

MultiWave: Multiresolution Deep Architectures through Wavelet Decomposition for Multivariate Time Series Prediction

Iman Deznabi
Madalina Fiterau

Manning College of Information & Computer Sciences, University of Massachusetts Amherst, MA, USA

IMAN@CS.UMASS.EDU
MFITERAU@CS.UMASS.EDU

Abstract

The analysis of multivariate time series data is challenging due to the various frequencies of signal changes that can occur over both short and long terms. Furthermore, standard deep learning models are often unsuitable for such datasets, as signals are typically sampled at different rates. To address these issues, we introduce MultiWave, a novel framework that enhances deep learning time series models by incorporating components that operate at the intrinsic frequencies of signals. MultiWave uses wavelets to decompose each signal into sub-signals of varying frequencies and groups them into frequency bands. Each frequency band is handled by a different component of our model. A gating mechanism combines the output of the components to produce sparse models that use only specific signals at specific frequencies. Our experiments demonstrate that MultiWave accurately identifies informative frequency bands and improves the performance of various deep learning models, including LSTM, Transformer, and CNN-based models, for a wide range of applications. It attains top performance in stress and affect detection from wearables. It also increases the AUC of the best-performing model by 5% for in-hospital COVID-19 mortality prediction from patient blood samples and for human activity recognition from accelerometer and gyroscope data. We show that MultiWave consistently identifies critical features and their frequency components, thus providing valuable insights into the applications studied.

Data and Code Availability We are evaluating our models on three publicly available datasets, the Wearable Stress and Affect Detection (WE-SAD) (Schmidt et al., 2018), and MHEALTH dataset (Banos et al., 2014, 2015) which both can be downloaded from the UCI Machine Learning Repository, and the COVID-19 dataset Yan et al. (2020). The

code is available in <https://github.com/Information-Fusion-Lab-Umass/MultiWave>.

Institutional Review Board (IRB) This research does not require IRB approval.

1. Introduction

Multivariate time series prediction and forecasting are critical tasks in healthcare, as they enable the analysis and prediction of patient outcomes and the allocation of resources based on historical data. These applications play a key role in improving patient care and outcomes and are increasingly important in a healthcare industry that is becoming more data-driven (Wiens and Shenoy, 2018). However, the final prediction in these applications can depend on many factors, such as information at different frequencies, as well as long-term and short-term changes in input signals. Moreover, in many tasks, observations come from multiple sources and are often collected at various sampling rates. Here, we propose *a model-agnostic approach that can leverage temporal dependencies at different frequencies and scales* in multivariate time series data. This approach is adept at handling multirate time series data, which refers to data collected at multiple sampling rates, through the use of multilevel discrete wavelet decomposition.

We have focused on three different applications of multivariate and multirate time series prediction: (1) stress and affect prediction using wearable devices, (2) COVID-19 in-hospital mortality prediction from blood samples collected over time, and (3) human activity recognition using wearable devices. The dataset corresponding to the first application is the Wearable Stress and Affect Detection (WE-SAD) dataset (Schmidt et al., 2018). This dataset includes physiological response data collected from a chest and a wrist-worn device from 15 subjects who were exposed to a baseline session, an amusement ses-

sion, and a stress session. Our goal was to predict to which session the subjects were exposed on the basis of the physiological data collected. For COVID-19 in-hospital mortality prediction, we used the dataset provided by Yan et al. (2020). This dataset includes 74 biomarkers collected from 375 patient blood samples between January 10 and February 18, 2020, at Tongji Hospital in Wuhan, China. Our objective was to predict in-hospital mortality based on the collected time series of biomarkers. For human activity recognition, we used the Mobile Health (MHEALTH) dataset (Banos et al., 2014, 2015), which contains wearable sensor data from 10 volunteers performing 12 different physical activities.

Deep learning-based methods that are introduced into time series analysis, such as Recurrent Neural Networks (RNN) (Williams and Zipser, 1989), Convolutional Neural Networks (CNN) (Zheng et al., 2016) and more recently transformers (Wen et al., 2022) achieve state-of-the-art results in these and similar applications (Lai et al., 2022; Tipirneni and Reddy, 2021; Huang et al., 2022; Sun et al., 2021; Dzieżyc et al., 2020). However, they have three notable shortcomings in handling multivariate time series data. First, these methods utilize only the information in the time domain and overlook the information available in the frequency domain of the signals. Second, they lack transparency in identifying the critical signals and subsignals that are relevant to the task at hand. Finally, they are unable to process signals that have been collected with varying frequencies (multirate signals) in their original form, which can lead to the introduction or loss of important temporal dependencies during up-sampling or down-sampling to a single rate (Che et al., 2018a; Tipirneni and Reddy, 2021; Che et al., 2018b).

We propose a new, model-agnostic approach to address these shortcomings. This framework employs discrete wavelet decomposition to segment signals into various frequency components (*subsignals*). After eliminating irrelevant subsignals, it organizes subsignals with similar frequencies into individual time series models. Finally, the output from these models is combined to generate a prediction. This framework brings the following improvements to multivariate time series modeling: 1) Uses the information available in both the time and frequency domains. 2) Provides unique insight into which frequencies of the signals are important for a given task. 3) Makes model-agnostic improvements to time

series models.

Our results demonstrate that MultiWave enhances the performance of multivariate time series prediction models in wearable stress and affect prediction, COVID-19 mortality prediction, and human activity recognition. Moreover, MultiWave’s capability to reveal useful signals and frequency components results in improved adaptability of deep learning models in these applications.

2. Related work

Multirate time series classification Methods in time series analysis that are based on deep learning, such as Recurrent Neural Networks (RNN) (Williams and Zipser, 1989), Convolutional Neural Networks (CNN) (Zheng et al., 2016), and more recently, Transformers (Wen et al., 2022), have demonstrated state-of-the-art results in numerous applications (Lai et al., 2022; Huang et al., 2022). However, as mentioned earlier, they fall short when handling multirate time series data. There are several approaches proposed to solve this problem, such as Che et al. (2018a); Tipirneni and Reddy (2021) and Chang et al. (2023), which can inherently handle irregularly sampled time series data and thus can model multirate time series data without aligning the signals (Sun et al., 2020; Shukla and Marlin, 2020). There are other methods, such as Che et al. (2018b); Armesto et al. (2008); Safari et al. (2014) that use architectures specifically developed for multirate time series data. All these approaches only consider the information available in the time domain of the time series data, while MultiWave is specifically developed for multirate data and is able to leverage the information in the frequency domain.

Frequency analysis of time series Frequency analysis of time series is an extensively studied subject in the signal processing community. Methods such as the discrete Fourier transform (Bracewell and Bracewell, 1986), the discrete wavelet transform (Daubechies, 1992), and the Z-Transform (Foster, 1996) have been used to analyze time series. For deep learning models, similar methods have been used in the preprocessing steps (Cui et al., 2016; Yuan et al., 2017; Song et al., 2021; Salim et al., 2023) or as part of neural networks (Koutnik et al., 2014; Lee-Thorp et al., 2021). Most of these models focus on univariate time series data and cannot be used directly on multivariate and multirate time series data.

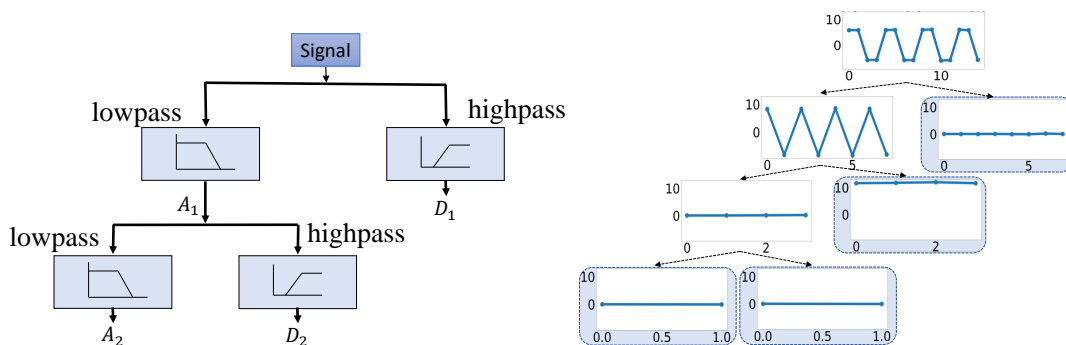


Figure 1: The process of multi-level discrete wavelet decomposition which uses lowpass and highpass filters to decompose signals. The image on the left depicts this process, while the image on the right showcases the decomposition of a signal using the Haar wavelet. As observed, the resulting signals demonstrate a value of zero, except for the signal component that corresponds to the true frequency of the original signal.

