

{COMPANYNAME}11K: AN UNSUPERVISED REPRESENTATION LEARNING DATASET FOR ARRHYTHMIA SUBTYPE DISCOVERY

Anonymous authors

Paper under double-blind review

ABSTRACT

We release the largest public ECG dataset of continuous raw signals for representation learning containing 11 thousand patients and 2 billion labelled beats. Our goal is to enable semi-supervised ECG models to be made as well as to discover unknown subtypes of arrhythmia and anomalous ECG signal events. To this end, we propose an unsupervised representation learning task, evaluated in a semi-supervised fashion. We provide a set of baselines for different feature extractors that can be built upon. Additionally, we perform qualitative evaluations on results from PCA embeddings, where we identify some clustering of known subtypes indicating the potential for representation learning in arrhythmia sub-type discovery.

1 INTRODUCTION

Arrhythmia detection is presently performed by cardiologists or technologists familiar with ECG readings. Recently, supervised machine learning has been successfully applied to perform detection of certain types of arrhythmia (Hannun et al., 2019; Yildirim et al., 2018; Mincholé & Rodriguez, 2019; Porumb et al., 2020).

However, there may be ECG anomalies that warrant further investigation because they do not fit the morphology of presently known arrhythmia. While cardiologists are able to see these differences, it is hard to conclude that they are real by finding the same anomalous signal across multiple timepoints and patients without a data driven approach. This motivates the representation learning potential of this data.

Our data is collected by the {DEVICENAME}TM, a single-lead heart monitor device from {COMPANYNAME}¹. The raw signals were recorded with a 16-bit resolution and sampled at 250Hz with the {DEVICENAME}TM in a modified lead 1 position. The wealth of data this provides us can allow us to improve on the techniques currently used by the medical industry to process days worth of ECG data, and perhaps to catch anomalous events earlier than currently possible.

The ethics institutional review boards at UNIVERISTY approved the study and release of data.

1.1 OBJECTIVE

We want to improve the state-of-the-art of automated arrhythmia detection via representation learning. Ideally, this representation should preserve as much information about the underlying true heart function as possible. Such representations and learned feature extractors can improve downstream tasks which require more complicated features than what is typically extracted to predict major cardiac issues. More concretely, we are proposing a semi-supervised challenge on ECG data.

While an objective method to evaluate such a representation would be to measure its performance on tasks of interest, the way to perform best on such an evaluation would be to directly run a supervised learning task on those objectives. However, in certain circumstances, like training a neural network, for example, doing so results in a loss of information about the input (Tishby & Zaslavsky,

¹We've anonymised the company and device name during review

2015). The process may remove information vital to the discovery of new sub-types. We will see an example of this in §5.2.

Extracting features which can predict outcomes of not just arrhythmia is also an existing field of study (Lerma & Glass, 2016; Karpagachelvi et al., 2010), and can benefit from learned feature extractors based on this data.

2 RELATED WORK

One of the first open dataset of ECG signals was the MIT-BIH dataset, created in 1979 (Moody & Mark, 2001). They “expected that the availability of a common database would foster rapid and quantifiable improvements in the technology of automated arrhythmia analysis.” The MIT-BIH is still in use today with just 47 subjects. However, Shah & Rubin (2007); Guglin & Thatai (2006) found that computer predictions during that time were fraught with errors.

Later, data collection efforts improved leading to the creation of many small specific datasets (Goldberger et al., 2000). The MIMIC-III Waveform Database (Johnson et al., 2016) contains 67,830 waveform records from 30,000 ICU patients. These samples are at a higher sampling rate and with more leads. However, they are only recorded for short periods of time. The ECG-VIEW II dataset (Kim et al., 2017) aims to be a freely available dataset of ECG records together with clinical data for 461,178 patients. However, instead of raw signals, only beat information is included: RR interval, PR interval, QRS duration, etc (See Figure 2 for what the letters P, Q, R, S refer to in a beat signal.) The STAFF III Database (Pablo Martínez et al., 2017) contains 104 patients under an acutely induced myocardial ischemia. This includes pre, during, and post catheter insertion.

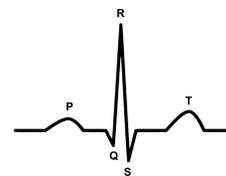


Figure 2: QRS regions for an ECG heartbeat signal.

