

Explaining Deep Classification of Time-Series Data with Learned Prototypes

Alan H. Gee^{*,1,2}, Diego Garcia-Olano^{*,1,2}, Joydeep Ghosh¹ and David Paydarfar²

¹Electrical and Computer Engineering, The University of Texas at Austin

²Neurology, Dell Medical School, The University of Texas at Austin

{alangee, diegoolano}@utexas.edu, ghosh@ece.utexas.edu, david.paydarfar@austin.utexas.edu

Abstract

The emergence of deep learning networks raises a need for explainable AI so that users and domain experts can be confident applying them to high-risk decisions. In this paper, we leverage data from the latent space induced by deep learning models to learn stereotypical representations or “prototypes” during training to elucidate the algorithmic decision-making process. We study how leveraging prototypes effect classification decisions of two dimensional time-series data in a few different settings: (1) electrocardiogram (ECG) waveforms to detect clinical bradycardia, a slowing of heart rate, in preterm infants, (2) respiration waveforms to detect apnea of prematurity, and (3) audio waveforms to classify spoken digits. We improve upon existing models by optimizing for increased prototype diversity and robustness, visualize how these prototypes in the latent space are used by the model to distinguish classes, and show that prototypes are capable of learning features on two dimensional time-series data to produce explainable insights during classification tasks. We show that the prototypes are capable of learning real-world features - bradycardia in ECG, apnea in respiration, and articulation in speech - as well as features within sub-classes. Our novel work leverages learned prototypical framework on two dimensional time-series data to produce explainable insights during classification tasks.

1 Introduction

Despite the recent surge of machine learning, adoption of deep learning models in decision critical domains, such as healthcare, has been slow because of limited transparency and explanations in black-box algorithms. This observation points to the critical need for black-box models to offer interpretable, faithful explanations of their decisions so that practitioners in high-risk domains can trust model outputs and leverage their results. One such high-risk domain is treating

preterm infants (~10% of births worldwide) in the neonatal intensive care unit (NICU).

A common disorder observed in majority of preterm infants is recurrent episodes of apnea (cessation of breathing) and bradycardia (slowing of heart rate). Both of these spontaneous events may cause end organ damage related to hypoxemia (low oxygenation of blood) and ischemia (reduced blood flow) [Martin and Wilson, 2012]. Early detection of apnea and bradycardia can help prevent hypoxic-ischemic injury in tissue with high-metabolic demands [Schmid *et al.*, 2015; Pichler *et al.*, 2003] and prevent the cascade into intermittent hypoxia, which leads to complications of retinopathy, developmental delays, and neuropsychiatric disorders [Williamson *et al.*, 2013; Poets *et al.*, 2015; Di Fiore *et al.*, 2015]. Leveraging explainability in deep neural network classification of these time series can reveal complex morphological and physiological features that clinicians may not readily see. Thus, machine learning algorithms need transparency in their decision-making process to highlight subtle patterns. One such technique in deep explainability is prototypes, a case-based reasoning technique.

Prototypes are representative examples, learned in-process during model training, that describe influential data regions in latent representations and provide insight into aggregated features across training data that are utilized by the model for classification. In contrast to post-hoc explainability, which trains a secondary model to infer decision reasoning from a primary model by only leveraging inputs and outputs, in-process explainable methods offer faithful explanations of a primary model’s decisions [Rudin, 2018]. So, users who employ prototypes can confidently gain direct insight into the decisions algorithms are making for classification tasks.

On data with unclear class boundaries, in-process methods can misbehave. For example when the model in [Li *et al.*, 2017] is applied to the MNIST dataset, the prototypes easily separate in the latent space because the latent data representation is separable and well-structured (Fig 1). However, when class boundaries and features do not form distinguishable clusters, learned prototypes become archetypes (extreme corner cases) that exist near the convex hull of the data in the latent space (Fig. 4). This phenomenon yields prototypes that represent extreme class types (i.e. archetypes) and can underperform on classifying data in overlapping class regions.

In this work, we provide a deep classification method with

*Equal Contribution

Copyright © 2019 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

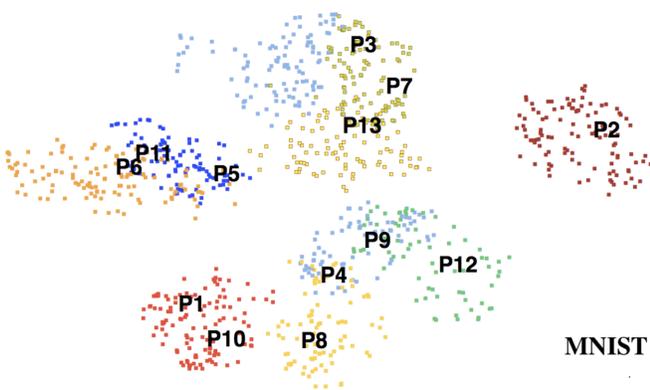


Figure 1: Learned prototypes of handwritten digits (MNIST) using the architecture from [Li *et al.*, 2017]. While colors represent the handwritten digits 0-9, the labels represent the learned prototypes. Because the latent representation of MNIST cluster distinctly, the prototypes are diverse. This may not be true when classes overlap