Wavelet decomposition Wavelet decompositions (Daubechies, 1992) are well-known methods for capturing information in time series in both the time and frequency domains. They have been used successfully as a preprocessing step for neural networks (Liu et al., 2013; Wang et al., 2020a; Alhnaity et al., 2021; Kim et al., 2021; Althelaya et al., 2021; Zucatelli et al., 2021) and as an integral part of them (Subasi et al., 2006; Zhang et al., 1995; Wang et al., 2018; Guo et al., 2022; Li et al., 2021; Zhuang et al., 2022; Zhou et al., 2022). Wang et al. (2018) proposes the methodology closest to our approach, implementing a trainable wavelet decomposition framework, which can be trained with the rest of the network. Although similar to our method, this paper uses wavelet decomposition to extract frequency-based information from datasets and models them using different model components; it is not suitable for utilization with multivariate and multirate time series data and does not incorporate the use of feature masks to eliminate frequency components that are not useful within the framework. Furthermore, their approach suggests a specific model to be employed, whereas our methodology is model-independent and can be integrated with any time series model.

3. Background

3.1. Notation

We denote multivariate and multirate time series data with m signals collected before time T as a set of signals $X^{1:T} = \{x_1^{1:T}, x_2^{1:T}, \dots, x_m^{1:T}\}$ where each signal is collected at sampling rates

$\{f_s(x_1), f_s(x_2), \dots, f_s(x_m)\}$. The length of each signal is proportional to its collected rate $Len_i = T \times f_s(x_i)$. Given $X^{1:T}$, we want to predict a label y , which can be continuous (regression) or discrete (classification). In the rest of the paper, we will remove the time indication for the signals and show the set of signals as X and the signal i as x_i .

3.2. Multilevel Discrete Wavelet Decomposition

We use wavelet decomposition to break down the signals into different frequencies. Multilevel discrete wavelet decomposition can extract multilevel time-frequency features from a time series by iteratively applying low-pass and high-pass filters derived from wavelets to the signal. The formula for this decomposition is given below:

$$x(t) = \sum_k A_{L,k} \phi_{L,k}(t) + \sum_k D_{L,k} \Psi_{L,k}(t) + \sum_k D_{L-1,k} \Psi_{L-1,k}(t) + \dots + \sum_k D_{1,k} \Psi_{1,k}(t)$$

$\Psi_{s,\tau}$ is the mother wavelet with scale s and time τ and ϕ is the father wavelet. This multilevel wavelet decomposition converts the input signal $x(t)$ into signals $A_L = \bigcup_k A_{L,k}$, which is a coarse general approximation of the signal (low frequency) and the detail coefficients $D_L = \bigcup_k D_{L,k}, D_{L-1} = \bigcup_k D_{L-1,k}, \dots, D_1 = \bigcup_k D_{1,k}$ that influence the function on various scales. Figure 1 depicts this decomposition. To simplify the notation, we show the decomposition of a signal x as a set $S(x) = \{D_1, D_2, \dots, D_L, A_L\}$ that includes signals retrieved when decomposing $x(t)$ at L levels.

This notation allows us to denote D_1 as $S_1(x)$, D_2 as $S_2(x)$, \dots , D_L as $S_L(x)$, and A_L as $S_{L+1}(x)$.

Many different wavelets are introduced in the literature, such as Haar, Daubechies, and Biorthogonal (Haar, 1909; Cohen et al., 1992; Daubechies, 1992). Our framework is independent of the type of wavelet used.

4. MultiWave

4.1. Signal Decomposition and Frequency Grouping

The overall structure of the framework is illustrated in Figure 2 for the case where the signals x_1 and x_2 are collected at frequencies of $64Hz$ and $30Hz$, respectively. The framework uses discrete wavelet decomposition to decompose each signal into different frequency groups. For a set of m input signals, $X = \{x_1, x_2, \dots, x_m\}$, a set of decomposed signals $\mathcal{S}(X) = \{S(x_1), S(x_2), \dots, S(x_m)\}$, will be generated. If all signals are sampled at the same rate, all elements of $\mathcal{S}(X)$ will have the same number of levels (denoted as L). In this case, all signals will have the same frequency at each level of decomposition; thus, all frequency components at the same level are grouped into a model component, denoted Φ_j . The input of the component j is $I_j = \{S_j(x_1), \dots, S_j(x_m)\}$ and the number of components is equal to the level of wavelet decomposition L . The outputs of all models are concatenated in the end and passed through a fully connected layer to obtain the final output.

$$\hat{y} = FC(\Phi_1(I_1) \oplus \Phi_2(I_2) \oplus \dots \oplus \Phi_L(I_L)) \quad (1)$$

Here, \oplus shows the concatenation operation. If the signals are collected with different sampling rates, we group the subsignals with frequencies that are close to each other together into a model component. In this case, the number of model components is equal to the maximum level of signal decomposition $L_{max} = \max(\{L_1, L_2, \dots, L_m\})$, which is determined by the signal with the highest sampling rate. Here, L_i represents the decomposition level for signal i . To better illustrate this, assume, without loss of generality, that we have two signals $m = 2$ where signal x_2 is collected twice as often as signal x_1 , $f_s(x_2) = 2 \times f_s(x_1)$. In this case, $L_2 = L_1 + 1$. Since the frequencies are reduced by half at each level of wavelet decomposition, we have:

$$f_s(S_j(x_1)) = f_s(S_{j+1}(x_2))$$

To ensure that the signals have consistent rates in each component, the input for the first component of the model is defined as $I_1 = \{S_1(x_2)\}$. The input for the subsequent components, j , is expressed as: $I_j = \{S_{j-1}(x_1), S_j(x_2)\}$. The difference between the number of levels between the decomposition of two signals is $\log(\frac{f_s(x_2)}{f_s(x_1)})$, so if this proportion is not a power of 2, the shorter signal at each component level, that receives both signals, is oversampled to the closest power of 2 to match the corresponding level of the other signal. For instance, if $f_s(x_2) = 6 \times f_s(x_1)$ in the above example, the input will be:

$$I_1 = \{S_1(x_2)\}, I_2 = \{S_2(x_2)\}, I_3 = \{S_3(x_2), S_1(x_1)\} \\ \dots, I_{L_2} = \{S_{L_1}(x_1), S_{L_2}(x_2)\}$$

since, for the component Φ_3 , $\frac{f_s(S_3(x_2))}{f_s(S_1(x_1))} = 1.5$, $S_1(x_1)$ is oversampled by a proportion of 1.5. Note that the oversampling proportion will always be less than 2. Although this oversampling may lead to the same issues mentioned earlier, using the proposed method will considerably decrease the degree of oversampling, and the issues mentioned earlier would not occur to the same extent. Figure 3 shows how the decomposition and grouping procedures work for two signals with different sampling rates.

Using this approach, the components trained on lower frequencies learn long-term changes in the data, while the faster frequency components learn short-term fast-changing trends in the data. Furthermore, since the signals are grouped with respect to their sampling rates, the signals that are input into each component have similar frequencies, which significantly reduces the amount of oversampling.

