NapTune: Efficient Model Tuning for Mood Classification using Previous Night's Sleep Measures along with Wearable Time-series

Debaditya Shome, Nasim Montazeri Ghahjaverestan, Ali Etemad
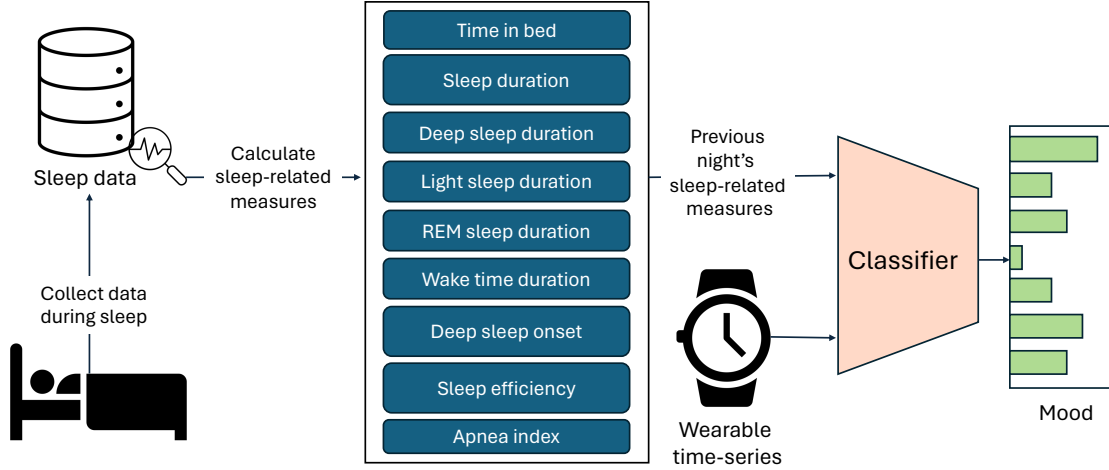Queen's University
Canada

Figure 1: Overview of our proposed pipeline.

ABSTRACT

Sleep is known to be a key factor in emotional regulation and overall mental health. In this study, we explore the integration of sleep measures from the previous night into wearable-based mood recognition. To this end, we propose NapTune, a novel prompt-tuning framework that utilizes sleep-related measures as additional inputs to a frozen pre-trained wearable time-series encoder by adding and training lightweight prompt parameters to each Transformer layer. Through rigorous empirical evaluation, we demonstrate that the inclusion of sleep data using NapTune not only improves mood recognition performance across different wearable time-series namely ECG, PPG, and EDA, but also makes it more sample-efficient. Our method demonstrates significant improvements over the best baselines and unimodal variants. Furthermore, we analyze the impact of adding sleep-related measures on recognizing different moods as well as the influence of individual sleep-related measures.

KEYWORDS

Mood Recognition, Affective Computing, Sleep, Deep Learning

1 INTRODUCTION

Affect recognition is a pivotal area in human-computer interaction, enabling a range of real-world applications [5] ranging from healthcare to user experience. Specifically, in healthcare, it offers valuable insights into mental health treatments and stress management [50].

In education, it can assess students' engagement and stress levels, tailoring learning experiences accordingly [25]. In the workplace, affect recognition can enhance employee well-being by identifying stress or dissatisfaction [36]. And finally, for user experience, affective computing paves the way for more intuitive and responsive products, from smart homes that adjust environments based on the occupants' moods to entertainment systems that adapt to user reactions [27].

Wearable devices have been widely adopted for affect recognition due to their ability to continuously and unobtrusively monitor physiological signals [45]. In particular, wearable signals such as electrocardiography (ECG) [17, 41], photoplethysmography (PPG) [31], and Electrodermal Activity (EDA) [46] have been the focus of numerous studies within this field. ECG measures the heart's electrical activity and can therefore capture changes in the emotional state of users. PPG detects blood volume changes throughout the body, offering insight into the physiological responses to emotions. EDA measures the electrical conductance of the skin, which dynamically fluctuates with sweat gland activity which is influenced by sympathetic nervous system activations. With advances in machine learning, these signals have been used to develop models capable of accurately recognizing affective states in real-time [22].

Mood and emotion are two common affect categories that are sometimes used interchangeably. In contrast to emotion, mood is an affective category that lasts for a long time, even up to several hours or days, and is thus harder to recognize from real-time data

[53]. While real-time wearable signals provide valuable information about a subject's affect state, they sometimes do not encapsulate historical information or external factors that may have a strong impact on mood. In particular, sleep-related measures stand out as such a salient external factor known to influence mood [42, 51]. These measures may include 'time in bed', 'sleep duration', 'deep sleep duration', 'light sleep duration', 'REM sleep duration', 'wake time', 'deep sleep onset', 'sleep efficiency', and 'Apnea index'. As a critical component of daily life, sleep impacts emotional regulation and mental health, thereby shaping subsequent affective states. The relationship between sleep patterns and mood dynamics is well-studied in the literature [58, 24, 29], with sleep disturbances known to aggravate stress and irritability [56]. Conversely, a pattern of restorative sleep is often correlated with enhanced mood stability and cognitive functioning [32].

A thorough investigation of the literature in this area demonstrates that despite the relationship between sleep measures and mood being well-established in prior works, previous night's sleep measures have not been incorporated into automated mood recognition models as a complementary modality to wearable signals. Moreover, what makes this issue even more challenging is the scarcity of public datasets that contain the previous night's sleep measures *paired* with wearable time-series from the *same users* over a period of time. This in turn inhibits the use of standard training frameworks from being effective in enabling a multimodal setup for both sleep and wearable time-series simultaneously.

In this paper, to address the challenges above, we propose Nap-Tune, a framework that integrates previous night's sleep measures as a complementary source of information into mood recognition from various wearable time-series (ECG, PPG, and EDA). First, our model uses an unimodal Transformer-based encoder to extract effective representations from the wearable time-series. Next, to enable our model to incorporate and analyze the previous night's sleep-related measures as an additional input, we freeze the Transformer encoder and add a small number of learnable parameters into each layer, which we then train for mood recognition. This enables our model to effectively utilize sleep-related measures alongside wearable time-series for mood recognition, with little requirement for paired training data. Extensive experiments demonstrate that incorporating previous night's sleep-related measures leads to an increase in mood recognition performance with up to 8% in F1 score. In summary, we make the following contributions.