explainable insights for health time-series data. We introduce a prototype diversity penalty that explicitly accounts for prototype clustering and encourages the model to learn more diverse prototypes. These diverse prototypes will help focus on areas of the latent space where class separation is most difficult and least defined to improve classification accuracies. We show the utility of this approach on three tasks in two-dimensional time-series classification: (1) bradycardia from ECG; (2) apnea from respiration; and (3) spoken digits from audio waveforms. The two-dimensional representation of time-series provides an interpretable method for domain experts (e.g. clinicians) to understand the evolution of clinically relevant features based on visible phenotypes in time-series data. Our work enables a closed-loop collaboration between experts and machine learning algorithms to accelerate the efficacy of outcome predictions. The learning algorithms can find nuance features through development of explainable prototypes, and the experts can fine-tune the algorithms by providing feedback through the regularization of the diversity penalty. This is especially important for clinician experts who need explainability in black-box models to understand and diagnose different pathological mechanisms. To the best of our knowledge this is the first application of prototypes and latent space analysis for health time-series data that could help reveal clinically relevant and explainable phenotypes to improve the baseline for standard of care with automatic monitoring and detection.

1.1 Relevant Work

Explainable methods [Ribeiro *et al.*, 2016; Caruana *et al.*, 2015; Zhou *et al.*, 2015] have largely focused on labeled image and tabular data sets where classes are clearly separable and less so on time-series data in general. Recent work has focused on using prototypes to provide in-process explainability of classification models, either by learning meaningful pixels in the entire image [Li *et al.*, 2017] or by applying attention through the use of sub-regions or patches over an image [Chen *et al.*, 2018]. Class attention maps (CAMs) pro-

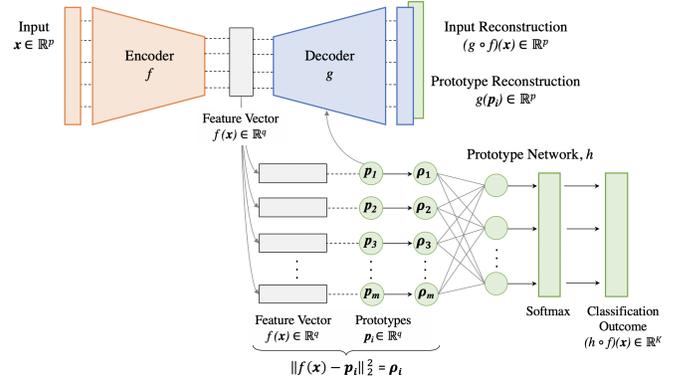


Figure 2: Prototype Architecture from [Li *et al.*, 2017]

vide probability maps to highlight areas of images that lead to a certain prediction [Zhou *et al.*, 2015], but do not give examples of prototypical examples of the data or explanations of how the training data relates to the end result. We focus on the former work [Li *et al.*, 2017] for example-based explainability where the generation of prototypes are intended to look like global representations of the training data.

Time-series classification on 1-D data with deep neural networks is a rapidly growing field, with almost 9,000 deep learning models [Fawaz *et al.*, 2018; Pons *et al.*, 2017; Faust *et al.*, 2018; Goodfellow *et al.*, 2018]. One such example leverages global average pooling to produce CAMs to provide explainability for a deep CNN to classify atrial fibrillation in ECG data [Goodfellow *et al.*, 2018]. However, the number of available healthcare datasets, specifically ECG waveforms, is limited [Fawaz *et al.*, 2018]. Within this context, time-series classification on ECG waveforms has been done on a small scale, typically with single beat or short-duration (10 s) arrhythmia classification [Faust *et al.*, 2018; Yildirim *et al.*, 2018].

2 Methods

2.1 Time-Series Explanation via Prototypes

We adopt the autoencoder-prototype architecture from [Li *et al.*, 2017]. Let $\mathcal{X} = (x_i, y_i)_i^n$ be the training set with $x_i \in \mathbb{R}^p$ and class labels $y_i \in \{1, \dots, K\}$ for each training point $i \in \{1, \dots, n\}$. The front-end autoencoder network learns a lower-dimension latent representation of the data with an encoder network, $f : \mathbb{R}^p \rightarrow \mathbb{R}^q$. The latent space is then projected back to the original dimension using a decoder function, $g : \mathbb{R}^q \rightarrow \mathbb{R}^p$. The latent representation, $f(x)$ is also passed to a feed-forward prototype network, $h : \mathbb{R}^q \rightarrow \mathbb{R}^K$, for classification. The prototype network learns m prototype vectors, $p_1, p_2, \dots, p_m \in \mathbb{R}^q$ using a four-layer fully-connected network over the latent space that learns a probability distribution over the class labels y_i (Fig 2). The learned prototypes can then be decoded using g and examined to infer what the network has learned. The choice of m is determined *a priori*, with larger values allowing for higher throughput and model capacity.

We improve prior work by adding a penalty for learned prototypes in the objective function of the above network to

increase prototype diversity and coverage of the data in latent representations. To align with the minimization of the objective function, this new prototype diversity penalty needs to be (1) small when distances between prototypes are far apart, and (2) large when distances between prototypes are close in distance. We can evaluate the feasibility of a set of prototypes by considering the distance of the two closest prototypes across all prototype combinations. So, we consider the average minimum squared L_2 distance between any two prototypes, p_i, p_j for our loss function. To achieve the desired property above, we take the inverse of this average distance:

$$PDL(p_1, \dots, p_m) = \frac{1}{\log\left(\frac{1}{m} \sum_{j=1}^m \min_{i>j \in [1, m]} \|p_i - p_j\|_2^2\right) + \epsilon} \quad (1)$$

