2024 Third International Conference on Artificial Intelligence, Computational Electronics and Communication System (AICECS) | 979-8-3503-9124-4/24/$31.00 ©2024 IEEE | DOI: 10.1109/AICECS63354.2024.10956646

Leveraging Deep Learning for Fever Temperature Analysis and Pattern Recognition

Anirudh Prabhakaran
Dept. of ECE
NITK, Surathkal
Mangalore, India
apml2304@gmail.com

Sumam David S.
Dept. of ECE
NITK, Surathkal
Mangalore, India
sumam@nitk.edu.in

Deepu Vijayasenan
Dept. of ECE
NITK, Surathkal
Mangalore, India
deepuv@nitk.edu.in

Chakrapani Mahabala*
Dept. of Medicine
Kasturba Medical College
Mangalore, India
chakrapani.m@manipal.edu

Pradeepa Dakappa
Dept. of Pharmacology
Siddaganga Medical College and Research Institute
Tumkur, India
pradeepahd@nanjappahospitals.com

***Abstract*—Tympanic temperature is one of the most fundamental indicators for the diagnosis of diseases. Due to its importance, using patients' temperature data to aid in the diagnostic process would be beneficial. This work uses temperature data collected from various patients to classify diseases. We consider dengue, tuberculosis, and non-infectious and non-tubercular bacterial diseases. Extracting essential features from the temperature data is necessary so that the downstream layers only have to consider important features, not miscellaneous information. This feature extraction is done using two methods - Convolution Neural Networks and Autoencoders. We introduce three models for Explainable Temperature Analysis - ExTemp-Conv-SM, ExTemp-Conv-LG and ExTemp-Auto. We achieve a classification accuracy of 70% over these four disease classes. We also use explainable AI tools, like GradCAM, to identify distinguishing patterns in temperature fluctuations that can characterize diseases. We generate such patterns for all four diseases under consideration. We note that the patterns generated for dengue and tuberculosis match the findings in biological observation studies. We hope that the methods in this paper can be leveraged for other diseases and used to aid the diagnostic process.**

***Index Terms*—machine learning, deep learning, temperature analysis, fever analysis, explainable AI, medical machine learning**

I. Introduction

Integrating advanced technologies has become a pivotal force for timely intervention and precise diagnostics in the rapidly evolving healthcare landscape. While conventional imaging modalities such as X-rays, MRI, and CAT scans offer intricate diagnostic insights, their complexity and cost often impede widespread accessibility. Conversely, temperature data, easily obtainable and reflective of disease progression, presents a promising avenue for diagnostic innovation. However, relying solely on temperature data for disease prediction confronts challenges rooted in the complex interplay of individual patient physiology and disease manifestation patterns.

This paper embarks on a multifaceted exploration to address these challenges through innovative methodologies and analytical paradigms. We aim to harness temperature data as a robust diagnostic tool while acknowledging its inherent limitations. Central to our endeavour are two interconnected avenues of investigation, each poised to illuminate novel insights into disease diagnosis and prognostication.

Firstly, we propose using modern machine learning and deep learning models explicitly tailored for temperature data analysis. Departing from traditional statistical approaches, we leverage advances like CNNs and autoencoders to make and train effective models. By re-framing this problem as a multi-class classification task, we aim to enhance the discriminative power of our models and extend the scope of analysis to encompass a diverse spectrum of diseases.

Secondly, we recognize the untapped potential of temperature data as a rich source of diagnostic information beyond its conventional role as a single-dimensional metric. Through Explainable AI techniques, we seek to identify characteristic signatures indicative of the onset or progression of specific diseases. By identifying the underlying features prioritized by our models in the diagnostic decision-making process, we aim to not just rely on classification but also develop a deeper understanding of disease dynamics.

The paper is organised as follows:

- Section 1 introduces the problem statement and its intrinsic difficulties. We outline our strategy for exploring this problem.
- Section 2 reviews related work in this field, establishing baselines and providing insights into the problem based on previous work.
- Section 3 describes our methods. It is broken down into various parts, each dealing with information about datasets, models used, and motivations for these choices.

* Corresponding Author. E-mail: chakrapani.m@manipal.edu

979-8-3503-9124-4/24/$31.00 ©2024 IEEE

- Section 4 presents the results of our experiments. We divide this into two subsections, separately reporting on the classification and temperature analysis tasks.
- Section 5 discusses the outcomes of our study. We also mention exciting ideas to explore in the future.