4.2. Masking frequency components

Not all the frequency components of all signals are important for the prediction of the final label. To filter these frequencies, we introduce a learnable mask for each frequency component of the signals. We use ReLU activations (Nair and Hinton, 2010) on the weights of the model to mask the non-informative components by setting the mask to zero so the input of the component i in the model is defined as:

$$I_j = \{ReLU(W_j^{(1)})S_j(x_1), ReLU(W_j^{(2)})S_j(x_2), \\ \dots, ReLU(W_j^{(m)})S_j(x_m)\}$$

To force the model to use a sparse mask for the feature components, we add the ℓ_1 norm of the weights

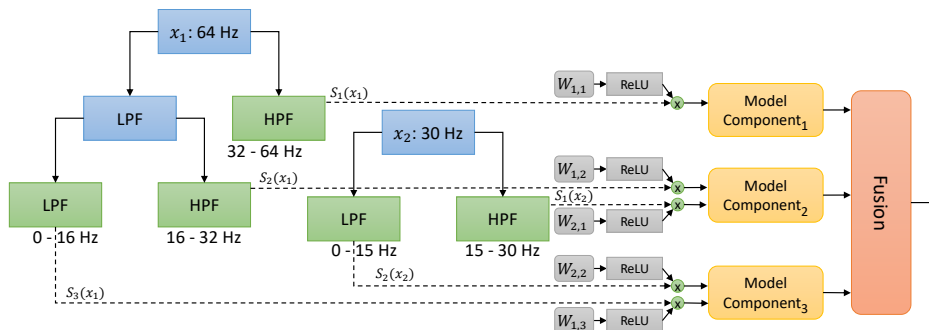


Figure 2: The structure of MultiWave with two signal inputs: one sampled at 64Hz (x_1) and the other at 30Hz (x_2). Utilizing wavelet decomposition, the signals are decomposed into their corresponding frequency components using low-pass filters (LPF) and high-pass filters (HPF), which are subsequently grouped into distinct model components. Notably, this architecture is model-agnostic, thereby enabling its applicability to any model designed to process multivariate time series data.




Figure 3: The decomposition and grouping of two signals sampled at 64Hz (x_1) and the other one at 30Hz (x_2). Three subsignals are extracted from x_1 through two levels of discrete wavelet decomposition, while x_2 is decomposed into two components in a single level. The subsignals are then organized into groups based on their frequency ranges.

to the final loss of the model:

$$Loss = \mathcal{L} + \alpha l_1(ReLU(W))$$

where \mathcal{L} is the normal loss of the model, which is usually defined by Mean-Squared-Error (MSE) for regression models and Cross-Entropy-Loss for classification models. α is a hyperparameter that determines the weight of the regularization term and sets a trade-off between minimizing the mask weights and the loss of the base model.

4.3. Final fusion of components

We conducted experiments using various fusion techniques such as attention, mean, weighted average, ensemble methods, transformer fusion, and others like GradBlend (Wang et al., 2020b) and Efficient Low-

rank Multimodal Fusion with Modality-Specific Factors (Liu et al., 2018; Du et al., 2018). However, we discovered that concatenating and using a fully connected layer produced the best results. This may be because of the unique structure of the input data, where each input provides relevant information for the target and needs to be combined to produce the output. Additionally, each input’s contribution to the final prediction may not be dependent on the other inputs. It is worth noting that MultiWave is not limited to any particular fusion technique, and different methods may work better for different datasets.

4.4. Model training

In our experiments, we used the Adam optimizer (Kingma and Ba, 2014) to train our models. Additionally, the feature masks of the baseline model, which takes in the original signals $X^{1:T} = x_1^{1:T}, x_2^{1:T}, \dots, x_m^{1:T}$ as input, were incorporated as a supplementary component in the fusion model. This allows MultiWave to revert to the baseline model (using a form of early fusion for frequencies) if the incorporation of frequency components does not lead to improvement, while still having the option to utilize frequency components for enhanced performance. We also used weights and biases (Biewald, 2020) to track and record the experiments, and Pytorch (Paszke et al., 2019) to implement and train our models.

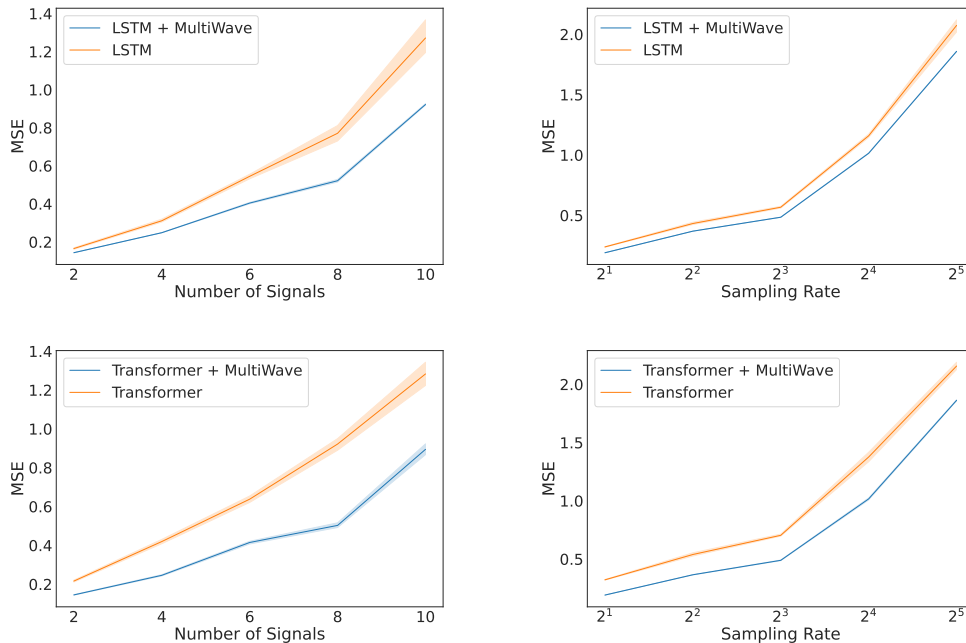


Figure 4: The Mean Squared Error (MSE) results for a synthetic dataset, with two columns displaying the effects of varying the number of signals and the difference between the signal sampling rates, respectively. The first row of the figure showcases the results for the LSTM model, while the second row illustrates the results for Transformers. Notably, the utilization of Multiwave results in improved performance of the models. This improvement can be attributed to Multiwave’s ability to identify and filter out irrelevant frequency components of the signals, consequently reducing the need for upsampling and mitigating errors caused by imputation.

5. Experiments

In this section, we evaluate the performance of MultiWave in synthetically generated and real-world datasets. We used Haar wavelet functions (Haar, 1909) throughout the experiments and, as mentioned in Section 4.3, we used a concatenation and a fully connected layer for the fusion of the output of model components.

5.1. Synthetically generated data

To determine the effectiveness of MultiWave in handling signals with different frequencies and sampling rates, we generated synthetic data. The generated data consist of multiple square signals, each with a different frequency and amplitude. The amplitude of each signal is randomly selected from a uniform distribution ranging from 0-10. We then add uniform noise to the input data with an amplitude of 3. The

label y is the sum of the amplitudes of the generated signals before the noise is added. The task is to predict the label y from given time series data. Figure 4 shows the results of the two experiments. For each experiment, we ran each setting five times with different random seeds and displayed the mean and standard deviation of these experiments. In the first experiment, two square wave signals were used as input with frequencies of 1Hz and 2Hz, sampled at a rate of 128Hz for a duration of 1 second. Then, additional signals were iteratively added with frequencies of 4Hz, 4Hz, 8Hz, 8Hz, 16Hz, 16Hz, 32Hz, and 32Hz, respectively, all sampled at the same rate of 128Hz. The MultiWave model was used with both LSTM and Transformer components, and the results are shown in the first column of Figure 4. The MultiWave model was able to improve performance steadily and reliably with the addition of new signals, owing to its

ability to identify significant frequency components, eliminate noise, and make more precise predictions.

In the second experiment, two square wave signals were generated with frequencies of 2Hz and 4Hz, sampled at rates of 64Hz and 128Hz, respectively, for 1 second. The sampling rate of the first signal was then reduced to 32Hz, 16Hz, 8Hz, and 4Hz. The results for the LSTM and Transformer models, with and without the use of MultiWave, are shown in the second column of Figure 4. The MultiWave model consistently improved the performance of the baseline models by grouping signals with similar sampling rates, which reduced the need for imputation and the potential for errors that can affect the accuracy of the model predictions.

Synthetic data experiments with frequencies other than the powers of two are in Appendix Section D.

