid base for further development of personalized models. To the best of our knowledge, this is the first study to address such a problem.

REFERENCES:

- [1] M. Gjoreski et al., "Chronic heart failure detection from heart sounds using a stack of machine-learning classifiers," in 2017 International Conference on Intelligent Environments (IE). IEEE, 2017, pp. 14-19.
- [2] J. Voigt et al., "A reevaluation of the costs of heart failure and its implications for allocation of health resources in the United States," *Clinical cardiology*, vol. 37, no. 5, pp. 312-321, 2014.
- [3] G. D. Clifford et al., "Classification of normal/abnormal heart sound recordings: the PhysioNet/Computing in Cardiology Challenge 2016," in 2016 Computing in Cardiology Conference (CinC). IEEE, 2016, pp. 609-612.
- [4] X. Jiang, Y. Pang, X. Li, and J. Pan, "Speed up deep neural network based pedestrian detection by sharing features across multi-scale models," *Neurocomputing*, vol. 185, pp. 163-170, 2016.
- [5] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in International Conference on Neural Information Processing Systems, 2012, pp. 1097-1105. [6] C. Szegedy et al., "Inception-v4, inception-ResNet and the impact of residual connections on learning," in Thirty-First AAAI Conference on Artificial Intelligence, 2017.

- [7] T. Young et al., “Recent trends in deep learning based natural language processing,” *IEEE Computational intelligence magazine*, vol. 13, no. 3, pp. 55-75, 2018.
- [8] Y. Bengio et al., “A neural probabilistic language model,” *Journal of machine learning research*, vol. 3, no. Feb, pp. 1137–1155, 2003.
- [9] S. Amiriparian et al., “Recognition of echolalic autistic child vocalisations utilising convolutional recurrent neural networks,” in *Interspeech*, 2018, pp. 2334-2338.
- [10] S. Amiriparian et al., “Bag-of-Deep-Features: Noise-robust deep feature representations for audio analysis,” in *International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2018, pp. 1-7.
- [11] M. Gjoreski et al. “Deep affect recognition from R-R intervals,” in *Proceedings of the 2017 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2017 ACM International Symposium on Wearable Computers*. ACM, 2017, pp. 754-762.
- [12] H. P. Martinez, Y. Bengio, G. N. Yannakakis, “Learning deep physiological models of affect,” *IEEE Computational intelligence magazine*, vol. 8, no. 2, pp. 20-33, 2013.
- [13] M. Gams, “Weak intelligence: Through the principle and paradox of multiple knowledge,” *Advances in computation: Theory and practice*, Volume 6, Nova science publishers, inc., NY, ISBN 1-56072-898-1, pp. 245, 2001.
- [14] M. Gjoreski et al., “Classical and deep learning methods for recognizing human activities and modes of transportation with smartphone sensors,” in *Information Fusion*. 2019, under revision.
- [15] M. Gjoreski et al., “Toward early detection and monitoring of chronic heart failure using heart sounds,” *Intelligent Environments*, pp. 336- 343, 2019.
- [16] S. Schmidt et al., “Segmentation of heart sound recordings from an electric stethoscope by a duration dependent hidden-markov model,” in *2008 Computers in Cardiology*. IEEE, 2008, pp. 345-348. [17] D. B. Springer, L. Tarassenko, and G. D. Clifford, “Logistic regression-HSSM-based heart sound segmentation,” *IEEE Transactions on Biomedical Engineering*, vol. 63, no. 4, pp. 822-832, 2015.

- [18] M. N. Homsy et al., "Automatic heart sound recording classification using a nested set of ensemble algorithms," in 2016 Computing in Cardiology Conference (CinC). IEEE, 2016, pp. 817-820.
- [19] F. Plesinger et al., "Heart sounds analysis using probability assessment," *Physiological measurement*, vol. 38, no. 8, p. 1685, 2017.
- [20] C. Potes et al., "Ensemble of feature-based and deep learning-based classifiers for detection of abnormal heart sounds," in 2016 Computing in Cardiology Conference (CinC). IEEE, 2016, pp. 621-624. [21] M. Zabihi et al., "Heart sound anomaly and quality detection using ensemble of neural networks without segmentation," in 2016 Computing in Cardiology Conference (CinC). IEEE, 2016, pp. 613- 616. [22] E. Kay and A. Agarwal, "DropConnected neural networks trained on time-frequency and inter-beat features for classifying heart sounds," *Physiological measurement*, vol. 38, no. 8, p. 1645, 2017.
- [23] I. J. D. Bobillo, "A tensor approach to heart sound classification," in 2016 Computing in Cardiology Conference (CinC). IEEE, 2016, pp. 629-632
- [24] J. Rubin et al., "Classifying heart sound recordings using deep convolutional neural networks and mel-frequency cepstral coefficients," in 2016 Computing in Cardiology Conference (CinC). IEEE, 2016, pp. 813-816.
- [25] S. Choi, Z. Jiang, "Comparison of envelope extraction algorithms for cardiac sound signal segmentation," *Expert Systems with Applications*, vol. 34, no. 2, pp. 1056-1069, 2008.
- [26] F. Eyben, F. Wening, F. Gross, B Schuller, "Recent Developments in openSMILE, the Munich Open-Source Multimedia Feature Extractor," In *Proc. ACM Multimedia (MM)*, Barcelona, Spain, ACM, ISBN 978-1-4503-2404-5, pp. 835-838, October 2013.
- [27] A. Gradišek et al., "Predicting species identity of bumblebees through analysis of flight buzzing sounds," *Bioacoustics*, pp. 63-76, 2017.