II. Related Works

Most previous work classifying diseases based on temperature data has focused on statistical techniques. Cuesta-Frau et al. (2019) [1] used Sample Entropy to classify fever patterns. They use one single extracted entropy feature to perform classification. Cuesta-Frau et al. (2020) [2] use Slope Entropy (2019) [3]. This is a novel method of estimating time series complexity based on symbolic patterns and amplitude information. There have also been attempts at classification using machine learning techniques. Dakappa et al. (2017) [4] tried to tackle this problem using quadratic support vector machines. Dakappa et al. (2018) [5] introduce artificial neural networks for classifying temperature patterns into two classes-infectious and non-infectious diseases.

There has been intense study on temperature patterns for various diseases as well. In particular, dengue and tuberculosis were studied over 24 hours. Dakappa et al. (2018) [6] observed the 24-hour continuous tympanic temperature pattern for 15 patients with dengue fever. They noted that a tri-phasic fever pattern was seen among most patients. Similarly, Dakappa et al. (2019) [7] studied 24-hour continuous tympanic temperature in tuberculosis on 81 patients and noted that most patients exhibited a slow temperature elevation, followed by a slow temperature fall.

For our models, we mainly look into those based on convolutional models and autoencoders. 1D convolutional models are typically used for supervised tasks, such as classification or regression, where the aim is to extract relevant patterns or features from sequential data using convolutional layers. These models excel at detecting local dependencies in the data by applying convolution operations along one dimension - in our case, time. Autoencoders are primarily used for unsupervised learning, focusing on encoding input data into a compressed latent space representation and then decoding it to reconstruct the original input. Their primary use cases include dimensionality reduction, anomaly detection, and data denoising. From the autoencoder, we are using the encoder part in our model, which will be useful for dimensionality reduction.

We break our study into two parts. In the first part, we aim to create new models. We try different model architectures to identify which ones perform best. Once experiments are completed, we finalize the models with the highest performance. The second part focuses on understanding the model's decision process. With the advance in Explainable AI techniques, deep learning models are no longer a black box. We apply GradCAM to our models for different disease classes and find what patterns encourage the model to make a specific prediction.

III. Methods

A. Dataset

We have obtained the data from Kasturba Medical College, Mangalore, India with ethics clearance number IEC KMC MLR 01-14/13. The dataset contains records from 185 patients diagnosed with nine different diseases. We are considering only four of these: Dengue, Non-Infectious Diseases, Non-Tubercular Bacterial Infection and Tuberculosis. This is done considering the availability of data.

TABLE I
Counts of classes provided in the data set.

Disease	Counts
Dengue	47
Leptospirosis	15
Malaria	16
Malignancy	7
Non-Infectious Diseases	28
Non-Tubercular Bacterial Infection	37
Pyogenic Sepsis	2
Thyroiditis	1
Tuberculosis	32

The preprocessing steps for the dataset are straightforward. We extract the data from the classes of interest to us. After creating this new data, we perform a simple normalization of the data to bring all the values into a comparable range. We use `MinMaxScaler` to bring the values between 0 and 1. This data is then split into training and testing sets. We use a split of 80% of the data as training data and the remaining 20% as the test set.

TABLE II
Number of samples for each data set.

Type	Counts
Train	115
Test	29

The data is provided as a 2D data set, with each record having 1440 entries. This corresponds to the number of minutes in 24 hours, as the temperature was recorded every minute for this duration. We provide some sample plots for reference.

B. Models

Based on our experiments, we note that convolution-based models outperform the others. This was because capturing the minor temporal differences in the temperature data is essential, rather than using the entire data in one go for classification. Convolution performs this feature extraction from temporal variations very well.