More recently, single-lead wearable devices provided much larger amounts of data than before. As these devices could be worn for throughout the day, over a period of a couple of weeks, machine learning had much more data to work with. Rajpurkar et al. (2017) created an annotated training dataset of ECG signals consisting of 30,000 patients (Turakhia et al., 2013). The authors’ approach, and the follow up work claim that their automated models perform at the level of trained cardiologists (Hannun et al., 2019). However, their data has not been made publicly available.

3 PRIVACY CONCERNS: HEARTBEATS AS BIOMETRICS

There are attempts to use ECG signal data as a biometrics to identify someone. This brings risk of re-identification from our anonymized signal data. A paper by Salloum & Kuo (2017) claims high performance but the evaluation does not seem very controlled. A company called Nymi also aims to use a wearable ECG as a method to authenticate users.

When viewed in the context of other literature, the claim that ECG is a reliable method of authentication seems to be diminished. For example, Song et al. (2017) explores alternative ways to sense

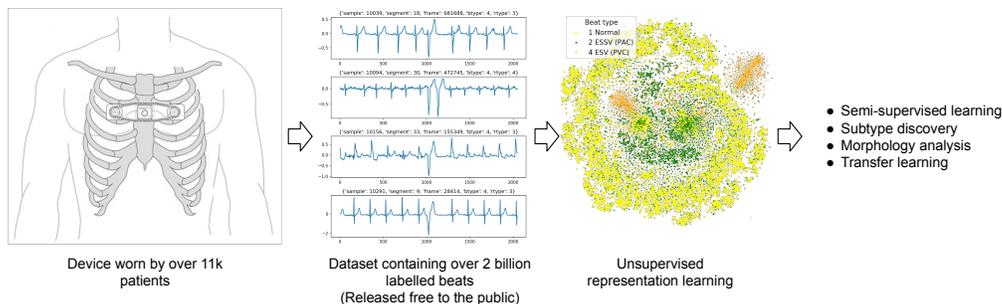


Figure 1: Overview of the project.

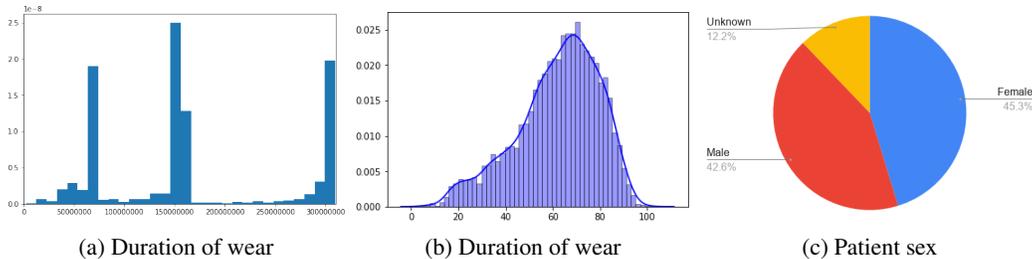


Figure 3: Demographics of the patients in the data.

cardiac motion (which they say is the identifying aspect), stating that ECG “biosignals are not related to cardiac motion, in which case indirect or incomplete cardiac characterization will compromise the advantages of cardiac motion as a biometric”.

Israel & Irvine (2012) state “Unlike fingerprint and face, the heartbeat data could contain health-related information as well as the personal identification information. This suggests a need for greater care in the collection, storage, and transmission of such data.” Additionally, they say that ECG has several limitations that must be overcome before they can be used as a biometric. Specifically, that (1) it requires a sufficient number of samples to identify an individual because the signal does not contain much information, (2) the expression of environmental variables on the heartbeat data is unique to the individual, (3) A target’s emotional state also requires intra-individual normalisation, and (4) a change in the contact location can reduce the ability to identify someone.

4 {COMPANYNAME} 11K DATASET

The dataset is processed from data provided by 11,000 patients who used the {DEVICENAME}TM device predominantly in Ontario, Canada, from various medical centers. While the device captures ECG data for up to two weeks, the majority of the prescribed duration of wear was one week. Figure 3a shows the distribution over duration of wear in the unprocessed data.

It should be noted that since the people who wear the device are patients, the dataset does not represent a true random sample of the global population. For one, the average age of the patient is 62.2 ± 17.4 years of age. Furthermore, whereas the {DEVICENAME}TM can be worn by any patient, it is mostly used for third line exam so that the majority of records in the dataset exhibit arrhythmias. No particular effort has been done on patient selection except data collection has been conducted over years 2017 and 2018. Figure 3c shows the distribution over age and gender.