(1) We propose NapTune, an efficient tuning framework for adapting a frozen wearable time-series encoder to utilize the previous night's sleep-related measures as an additional input modality. Our framework demonstrates strong performances for mood classification based on sleep-related measures and wearable time-series, outperforming various multimodal baselines.

(2) For training, our framework requires minimal data in the form of wearable time-series paired with previous night's sleep measures from the same users, overcoming the problem of insufficient paired training data.

(3) We study the effect of the previous night's sleep-related measures in aiding mood recognition from wearable time-series. Particularly, we train and evaluate models for mood recognition with ECG, PPG, and EDA, with and without the use of sleep measures.

Our findings show an increase of 9% to 11% in F1 when utilizing the sleep-related measures as an additional modality.

2 RELATED WORK

In this section, we review three key areas relevant to our work. First, we review affect recognition studies from wearable time-series. Next, we present a review of prior works focusing on the relationship between sleep and affect. Finally, we conclude this section with a thorough review of recent works on prompt-tuning, a technique that is central to our proposed framework.

2.1 Wearable affect recognition

Several recent studies have demonstrated the potential of wearable time-series for affect recognition. In [60], a framework for personalized and generalized mood recognition in a workplace setting was proposed using smartphone-based wearable sensors including ECG, PPG, 3-axis acceleration, and skin temperature. Several ML algorithms were evaluated to find a bagged ensemble of decision trees as the best performing solution. In [63], a novel pipeline using popular wearables was developed to recognize daily activities and a regression model was trained for mood assessment. The method was tested in a real-world study with 18 users over 93 user-days and achieved mood inference with a mean absolute error of 0.24 π radians on the Circumplex Model of Affect. In [43], a self-supervised learning method was introduced for emotion recognition using ECG time-series, that achieves comparable or superior results to fully-supervised methods on the SWELL [28] and AMIGOS [37] datasets. In [12], a study was conducted on combining smartphone and smartwatch data to assess mood. The results demonstrated the significance of temporal features and heart rate obtained from such wearable time-series. In [62], an emotion recognition system was developed using spectrogram representations of EDA time-series with a combination of CNN and Bi-GRU for feature learning and classification. The model was evaluated on the AMIGOS [37] dataset demonstrating high classification accuracies of 83.4% for arousal and 81.2% for valence.

2.2 Sleep and affect

Although the aim of this paper is to study the impact of sleep as an additional modality to improve mood recognition, in this subsection, we review the prior works that study the overall link between sleep and affect, i.e., how sleep impacts affect and vice-versa. In [51], a conceptual review was conducted to explore the relationship between sleep and affect. A granular framework was proposed to deconstruct sleep and affect into three dimensions: domains, methods, and timescales. This framework was applied to systematically review empirical studies from PubMed, focusing on associations between various aspects of sleep and affect. The study found evidence of links between sleep disturbances, sleep duration, and affect, but noted that evidence was often inconclusive or sparse for other aspects. In [29], the reciprocal relationships between daily sleep and mood were examined, focusing on real-world studies. Electronic databases were utilized to collect studies that investigated daily associations between sleep and mood using ambulatory diary techniques. The findings of this study supported a reciprocal relationship between subjective sleep variables (quality,

duration, latency) and daytime affective states, emphasizing the potential clinical importance of sleep disturbance in predicting and preventing psychopathology.

In [24], a 2-week study was conducted to explore the interplay between daily affect and sleep in young women. Using daily sampling, they investigated how variations in positive and negative affect influenced self-reported sleep-onset latency, sleep duration, and sleep quality. The study revealed significant associations between sleep and emotions, where sadness and serenity were found to be strong predictors of sleep-related measures, and better sleep quality predicted greater happiness the next day. In [8], the interplay of neuroticism, affect, and hyperarousal with perceived sleep quality was explored. Using an online survey, the impact of these factors on sleep quality was assessed among 498 Italian participants. The study found that neuroticism was the primary personality predictor of poor sleep quality. Additionally, hyperarousal and positive affect also significantly predicted good sleep quality. In [30], a systematic review and meta-analysis was conducted to evaluate the causal impact of emotions on sleep. 31 experimental studies focusing on the effects of emotion inductions on various sleep-related measures were analyzed. The study found a moderately significant effect of emotion inductions on delayed sleep onset latency, but no consistent impacts on other sleep-related measures.

2.3 Prompt tuning

Prompt tuning was first introduced as a method for adapting large pre-trained language models to specific tasks using soft prompts learned through backpropagation [33]. This technique allows a large frozen model to be used across multiple tasks without significant computational overhead. Inspired by the success of this approach in large language models, prompt tuning has since been further adopted for other domians including vision and time-series. In [21], Visual Prompt Tuning (VPT) was introduced as an efficient method for adapting large-scale Transformer models in vision tasks. VPT incorporates a small amount of trainable parameters into the input space while keeping the model backbone frozen, leveraging the advantages of large pre-trained models without the need for extensive retraining. Similarly, prompt-tuning was recently adopted for several time-series tasks demonstrating promising results. In [6], a novel framework named TEMPO was proposed as a prompt-based generative pre-trained Transformer for time-series forecasting. TEMPO addresses the challenges in time-series forecasting by integrating decomposition of complex interactions between trend, seasonal, and residual components, and by introducing selection-based prompts to adapt to non-stationary time-series. The framework demonstrated superior performance over state-of-the-art methods in various benchmark datasets. In [59], PromptCast, a novel approach to time-series forecasting utilizing a prompt-based paradigm, was introduced. It transforms numerical data into prompts, leveraging pre-trained language models for forecasting.

In [34], PT-Tuning was introduced as an approach that enhances time-series forecasting by bridging the gap between masked reconstruction and forecasting through prompt token tuning. The unique aspect of PT-Tuning is its integration of a few trainable prompt tokens, while all other pre-trained parameters remain frozen. This effectively addresses the issue of task difficulty variation. The approach demonstrated remarkable performance improvements over

other methods in extensive experiments with real-world datasets. In [55], a method called POND was introduced for multi-source time-series domain adaptation using prompt tuning. This method aims to tackle three challenges: the exploitation of domain-specific information for domain adaptation, the capturing of dynamic domain-specific information, and the assessment of learned domain-specific data. POND uses prompts to capture both common and specific information across domains, introduces a conditional module for generating dynamic prompts for each source domain, and employs specific criteria for effective prompt selection. In [9], SpeechPrompt v2 was introduced, which is a prompt tuning framework for various speech classification tasks. This approach uses generative spoken language models with trainable prompt vectors, allowing efficient adaptation to different languages and tasks with minimal parameter updates. The framework introduces a learnable verbalizer, enhancing its adaptability and performance across a wide range of speech classification tasks, including emotion recognition and language identification, while maintaining computational efficiency.