Based on this observation, we decided to use a 1D convolution-based model. Previously, Singstad et al. [8] have proven that 1D Convolution Neural Networks in medical applications perform well. We provide two models, named **ExTemp-Conv-SM** and **ExTemp-Conv-LG**. Both are based on convolution blocks, followed by a fully connected neural network to perform classification. The difference between the

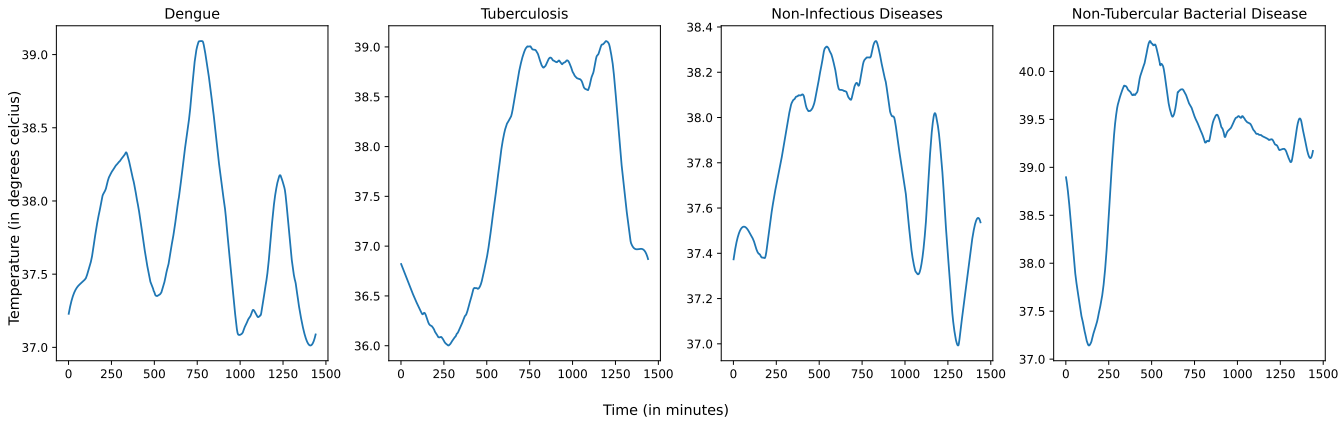

Fig. 1. Sample Temperature Patterns for Dengue, Tuberculosis, Non-Infectious Diseases and Non-Tubercular Bacterial Diseases. From the visualization of the data points provided, there is quite a fluctuation in the patterns of diseases like Dengue and Tuberculosis. Compared to that, temperature graphs for Non-Tubercular Bacterial Diseases and Non-Infectious Diseases are comparatively smoother, although they have fluctuations.

two models comes from slight differences in architecture, which leads to a difference in the number of parameters for each model. Different numbers of convolution layers are used to properly experiment with extracting low and high level features. We also note that data has to be slightly modified for these two models. Instead of using the standard 2D data provided, we need to reshape the data to be used by convolution. This is because the convolution block expects data to be $(num\_channels, num\_features)$ instead of our current shape of data, which is $(num\_features)$. We have reshaped the data so that the number of channels in the input data is 1, i.e. reshaping leads to the data input of shape $(1, 1440)$.

There has also been previous work to denote that convolution blocks act as an effective method for dimensionality reduction [9]. We explore this idea more and look at other promising methods for dimensionality reductions. Hinton et al. [10] show that autoencoders can be effectively used for feature extraction and dimensionality reduction. Using the same thought process, we replace the convolution blocks in our proposed network with the encoder block of an autoencoder. This motivates the third model - **ExTemp-Auto**.

TABLE III
DETAILS ABOUT EXTEMP FAMILY OF MODELS

Parameter	**ExTemp-Conv-SM**	**ExTemp-Conv-LG**	**ExTemp-Auto**
d_model	(1, 1440)	(1, 1440)	(1440,)
Layers	24	14	10
Feature Layers	15	5	2
Classifier Layers	8	8	8
Feature Params	1,464	40	17,556
Classifier Params	369,236	745,892	368,896
Total Params	370,700	745,932	386,452

One-dimensional Convolutional Neural Networks (1D CNNs) [11] is a variant of the traditional CNN architecture tailored for processing sequential data such as time series, signals, and text. They employ convolutional layers that slide one-dimensional kernels along the input sequence, capturing local patterns and dependencies. 1D CNNs excel at learning meaningful representations from sequential data, automatically capturing patterns and temporal dependencies. Their ability to learn hierarchical features and exploit local correlations in sequential data has led to widespread adoption across various domains, highlighting their versatility and effectiveness in modelling and analyzing sequential information.