The data is analysed by {COMPANYNAME}'s technologists who performed annotation on proprietary analysis tools. When the data is first extracted from the device, beat detection is performed automatically. A first technologist looks at the record as soon as possible to quickly send a feedback on the severity of the case. A second technologist then analyses the record labelling beats and rhythms performing a full disclosure analysis i.e. he / she sees the whole recording. The types of labels are described in more detail in Section 5.1. Finally, the analysis is approved by a senior technologist.

To prepare the data, we segment each patient record into segments of $2^{20} + 1$ signal samples (≈ 70 minutes). From this, we randomly select 50 of the segments and their respective labels from the list of segments. The goal here is to reduce the size of the dataset while maintaining a fair representation of each patient. In the training data we remove the labels for 80% of the patients. For the remaining 20%, half will be kept for the semi-supervised task, while another half will remain as test data for evaluation. Further details of nomenclature and statistics of the unprocessed and processed data can be found in Table 1.

We describe in further detail the different levels of hierarchy we have separated the data into:

Patient level (3-14 days) At this level, the data can capture features which vary in a systematic way and not isolated events, like the placement of the probes or patient specific noise.

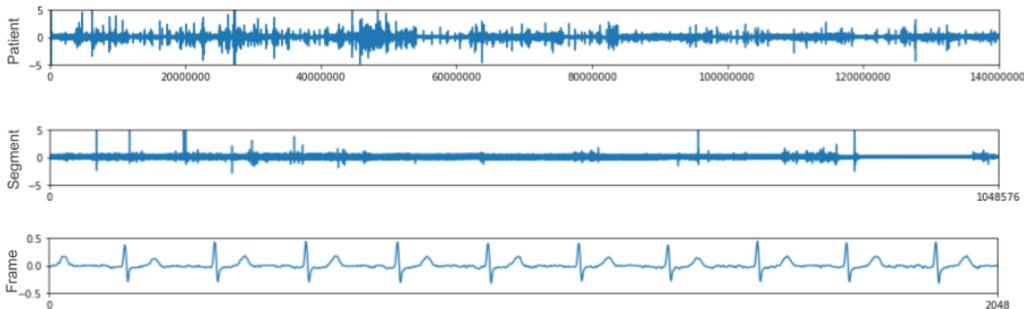


Figure 4: ECG data at different levels of the hierarchy. From top to bottom, a full patient record, a segment, and a frame.

Segment level (approximately 1 hour) A cardiologist can look at a specific segment and identify patterns which indicate a disease while ignoring noise from the signal such as a unique signal amplitude. Looking at trends in the segment help to correctly identify arrhythmia as half an hour provides the necessary context to observe the stress of a specific activity.

Frame level (approximately 8 seconds) At this level, the data can capture features about the beat as well as the rhythm.

While we have provided baselines only for frame-level features in this paper, we believe that processing the data with these levels of hierarchy results in some grouping information that could be leveraged to attain better results.

5 UNSUPERVISED REPRESENTATION LEARNING TASK

While the processed data includes labelled beat and arrhythmia information, we propose an *unsupervised* representation learning challenge to the community.

The goal of this data is to develop unsupervised representations of the ECG signal which can aid in two aspects:

1. Improve the performance of supervised tasks by using the learned representations.
2. Identify unknown subtypes of disease by studying the clustering of the representations.

These issues are addressed in quantitative and qualitative evaluations in the next two sections. The focus of this section studies the frame level embeddings which are typically enough for cardiologists to interpret.

Term	Definition	Statistic	# (units)
Segment	Fixed length contiguous region of a signal.	Number of Patients	11,000
		Number of labeled beats	2,774,054,987
Sample	As used in signal processing: A scalar value representing the amplitude of the signal in time.	Sample Rate	250Hz
		Frame size	$2^{11} + 1 = 2,049$ samples
		Segment size	$2^{20} + 1 = 1,048,577$ samples
		Total number of frames	1,084,314
Event	A specific arrhythmia occurring.	Total number of segments	542,157
		Dataset Size	271.27GB