The logarithm function tapers large distances so that the penalty does not quickly vanish, and the ϵ term is for numeric stability. By taking the inverse of the log of the prototype distances, we penalize prototypes that are close in distance while making sure the minimum distance between prototypes does not get too large. This prototype diversity loss (PDL) promotes coverage over the latent space. We update the objective function to:

$$\begin{aligned} \mathcal{L}((f, g, h), X) = & E(h \circ f, X) + \lambda_R R(g \circ f, X) \\ & + \lambda_1 R_1(p_1, \dots, p_m, X) \\ & + \lambda_2 R_2(p_1, \dots, p_m, X) \\ & + \lambda_{pd} PDL(p_1, \dots, p_m) \end{aligned} \quad (2)$$

where E is the classification (cross entropy) loss, R is the reconstruction loss of the autoencoder (*i.e.* L_2 norm), and R_1 and R_2 are the loss terms that relate the distances of the feature vectors to the prototype vectors in latent space [Li *et al.*, 2017]:

$$R_1(p_1, \dots, p_m, X) = \frac{1}{m} \sum_{j=1}^m \min_{i \in [1, m]} \|p_j - f(x_i)\|_2^2, \quad (3)$$

$$R_2(p_1, \dots, p_m, X) = \frac{1}{n} \sum_{i=1}^n \min_{j \in [1, m]} \|f(x_i) - p_j\|_2^2 \quad (4)$$

The minimization of the R_1 loss term promotes each prototype vector to learn one of the encoded training examples, while the minimization of R_2 loss promotes encoded training examples to be close to one of the prototypes. This balance gives meaningful pixel-to-pixel representations between the prototypes and training data.

We train our models with a randomly shuffled batch size of 100 (ECG, Speech) and 125 (Respiration). We parameterize the number of prototypes (see supplement) and the regularization term λ_{pd} for the classification tasks while keeping the other hyperparameters as in [Li *et al.*, 2017].

2.2 Prototype Diversity Score

We adopt a version of the group fairness metric presented in [Mehrotra *et al.*, 2018] and refer to it as the prototype diver-

sity score, Ψ :

$$\Psi = \frac{1}{Z} \sum_{i=1}^t \sqrt{|\phi_i|} \quad (5)$$

where $\phi_i, i \in \{1, \dots, t\}$ is defined for a specific metric and Z is the normalization constant. For the neighbor diversity metric Ψ_N , ϕ_i is the set of prototypes that have nearest neighbor i and Z is the number of prototypes m . For the class diversity metric Ψ_C , ϕ_i is the set of prototypes that are from class i and Z is the number of classes K . Higher scores will occur when prototypes have more unique elements. Thus, $\max(\Psi_D) = 1$.

2.3 Datasets

The neonatal intensive care unit (NICU) dataset is composed of two sources: (1) ECG and Respiration waveforms from PhysioNet’s PICS database [Gee *et al.*, 2017; Goldberger *et al.*, 2000]; and (2) ECG waveforms (500 Hz, Intellivue MP450) collected from a preterm infant over their entire stay (~10 weeks) at Seton Medical Center Austin. The inclusion of (2) helps supplement the ECG events from (1). The image data used in this study are made publicly available¹.

The inter-breath intervals (IBIs) from the respiration were extracted using a standard peak finder. The respiration signals were clipped into 60 second segments that were normalized to zero-mean, unit variance. The R-R intervals for the ECG of the NICU dataset were extracted using a Morlet wavelet transformation of the ECG signal. An open-source peak finder was applied to the wavelet scale range (0.01 to .04 scales) related to QRS complex formation in the spectrogram. The ECG waveforms were clipped at 15 seconds with the event in the middle. All ECG segments were band-pass filtered from 3 to 45 GHz, scaled to zero-mean, unit-variance, and scaled to the median QRS complex amplitude. Images were then captured to mimic what a clinician would see upon investigation of an ECG signal. Waveforms with no

¹<https://physionet.org/physiobank/database/picsdb>

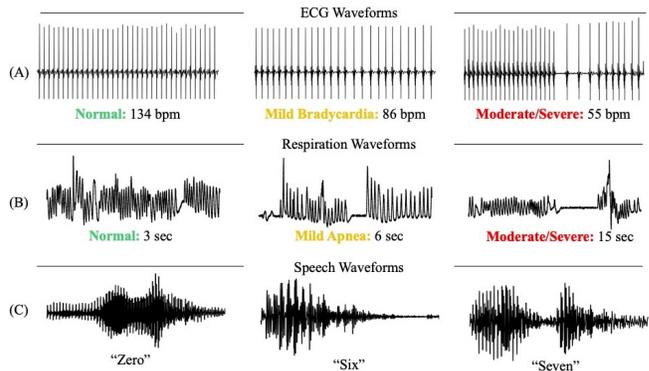


Figure 3: Examples of waveforms for each task: (A) Electrocardiogram (ECG) waveforms related to bradycardia classification, (B) Respiration waveforms related to apnea classification, and (C) Speech waveforms for a particular a speaker (Jackson). For (A) and (B) we classify the segments based on severity (*i.e.* time difference between peaks), and for (C) we classify based on digit class.