5.2. The Wearable Stress and Affect Detection (WESAD)

Wearable Stress and Affect Detection (WESAD) (Schmidt et al., 2018) is a publicly available multi-modal dataset for stress and affect detection. WESAD contains physiological response data from 15 subjects during three sessions of baseline, amusement, and stress. The baseline session is 20 minutes, where the subject is doing a neutral reading task, the amusement session is watching a set of funny videos for 392 seconds, and the stress session is when the subject is exposed to the Trier Social Stress Test (Kirschbaum et al., 1993) for 10 minutes. During these sessions, various physiological measurements are taken using a chest-worn device and a wrist-worn device. Measurements include blood volume pulse (BVP), electrocardiogram (ECG), electrodermal activity (EDA), electromyogram (EMG), respiration (RESP), temperature (TEMP), and accelerometer (ACC). The chest-worn device collects data at a frequency of 700 Hz, while the wrist-worn device collects data at frequencies of 64 Hz, 32 Hz, and 4 Hz. We followed the data preprocessing method described in Dzieżyc et al. (2020) but used sampling rates that are powers of 2 to obtain more consistent signals. More details on the preprocessing of the signals in this dataset are given in the Supplementary Section A.1.

The results achieved in this dataset are shown in Table 1. To our knowledge, the fully convolutional per channel neural network (FCN-PC) described in the paper Dzieżyc et al. (2020) is the state-of-the-art

deep learning model for this dataset in these prediction settings, and MultiWave brings more than 5% improvement to this model. MultiWave improves the performance of all baseline models since it allows the model components to learn short-term and long-term changes and can combine these multirate signals without the need for alignment and imputation. Table 2 shows the signals and their frequency components consistently selected by the MultiWave model. The lower frequencies of the accelerometer data are selected, which is consistent with previous research (Tatler et al., 2018; Rooney et al., 2022; Alhassan and Robinson, 2010). Other physiological data were also consistently selected at lower frequencies, which aligns with earlier research that highlights the significance of summary statistics for this type of data (Rooney et al., 2022). More details about these experiments are given in the Supplementary Section A. Additionally, to investigate whether the performance improvements of the models are solely attributable to the inclusion of frequency features, we constructed a naive model that utilizes both FFT features and the signal itself. The architecture of this model is illustrated in Appendix Section E. We used the same model component for both FFT features and the input signal and concatenated the results before passing them through a 2-layer fully connected network. The experimental results are presented in Table 3.

5.3. COVID-19

The COVID-19 dataset (Yan et al., 2020) is a publicly available dataset that contains 74 indicators of 375 patient blood samples from 10 January to 18 February 2020 at Tongji Hospital, Wuhan, China. These indicators are collected in irregular time intervals, and sampling rates range from 0 to 6 per day. We are interested in the task of predicting in-hospital mortality given the time series of biomarkers. The duration of hospital stays for patients varies from 2 hours to 35 days.

To process this dataset, we sampled the features with different rates ranging from 1 to 8-day intervals (more details in the Supplementary Section B). If multiple values are recorded for a feature in the determined rate, we use the last recorded value. We fill in the missing values by linear interpolation. If a feature is completely missing for a patient, we impute it using that feature’s mean value (i.e., the average of the recorded values for all patients in the training set). To evaluate the capability to select features in

Dataset	Model	AUC without MultiWave	AUC with MultiWave
WESAD	LSTM	0.822 ± 0.04	0.828 ± 0.04
	CNN-Attn	0.831 ± 0.03	0.877 ± 0.03
	CNN-LSTM	0.807 ± 0.04	0.839 ± 0.04
	Transformer	0.805 ± 0.04	0.824 ± 0.03
	FCN	0.805 ± 0.04	0.833 ± 0.05
	FCN-PC	0.860 ± 0.04	0.904 ± 0.03
COVID-19 0 days ahead	LSTM	0.983 ± 0.008	0.989 ± 0.004
	CNN-Attn	0.978 ± 0.012	0.979 ± 0.02
	CNN-LSTM	0.979 ± 0.010	0.981 ± 0.012
	Transformer	0.980 ± 0.007	0.984 ± 0.008
COVID-19 12 days ahead	LSTM	0.977 ± 0.007	0.979 ± 0.006
	CNN-Attn	0.967 ± 0.008	0.967 ± 0.013
	CNN-LSTM	0.961 ± 0.012	0.962 ± 0.009
	Transformer	0.969 ± 0.01	0.972 ± 0.01

Table 1: The AUC performance of MultiWave on real-world datasets WESAD and COVID-19, both with and without the use of MultiWave (when the model is used as Φ component in MultiWave architecture). CNN-Attn is a Convolutional Neural Network (CNN) model followed by an attention layer; CNN-LSTM is a CNN model followed by a Long Short-Term Memory (LSTM) layer; FCN is a fully convolutional neural network; and FCN-PC is a per-channel FCN model, as described in [Dzieżyc et al. \(2020\)](#).

Frequency component	Features
0 – 4 Hz	Chest acceleration Y, Chest acceleration Z, Wrist acceleration X, Chest ECG, Wrist BVP, Wrist EDA, Wrist temperature
4 – 8 Hz	Chest ECG

Table 2: features with non-zero mask weights after the training procedure on the WESAD dataset. These features have been consistently selected in 5 separate runs of the training, demonstrating their significance for the predictive task. This aligns with the domain expertise and supports the ability of MultiWave to automatically identify relevant feature dependencies, as well as determining the essential subsignals among different frequencies.

Component	AUC Simple	AUC w/ FFT	AUC w/ MultiWave
LSTM	0.822 ± 0.04	0.816 ± 0.03	0.828 ± 0.04
CNN-Attn	0.831 ± 0.03	0.816 ± 0.05	0.877 ± 0.03
CNN-LSTM	0.807 ± 0.04	0.835 ± 0.04	0.839 ± 0.04
Transformer	0.805 ± 0.04	0.813 ± 0.05	0.824 ± 0.03
FCN	0.805 ± 0.04	0.800 ± 0.05	0.833 ± 0.05
FCN-PC	0.860 ± 0.04	0.884 ± 0.04	0.904 ± 0.03

Table 3: This table compares the AUC results of models using FFT features versus those using MultiWave on the WESAD dataset. Some model components struggle to handle FFT values, and since we’re adding complexity to the model, there may not be sufficient data for the model with FFT features to exclusively learn from time domain features. Additionally, because we’re splitting the parameters between the two components, there are cases where adding FFT features results in worse performance than the baseline model.

Frequency component	Features
$0 - \frac{1}{16}$ days	High sensitivity C-reactive protein, Glucose
$\frac{1}{16} - \frac{1}{8}$ days	Lactate dehydrogenase
$\frac{1}{8} - \frac{1}{4}$ days	D-D dimer
$\frac{1}{4} - \frac{1}{2}$ days	-
$\frac{1}{2} - 1$ days	(%)lymphocyte, High sensitivity C-reactive protein

Table 4: The features with nonzero mask weights in the COVID-19 dataset at the end of the training. These features were consistently selected in 5 different runs of the training procedure, indicating that they are informative for the prediction task, which is consistent with domain expertise. MultiWave was able to automatically determine these dependencies, as well as determine which subsignals of different frequencies are relevant, as shown in Figure 5