(a) Glossary of terms

(b) Dataset Statistics

Table 1: Reference tables

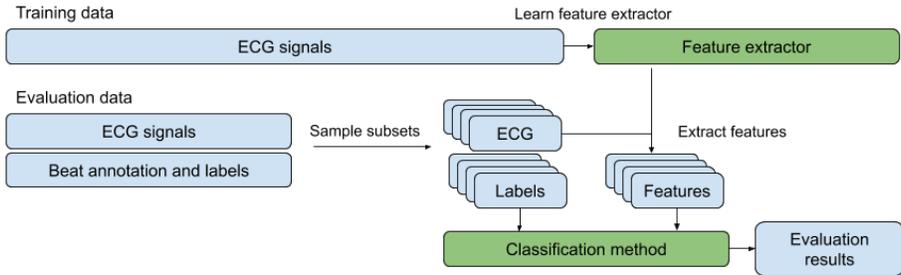


Figure 5: Diagram detailing the training and evaluation pipeline for the representation learning task. We have provide different methods in this paper for the blocks colored in green.

Beat labels	Count	Rhythm Labels	Count
Normal	174,249	NSR (Normal Sinusal Rhythm)	261,377
Premature Atrial Contractions	58,780	AFib (Atrial Fibrillation)	13,056
Premature Ventricular contractions	44,835	AFlutter (Atrial Flutter)	3,330

(a) Beat labels in the test set

(b) Rhythm labels in the test set

Table 2: Label counts in the test subset (patients 9000-10999). Each frame has a label. Only 2 types of labels are provided. Only these meaningful labels are used for evaluation and presented to the classifier.

5.1 QUANTITATIVE EVALUATION

For the quantitative evaluation we will benchmark common unsupervised algorithms in a semi-supervised setting to establish base quality. We make all code and models public in order to facilitate reproducibility and future work [GITHUB](#).

The evaluation consists of predicting the beat and rhythm for each frame in a hold out set (samples id’s > 10,000). The beat task is to predict if a frame contains all normal beats or contains at least one premature ventricular contractions (PVC) or premature atrial contraction (PAC) anywhere in a frame. Classifying a beat alone regardless of its surrounding beats can be challenging as, for example, a PAC is an abnormal beat only because it appears too soon and disrupts the rhythm (frequency). Furthermore, a PAC beat has the same shape as a normal beat, so taken alone, you can nearly not make the difference with a normal beat. The model will need to construct features about the nearby beats as well.

The second task is to predict the rhythm type given a frame. For a given frame the classification method must predict if the rhythm is normal, atrial fibrillation (AFib)², or atrial flutter, based on the input representation. AFib is indicated by irregular RR intervals, no distinct P waves and usually variable intervals between two atrial activations (Vollmer et al., 2018). Flutter appears as a saw-tooth pattern of R waves. Both require a representation that will compose a representation showing the difference between beats over time.

Both these tasks used in a supervised classification problem as a proxy for evaluating the usefulness of extracted features for detecting such events.

Figure 5 shows the pipeline for our evaluation method. Code to perform this evaluation in a consistent fashion is made available online for replicating the results and implementing new methods. The bulk of the training data does not come with beat annotation and labels, and can be used to train or fit a feature extraction method. The evaluation consists of sampling N frames from the test set and computing representations using the *feature extractor*. 50% of the data is then used to train a

²AFib is a controversial rhythm as cardiologists do not agree on the minimum duration. 8 second frames might not be sufficient to make such a decision.

Model	KNN				MLP			
	$N = 1000$		$N = 20000$		$N = 1000$		$N = 20000$	
	Beat	Rhythm	Beat	Rhythm	Beat	Rhythm	Beat	Rhythm
Random	0.33±0.02	0.33±0.00	0.33±0.01	0.33±0.01	0.33±0.01	0.33±0.01	0.33±0.02	0.33±0.00
Raw Sequence	0.44±0.03	0.33±0.00	0.61±0.01	0.34±0.00	0.54±0.03	0.38±0.08	0.67±0.01	0.33±0.01
PCA R^{100}	0.50±0.04	0.33±0.01	0.65±0.01	0.34±0.01	0.55±0.04	0.36±0.07	0.67±0.01	0.33±0.00
PCA R^{50}	0.51±0.03	0.33±0.00	0.64±0.01	0.34±0.00	0.55±0.04	0.34±0.06	0.64±0.01	0.33±0.00
PCA R^{10}	0.46±0.02	0.34±0.01	0.52±0.01	0.34±0.00	0.47±0.03	0.40±0.06	0.50±0.01	0.33±0.00
FFT	0.48±0.02	0.37±0.03	0.53±0.01	0.36±0.01	0.50±0.04	0.41±0.09	0.54±0.01	0.33±0.00
Periodogram	0.43±0.02	0.35±0.03	0.47±0.01	0.36±0.01	0.49±0.03	0.44±0.10	0.53±0.01	0.33±0.00
BioSPPy mean beat	0.35±0.02	0.34±0.01	0.38±0.01	0.39±0.01	0.40±0.03	0.34±0.08	0.40±0.01	0.33±0.00
AE (Random init)	0.41±0.03	0.33±0.00	0.53±0.01	0.34±0.00	0.45±0.02	0.36±0.08	0.55±0.00	0.33±0.00
AE	0.51±0.04	0.34±0.01	0.64±0.01	0.34±0.01	0.56±0.02	0.38±0.07	0.66±0.01	0.33±0.00