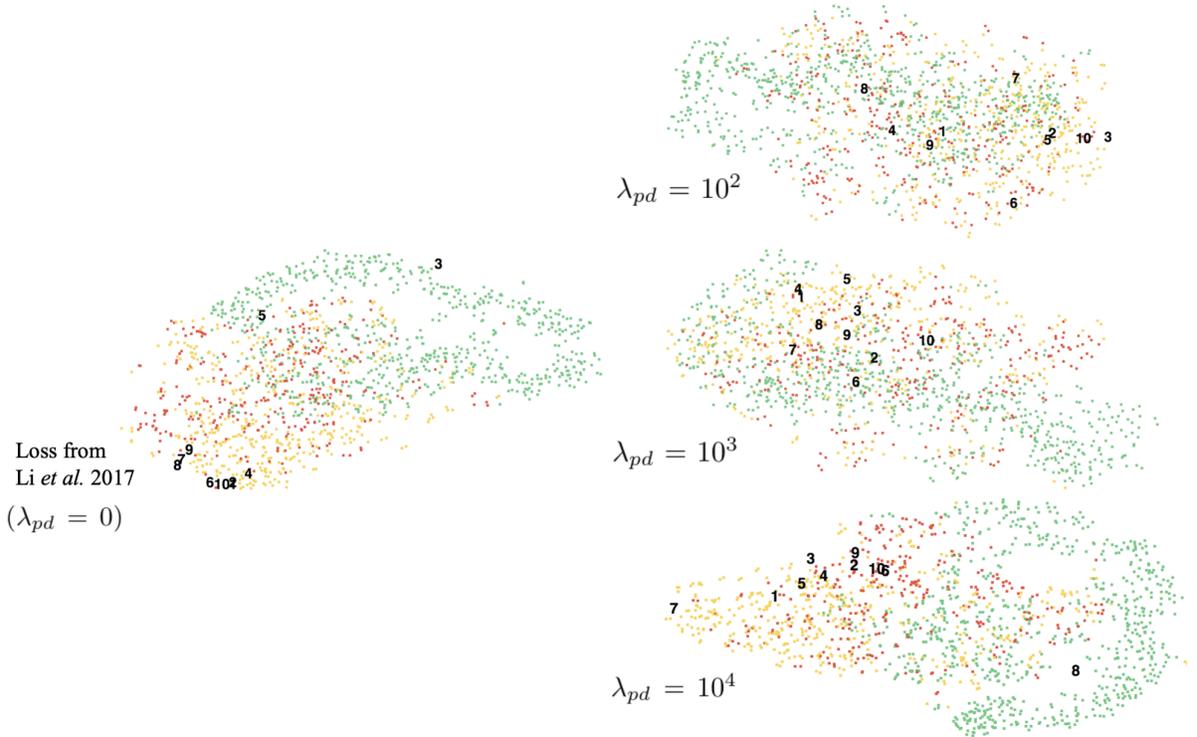


Figure 4: Effect of loss regularization on the latent space and spread of prototypes for the NICU classification task using 10 prototypes with $\lambda_{pd} = 0$ (baseline) and $\lambda_{pd} = 10^3$. The second and third dimensions of a t-SNE projection on each space shows prototypes with more coverage and diversity in the latter case.

visibly distinguishable QRS complexes or respiratory peaks were discarded because these waveforms are too obscure for even a clinician expert to evaluate.

Class breakdowns for bradycardia in the ECG signal follow clinical thresholds [Perlman and Volpe, 1985]: $X_{ECG} = \{ \text{normal} (>100 \text{ bpm}): 1039, \text{mild} (100-80 \text{ bpm}): 634, \text{moderate} (80-60 \text{ bpm}): 306, \text{severe} (<60 \text{ bpm}): 132 \}$. Moderate and severe events were combined into a single class. The class breakdown for apneas in respiration are: $X_{RESP} = \{ \text{normal} (1-3 \text{ s}): 1939, \text{mild} (4-6 \text{ s}): 1921, \text{moderate/severe} (> 6 \text{ s}): 1487 \}$.

The Free Spoken Digit Dataset [Jackson *et al.*, 2018] consists of 2000 audio clips (8 kHz) of four speakers repeating the digits 0 through 9, 50 times each. Each segment was normalized to zero-mean, unit-variance and clipped for white space (Fig. 3). This data can be thought of as “spoken MNIST”. We perform speaker classification and digit classification within a speaker.

2.4 Visualization of Latent Space

We use PCA to reduce the latent space vectors to a dimension of 500, which retains 98% of the variability. We then calculate the cosine similarity between these 500 dimensional vectors to produce a similarity matrix and use t-distributed stochastic neighbor embedding (t-SNE) from [Van der Maaten and Hinton, 2008] to reduce the 500 x 500 similarity matrix down to three dimensions for visualization

purposes. This technique calculates the KL-Divergence between the higher-order dimensional latent space and the lower dimensional space used to represent the former visually. This approach is non-deterministic so the global position in the lower space is uninformative and instead proximity to neighbors is the key insight to gain. Additionally while the first two dimensions of the projection show the general spread of information, the second and third dimensions maybe useful for visualizing within group information. Thus, we use the second and third dimensions for our visualizations.

3 Results

3.1 Classification of ECG with 2-D Prototypes

We test our prototype implementation on ECG waveforms related to bradycardia using the NICU data for a 3-class classification task using 10 prototypes. We treat the input waveforms as 2-D images and use a four-layer autoencoder to learn complex representations over the data.