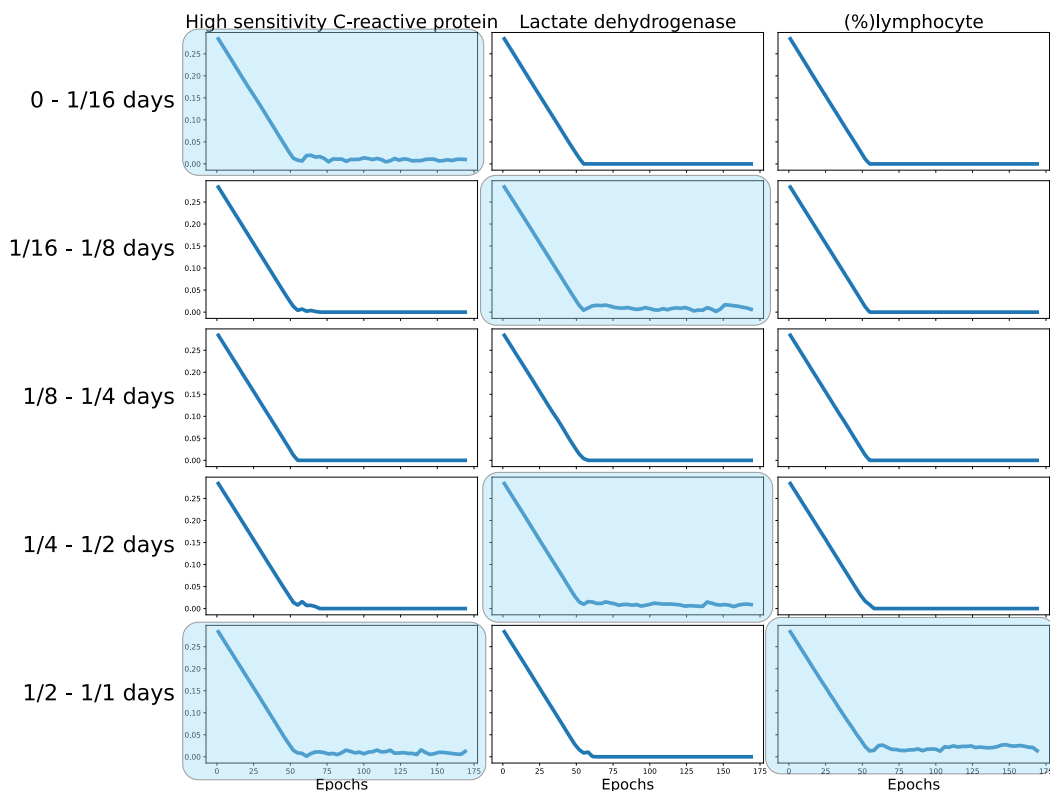


Figure 5: The mask weights over training epochs for the three most important features for different frequency components in COVID-19 dataset. Masks that are not zeroed out are highlighted.

our model, we used all 74 features. Because of this, we were unable to use the original test set provided with the dataset as it only contains 3 features. Therefore, to evaluate our models, we separated 100 patients from the original dataset and used 50 patients for validation and 50 patients to test our models.

The first group of experiments aims to predict mortality at the time of hospital discharge (0 days ahead of the prediction) by considering all patient data leading up to discharge. The second group of experiments is focused on predicting mortality 12 days ahead of hospital discharge (12 days ahead of prediction), incorporating patient data up to that point.

The average AUC results for five runs of each experiment are included in Table 1. For this dataset, we used the LSTM, CNN-Attn, CNN-LSTM, and Transformer models, and the results are reported with and without the inclusion of the MultiWave framework. Consistent with the results reported in Sun et al. (2021), LSTM-based models achieve the best results, and MultiWave brings consistent improvements to all baseline models.

Yan et al. (2020) reports three features as the most important characteristics in the prediction of hospital mortality in this dataset: lactic dehydrogenase, lymphocytes, and high-sensitivity C-reactive protein. To determine whether MultiWave can recognize the correct features in predicting the target, we looked at non-zero masked values in each model component for five different runs, and we show the common ones in Table 4. As can be seen, the model consistently selects these three features as important predictors, and the model has identified important subsignals within these features. These features are also reported to be important by Gao et al. (2021), Sun et al. (2021) and Jiang et al. (2021). In Figure 5 we show how the mask weights for these features across the components are changed over the training epochs of the model for one run.

5.4. MHEALTH

The MHEALTH (Banos et al., 2014, 2015) dataset is a publicly available dataset for health monitoring applications that contains data from 10 volunteers with diverse profiles while performing 12 different physical activities. The dataset was collected using Shimmer2 wearable sensors (Burns et al., 2010), which were placed on the subject’s chest, right wrist, and left ankle and attached using elastic straps. The dataset includes various physiological measurements, including acceleration, rate of turn, and magnetic field orientation, recorded at a sampling rate of 50 Hz. In addition, the chest sensor provides 2-lead ECG measurements.

In this dataset, we are trying to recognize different activities and classify them into 13 classes, which include 12 activities and the absence of any activity. To do this, we separated our dataset into two sets: a test set consisting of data from two subjects and a training set consisting of data from the other eight subjects. We only used acceleration and gyroscope data for this experiment (for more details, see Supplementary Section C).

We conducted two different experiments on this dataset, one with the original uniformly sampled data, and for the second experiment, we resampled the features with frequencies of $1/3, 1/5, 1/7, \dots, 1/25$ to obtain a multirate dataset. Because the classes are balanced, and following prior research (Demrozi et al., 2020; Chen et al., 2021) we report the classification accuracy in this dataset in Table 5. MultiWave brings consistent improvement to all baseline models, and the LSTM model with the inclusion of MultiWave achieved the best accuracy in the uniformly sampled dataset in which MultiWave brings an improvement of 9.8%. In the multirate version of the dataset the CNN-LSTM model with MultiWave achieved the best accuracy in which MultiWave brought an improvement of 5.6%. Table 6 shows the signals and their frequency components that were consistently selected by the MultiWave model. Accelerometer data are selected with different frequencies, consistent with previous research showing the importance of accelerometer data in activity recognition (Stisen et al., 2015; Ordóñez and Roggen, 2016).

5.5. The effect of frequency masks

The outcomes displayed in Table 7 show how incorporating frequency masks impacts the performance of MultiWave on the three datasets. The inclusion of component masks enhances MultiWave’s performance in all cases, indicating that the use of frequency masks leads to improved performance by eliminating unnecessary frequency components from the signals.

6. Conclusions

In this paper, we presented a new framework called MultiWave that augments any deep learning time series model with components that operate at different frequencies of signals using wavelet decomposition. We further improved MultiWave by introducing frequency masks, which remove noninformative frequency components of signals from the component inputs. We show that this framework improves the performance of time series models in synthetic datasets, as well as three real-world datasets for wearable affect and stress detection, COVID-19 prediction, and human activity recognition.

Dataset	Model	Accuracy without MultiWave	Accuracy with MultiWave
MHEALTH	LSTM	88.90% \pm 4.48	97.59% \pm 0.68
	CNN-Attn	92.90% \pm 2.69	95.50% \pm 2.11
	CNN-LSTM	92.04% \pm 3.33	96.48% \pm 1.20
	Transformer	95.83% \pm 3.53	96.14% \pm 1.14
MHEALTH Multirate	LSTM	81.84% \pm 3.68	92.34% \pm 0.69
	CNN-Attn	90.24% \pm 2.70	94.03% \pm 1.34
	CNN-LSTM	90.19% \pm 3.94	95.28% \pm 0.41
	Transformer	91.16% \pm 3.86	91.52% \pm 1.63

Table 5: The accuracy performance of MultiWave on the MHEALTH dataset both with and without the use of MultiWave (when the model is used as Φ component in MultiWave architecture). The first part of the table shows the results for the original dataset (uniformly sampled data), and the second part shows the results when the dataset is resampled to get a multirate dataset.

Frequency component	Features
0 – 3.125 Hz	Left-ankle acceleration (X-axis, Y-axis, Z-axis), Right-lower-arm acceleration (X-axis, Y-axis, Z-axis)
3.125 – 6.25 Hz	Left-ankle acceleration (Y-axis, Z-axis), Right-lower-arm acceleration (Y-axis)
6.25 – 12.5 Hz	Left-ankle acceleration (Y-axis, Z-axis)
12.5 – 25 Hz	Right-lower-arm acceleration (Y-axis)
25 – 50 Hz	Left-ankle acceleration (X-axis), Right-lower-arm acceleration (Y-axis)

Table 6: Features with nonzero masks after the training procedure on the MHEALTH dataset. These features have been consistently selected in 5 separate training runs, demonstrating their significance for the prediction task.

Dataset	AUC w/o MultiWave	AUC w/ MultiWave No Mask	AUC w/ MultiWave
WESAD	0.831 \pm 0.03	0.866 \pm 0.06	0.877 \pm 0.03
COVID-19 0 d ahead	0.980 \pm 0.007	0.984 \pm 0.01	0.984 \pm 0.008
COVID-19 12 d ahead	0.969 \pm 0.01	0.964 \pm 0.02	0.972 \pm 0.01
	Acc w/o MultiWave	Acc w/ MultiWave No Mask	Acc w/ MultiWave
MHEALTH	88.90% \pm 4.48	96.68% \pm 2.99	97.59% \pm 0.68

Table 7: The AUC results with and without the inclusion of frequency masks. We used the CNN-Attn model as a component for the WESAD dataset and the Transformer model for the COVID-19 dataset. We also report the accuracy results for MHEALTH dataset with the LSTM model. We selected these models for the ablation study because MultiWave brought the greatest improvement to the performance of the baseline models, which made it easier to reliably quantify the benefit brought by the frequency splitting and the benefit brought by the sparsity-inducing masks.