Table 3: Performance on a semi-supervised task computed as balanced accuracy. Given a random subset of labels from the training set predict the labels in the test set. Evaluated over 10 random subsets.

classification method and then evaluated on the held out 50%. Two classification models are used: (1) A k-nearest neighbors (KNN) method with $k = 3$, and (2) an MLP method, which consists of 4 layers of dimensions 1024, 1024, 512, and 512. The MLP model was trained for 10 epochs with Adam optimizer. We applied dropout (Srivastava et al., 2014) to prevent overfitting.

Each representation is learned without knowledge of the tasks — the feature extraction model is not updated during training of the classifier. We provide the evaluation results for the following baseline feature extraction methods:

Principal Components Analysis (PCA) computes the principal components from 30k examples from the training data. Then projects the test data onto 100, 50, or 10 principal components.

Fast Fourier Transform (FFT) computes a Fourier transform representing the magnitude of frequencies between 1Hz and 125Hz (Cooley & Tukey, 1965).

Periodogram computes an estimate of the power spectral density using Welch’s method (Welch, 1967).

BioSPPy identifies each beat using the detection algorithm by Carreiras et al. (2015) and computes the mean and standard deviation then concatenates them together to form the representation.

Autoencoder (Hinton, 1990) comprises of 2 MLPs, an encoder with an input size of 2049, a hidden layer of dimension 200, and a bottleneck representation of 100 dimensions. The decoder has the same architecture in reverse. There are residual connections before each non-linearity, and a batch normalization (Ioffe & Szegedy, 2015) is performed at the bottleneck layer. The model is trained for 3 epochs with Adam (Kingma & Ba, 2014) at a learning rate 10^{-4} with the L2 loss.

Our hope is that evaluation using a semi-supervised setting on high-level abnormality labels (e.g. premature atrial contraction, premature ventricular contraction) is a sufficient proxy for the quality of a representation — that these representations will prove useful for discovering unknown disease subtypes. Two models are used to evaluate the representations. We utilize small numbers of samples ($N = 1000$ and $N = 20000$) for evaluation to simulate the situation where a small cohort of patients is augmented using the unlabelled data we provide. Balanced accuracy is used to compute performance because there is a large imbalance between classes. If a model is to predict the same class for all samples the maximum balanced accuracy will be 0.33. We expect that this also becomes a source of noise at $N = 1000$ because an underrepresented class has a large impact in the performance if random predictions get a few samples right by chance.

The results are shown in Table 3. Currently autoencoders are not able to perform as well as we expected. PCA is able to perform the best at beat detection when using the KNN model while the MLP is able to predict better using the raw signal. One surprise is that rhythm detection is difficult. It is possible that, because the Periodogram and FFT captures periodicity in the signal, it performs better than the other feature extraction methods. Work by Vollmer et al. (2018) has shown that it is possible in a supervised setting.