3 METHOD

3.1 Problem and solution overview

We aim to develop a framework to enable a pre-trained wearable time-series encoder to utilize the users' previous night's sleep-related measures as an auxiliary input to enhance mood recognition given a small training dataset of paired data. Inspired by recent advances in prompt tuning in large language foundation models, we propose an efficient tuning solution to address this problem. Our method includes two stages. First, we train a Transformer-based wearable time-series encoder using self-supervised pre-training over a large corpus of unlabeled unimodal wearable time-series. Next, we freeze the weights of this pre-trained encoder and add a linear projection to include sleep-related measures, and additional learnable parameters in each Transformer layer for efficient training without a high computational overhead. Figure 2 depicts an overview of our approach.

3.2 Proposed approach

3.2.1 Wearable time-series encoder. Inspired by [2] we adopt a CNN-Transformer model as an encoder to learn the wearable time-series. The detailed architecture of this encoder is presented in Figure 3. As shown in the figure, the encoder consists of a convolutional projection module with a series of ω convolution blocks. The initial convolutional layer uses a kernel size of k_l (long kernel) and a stride of s_l (long stride). This layer is specifically designed to better capture the global patterns in the input time-series. The output from this layer undergoes normalization using GroupNorm, followed by GELU activation. The subsequent convolutional layers follow a similar pattern where each of these layers use smaller kernels of size k_s (short kernel) and a stride of s_s (short stride), which focus on extracting more detailed local features from the input time-series. We apply LayerNorm to the extracted convolutional features and apply a GELU-activated linear layer to project the output of this convolutional projection module into a $L \times D$ dimensional time-series embedding $z_{time-series}$. Next, we represent
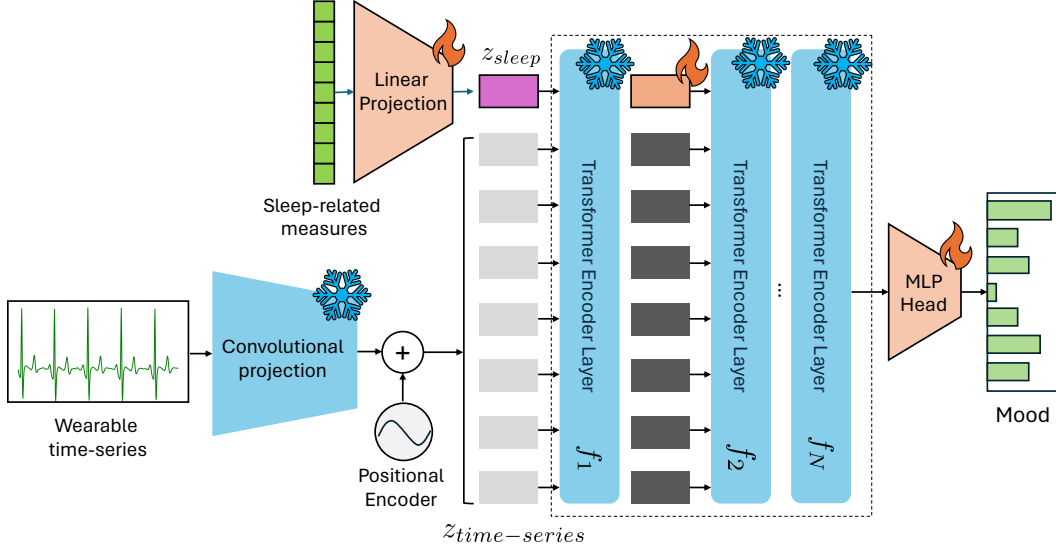
Figure 2: Our proposed NapTune framework.

$z_{time-series}$ as a set of L tokens each with a dimension D, as

$$z_{time-series} = \{z_i \in \mathbb{N} | 1 \le i < L\}. \tag{1}$$

Next, in order to preserve sequential information we use a Positional Encoder. Given a time index pos and dimension index i, the Positional Encoder generates a positional embedding Π defined as

$$\Pi(pos, 2i) = \sin\left(\frac{pos}{10000^{2i/D}}\right), \Pi(pos, 2i+1) = \cos\left(\frac{pos}{10000^{2i/D}}\right) \tag{2}$$

Lastly, the model consists of a Transformer encoder, which we refer to as f, that consists of N layers with an embedding dimension of D. Each layer consists of a multi-head attention module, Layer-Norm, and GELU-activated feedforward layers, as shown in Figure 3. Each multi-head attention module performs scaled-dot product attention in parallel over multiple attention heads by utilizing the layer's input representation as query, key, and value. Specifically, the scaled-dot product attention operation involves calculating the dot product between scaled query (Q) and key (K) matrices, which is multiplied with a pre-defined mask of attention weights. Then, the softmax operation is applied to derive the attention weights and subsequently multiplied with the value matrix (V) to obtain the final attention output represented as

$$\text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{Q \cdot K^T}{\sqrt{d}} * \text{Mask}\right) V, \tag{3}$$

where, d represents the dimension of K, Q, and V, which is used as a scaling factor to prevent vanishing gradients.

3.2.2 Pre-training. To pre-train the wearable time-series encoder f to learn strong representations from unlabeled time-series, we pre-train this module using SimCLR [11]. SimCLR is a widely used contrastive self-supervised learning framework [11, 49]. Given each input time-series x, we create two views x_1 and x_2 using two distinct stochastic augmentations t_1 and t_2. First, the encoder backbone

extracts embedding h from the augmented views of each time-series. Next, a projector network g with a single linear layer projects the embedding into $z = g(h)$. In a minibatch of M samples resulting in $2M$ pairs, we obtain one positive pair (p, q), and the rest $2(M-1)$ as negative pairs. We use the contrastive loss function [48], which can be defined as

$$\ell_{p,q} = -\log \frac{\exp(\text{sim}(z_p, z_q)/\tau)}{\sum_{k=1}^{2M} \mathbb{1}_{k \ne i} \exp(\text{sim}(z_p, z_k)/\tau)}, \tag{4}$$

where sim indicates pairwise dot product similarity and τ is the temperature parameter. Both t_1 and t_2 are sampled from the same family of augmentations which are suitable for temporal time-series, as described below.