We observe more diverse prototypes and comparable or better test accuracy with our model $93.1 \pm 0.4\%$ compared with $92.1 \pm 0.1\%$ from the baseline model in [Li *et al.*, 2017] (Table 1). Both models perform well on the classification of the normal class, as expected since normal waveforms have near-constant phase. Both models additionally have difficulty separating between the mild and moderate/severe classes, often confusing the classification between these two (see supplement). This behavior is expected since data near these

ECG: Bradycardia			
λ_{pd}	Accu.	Ψ_N	Ψ_C
0	92.1 \pm 0.1%	0.83 \pm 0.04	0.78 \pm 0.19
500	92.7 \pm 1.0 %	0.86 \pm 0.07	0.89 \pm 0.19
1e3	92.4 \pm 1.3%	0.87 \pm 0.11	0.89 \pm 0.19
2e3	93.1 \pm 0.4%	0.90 \pm 0.04	1.00 \pm 0.00

Respiration: Apnea			
λ_{pd}	Acc.	Ψ_N	Ψ_C
0	81.4 \pm 3.6%	0.96 \pm 0.07	1.00 \pm 0.00
500	82.3 \pm 3.8%	0.94 \pm 0.09	1.00 \pm 0.00
1e3	77.1 \pm 0.6 %	1.00 \pm 0.00	1.00 \pm 0.00
2e3	80.2 \pm 2.5 %	0.97 \pm 0.04	0.84 \pm 0.23

Table 1: Diversity score for neighbors Ψ_N and class Ψ_C . We report Ψ 's related to the epoch with the highest test accuracy. Our model, $\lambda_{pd} > 0$, returns better accuracies and diversity scores (bold) than the baseline model, which is row $\lambda_{pd} = 0$, across ECG and Respiration waveforms. (Model details: 3-class, 10-prototypes, learning rate = 0.002).

two class boundaries are difficult to discern, even for domain experts, due to events existing in both classes with possible subtle time differences in cardiac firing. Our model also improves prototype diversity (Table 1) over the baseline model. This result suggests that the prototype diversity loss encourages exploration, through learning diverse prototypes, within the data represented in the latent space. As a result, our model finds more helpful features and prototypes and thus, improves classification results.

Because prototypes are generated during training, we infer features that the algorithm utilized to classify waveforms at different points during training (Fig 5). For example, by epoch 100, we see that some of the prototypes exhibit global morphological features of the normal waveform class after random initialization at epoch 0. As training progresses, we observe other complex phenotypes emerging: one prototype learns that large gaps in cardiac firings are important for identifying severe cases and another prototype learns the consistent pattern of spikes are important for mild cases. Since the mild class shares mixed features of both normal and positive events, it is not surprising that more prototypes are needed in this class to learn subtleties of the class features (see supplement). Thus, prototypes highlight waveform structures that the algorithm deemed as important when trying to learn the classification of bradycardia. This finding aligns with the idea of clinicians using visible features present in a bradycardia (i.e. the increasing distance between QRS complexes) to decide whether or not a bradycardia exists in an image.

We compare the latent space of [Li *et al.*, 2017] to the latent space of our model with prototype diversity loss via t-SNE projections, where proximity in 2-D space suggests that points are "close" in distance in the original latent space. We represent the learned prototypes by mapping each prototype to its nearest neighbor (Fig 4). We find that by increasing our loss term, PDL , our model increases the local coverage of the prototypes compared with the baseline model (i.e. $\lambda_{pd} = 0$). However, if we regularize our loss term too much (i.e. $\lambda_{pd} > 10^4$), we begin to introduce clustering of prototypes and diversity suffers. Thus with the additional prototype distance penalty, we achieve higher diversity scores and classification accuracies for various hyperparameters (Fig 9).

3.2 Case Study with Prototypes: Exploring ECG Morphology and Bradycardia Classification.

We observe that ECG events in a local neighborhood share similar QRS complex morphology, despite having different

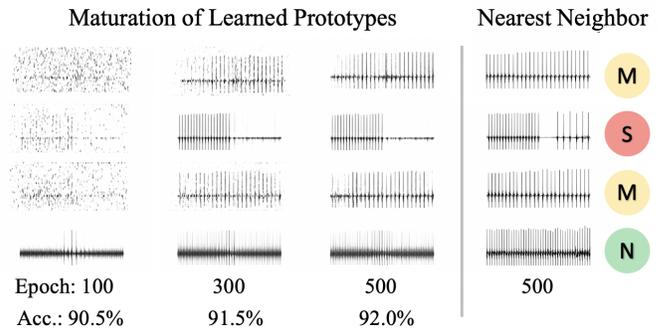


Figure 5: Prototype evolution with in-process explainability over training time. High level features are easily learned in early epochs of training, while more complex features are developed over time. The final nearest neighbors are depicted on the right. The prototypes correspond to a subset of the $\lambda_{pd} = 10^3$ latent space cloud in Figure 4. Model details: 3-class, 10-prototypes.

class labels and cardiac firing periods (Fig. 6, bottom). Even though we did not impose a class constraint, we observe that the algorithm found two separate features within the moderate/severe class that were important in the classification task (i.e. prototypes 2 and 10 shown at the top of the (Fig 6). These two prototypes explore two different cardiac timings as prototype 2 exhibits a progressive delay in cardiac firing, while prototype 10 exhibits a large spontaneous delay. The incorporation of the prototype diversity loss encouraged this exploration of the latent space. These results suggest that there are physiologic dependencies (i.e. clustering based on cardiac morphology and function) that can be learned using our model to investigate physiological phenomena, and possibly applied to other clinical areas, like cardiac ischemia or apnea of prematurity in respiration - both exhibit visible, abnormal waveform behavior. This work provides a visualization tool for clinician experts to evaluate different morphological of physiological time-series data².