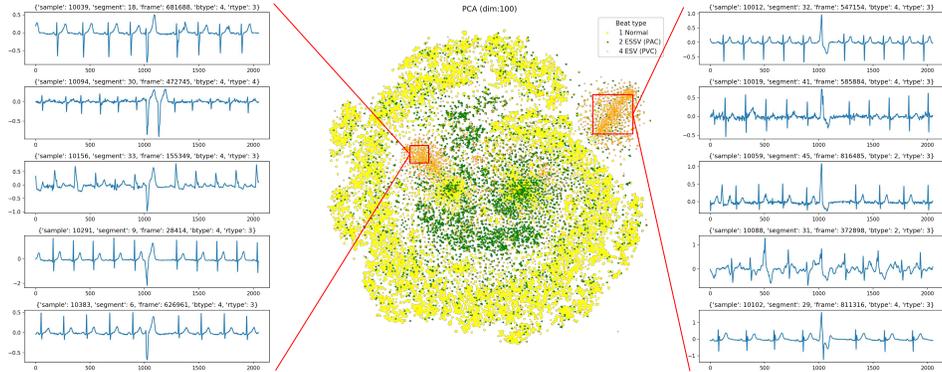


Figure 6: An analysis of the specific clusters resulting from the PCA features of 100 dimensions visualized with a t-SNE. 40,000 example frames were randomly sampled from the test data.

The results also shows the issues with using MLPs as a classification method for this task. MLPs typically requires more data points for training, and this issue shows up in the $N = 1000$ case, where there is a higher variance in the accuracy for each subset. The effect is even larger in rhythm classification, where the classes are imbalanced, resulting in huge variations in the *balanced accuracy*. When more data is available ($N = 20000$), variance is lower. As the ultimate purpose of this task is to learn better representation of the ECG signal, having a powerful parametric models like an MLP that works well only on higher instance counts may be offloading the representation learning to the classification method, which, as we alluded to before, is not favourable in our setting.

5.2 QUALITATIVE EVALUATION

Medical literature has discussed multiple types of PVC (Kanei et al., 2008; Phibbs, 2006). PVCs can be monomorphic or multimorphic (have different morphologies). Additionally, PVCs can also be multifocal and manifest in a different shape. When arising from the right ventricle, it has a dominant S wave in V1 lead. Whereas it has a dominant R wave if generated from the left ventricle Phibbs (2006).

We investigate the clustering of the signals by looking at the PCA encoding of 40,000 frames using a t-SNE in Figure 6. The plots clearly show two clusters of PVC that we can interpret as two different morphologies of this arrhythmia. We note that these are easy to see because of the different colors we use to highlight the points, but there seems to be remaining clusters that have not been analysed. The correlation between having two clusters for PVCs and PVCs being multimorphic aspect may be of interest to medical researchers to further explore clusters in this space created by different feature extractors.

Many other encoding methods, shown in Figure 7, also show clustering related to PVC and PAC. Notably FFT and BioSppy do not break the PVCs into two clusters. Although we can observe rhythm having some grouping it does not appear significant in the quantitative evaluation.

Such analysis is similar to what is done by Kachuee et al. (2018). However, in that work the features were constructed using a supervised task.

6 CONCLUSION

Single-lead heart monitors like the {DEVICENAME}TM are increasingly common, and have the potential for cardiologists to learn much more about arrhythmia and related heart diseases. However, this amount of data means manual analysis is no longer practical.

Machine learning has been widely deployed in the medical field by training a model to predict the right diagnosis based on human expert labels. Supervised learning serves well as an assistant in medical field; however, it hardly provides information beyond human knowledge. Additionally, certain human body signals can be very complex and imply non-linear features that cannot be easily

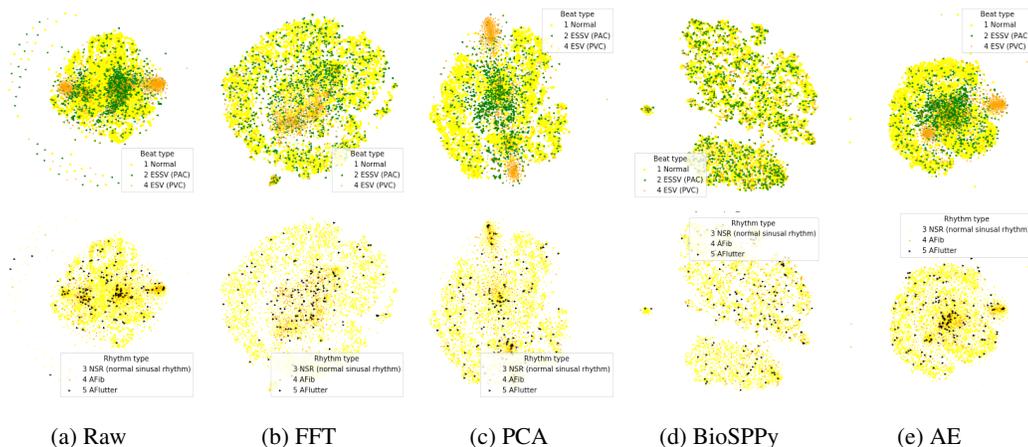


Figure 7: t-SNE plot of embeddings produced by different frame-level encoders. Colors represent the three basic labels. Each plot is computed using the same 20,000 frame examples encoded using each method and then having a t-SNE applied.