Time Warping. In this augmentation, we split the time-series x into r segments $x = [x_1, ..., x_r]$. We randomly select $\frac{r}{2}$ segments and apply 1D interpolation-based time-warping to stretch them by a factor of $\sigma\%$, while squeezing the remaining segments by the same factor. Finally, we concatenate the segments and apply zero padding if the resulting length is an odd number.

Gaussian Noise. In this augmentation, gaussian noise is added to the time-series x. The noise standard deviation is randomly selected within a predetermined range, defined by the product of the time-series's standard deviation and specified minimum and maximum signal-to-noise ratios (SNR). Gaussian noise is then generated according to this noise standard deviation and added to x, resulting in a controlled level of random noise added to the time-series.

Random Scaling. This augmentation involves altering the amplitude of the time-series x. It is achieved by multiplying x with a scaling factor α, where α is a positive value randomly chosen from a predefined range.
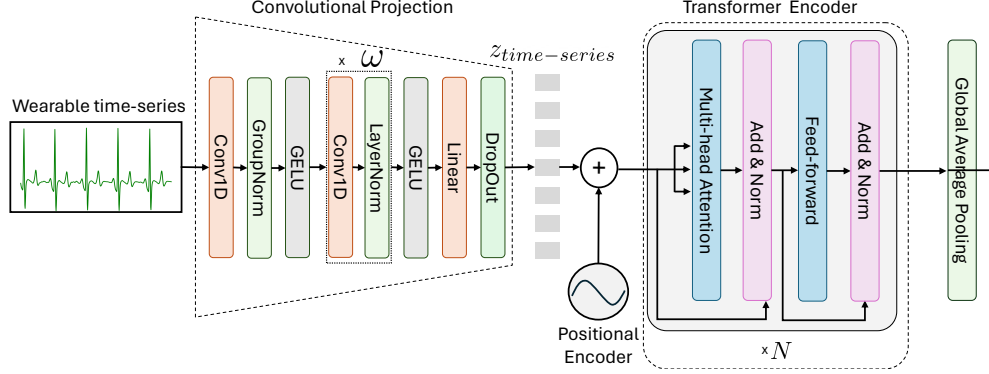
Figure 3: The architecture of the wearable time-series encoder.

3.3 Tuning

Once the wearable time-series encoder is pre-trained, we freeze the weights of the convolutional projection module and the transformer encoder f. We add a linear projection module to encode sleep-related measures into a D-dimensional representation z_{sleep}. Using the wearable time-series encoder, we obtain $z_{time-series}$ using Equation 1. Next, we concatenate z_{sleep} with $z_{time-series}$ which results in z_0, a D-dimensional set of $(L + 1)$ as input tokens for f. Lastly, we add learnable parameters into each transformer layer of f which are updated using backpropagation. Here, f has N Transformer encoder layers which can be represented as

$$z_{n+1} = \begin{cases} f_1(z_1), & \text{if } n = 1 \\ f_n(\text{concat}[P_n, z_n]), & \text{otherwise,} \end{cases} \quad (5)$$

where P_n refers to a set of unique learnable parameters at layer n. We train the model using the Binary Cross Entropy (BCE) loss function for our downstream task of multi-label mood classification. Specifically, BCE loss can be defined as

$$\text{BCE}(p, \tilde{p}) = -\frac{1}{C} \sum_{j=1}^{C} [p_j \log \tilde{p}_j + (1 - p_j) \log (1 - \tilde{p}_j)], \quad (6)$$

where, C is the number of classes, and p and $\tilde{p}$ represent the actual mood labels and the predicted logits by our model, respectively.

4 EXPERIMENT SETUP

4.1 Datasets

Here, we describe the details of the datasets used in this paper. First, we use a collection of datasets for pre-training the wearable time-series, as well as a downstream mood classification dataset which we use for evaluating our proposed method.

4.1.1 Pre-training datasets. The following datasets are used for pre-training the wearable time-series encoders.

WESAD [44]. This dataset contains 24 hours of Lead-II ECG and PPG data, sampled at 700 Hz and 64 Hz respectively, from 15 participants. Annotations in the dataset include stress and various affect states, providing insights into emotional and stress-related physiological responses. We use this dataset to pre-train encoders for ECG, PPG, and EDA.

MIMIC AFib [1]. This datasets is part of the larger MIMIC-III waveform collection [23] dataset. This subset includes data from 35 critically ill adults, including 19 with Atrial Fibrillation (AFib). Both ECG and PPG time-series are present, sampled at a frequency of 125 Hz.

We use this dataset to pre-train ECG and PPG encoders.

CAPNO [26]. This dataset features approximately 5.6 hours of Lead-II ECG and PPG data, recorded at a 300 Hz sampling rate, from 42 subjects under clinical supervision. It extends the range of physiological data to include medically monitored scenarios. We use this dataset to pre-train both ECG and PPG models.

BIDMC [39]. Collected from 53 ICU patients, the dataset contains nearly 7 hours of ECG (across multiple leads including Lead II, V, and AVR) and PPG recordings, each sampled at 125 Hz. We use this dataset to pre-train encoders for ECG and PPG.

DALIA [40]. Comprising approximately 35 hours of data, this dataset includes PPG and Lead-II ECG recordings from 15 subjects performing everyday activities. The sampling rates for ECG and PPG are 700 Hz and 64 Hz, respectively.

In our study, we use this dataset to pre-train ECG and PPG encoders.

Fekri Azgomi et al. [13]. This dataset includes EDA signals collected using an Empatica E4 device during cognitive tasks involving auditory, gustatory, and olfactory stimulation. It is designed to explore the modulation of cognitive states and its effects on physiological responses. We utilize this to pre-train the EDA encoder.