References

- Sofiya Alhassan and Thomas N Robinson. Defining accelerometer thresholds for physical activity in girls using roc analysis. *Journal of Physical Activity and Health*, 7(1):45–53, 2010.
- Bashar Alhnaity, Stefanos Kollias, Georgios Leontidis, Shouyong Jiang, Bert Schamp, and Simon Pearson. An autoencoder wavelet based deep neural network with attention mechanism for multi-step prediction of plant growth. *Information Sciences*, 560:35–50, 2021.
- Khaled A Althelaya, Salahadin A Mohammed, and El-Sayed M El-Alfy. Combining deep learning and multiresolution analysis for stock market forecasting. *IEEE Access*, 9:13099–13111, 2021.
- Leopoldo Armesto, Josep Tornero, and Markus Vincze. On multi-rate fusion for non-linear sampled-data systems: Application to a 6d tracking system. *Robotics and Autonomous Systems*, 56(8):706–715, 2008.
- Oresti Banos, Rafael Garcia, Juan A Holgado-Terriza, Miguel Damas, Hector Pomares, Ignacio Rojas, Alejandro Saez, and Claudia Villalonga. mhealthdroid: a novel framework for agile development of mobile health applications. In *Ambient Assisted Living and Daily Activities: 6th International Work-Conference, IWAAL 2014, Belfast, UK, December 2-5, 2014. Proceedings 6*, pages 91–98. Springer, 2014.
- Oresti Banos, Claudia Villalonga, Rafael Garcia, Alejandro Saez, Miguel Damas, Juan A Holgado-Terriza, Sungyong Lee, Hector Pomares, and Ignacio Rojas. Design, implementation and validation of a novel open framework for agile development of mobile health applications. *Biomedical engineering online*, 14(2):1–20, 2015.
- Lukas Biewald. Experiment tracking with weights and biases, 2020. URL <https://www.wandb.com/>. Software available from wandb.com.
- Ronald Newbold Bracewell and Ronald N Bracewell. *The Fourier transform and its applications*, volume 31999. McGraw-Hill New York, 1986.
- Adrian Burns, Barry R Greene, Michael J McGrath, Terrance J O’Shea, Benjamin Kuris, Steven M Ayer, Florin Stroiescu, and Victor Cionca. Shimmer™—a wireless sensor platform for noninvasive biomedical research. *IEEE Sensors Journal*, 10(9):1527–1534, 2010.
- Ping Chang, Huayu Li, Stuart F Quan, Janet Roveda, and Ao Li. Tdstf: Transformer-based diffusion probabilistic model for sparse time series forecasting. *arXiv preprint arXiv:2301.06625*, 2023.
- Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. Recurrent neural networks for multivariate time series with missing values. *Scientific reports*, 8(1):1–12, 2018a.
- Zhengping Che, Sanjay Purushotham, Guangyu Li, Bo Jiang, and Yan Liu. Hierarchical deep generative models for multi-rate multivariate time series. In *International Conference on Machine Learning*, pages 784–793. PMLR, 2018b.
- Kaixuan Chen, Dalin Zhang, Lina Yao, Bin Guo, Zhiwen Yu, and Yunhao Liu. Deep learning for sensor-based human activity recognition: Overview, challenges, and opportunities. *ACM Computing Surveys (CSUR)*, 54(4):1–40, 2021.
- Albert Cohen, Ingrid Daubechies, and J-C Feauveau. Biorthogonal bases of compactly supported wavelets. *Communications on pure and applied mathematics*, 45(5):485–560, 1992.
- Zhicheng Cui, Wenlin Chen, and Yixin Chen. Multi-scale convolutional neural networks for time series classification. *arXiv preprint arXiv:1603.06995*, 2016.
- Ingrid Daubechies. *Ten lectures on wavelets*. SIAM, 1992.
- Florenc Demrozi, Graziano Pravadelli, Azra Bihorac, and Parisa Rashidi. Human activity recognition using inertial, physiological and environmental sensors: A comprehensive survey. *IEEE access*, 8: 210816–210836, 2020.
- Yunshu Du, Wojciech M Czarnecki, Siddhant M Jayakumar, Mehrdad Farajtabar, Razvan Pascanu, and Balaji Lakshminarayanan. Adapting auxiliary losses using gradient similarity. *arXiv preprint arXiv:1812.02224*, 2018.
- Maciej Dzieżyc, Martin Gjoreski, Przemysław Kazienko, Stanisław Saganowski, and Matjaž

- Gams. Can we ditch feature engineering? end-to-end deep learning for affect recognition from physiological sensor data. *Sensors*, 20(22), 2020. ISSN 1424-8220. doi: 10.3390/s20226535. URL <https://www.mdpi.com/1424-8220/20/22/6535>.
- Grant Foster. Wavelets for period analysis of unevenly sampled time series. *The Astronomical Journal*, 112:1709–1729, 1996.
- Ya-dong Gao, Mei Ding, Xiang Dong, Jin-jin Zhang, Ahmet Kursat Azkur, Dilek Azkur, Hui Gan, Yuan-li Sun, Wei Fu, Wei Li, et al. Risk factors for severe and critically ill covid-19 patients: a review. *Allergy*, 76(2):428–455, 2021.
- Wenhui Guo, Guixun Xu, Baodi Liu, and Yanjiang Wang. Hyperspectral image classification using cnn-enhanced multi-level haar wavelet features fusion network. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5, 2022.
- Alfred Haar. *Zur theorie der orthogonalen funktionensysteme*. Georg-August-Universitat, Göttingen., 1909.
- Lei Huang, Feng Mao, Kai Zhang, and Zhiheng Li. Spatial-temporal convolutional transformer network for multivariate time series forecasting. *Sensors*, 22(3):841, 2022.
- Nan Jiang, Zhijun Li, Bo Yang, Mengdi Jin, Yaoyao Sun, Yang He, Yang Liu, Yueying Wang, Daoyuan Si, Piyong Ma, et al. Peripheral inflammatory cytokines and lymphocyte subset features of deceased covid-19 patients. *BioMed Research International*, 2021, 2021.
- Jusong Kim, Xiaoli Wang, Chollyong Kang, Jinwon Yu, and Penghui Li. Forecasting air pollutant concentration using a novel spatiotemporal deep learning model based on clustering, feature selection and empirical wavelet transform. *Science of The Total Environment*, 801:149654, 2021.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Clemens Kirschbaum, Karl-Martin Pirke, and Dirk H Hellhammer. The ‘trier social stress test’—a tool for investigating psychobiological stress responses in a laboratory setting. *Neuropsychobiology*, 28(1-2):76–81, 1993.
- Jan Koutnik, Klaus Greff, Faustino Gomez, and Jürgen Schmidhuber. A clockwork rnn. In *International Conference on Machine Learning*, pages 1863–1871. PMLR, 2014.
- Chi Qin Lai, Haidi Ibrahim, Aini Ismafairus Abd Hamid, and Jafri Malin Abdullah. Lstm network as a screening tool to detect moderate traumatic brain injury from resting-state electroencephalogram. *Expert Systems with Applications*, page 116761, 2022.
- James Lee-Thorp, Joshua Ainslie, Ilya Eckstein, and Santiago Ontanon. Fnet: Mixing tokens with fourier transforms. *arXiv preprint arXiv:2105.03824*, 2021.
- Yiqun Li, Songjian Chai, Zhengwei Ma, and Guibin Wang. A hybrid deep learning framework for long-term traffic flow prediction. *IEEE Access*, 9:11264–11271, 2021.
- Hui Liu, Hong qi Tian, Di fu Pan, and Yan fei Li. Forecasting models for wind speed using wavelet, wavelet packet, time series and artificial neural networks. *Applied Energy*, 107:191–208, 2013. ISSN 0306-2619. doi: <https://doi.org/10.1016/j.apenergy.2013.02.002>. URL <https://www.sciencedirect.com/science/article/pii/S0306261913001104>.
- Zhun Liu, Ying Shen, Varun Bharadhwaj Lakshminarasimhan, Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. Efficient low-rank multimodal fusion with modality-specific factors. *arXiv preprint arXiv:1806.00064*, 2018.
- Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 807–814, 2010.
- Francisco Javier Ordóñez and Daniel Roggen. Deep convolutional and lstm recurrent neural networks for multimodal wearable activity recognition. *Sensors*, 16(1):115, 2016.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner,

- Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- Sydney R Rooney, Evan L Reynolds, Mousumi Banerjee, Sara K Pasquali, John R Charpie, Michael G Gaies, and Gabe E Owens. Prediction of extubation failure in the paediatric cardiac icu using machine learning and high-frequency physiologic data. *Cardiology in the Young*, 32(10):1649–1656, 2022.
- Sajjad Safari, Faridoon Shabani, and Dan Simon. Multirate multisensor data fusion for linear systems using kalman filters and a neural network. *Aerospace Science and Technology*, 39:465–471, 2014.
- Djebbouai Salim, Souag-Gamane Doudja, Ferhati Ahmed, Djoukbala Omar, Dougha Mostafa, Bense-lama Oussama, and Hasbaia Mahmoud. Comparative study of different discrete wavelet based neural network models for long term drought forecasting. *Water Resources Management*, pages 1–20, 2023.
- Philip Schmidt, Attila Reiss, Robert Duerichen, Claus Marberger, and Kristof Van Laerhoven. Introducing wesad, a multimodal dataset for wearable stress and affect detection. In *Proceedings of the 20th ACM international conference on multimodal interaction*, pages 400–408, 2018.
- Satya Narayan Shukla and Benjamin M Marlin. A survey on principles, models and methods for learning from irregularly sampled time series. *arXiv preprint arXiv:2012.00168*, 2020.
- Donghwan Song, Adrian Matias Chung Baek, and Namhun Kim. Forecasting stock market indices using padding-based fourier transform denoising and time series deep learning models. *IEEE Access*, 9: 83786–83796, 2021.
- Allan Stisen, Henrik Blunck, Sourav Bhattacharya, Thor Siiger Prentow, Mikkel Baun Kjærgaard, Anind Dey, Tobias Sonne, and Mads Møller Jensen. Smart devices are different: Assessing and mitigating mobile sensing heterogeneities for activity recognition. In *Proceedings of the 13th ACM conference on embedded networked sensor systems*, pages 127–140, 2015.
- Abdulhamit Subasi, Mustafa Yilmaz, and Hasan Riza Ozcalik. Classification of emg signals using wavelet neural network. *Journal of Neuroscience Methods*, 156(1):360–367, 2006. ISSN 0165-0270. doi: <https://doi.org/10.1016/j.jneumeth.2006.03.004>. URL <https://www.sciencedirect.com/science/article/pii/S0165027006001440>.
- Chenxi Sun, Shenda Hong, Moxian Song, and Hongyan Li. A review of deep learning methods for irregularly sampled medical time series data. *arXiv preprint arXiv:2010.12493*, 2020.
- Chenxi Sun, Shenda Hong, Moxian Song, Yanxiu Zhou, Yongyue Sun, Derun Cai, and Hongyan Li. Te-esn: Time encoding echo state network for prediction based on irregularly sampled time series data. *arXiv preprint arXiv:2105.00412*, 2021.
- Jack Tatler, Phillip Cassey, and Thomas A. A. Prowse. High accuracy at low frequency: detailed behavioural classification from accelerometer data. *Journal of Experimental Biology*, 221(23), 11 2018. ISSN 0022-0949. doi: 10.1242/jeb.184085. URL <https://doi.org/10.1242/jeb.184085>. jeb184085.
- Sindhu Tipirneni and Chandan K Reddy. Self-supervised transformer for multivariate clinical time-series with missing values. *arXiv preprint arXiv:2107.14293*, 2021.
- Hongyu Wang, Wenrui Ding, Duona Zhang, and Baochang Zhang. Deep convolutional neural network with wavelet decomposition for automatic modulation classification. In *2020 15th IEEE Conference on Industrial Electronics and Applications (ICIEA)*, pages 1566–1571. IEEE, 2020a.
- Jingyuan Wang, Ze Wang, Jianfeng Li, and Junjie Wu. Multilevel wavelet decomposition network for interpretable time series analysis. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2437–2446, 2018.
- Weiyao Wang, Du Tran, and Matt Feiszli. What makes training multi-modal classification networks hard? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12695–12705, 2020b.

- Qingsong Wen, Tian Zhou, Chaoli Zhang, Weiqi Chen, Ziqing Ma, Junchi Yan, and Liang Sun. Transformers in time series: A survey. *arXiv preprint arXiv:2202.07125*, 2022.
- Jenna Wiens and Erica S Shenoy. Machine learning for healthcare: on the verge of a major shift in healthcare epidemiology. *Clinical Infectious Diseases*, 66(1):149–153, 2018.
- Ronald J Williams and David Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1(2):270–280, 1989.
- Li Yan, Hai-Tao Zhang, Jorge Goncalves, Yang Xiao, Maolin Wang, Yuqi Guo, Chuan Sun, Xiuchuan Tang, Liang Jing, Mingyang Zhang, et al. An interpretable mortality prediction model for covid-19 patients. *Nature machine intelligence*, 2(5):283–288, 2020.
- Ye Yuan, Guangxu Xun, Kebin Jia, and Aidong Zhang. A multi-view deep learning method for epileptic seizure detection using short-time fourier transform. In *Proceedings of the 8th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, pages 213–222, 2017.
- Jun Zhang, Gilbert G Walter, Yubo Miao, and Wan Ngai Wayne Lee. Wavelet neural networks for function learning. *IEEE transactions on Signal Processing*, 43(6):1485–1497, 1995.
- Yi Zheng, Qi Liu, Enhong Chen, Yong Ge, and J Leon Zhao. Exploiting multi-channels deep convolutional neural networks for multivariate time series classification. *Frontiers of Computer Science*, 10(1):96–112, 2016.
- Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *International Conference on Machine Learning*, pages 27268–27286. PMLR, 2022.
- Yufan Zhuang, Zihan Wang, Fangbo Tao, and Jingbo Shang. Waveformer: Linear-time attention with forward and backward wavelet transform. *arXiv preprint arXiv:2210.01989*, 2022.
- PJ Zucатели, EGS Nascimento, AÁB Santos, AMG Arce, and DM Moreira. An investigation on deep learning and wavelet transform to nowcast wind power and wind power ramp: A case study in brazil and uruguay. *Energy*, 230:120842, 2021.

Appendix A. WESAD

A.1. Preprocessing

We follow [Dziężyc et al. \(2020\)](#) to preprocess the data and use a similar technique to down-sample the signals. However, instead of the frequencies used in the article, we used frequencies that are powers of 2 given in [Table 8](#). We used a five-fold data splitting technique to train our model. For each fold, we trained the model on three folds, validated it on one fold, and tested it on the remaining fold. We do this five times with random seed values of 123 - 127 for all the models, and the mean and standard deviation are reported in the paper.

A.2. Hyperparameters

We perform a grid search on the validation data to select the hyperparameters for the baseline model. Then, we set the sizes of the MultiWave components so that they have a similar number of parameters as the baseline model. The hyperparameters selected for the WESAD dataset are given in [Table 10](#).

Appendix B. COVID-19 dataset

B.1. Preprocessing

This is an irregularly sampled dataset, we resample features based on their overall rate in all training data according to [Table 9](#). If a feature is missing at the time of sampling, we used the closest available value. If a feature is completely missing for a patient, we fill the values with the average of that feature among all the training samples. We also normalize the features to values between 0 and 1 using a min-max scaler.

B.2. Hyperparameters

We use the validation data to select the hyperparameters for the baseline model using grid search, and then for MultiWave components, we set the model sizes such that the number of parameters is approximately the same between the two models. We use data from 50 subjects for validation. The hyperparameters selected for MultiWave in this dataset are given in [Table 11](#).