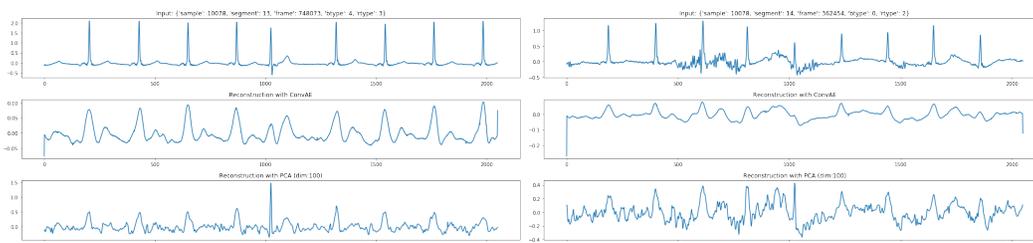


Figure 8: Reconstructions using the AE and PCA¹⁰⁰. Two samples are shown, one for each column. The input is shown on the top followed by the AE and then PCA.

identifiable manually. At present, representation learning methods have a potential in disentangling complex features, and potentially, unveil new signal structures of certain diseases which can correlate with clinical presentations.

By releasing this dataset, we believe that we can leverage unsupervised representation learning expertise to not only help to enable training models with lower number of samples, but potentially find new diseases and identify patterns associated with them.

We have proposed an evaluation pipeline for learning a feature extractor and evaluating extracted features using known arrhythmia as a proxy to measure the usefulness of the features. In addition, we have provided baseline results for *frame*-level representations under different feature extraction methods. Our data preparation makes a three level hierarchy available — the *segment* and *patient* level grouping of data. While we did not provide baselines that exploit this, future work that can take advantage of this context to extract better representations, and perhaps, find more interesting structure in the representation space. We also believe that this dataset can serve as a benchmark in other areas of machine learning, such as anomaly and outlier detection and hierarchical sequence modelling.

REFERENCES

- Carlos Carreiras, Ana Priscila Alves, André Lourenço, Filipe Canento, Hugo Silva, Ana Fred, and Others. BioSPPy: Biosignal Processing in Python, 2015.
- James W Cooley and John W Tukey. An Algorithm for the Machine Calculation of Complex Fourier Series. *Mathematics of Computation*, 1965. doi: 10.2307/2003354.