4.1.2 Downstream Dataset. For our experiments on the downstream task of mood classification, we require a datasets that contains both wearable time-series data with mood labels, paired with previous night's sleep-related measures from the same subjects. To our knowledge, **ECSMP dataset [14]** is the only dataset that contains such information. It comprises multiple physiological time-series, including ECG, PPG, and EDA collected from 89 healthy college students. The sampling rates of ECG, PPG, and EDA are 512 Hz, 64 Hz, and 4 Hz respectively. These signals have been collected during various states, such as resting, emotional induction and recovery, and cognitive assessments. Additionally, the dataset consists of multiple self-reported questionnaires, among which we utilize the Profile of Mood States (POMS) as annotations for mood classification. It includes seven mood labels represented by scores

to categorize tension, anger, fatigue, depression, vigor, confusion, and esteem. We convert the scores to binary labels and formulate this task as a 7-way multi-label classification problem.

The dataset also contains sleep-related measures estimated from the ECG collected during the previous night's sleep from each subject. The authors used the cardiopulmonary coupling analysis algorithm [52] to estimate the sleep-related measures from the ECG. These sleep-related measures include:

- *Time in bed:* This measures the total number of hours that the subject spent in bed.
- *Sleep duration:* This refers to the total number of hours the subject has slept.
- *Deep sleep duration:* This refers to the total number of hours that the subject has been in deep sleep, the most restorative stage of sleep.
- *Light sleep duration:* This refers to the total number of hours that the subject has been in light sleep, the transition period between wakefulness and deep sleep.
- *REM sleep duration:* This measures the total number of hours in Rapid Eye Movement (REM) sleep, the sleep stage characterized by rapid eye movements, vivid dreaming, and high brain activity, resembling wakefulness.
- *Wake time:* This measures the total number of hours the subject has spent in bed without sleeping.
- *Deep sleep onset:* This indicates the number of hours needed by the subject to transition from light sleep to deep sleep.
- *Sleep efficiency:* This refers to the ratio of sleep duration to the total time spent in bed.
- *Apnea index:* This refers to the average number of apnea occurrences per hour of sleep.

4.2 Data pre-processing

We apply standard pre-processing steps [38, 35] for ECG, PPG, and EDA time-series. The ECG signals are processed using a high-pass Butterworth filter with a 0.5 Hz cut-off frequency. For the PPG time-series, we apply a band-pass Butterworth filter with frequencies ranging from 0.5 to 8 Hz. We skip the filtering of EDA time-series due to the low sampling frequency of 4 Hz. Z-score normalization is applied on each type of time-series to adjust for individual subject-specific differences. After normalization, the time-series are scaled to fit within a [-1, 1] range using min-max scaling. The final step in our pre-processing is segmenting the ECG, PPG, and EDA time-series into 10-second windows.

4.3 Baselines

Here we describe the models that we use as baselines for comparison to our proposed method. To our knowledge, no prior work has performed mood classification with the aid of the previous night's sleep measures on this dataset. As a result, we create several baselines based on state-of-the-art wearable time-series encoders and pre-train each encoder in the same manner that our own model's encoders are pre-trained. To create these baselines, we use an encoder/projection model for the sleep measures and perform late-stage fusion with the representations of the wearable time-series encoders. For the sleep encoder models, we use a GELU-activated 2-layer linear projection module for the Transformer-based baseline,

while we use a ReLU-activated 2-layer linear projection module for others. Finally, for late-stage fusion, we combine the sleep embeddings with the global-average pooled output embeddings from the wearable time-series encoder and pass the outcome to a sigmoid-activated linear layer to predict mood logits. To ensure that the best possible performances are obtained from these baselines, we optimize the hyper-parameters to the best of our ability. Following we describe the architectures and hyper-parameters of these encoders.

Chen et al. [10]. This model employs a VGG-19 [47] backbone which takes 2D Short-Time Fourier Transform (STFT) representations of the wearable time-series as inputs. The network comprises a series of convolutional layers, initially with 64 filters, doubling at each stage following a max-pooling layer, up to layers with 512 filters. Each convolutional block consists of layers in configurations of two or four, paired with max pooling. Convolutional layers use 3x3 kernels with batch normalization and ReLU activation.

Zihlmann et al.[64] This model follows a Convolutional Recurrent neural network (CRNN) architecture that takes wearable time-series converted into logarithmic spectrograms as input. It comprises 3 convolutional blocks, with each applying a set of 5x5 convolutional filters followed by batch normalization and ReLU activation. The channels in these layers increase with each block. A 3-layer bidirectional LSTM processes the outputs from the convolutional layers, providing an aggregated feature vector.

Hong et al. [19] We adopt ResNet1D [19] as our baseline, an adaptation of the ResNet model [18] for 1D time-series. The network is comprised of 34 blocks. Each block consists of convolution layers with a kernel size of 64 and a stride of 2, followed by batch normalization ReLU activation, and a dropout of 20%. We use 64 as the number of base filters, which doubles at specified intervals across the network.

Behinaein et al. [2] We use the same pre-trained wearable signal encoder based on Transformer-based architecture used for our proposed method as a baseline.

4.4 Evaluation Protocol

We use a cross-subject 3-fold cross-validation scheme for evaluation. This involves dividing all the subjects into three distinct groups (folds), with each fold being iteratively used as the test set while the remaining two serve as the training set. Due to the inherent imbalance in the mood states represented in the dataset, we report both weighted F1 score and accuracy.

To obtain the final results, we generally perform linear evaluation unless otherwise specified. For linear evaluation, we use transfer learning, where we first freeze the weights of the pre-trained wearable time-series encoders. Then, we train a linear classifier over the learned representations for mood recognition.

4.5 Implementation Details

We use 4×NVIDIA A100 GPUs for self-supervised pre-training using a batch size of 2048, with AdamW optimizer and a learning rate of 1×10^{-4} for 150 epochs. For downstream experiments, we use 2×NVIDIA A100 GPUs with a batch size of 1024. We use AdamW optimizer with CosineWarmup, starting from a base learning rate of 1×10^{-5} for 30 epochs. To ensure reproducibility we present all the hyper-parameters used in our work, in Table 1.

Table 1: The details of our wearable time-series encoder.

Module	Parameter	Value
Convolutional Projection	Activation	GELU
	k_l	10
	s_l	5
	k_s	3
	s_s	2
	ω	6
	Linear dim.	512
	DropOut	0.2
Linear Projection	Activation	GELU
	Linear 1 dim.	128
	Linear 2 dim.	512
Transformer Encoder	Activation	GELU
	N	6
	D	512
	heads	4
Classification Block	Activation	GELU
	Number of layers	2
	Dropout rate	0.5

Table 2: Performance of our proposed method.