3.3 Classification of Apnea in Respiration

Apnea of prematurity is common among preterm infants, and is visually apparent as a pause of inhalation and exhalation (i.e. absence of sinusoidal behavior) in the respiratory signal. We next test our prototype implementation on respiration

²<https://github.com/alangee/ijcai19-ts-prototypes>

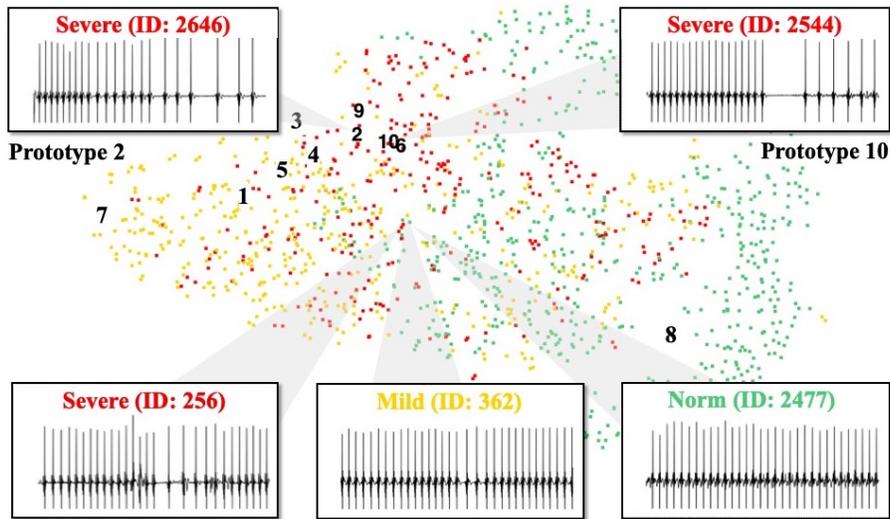


Figure 6: Learned prototypes showcase the diversity of features that are important for understanding ECG morphology while classifying bradycardia events. (10 -prototypes, $\lambda_{pd} = 10^4$).

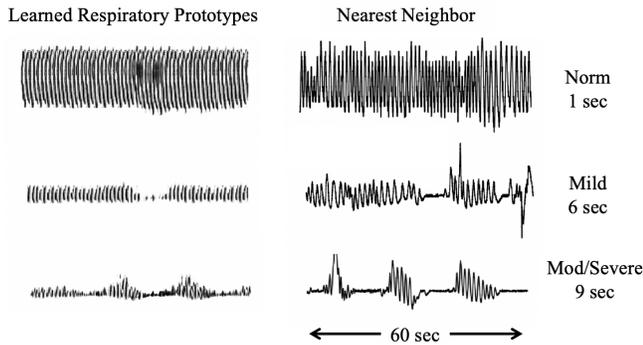


Figure 7: Learned prototypes showcase the diversity of features across classes that are important for understanding respiration morphology while classifying apnea events. For this classification task, we observe a variety of prototypes (at epoch 500) that learn various cases with cessation of breathing (6 and 9 second gaps) and the global features within the segment that are important for the model’s classification. (8 -prototypes, $\lambda_{pd} = 500$).

waveforms that are related to apnea in a 3-class classification task. We treat the input waveforms as 2-D images again, since clinicians evaluate apneas through visual inspection of the respiration signal.

We observe more diverse prototypes and comparable or better test accuracy with our model $82.3 \pm 3.8\%$ compared with $81.4 \pm 3.6\%$ from the baseline model, and with overall unique nearest neighbors ($\Psi_N = 1$) and class diversity ($\Psi_C = 1$) (Table 1). Both models have difficulty separating between the event classes because data near these two class boundaries are difficult to visually discern (i.e. 6 second gap versus 7 second gap) and have common behavior with regular respiratory function that is found in the normal class. We find that the addition of a prototype diversity loss maintains or improves performance and yields more diverse prototypes (Table 1).

We also note that the algorithm is able to discern physi-

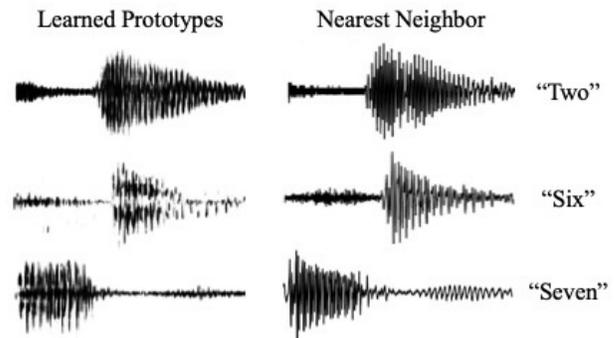


Figure 8: Learned prototypes from audio waveforms of spoken digits by Nicolas from the FSDD ($\lambda_{pd} = 500$).

ological examples and generate learned prototypes that distinctly relate to physiological behavior. For example, in Fig. 7, we see that algorithm finds segments that are related to periodic breathing of 9 second duration (moderate/severe). These segments are physiologically different from normal apneas of 6 seconds (mild), and clearly different from normal breathing with periodicity of 1 second (Fig 7). In the set of eight learned prototypes, the algorithm finds three different classes easily, each with different respiratory phenomena, that are critical in the classifying various types of apneas.