- Ary L. Goldberger, Luis A. N. Amaral, Leon Glass, Jeffrey M. Hausdorff, Plamen Ch. Ivanov, Roger G. Mark, Joseph E. Mietus, George B. Moody, Chung-Kang Peng, and H. Eugene Stanley. PhysioBank, PhysioToolkit, and PhysioNet. 2000 *Circulation*. doi: 10.1161/01.CIR.101.23.e215.
- Maya E Guglin and Deepak Thatai. Common errors in computer electrocardiogram interpretation. *International journal of cardiology*, 2006.
- Awni Y. Hannun, Pranav Rajpurkar, Masoumeh Haghpanahi, Geoffrey H. Tison, Codie Bourn, Mintu P. Turakhia, and Andrew Y. Ng. Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network. *Nature Medicine*, 2019. doi: 10.1038/s41591-018-0268-3.
- Geoffrey E Hinton. Connectionist learning procedures. In *Machine learning*. 1990.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, 2015.
- Steven A. Israel and John M. Irvine. Heartbeat biometrics: a sensing system perspective. *International Journal of Cognitive Biometrics*, 2012. doi: 10.1504/ijcb.2012.046514.
- Alistair E.W. Johnson, Tom J. Pollard, Lu Shen, Li-wei H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. MIMIC-III, a freely accessible critical care database. *Scientific Data*, 2016. doi: 10.1038/sdata.2016.35.
- Mohammad Kachuee, Shayan Fazeli, and Majid Sarrafzadeh. ECG heartbeat classification: A deep transferable representation. 2018 In *International Conference on Healthcare Informatics*. doi: 10.1109/ICHI.2018.00092.
- Yumiko Kanei, Meir Friedman, Naomi Ogawa, Sam Hanon, Patrick Lam, and Paul Schweitzer. Frequent premature ventricular complexes originating from the right ventricular outflow tract are associated with left ventricular dysfunction. *Annals of Noninvasive Electrocardiology*, 2008.
- S. Karpagachelvi, M. Arthanari, and M. Sivakumar. ECG Feature Extraction Techniques - A Survey Approach. In *International Journal of Computer Science and Information Security*, 2010.
- Young-Gun Kim, Dahye Shin, Man Young Park, Sukhoon Lee, Min Seok Jeon, Dukyong Yoon, and Rae Woong Park. ECG-VIEW II, a freely accessible electrocardiogram database. *PloS one*, 2017. doi: 10.1371/journal.pone.0176222.
- Diederik Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *International Conference on Learning Representations*, 2014. doi: 10.1145/1830483.1830503.
- Claudia Lerma and Leon Glass. Predicting the risk of sudden cardiac death. 2016 *The Journal of Physiology*. doi: 10.1113/JP270535.
- Ana Mincholé and Blanca Rodriguez. Artificial intelligence for the electrocardiogram. 2019 *Nature Medicine*. doi: 10.1038/s41591-018-0306-1.
- G B Moody and R G Mark. The impact of the MIT-BIH arrhythmia database. *IEEE engineering in medicine and biology magazine : the quarterly magazine of the Engineering in Medicine & Biology Society*, 2001.
- Juan Pablo Martínez, Olle Pahlm, Michael Ringborn, Stafford Warren, Pablo Laguna, and Leif Sörnmo. The STAFF III Database: ECGs Recorded During Acutely Induced Myocardial Ischemia. In *Computing in Cardiology (CinC)*, 2017. doi: 10.13026/C20P4H.
- Brendan Phibbs. *Advanced ECG: boards and beyond*. Elsevier Health Sciences, 2006.
- Mihaela Porumb, Ernesto Iadanza, Sebastiano Massaro, and Leandro Pecchia. A convolutional neural network approach to detect congestive heart failure. *Biomedical Signal Processing and Control*, 2020. doi: 10.1016/J.BSPC.2019.101597.
- Pranav Rajpurkar, Awni Y Hannun, Masoumeh Haghpanahi, Codie Bourn, and Andrew Y Ng. Cardiologist-level arrhythmia detection with convolutional neural networks. *arXiv preprint arXiv:1707.01836*, 2017.

- Ronald Salloum and C.-C. Jay Kuo. ECG-based biometrics using recurrent neural networks. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2017 IEEE. doi: 10.1109/ICASSP.2017.7952519.
- Atman P. Shah and Stanley A. Rubin. Errors in the computerized electrocardiogram interpretation of cardiac rhythm. *2007 Journal of Electrocardiology*. doi: 10.1016/j.jelectrocard.2007.03.008.
- Chen Song, Feng Lin, Yan Zhuang, Wenyao Xu, Changzhi Li, and Kui Ren. Cardiac Scan: A Non-Contact and Continuous Heart-Based User Authentication System. In *International Conference on Mobile Computing and Networking*, 2017.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 2014.
- Naftali Tishby and Noga Zaslavsky. Deep learning and the information bottleneck principle. In *2015 Information Theory Workshop (ITW)*, 2015.
- Mintu P Turakhia, Donald D Hoang, Peter Zimetbaum, Jared D Miller, Victor F Froelicher, Uday N Kumar, Xiangyan Xu, Felix Yang, and Paul A Heidenreich. Diagnostic utility of a novel leadless arrhythmia monitoring device. *The American journal of cardiology*, 2013.
- Marcus Vollmer, Philipp Sodmann, Leonard Caanitz, Neetika Nath, and Lars Kaderali. Can Supervised Learning Be Used to Classify Cardiac Rhythms? In *2017 Computing in Cardiology Conference (CinC)*, volume 44, 2018. doi: 10.22489/cinc.2017.347-176.
- Peter D. Welch. The Use of Fast Fourier Transform for the Estimation of Power Spectra: A Method Based on Time Averaging Over Short, Modified Periodograms. *Transactions on Audio and Electroacoustics*, 1967. doi: 10.1109/TAU.1967.1161901.
- Özal Yıldırım, Paweł Pławiak, Ru-San Tan, and U Rajendra Acharya. Arrhythmia detection using deep convolutional neural network with long duration ecg signals. *Computers in biology and medicine*, 2018.