Modality	Method	F1	Acc.
	Hong et al. [19]	0.33	0.30
	Zihlmann et al. [64]	0.29	0.27
EDA + Sleep	Chen et al. [10]	0.31	0.30
	Behinaein et al. [2]	0.40	0.36
	Ours	**0.42**	**0.39**
	Hong et al. [19]	0.48	0.46
	Zihlmann et al. [64]	0.44	0.42
PPG + Sleep	Chen et al. [10]	0.52	0.50
	Behinaein et al. [2]	0.60	0.59
	Ours	**0.72**	**0.69**
	Hong et al. [19]	0.56	0.52
	Zihlmann et al. [64]	0.49	0.47
ECG + Sleep	Chen et al. [10]	0.50	0.46
	Behinaein et al. [2]	0.65	0.58
	Ours	**0.76**	**0.71**

Table 3: Comparison of mood classification using wearable time-series with and without sleep-related measures.

Modality	F1	Acc.
EDA	0.31	0.29
EDA + Sleep	0.42	0.39
PPG	0.63	0.57
PPG + Sleep	0.72	0.69
ECG	0.67	0.64
ECG + Sleep	0.76	0.71

5 RESULTS AND DISCUSSION

First, we compare our proposed method against several baselines using linear evaluation, and present the results in Table 2. We observe that our approach achieves the best performance for all three types of wearable time-series for both metrics (F1 and accuracy) by considerable margins. More specifically, for EDA, our method achieves the best results by 2% improvement in F1 and 3% improvement in accuracy over the best baseline [2]. For PPG, our method leads to a significant increase of 12% and 10% in F1 and accuracy respectively over the best baseline. With ECG, we achieve a similar improvement of 11% and 13% in F1 and accuracy, respectively. Out of the three modalities, NapTune performs best with ECG achieving an F1 of 0.76 and an accuracy of 0.71. We believe that the relatively lower performance when using EDA is due to the low sampling rate and the presence of noise.

Next, we examine the effect of using sleep-related measures on the performance of our method. To this end, we use the frozen wearable time-series encoder and perform linear evaluation without the rest of our method which uses the sleep-related measures. We present the results in Table 3, which shows that using sleep-related measures as a complimentary modality leads to improvements across all three wearable time-series. In ECG-based mood recognition, adding sleep-related measures provides a performance boost of 9% in F1 and 7% in accuracy. With PPG, we observe boosts of 9% and 12% in F1 and accuracy respectively. Lastly for EDA, we observe a similar trend where 11% and 10% improvements are obtained for in F1 and accuracy respectively.

We then investigate the performance breakdown for different classes of mood, with and without the sleep-related measures, and present the results in Figure 4. Overall, we observe that the addition of sleep boosts performance for the classification of every mood class regardless of the type of wearable time-series used. Next, it can be seen that using EDA (without sleep), tension and anger are classified with the highest F1 scores, whereas the addition of sleep results in the highest boosts for vigor, followed by esteem.

On the other hand, when using PPG (without sleep), the highest F1 scores are achieved for fatigue and esteem recognition, while the addition of sleep results in the highest boosts for depression and vigor. Lastly, using ECG (without sleep) leads to the highest F1 scores for vigor and esteem, whereas the highest improvements by adding sleep are achieved for depression and tension. These findings are aligned with our understanding of the physiological connections between sleep and depression, as documented in prior research [7, 61]. These studies have demonstrated that the quality of sleep from the previous night can significantly impact the onset of depression.

Furthermore, we analyze the impact of different training methods on mood recognition performance in Table 4, specifically full-finetuning, pre-training, and NapTune. For full-finetuning, we update all the weights of the pre-trained encoder with sleep-related measures paired with wearable signals as input, without adding any prompt parameters. We observe an improvement of 3% to 4% in the case of finetuning the ECG and PPG-based models across both F1 and accuracy, while an increase of 8% and 9% respectively in the case of the EDA-based model. Our intuition suggests that the pre-training of EDA is not as strong as PPG and ECG counterparts, particularly due to a lesser amount of pre-training data. This in turn results in such an improvement in performance with finetuning compared to NapTune which uses Linear protocol. Upon training

Figure 4: Impact of sleep on classification of different classes of mood with EDA (left), PPG (middle), and ECG (right).


**Figure 5: Ablation of individual sleep-related measures.**

Figure 6: Impact of the amount of data used for training with ECG+Sleep.

the model from scratch without any pre-trained weights, we observe a similar trend for all the modalities, with a performance drop ranging between 2% to 4% in both metrics.

Next, we study the contribution of individual sleep-related measures on mood recognition performance. Specifically, we ablate each sleep-related measure through masking at test time, and evaluate the model's performance. As shown in Figure 5, REM sleep duration has the highest impact on mood recognition, followed by deep sleep and sleep duration. These findings are well-aligned with our understanding of how REM and deep sleep influence emotion regulation. Several studies [15, 20, 57, 54, 16] have highlighted that during REM sleep the brain strengthens important emotional memories, which influences emotional regulation after waking up. Moreover, deep sleep has been highlighted as primarily responsible for consolidating and stabilizing emotional memories [4, 3].

Table 4: Performance of different training methods.

Modality	Variants	F1	Acc.
	Ours	0.42	0.39
EDA + Sleep	w/ finetuning	0.51 (↑ 9%)	0.47 (↑ 8%)
	w/o pre-training	0.40 (↓ 2%)	0.36 (↓ 3%)
	Ours	0.72	0.69
PPG + Sleep	w/ finetuning	0.75 (↑ 3%)	0.73 (↑ 4%)
	w/o pre-training	0.69 (↓ 3%)	0.67 (↓ 2%)
	Ours	0.76	0.71
ECG + Sleep	w/ finetuning	0.79 (↑ 4%)	0.74 (↑ 3%)
	w/o pre-training	0.72 (↓ 4%)	0.68 (↓ 3%)

Accordingly, both factors are highly influential in regulation of emotions and mood, which conforms with our observations from this experiment.