3.4 Spoken Digits Classification and Analysis

Speech abnormalities can be suggestive of underlying pathological dysfunction, and common features that clinicians visibly discern in waveforms to assess speech include cadence, prosody, and syllable articulation. To aid in speech feature detection, we assess our model on high-frequency audio waveforms of spoken digits (FSDD) from medically-normal individuals. These digits are treated as 2-D images for 4 class speaker and 10 digit classification tasks with 4 and 10 proto-

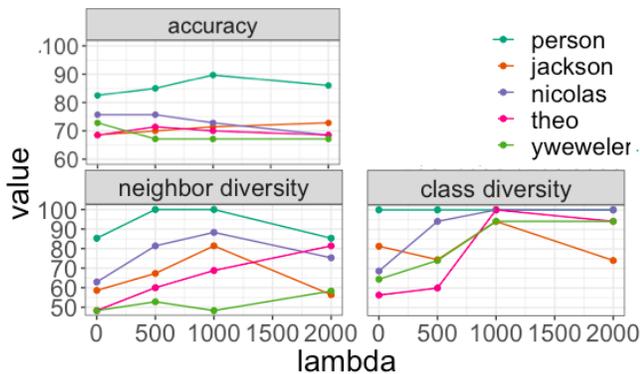


Figure 9: Accuracy and diversity metrics for the spoken digits experiments using the FSDD. We divide this dataset into two tasks: (1) classifying the person speaking and (2) classifying the digit spoken within each person.

types, respectively. The waveform envelope and syllables of these spoken digits are discernible to the eye (see “six” and “se-ven” in Fig 2) and, as such, make good candidates for our image-based explainability model. We demonstrate some of the learned prototypes in Fig. 8, which show representations the model finds useful in classifying digits for a given speaker. Experiments show that by varying regularization of the prototype diversity penalty, we observe slightly better or similar accuracies when compared to the baseline model (Fig. 9). With a fine-tuned λ_{pd} we can increase diversity of the prototypes and correspondingly see improved accuracy and data coverage (see supplement). For example, $\lambda_{pd} = 500$ gives a higher diversity score across all tasks, indicating prototypes with more unique nearest neighbors as compared with the baseline model (Fig 9).

Experiments show that increasing the depth of the network and fine-tuning the learning rate lead to both increased accuracy and diversity over all tasks. Similarly, recent data augmentation techniques in medical [Bahadori and Lipton, 2019] and speech recognition [Park *et al.*, 2019] domains could help further improve performance. The purpose of this work, however, is not to obtain the best performance on these tasks, but rather to show the utility of learned prototypes as faithful explanations of decisions made by a model.

4 Discussion

We presented a new autoencoder-prototype model that promotes diversity in learned prototypes by penalizing prototypes that are too close in squared L_2 distance in the latent space. The new term, $\lambda_{pd} PDL(p_1, \dots, p_m)$, in the loss function (Eq. 2) promotes prototype diversity while improving classification accuracy and prototype coverage of data represented in the latent space. These prototypes help explain which global features and representative segments in the training data are most useful for deep time-series classification. This in-process generation of prototypes offers explainable insights into deep classifiers.

Our model and results provide an important significance that previous works lack. Depending on the clinical context of the case, experts may want to either trivialize big differ-

ences in the time series features, or conversely accentuate nuanced differences in learned prototypes as clinically important signs of impending adverse outcomes. Therefore, our implementation offers a collaborative method for clinician experts to use their insight interactively with machine learning algorithms: increasing λ_{pd} promotes large observable differences in the prototypes, while decreasing λ_{pd} promotes diverse features and prototypes. In turn, our model enables a closed-loop feedback framework to accelerate phenotype discovery to lead clinicians to better-informed decision.

We evaluate the performance of our model on increasingly difficult physiological datasets to demonstrate the effect of λ_{pd} . The ECG signal is more robust against movement artifact and produces a cleaner signal for the 2-D visualization task, whereas the respiration signal, which is the resultant voltage change across diaphragm movement, is highly susceptible to signal artifact. Additionally, speech waveforms are compressed, high-frequency waveforms (kHz) which make it difficult to visibly extract high-resolution features. We find that our model allocates more prototypes to learn the intricacies of the more indistinguishable classes (i.e. mild and moderate/severe) that are hard for a human to discern, especially the mild cases because this class is a mixture and intermediary of the two extreme classes.

We observe, however, that the high number of loss terms creates a trade-off between prototype interpretability and model accuracy. For example, we observe that for a small number of prototypes, we achieve near-perfect prototype reconstruction but at the cost of classification accuracy. When the number of prototypes was large, we achieve higher accuracy but received noisy prototypes. In future implementations, we can replace the front-end autoencoder with a model that operates well on 1-D time series, like a recurrent neural network, to balance accuracy and prototype interpretability.

There has also been work on computing prototypical patches over 2-D images to generate explainable sub-features [Chen *et al.*, 2018]. Extending the idea of patches to 1-D time-series signals would allow for parsing the signal for sub-frequencies and features that could better explain how events are triggered. Nonetheless, the work presented in this paper provides a more robust prototype model to help explain algorithmic behavior and decision-making in deep time-series classification tasks with promising results in clinically relevant datasets.

Acknowledgement

The authors would like to thank Sinead Williamson and the three reviewers for providing helpful feedback and critical reviews of our work.