Appendix C. MHEALTH

C.1. Preprocessing

To prepare the dataset for analysis, several steps were taken. First, we resampled the data for the “no activity” category to 30000, bringing it in line with the other activities. Our experiment was designed to use only accelerometer and gyroscope data, which is more challenging and better simulates real-world conditions. We then eliminated duplicate values and removed any outliers falling below the 0.01 quantile or above the 0.99 quantile for each feature. Finally, we selected windows of size 100 and retained the original activity labels. To split the dataset into training and testing, we utilized data from subjects 9 and 10 for testing and the remaining subjects for training.

C.2. Hyperparameters

Similar to the previous datasets, we used the validation data to perform a grid search and select the hyperparameters for the baseline model. As for the MultiWave components, we adjust the model sizes to ensure that they have roughly the same number of parameters as the baseline model. The hyperparameters selected for the MHEALTH dataset are given in [Table 12](#).

Appendix D. Experiments with frequency components that are not powers of 2

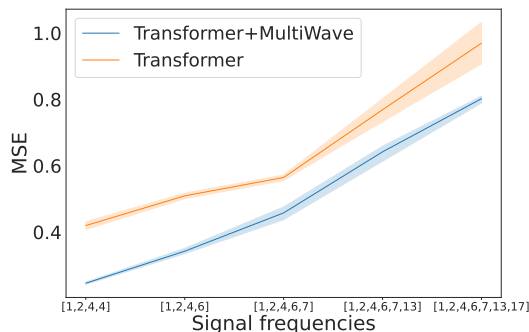


Figure 6: The Mean Squared Error (MSE) results for a synthetic dataset, on Transformer model with and without the addition of MultiWave. We add signals with frequencies 6Hz, 7Hz, 13Hz and 17Hz.

For this experiment, we repeated experiment 1 from section 5.1 using the Transformer model with varying frequencies, but we added signals with frequencies of 6 Hz, 7 Hz, 13 Hz, and 17 Hz, instead of using powers of 2. We began with signals having $1Hz$, $2Hz$, $4Hz$, $4Hz$ frequencies and changed the frequency of the last signal from 4 to 6 to observe how the models' performance would be affected, before adding the remaining signals. The results are displayed in figure 6, which demonstrates that despite the added signal frequencies not being powers of 2, MultiWave consistently improves performance.

Appendix E. Model with FFT features

To investigate whether the performance improvements of the models are solely attributable to the inclusion of frequency features, we constructed a naive model that utilizes both FFT features and the signal itself. The architecture of this model is illustrated in Figure 7. We used the same model component for both FFT features and the input signal, and concatenated the results before passing them through a 2-layer fully connected network. The experimental results are presented in Table 3.




Figure 7: Model with FFT features included

Signal	Original Sampling	Downsampled to
ECG RespiBAN	700 Hz	64 Hz
ACC RespiBAN	700 Hz	8 Hz
EMG RespiBAN	700 Hz	8 Hz
EDA RespiBAN	700 Hz	4 Hz
TEMP RespiBAN	700 Hz	4 Hz
Respiration RespiBAN	700 Hz	4 Hz
BVP Empatica	64 Hz	64 Hz
ACC Empatica	32 Hz	8 Hz
EDA Empatica	4 Hz	4 Hz
TEMP Empatica	4 Hz	4 Hz

Table 8: WESAD dataset feature frequencies

Signals	Sampled Rates
hemoglobin, Serum chloride, Prothrombin time, procalcitonin, eosinophils(%), Alkaline phosphatase, albumin, basophil(%), Total bilirubin, Platelet count, monocytes(%), indirect bilirubin, Red blood cell distribution width, neutrophils(%), total protein, Prothrombin activity, mean corpuscular volume, hematocrit, White blood cell count, mean corpuscular hemoglobin concentration, fibrinogen, Urea, lymphocyte count, Red blood cell count, Eosinophil count, Corrected calcium, Serum potassium, glucose, neutrophils count, Direct bilirubin, Mean platelet volume, RBC distribution width SD, Thrombin time, (%)lymphocyte, D-D dimer, Total cholesterol, aspartate aminotransferase, Uric acid, HCO ₃ ⁻ , calcium, Lactate dehydrogenase, platelet large cell ratio, monocytes count, PLT distribution width, globulin, γ -glutamyl transpeptidase, International standard ratio, basophil count(#), mean corpuscular hemoglobin, Activation of partial thromboplastin time, High sensitivity C-reactive protein, serum sodium, thrombocytocrit, glutamic-pyruvic transaminase,eGFR, creatinine	every day
antithrombin, Quantification of Treponema pallidum antibodies, HBsAg, HCV antibody quantification, Amino-terminal brain natriuretic peptide precursor(NT-proBNP), Fibrin degradation products, HIV antibody quantification, ESR	every 2 days
PH value	every 3 days
Interleukin 2 receptor, Interleukin 10, Interleukin 8, Tumor necrosis factor α , Interleukin 1 β , Interleukin 6	every 5 days
2019-nCoV nucleic acid detection	every 7 days
ferritin	every 8 days

Table 9: COVID dataset feature frequencies

Model	Hyper parameter	Value
LSTM	LSTM cell size	28
	Initial Mask Weight values	0.5
	Mask norm weight	0.05
	Number of Layers	1
	Learning rate	0.001
	Batch size	16
	patience	15
	Number of Components	6
	AddBaseline	True
CNN-Attn	CNN Kernel size	7
	Initial Mask Weight values	0.5
	Mask norm weight	0.1
	Number of Layers	2
	Learning rate	0.001
	Batch size	16
	patience	15
	Number of Components	6
	AddBaseline	True
CNN-LSTM	CNN Kernel size	7
	Initial Mask Weight values	0.5
	Mask norm weight	0.1
	Number of Layers	2
	Learning rate	0.001
	Batch size	16
	patience	15
	Number of Components	6
	AddBaseline	False
FCN	CNN Kernel size	7
	Initial Mask Weight values	0.5
	Mask norm weight	0.1
	Number of Layers	2
	Learning rate	0.001
	Batch size	16
	patience	15
	Number of Components	6
	AddBaseline	False

Table 10: The selected MultiWave hyperparameters for WESAD dataset

Model	Hyper parameter	Value
LSTM	LSTM cell size	19
	Initial Mask Weight values	0.3
	Mask norm weight	0.1
	Number of Layers	1
	Learning rate	0.001
	Batch size	16
	patience	25
	Number of Components	6
	Reset Model	False
CNN-Attn	CNN Kernel size	7
	Initial Mask Weight values	0.3
	Mask norm weight	0.1
	Number of Layers	2
	Learning rate	0.001
	Batch size	16
	patience	25
	Number of Components	6
	Reset Model	False
CNN-LSTM	CNN Kernel size	7
	Initial Mask Weight values	0.3
	Mask norm weight	0.1
	Number of Layers	2
	Learning rate	0.001
	Batch size	16
	patience	25
	Number of Components	6
	Reset Model	False
Transformer	Hidden Embedding Size	13
	Initial Mask Weight values	0.3
	Mask norm weight	0.1
	Number of Layers	4
	Number of Heads	3
	Learning rate	0.001
	Batch size	16
	patience	25
	Reset Model	False

Table 11: The selected MultiWave hyperparameters for COVID dataset

Model	Hyper parameter	Value
LSTM	LSTM cell size	16
	Initial Mask Weight values	1.5
	Mask norm weight	0.001
	Number of Layers	1
	Learning rate	0.001
	Batch size	32
	patience	15
	Number of Components	6
	Reset Model	False
CNN-Attn	CNN Kernel size	7
	Initial Mask Weight values	0.5
	Mask norm weight	0.01
	Number of Layers	2
	Learning rate	0.001
	Batch size	32
	patience	15
	Number of Components	6
	Reset Model	False
CNN-LSTM	CNN Kernel size	7
	Initial Mask Weight values	1.5
	Mask norm weight	0.001
	Number of Layers	1
	Learning rate	0.001
	Batch size	32
	patience	15
	Number of Components	6
	Reset Model	False
Transformer	Hidden Embedding Size	12
	Initial Mask Weight values	0.4
	Mask norm weight	0.5
	Number of Layers	1
	Number of Heads	3
	Learning rate	0.001
	Batch size	32
	patience	15
	Reset Model	False

Table 12: The selected MultiWave hyperparameters for MHEALTH dataset