Lastly, we analyze the trend of mood recognition performance based on the amount of training data used for all the methods. As shown in Figure 6, our proposed NapTune shows a much less steep decrease in performance when reducing training data, with up to a 62% F1 score when we use only 25% of the training data. In contrast, the other baselines show a much steeper descent, especially when reducing the training data from 50% to 25%.

6 CONCLUSION

We introduced NapTune, a prompt-tuning framework that involves adapting a frozen Transformer encoder, by adding lightweight prompt parameters into each transformer layer, to efficiently utilize the previous night's sleep-related measures as complementary information alongside wearable time-series data. Our results demonstrate that incorporating sleep data significantly improves mood classification, achieving up to an 11% increase in F1 scores over baseline methods that only use wearable data. Our proposed NapTune framework outperforms several state-of-the-art multimodal baselines by a minimum margin of 11%, 12%, and 2% for ECG, PPG, and EDA respectively. Future research directions may include studying other approaches to using pre-trained encoders, such as adaptors. Additionally, the notion of using generative models to fill the gap in instances where previous night's sleep measures may be missing, can be explored.

REFERENCES

[1] Syed Khairul Bashar, Eric Ding, Allan J Walkey, David D McManus, and Ki H Chon. 2019. Noise detection in electrocardiogram signals for intensive care unit patients. *IEEE Access*, 7, 88357–88368.

[2] Behnam Behinaein, Anubhav Bhatti, Dirk Rodenburg, Paul Hungler, and Ali Etemad. 2021. A transformer architecture for stress detection from ecg. In *Proceedings of the ACM International Symposium on Wearable Computers*, 132–134.

[3] Scott A Cairney, Simon J Durrant, Johan Hulleman, and Penelope A Lewis. 2014. Targeted memory reactivation during slow wave sleep facilitates emotional memory consolidation. *Sleep*, 37, 4, 701–707.

[4] Scott A Cairney, Simon J Durrant, Rebecca Power, and Penelope A Lewis. 2015. Complementary roles of slow-wave sleep and rapid eye movement sleep in emotional memory consolidation. *Cerebral Cortex*, 25, 6, 1565–1575.

[5] Rafael A Calvo and Sidney D'Mello. 2010. Affect detection: an interdisciplinary review of models, methods, and their applications. *IEEE Transactions on Affective Computing*, 1, 1, 18–37.

[6] Defu Cao, Furong Jia, Sercan O Arik, Tomas Pfister, Yixiang Zheng, Wen Ye, and Yan Liu. 2023. Tempo: prompt-based generative pre-trained transformer for time series forecasting. *arXiv preprint arXiv:2310.04948*.

[7] Chiara EG Castiglione-Fontanellaz and Leila Tarokh. 2023. Sleep and adolescent depression. *Clinical and Translational Neuroscience*, 8, 1, 3.

[8] Nicola Cellini, Katherine A Duggan, and Michela Sarlo. 2017. Perceived sleep quality: the interplay of neuroticism, affect, and hyperarousal. *Sleep Health*, 3, 3, 184–189.

[9] Kai-Wei Chang, Yu-Kai Wang, Hua Shen, Iu-thing Kang, Wei-Cheng Tseng, Shang-Wen Li, and Hung-yi Lee. 2023. Speechprompt v2: prompt tuning for speech classification tasks. *arXiv preprint arXiv:2303.00733*.

[10] Jianzhong Chen, Ke Sun, Yi Sun, and Xinxin Li. 2021. Signal quality assessment of ppg signals using stft time-frequency spectra and deep learning approaches. In *43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society*. IEEE, 1153–1156.

[11] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*. PMLR, 1597–1607.

[12] Anja Exler, Andrea Schankin, Christoph Klebsattel, and Michael Beigl. 2016. A wearable system for mood assessment considering smartphone features and data from mobile ecgs. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 1153–1161.

[13] Hamid Fekri Azgomi, Luciano R F. Branco, Md Rafiul Amin, Saman Khazaei, and Rose T Faghih. 2023. Regulation of brain cognitive states through auditory, gustatory, and olfactory stimulation with wearable monitoring. *Scientific reports*, 13, 1, 12399.

[14] Zhilin Gao, Xingran Cui, Wang Wan, Wenming Zheng, and Zhongze Gu. 2021. Ecsmp: a dataset on emotion, cognition, sleep, and multi-model physiological signals. *Data in Brief*, 39, 107660.

[15] Lisa Genzel, VI Spoormaker, BN Konrad, and M Dresler. 2015. The role of rapid eye movement sleep for amygdala-related memory processing. *Neurobiology of learning and memory*, 122, 110–121.

[16] Andrea N Goldstein and Matthew P Walker. 2014. The role of sleep in emotional brain function. *Annual review of clinical psychology*, 10, 679–708.

[17] Muhammad Anas Hasnul, Nor Azlina Ab Aziz, Salem Alelyani, Mohamed Mohana, and Azlan Abd Aziz. 2021. Electrocardiogram-based emotion recognition systems and their applications in healthcare—a review. *Sensors*, 21, 15, 5015.

[18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on cComputer Vision and Pattern Recognition*, 770–778.

[19] Shenda Hong, Yanbo Xu, Alind Khare, Satria Priambada, Kevin Maher, Alaa Aljiffry, Jimeng Sun, and Alexey Tumanov. 2020. Holmes: health online model ensemble serving for deep learning models in intensive care units. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1614–1624.

[20] Isabel C Hutchison and Shailendra Rathore. 2015. The role of rem sleep theta activity in emotional memory. *Frontiers in Psychology*, 6, 154468.

[21] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. 2022. Visual prompt tuning. In *European Conference on Computer Vision*. Springer, 709–727.

[22] Wenxiang Jiao, Michael Lyu, and Irwin King. 2020. Real-time emotion recognition via attention gated hierarchical memory network. In *Proceedings of the AAAI conference on Artificial Intelligence* number 05. Vol. 34, 8002–8009.

[23] Alistair EW Johnson et al. 2016. Mimic-iii, a freely accessible critical care database. *Scientific Data*, 3, 1, 1–9.

[24] David A Kalmbach, Vivek Pillai, Thomas Roth, and Christopher L Drake. 2014. The interplay between daily affect and sleep: a 2-week study of young women. *Journal of Sleep Research*, 23, 6, 636–645.