Autoencoders [10] is a type of neural network architecture employed in unsupervised learning tasks, primarily for dimensionality reduction, feature learning, and data generation. Combining an encoder and a decoder, autoencoders aim to reconstruct their input data as accurately as possible. The encoder compresses the input into a latent space representation, typically of lower dimensionality, while the decoder attempts to reconstruct the original input from this representation. Through this process, autoencoders learn to capture meaningful features and patterns in the data, effectively reducing its dimensionality while preserving important information.

C. Loss Calculation

The convolution-based models are trained using the standard cross entropy loss [12] that is used for other classification tasks. For the autoencoder-based model, we use a slightly different approach. Since two models are in play here, we use two corresponding loss computations and optimization criterion functions. We use the cross-entropy loss for the classification part of the model. However, for the autoencoder model, we use the Mean Squared error. We then summate these losses and apply our back-propagation techniques on this computed **Total Loss**. We define the Total Loss as the following:

$$\text{Total Loss} = \text{Cross-Entropy Loss} + \text{MSE} \quad (1)$$

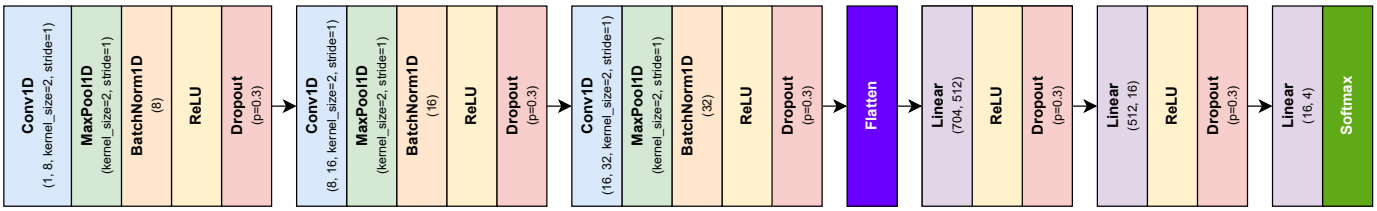

Fig. 2. ExTemp-Conv-SM Model

Fig. 3. ExTemp-Conv-LG Model

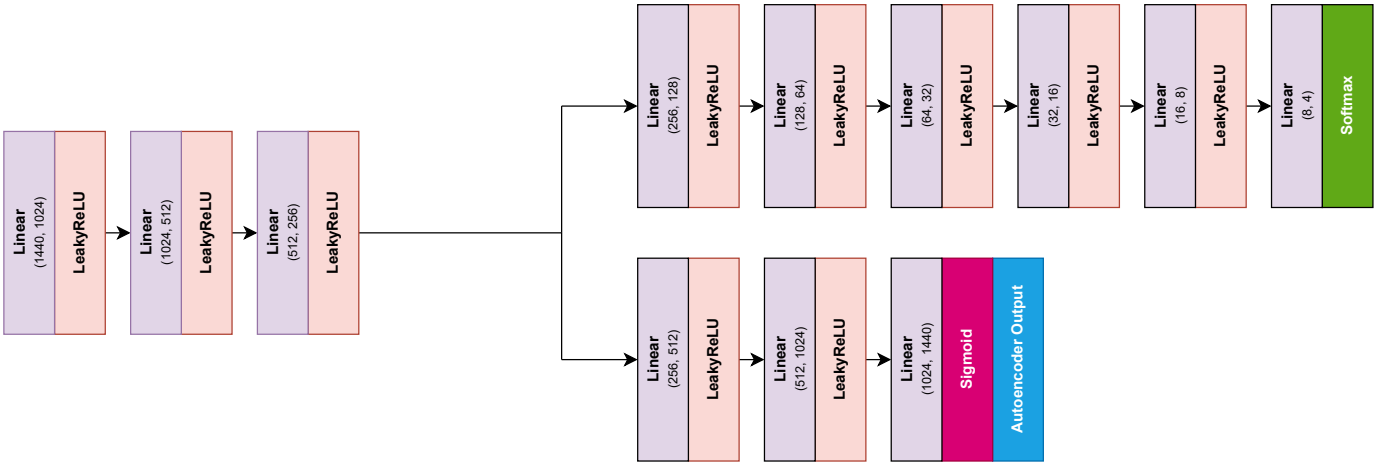

Fig. 4. ExTemp-Auto Model

Gradient-weighted Class Activation Mapping, also known as GradCAM [13], is one of the most popular Explainable AI techniques. It is a technique widely utilized in computer vision and deep learning to interpret the decisions made by convolutional neural networks (CNNs). By computing gradients of the target class score with respect to the feature maps of the last convolutional layer, GradCAM identifies the crucial regions within an input image that contribute most significantly to the model's prediction for a particular class. GradCAM enhances the interpretability and transparency of CNNs, aiding in validating and understanding their predictions. We leverage the `IntegratedGradients` function provided by Captum [14]. This module employs a few more optimizations to increase the computation speed of Class Activation Maps. It also helps quickly adapt the GradCAM code to a 1-dimensional use case, similar to what we used in this paper.

IV. Results

A. Classification

We propose three models for Explainable Temperature analysis [1] - **ExTemp-Conv-SM**, **ExTemp-Conv-LG** and **ExTemp-Auto**. Hyperparameters for these models can be found in (IV). These hyperparameters were discovered as ideal with the help of the `RayTune` [15] project. The models were trained using Adam [16] optimizer. As noted previously, loss functions are a cross-entropy loss for the convolution-based models. Total loss consists of cross-entropy and mean squared loss for the autoencoder-based models.