References

- [Bahadori and Lipton, 2019] Mohammad Taha Bahadori and Zachary Chase Lipton. Temporal-clustering invariance in irregular healthcare time series. *arXiv preprint arXiv:1904.12206*, 2019.
- [Caruana *et al.*, 2015] Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noemie Elhadad. Intelligible models for healthcare: Predicting pneumonia

- risk and hospital 30-day readmission. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '15*, pages 1721–1730, New York, NY, USA, 2015. ACM.
- [Chen *et al.*, 2018] Chaofan Chen, Oscar Li, Alina Barnett, Jonathan Su, and Cynthia Rudin. This looks like that: deep learning for interpretable image recognition. *CoRR*, abs/1806.10574, 2018.
- [Di Fiore *et al.*, 2015] J.M. Di Fiore, E Gauda, R.J. Martin, and P MacFarlane. Cardiorespiratory events in preterm infants: interventions and consequences. *Journal Of Perinatology*, 36(251), 2015.
- [Faust *et al.*, 2018] Oliver Faust, Yuki Hagiwara, Tan Jen Hong, Oh Shu Lih, and U Rajendra Acharya. Deep learning for healthcare applications based on physiological signals: A review. *Computer Methods and Programs in Biomedicine*, 161:1 – 13, 2018.
- [Fawaz *et al.*, 2018] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. Deep learning for time series classification: a review. *CoRR*, abs/1809.04356, 2018.
- [Gee *et al.*, 2017] A. H. Gee, R. Barbieri, D. Paydarfar, and P. Indic. Predicting bradycardia in preterm infants using point process analysis of heart rate. *IEEE Transactions on Biomedical Engineering*, 64(9):2300–2308, 2017.
- [Goldberger *et al.*, 2000] 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: Components of a new research resource for complex physiologic signals. *Circulation*, 101(23):e215–e220, June 2000.
- [Goodfellow *et al.*, 2018] Sebastian Goodfellow, Andrew Goodwin, Danny Eytan, Robert Greer, Mjaye Mazwi, and Peter Laussen. Towards understanding ecg rhythm classification using convolutional neural networks and attention mappings. In *Proceedings of Machine Learning for Healthcare, MLHC '18*, pages 2243–2251, 08 2018.
- [Jackson *et al.*, 2018] Zohar Jackson, César Souza, Yuxin Flaks, Jason; Pan, Hereman Nicolas, and Adhish Thite. Free spoken digit dataset (fsdd). 2018.
- [Li *et al.*, 2017] Oscar Li, Hao Liu, Chaofan Chen, and Cynthia Rudin. Deep learning for case-based reasoning through prototypes: A neural network that explains its predictions. *CoRR*, abs/1710.04806, 2017.
- [Martin and Wilson, 2012] Richard J. Martin and Christopher G. Wilson. Apnea of prematurity. pages 2923–2931, 2012.
- [Mehrotra *et al.*, 2018] Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lahmas, and Fernando Diaz. Towards a fair marketplace: Counterfactual evaluation of the trade-off between relevance, fairness & satisfaction in recommendation systems. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 2243–2251, 2018.
- [Park *et al.*, 2019] Daniel S Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D Cubuk, and Quoc V Le. SpecAugment: A simple data augmentation method for automatic speech recognition. *arXiv preprint arXiv:1904.08779*, 2019.
- [Perlman and Volpe, 1985] Jeffrey M. Perlman and Joseph J. Volpe. Episodes of apnea and bradycardia in the preterm newborn: Impact on cerebral circulation. *Pediatrics*, 76(3):333–338, 1985.
- [Pichler *et al.*, 2003] G. Pichler, B. Urlesberger, and W. Muller. Impact of bradycardia on cerebral oxygenation and cerebral blood volume using apnoea in preterm infants. *Physio. Measurement*, 24(3):671–680, 2003.
- [Poets *et al.*, 2015] Christian F. Poets, Robin S. Roberts, Barbara Schmidt, Robin K. Whyte, Elizabeth V. Asztalos, David Bader, Aida Bairam, Diane Moddemann, Abraham Peliowski, Yacov Rabi, Alfonso Solimano, and Harvey Nelson. Association between intermittent hypoxemia or bradycardia and late death or disability in extremely preterm infants. *JAMA*, 314(6):595–603, 08 2015.
- [Pons *et al.*, 2017] Jordi Pons, Oriol Nieto, Matthew Prockup, Erik M. Schmidt, Andreas F. Ehmann, and Xavier Serra. End-to-end learning for music audio tagging at scale. *CoRR*, abs/1711.02520, 2017.
- [Ribeiro *et al.*, 2016] Marco Túlio Ribeiro, Sameer Singh, and Carlos Guestrin. ”why should I trust you?”: Explaining the predictions of any classifier. *CoRR*, abs/1602.04938, 2016.
- [Rudin, 2018] Cynthia Rudin. Please stop explaining black box models for high stakes decisions. *CoRR*, abs/1811.10154, 11 2018.
- [Schmid *et al.*, 2015] M.B. Schmid, R.J. Hopfner, S. Lenhof, H.D. Hummler, and H. Fuchs. Cerebral oxygenation during intermittent hypoxemia and bradycardia in preterm infants. *Neonatology*, 107:137–146, 2015.
- [Van der Maaten and Hinton, 2008] L. Van der Maaten and G. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605, 2008.
- [Williamson *et al.*, 2013] James R. Williamson, Daniel W. Bliss, and David Paydarfar. Forecasting respiratory collapse: Theory and practice for averting life-threatening infant apneas. *Respiratory Physiology & Neurobiology*, 189(2):223 – 231, 2013.
- [Yildirim *et al.*, 2018] Ozal Yildirim, Pawel Plawiak, Ru-San Tan, and U. Rajendra Acharya. Arrhythmia detection using deep convolutional neural network with long duration ecg signals. *Computers in Biology and Medicine*, 102:411 – 420, 2018.
- [Zhou *et al.*, 2015] Bolei Zhou, Aditya Khosla, Àgata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. *CoRR*, abs/1512.04150, 2015.