TABLE IV
Ideal Hyperparameter values for ExTemp Models

Parameter	**Value**
Number of Epochs	1000
Learning Rate	0.00624
Dropout Rate	0.4

[1]Code can be found at https://github.com/anirudhprabhakaran3/extemp

The ExTemp-Conv-SM is called the ”small” model because of the lower number of parameters than the other convolution-based model. However, note that this model contains more convolution blocks for feature extraction. Their performance is summarized in the tables below.

TABLE V
SUMMARY OF TRAINING METRICS.

model	precision	recall	f1-score	accuracy
ExTemp-Conv-SM	0.98	0.97	0.97	0.97
ExTemp-Conv-LG	0.91	0.90	0.91	0.90
ExTemp-Auto	0.77	0.74	0.74	0.74

TABLE VI
SUMMARY OF TESTING METRICS.

model	precision	recall	f1-score	accuracy
ExTemp-Conv-SM	0.62	0.55	0.62	0.65
ExTemp-Conv-LG	0.60	0.57	0.55	**0.70**
ExTemp-Auto	0.56	0.52	0.51	0.52

We note that the ExTemp-Conv-LG model performs the best out of these three models in terms of the testing data. We also note that the performance of ExTemp-Auto on the test data is lower than that of the convolution-based models. It is likely due to the lack of data - a lot of data is required for the encoding model to learn proper and valid encodings. Another trend we note with this model is that the class with the least support especially suffers. We attribute this to the inability to properly learn the characteristics of that particular disease's temperature variations.

B. Temperature Analysis

Below, we present the results of the GradCAM analysis. We only report results from the ExTemp-Conv-LG model, which is the best-performing one. We plot the points that have strongly contributed to classifying a particular disease. We define ”strongly positively” contributed if the class activation map value is more than the mean positive class activation map value. We only consider the mean of the positive values, as data points might be negatively contributing toward classification. This means points that encourage the model not to consider the selected class. Although this might interest other analytical studies, we do not use this metric. We are more interested in finding which patterns contribute to the classification decision.

Examining the graphs for tuberculosis in (IV-B), it becomes evident that the model accurately captures the gradual rise in temperature, maintenance of this trend, followed by a gradual decline. This observed pattern aligns closely with the findings of Dakappa et al. [7], indicating a consistent temporal rhythm in temperature variations. Furthermore, the model demonstrates attentiveness to finer details, such as saddle points and other fluctuations within the overarching trend.

Here, for dengue, in (IV-B), the model accurately recognizes the temperature's cyclic nature, mirroring observed increases and decreases. These patterns correspond with the tri-phasic temperature phenomena outlined by Dakappa et al. [6]. The model can do this even with noisy dengue temperature data.

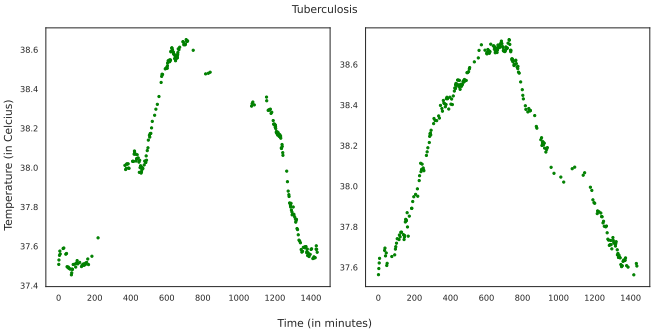

Fig. 5. Plots for Tuberculosis generated by ExTemp-Conv-LG

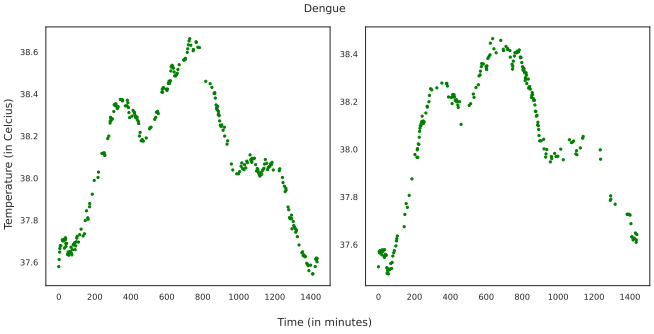

Fig. 6. Plots for Dengue generated by ExTemp-Conv-LG

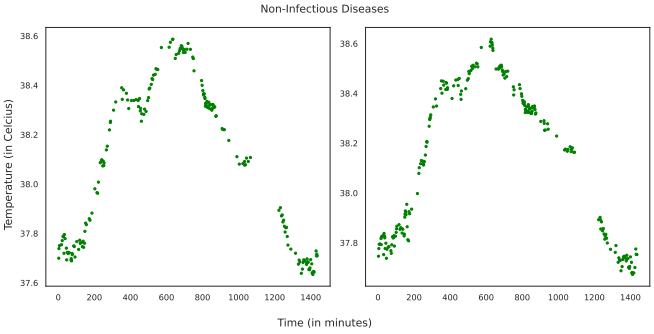

Fig. 7. Plots for Non-Infectious Diseases generated by ExTemp-Conv-LG

In this graph for non-infectious diseases in (IV-B), the temperature fluctuates mildly throughout the day, with a slight increase around mid-morning followed by a gradual decline. However, there is a sharp spike in temperature around noon, peaking in the afternoon before gradually returning to normal levels by evening.

Here, for non-tubercular bacterial diseases in (IV-B), the temperature starts to rise steadily around late morning, peaking in the early afternoon before gradually declining toward

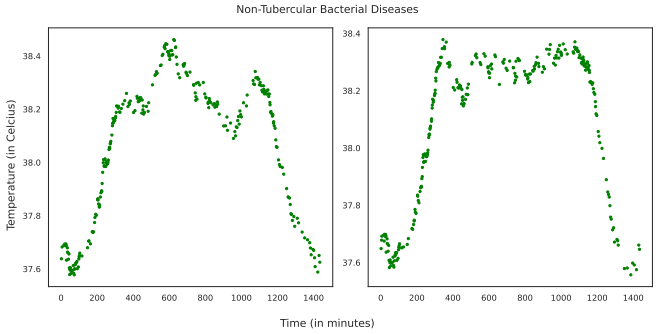

Fig. 8. Plots for Non-Tubercular Bacterial Diseases generated by ExTemp-Conv-LG

evening and becomes relatively stable across the mean range of temperature. The temperature rise is more stable at noon than the previous graph for the same class.

V. Discussion and Future Scope

In this paper, we explored two tasks. First, we looked into deploying deep learning algorithms and techniques for disease classification from 24-hour tympanic temperature data. We devised three new models for this use case - two based on convolution and one based on autoencoders. Secondly, we looked at leveraging explainable AI tools to understand how models made these decisions. Observing those trends improved our understanding of temperature patterns for various diseases.

In conclusion, this paper aims to contribute to the ever-growing field of medical machine learning. The results and outcomes of this paper aim to be an effective tool available in the arsenal for non-invasive diagnostics. These findings have important implications in diagnostic and clinical settings and can be used with other traditional differential diagnosis techniques and AI-enhanced medical procedures. A couple of improvements can be made to the model itself. One essential thing would be to gather more data to train our models better. We can also extend this model to other diseases not considered in this study, to analyse the behaviour and effectiveness of our models and methods.

Another important step would be to get a control set of data points corresponding to non-fever people. This will allow the model to learn the baselines, improving its understanding of the patterns. It would also extremely helpful if we could incorporate doctors' knowledge about various diseases to help in model prediction. Creation of a framework that allows us to include this knowledge is something to be worked on. Finally, we have implemented a simple Autoencoder in the autoencoder model. However, we can use more sophisticated techniques like Variational Autoencoders (VAE) and Convolutional Variational Autoencoders.

References

[1] D. Cuesta Frau, P. Miró Martínez, S. Oltra Crespo, A. Molina Picó, B. Vargas, P. González, C. Mahabala, and P. H. Dakappa, "Classification of fever patterns using a single extracted entropy feature: A feasibility study based on sample entropy," *Mathematical Biosciences and Engineering*, vol. 17, no. 1, pp. 235–249, 2019.

[2] D. Cuesta-Frau, P. H. Dakappa, C. Mahabala, and A. R. Gupta, "Fever time series analysis using slope entropy. application to early unobtrusive differential diagnosis," *Entropy*, vol. 22, no. 9, p. 1034, 2020.

[3] D. Cuesta-Frau, "Slope entropy: A new time series complexity estimator based on both symbolic patterns and amplitude information," *Entropy*, vol. 21, no. 12, p. 1167, 2019.

[4] P. H. Dakappa, K. Prasad, S. B. Rao, G. Bolumbu, G. K. Bhat, C. Mahabala, *et al.*, "A predictive model to classify undifferentiated fever cases based on twenty-four-hour continuous tympanic temperature recording," *Journal of Healthcare Engineering*, vol. 2017, 2017.

[5] P. H. Dakappa, K. Prasad, S. B. Rao, G. Bolumbu, G. K. Bhat, and C. Mahabala, "Classification of infectious and noninfectious diseases using artificial neural networks from 24-hour continuous tympanic temperature data of patients with undifferentiated fever," *Critical Reviews™ in Biomedical Engineering*, vol. 46, no. 2, 2018.

[6] P. H. Dakappa, S. B. Rao, G. Bhat, *et al.*, "Tri-phasic fever in dengue fever.," *Tropical Doctor*, vol. 48, no. 2, pp. 93–97, 2018.

[7] P. H. Dakappa, S. B. Rao, G. K. Bhat, and C. Mahabala, "Unique temperature patterns in 24-h continuous tympanic temperature in tuberculosis," *Tropical Doctor*, vol. 49, no. 2, pp. 75–79, 2019.

[8] B.-J. Singstad and C. Tronstad, "Convolutional neural network and rule-based algorithms for classifying 12-lead ecgs," in *2020 Computing in Cardiology*, pp. 1–4, IEEE, 2020.

[9] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller, "Striving for simplicity: The all convolutional net," *arXiv preprint arXiv:1412.6806*, 2014.

[10] G. E. Hinton and R. R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *Science*, vol. 313, no. 5786, pp. 504–507, 2006.

[11] S. Kiranyaz, O. Avci, O. Abdeljaber, T. Ince, M. Gabbouj, and D. J. Inman, "1d convolutional neural networks and applications: A survey," *Mechanical Systems and Signal Processing*, vol. 151, p. 107398, 2021.

[12] A. Mao, M. Mohri, and Y. Zhong, "Cross-entropy loss functions: Theoretical analysis and applications," in *International Conference on Machine Learning*, pp. 23803–23828, PMLR, 2023.

[13] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-cam: Visual explanations from deep networks via gradient-based localization," in *Proceedings of the IEEE International Conference on Computer Vision*, pp. 618–626, 2017.

[14] N. Kokhlikyan, V. Miglani, M. Martin, E. Wang, B. Alsallakh, J. Reynolds, A. Melnikov, N. Kliushkina, C. Araya, S. Yan, and O. Reblitz-Richardson, "Captum: A unified and generic model interpretability library for pytorch," 2020.

[15] R. Liaw, E. Liang, R. Nishihara, P. Moritz, J. E. Gonzalez, and I. Stoica, "Tune: A research platform for distributed model selection and training," *arXiv preprint arXiv:1807.05118*, 2018.

[16] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.