[25] Ashish Kapoor and Rosalind W Picard. 2005. Multimodal affect recognition in learning environments. In *Proceedings of the 13th Annual ACM International Conference on Multimedia*, 677–682.

[26] Walter Karlen. 2021. Capnobase ieee tbme respiratory rate benchmark. (2021).

[27] Harshit Kaushik, Tarun Kumar, and Kriti Bhalla. 2022. Isecurehome: a deep fusion framework for surveillance of smart homes using real-time emotion recognition. *Applied Soft Computing*, 122, 108788.

[28] Saskia Koldijk, Maya Sappelli, Suzan Verberne, Mark A Neerincx, and Wessel Kraaij. 2014. The swell knowledge work dataset for stress and user modeling research. In *Proceedings of the 16th International Conference on Multimodal Interaction*, 291–298.

[29] Monika Konjarski, Greg Murray, V Vien Lee, and Melinda L Jackson. 2018. Reciprocal relationships between daily sleep and mood: a systematic review of naturalistic prospective studies. *Sleep Medicine Reviews*, 42, 47–58.

[30] Zlatan Krizan, Nicholas A Boehm, and Caroline B Strauel. 2023. How emotions impact sleep: a quantitative review of experiments. *Sleep Medicine Reviews*, 101890.

[31] Min Seop Lee, Yun Kyu Lee, Dong Sung Pae, Myo Taeg Lim, Dong Won Kim, and Tae Koo Kang. 2019. Fast emotion recognition based on single pulse ppg signal with convolutional neural network. *Applied Sciences*, 9, 16, 3355.

[32] Ruth LF Leong and Michael WL Chee. 2023. Understanding the need for sleep to improve cognition. *Annual Review of Psychology*, 74, 27–57.

[33] Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 3045–3059.

[34] Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Lam Tam, Zhengxiao Du, Zhilin Yang, and Jie Tang. 2021. P-tuning v2: prompt tuning can be comparable to fine-tuning universally across scales and tasks. *arXiv preprint arXiv:2110.07602*.

[35] Dominique Makowski, Tam Pham, Zen J. Lau, Jan C. Brammer, François Lespinasse, Hung Pham, Christopher Schölzel, and S. H. Annabel Chen. 2021. NeuroKit2: a python toolbox for neurophysiological signal processing. *Behavior Research Methods*, 53, 4, 1689–1696. DOI: 10.3758/s13428-020-01516-y.

[36] Peter Mantello and Manh-Tung Ho. 2023. Emotional ai and the future of well-being in the post-pandemic workplace. *AI & Society*, 1–7.

[37] Juan Abdon Miranda-Correa, Mojtaba Khomami Abadi, Nicu Sebe, and Ioannis Patras. 2018. Amigos: a dataset for affect, personality and mood research on individuals and groups. *IEEE Transactions on Affective Computing*, 12, 2, 479–493.

[38] Jiapu Pan and Willis J Tompkins. 1985. A real-time qrs detection algorithm. *IEEE Transactions on Biomedical Engineering*, 3, 230–236.

[39] Marco AF Pimentel, Alistair EW Johnson, Peter H Charlton, Drew Birrenkott, Peter J Watkinson, Lionel Tarassenko, and David A Clifton. 2016. Toward a robust estimation of respiratory rate from pulse oximeters. *IEEE Transactions on Biomedical Engineering*, 64, 8, 1914–1923.

[40] Attila Reiss, Ina Indlekofer, Philip Schmidt, and Kristof Van Laerhoven. 2019. Deep ppg: large-scale heart rate estimation with convolutional neural networks. *Sensors*, 19, 14, 3079.

[41] Kyle Ross, Paul Hungler, and Ali Etemad. 2023. Unsupervised multi-modal representation learning for affective computing with multi-corpus wearable data. *Journal of Ambient Intelligence and Humanized Computing*, 14, 4, 3199–3224.

[42] Akane Sano, Z Yu Amy, Andrew W McHill, Andrew JK Phillips, Sara Taylor, Natasha Jaques, Elizabeth B Klerman, and Rosalind W Picard. 2015. Prediction of happy-sad mood from daily behaviors and previous sleep history. In *37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 6796–6799.

[43] Pritam Sarkar and Ali Etemad. 2020. Self-supervised learning for ecg-based emotion recognition. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*. IEEE, 3217–3221.

[44] Philip Schmidt, Attila Reiss, Robert Duerichen, Claus Marberger, and Kristof Van Laerhoven. 2018. Introducing wesad, a multimodal dataset for wearable stress and affect detection. In *Proceedings of the 20th ACM International Conference on Multimodal Interaction*, 400–408.

[45] Philip Schmidt, Attila Reiss, Robert Dürichen, and Kristof Van Laerhoven. 2019. Wearable-based affect recognition—a review. *Sensors*, 19, 19, 4079.

[46] Jainendra Shukla, Miguel Barreda-Angeles, Joan Oliver, Gora Chand Nandi, and Domenec Puig. 2019. Feature extraction and selection for emotion recognition from electrodermal activity. *IEEE Transactions on Affective Computing*, 12, 4, 857–869.

[47] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.

[48] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. 2020. Fixmatch: simplifying semi-supervised learning with consistency and confidence. *Advances in Neural Information Processing Systems*, 33, 596–608.

[49] Sahar Soltanieh, Javad Hashemi, and Ali Etemad. 2023. In-distribution and out-of-distribution self-supervised ecg representation learning for arrhythmia detection. *IEEE Journal of Biomedical and Health Informatics*.

[50] Sara Taylor, Natasha Jaques, Ehimwenma Nosakhare, Akane Sano, and Rosalind Picard. 2017. Personalized multitask learning for predicting tomorrow's mood, stress, and health. *IEEE Transactions on Affective Computing*, 11, 2, 200–213.

[51] Maia Ten Brink, Jessica R Dietch, Joshua Tutek, Sooyeon A Suh, James J Gross, and Rachel Manber. 2022. Sleep and affect: a conceptual review. *Sleep Medicine Reviews*, 101670.

[52] Robert Joseph Thomas, Joseph E Mietus, Chung-Kang Peng, and Ary L Goldberger. 2005. An electrocardiogram-based technique to assess cardiopulmonary coupling during sleep. *Sleep*, 28, 9, 1151–1161.

[53] Giuseppe Romano Tizzano, Matteo Spezialetti, and Silvia Rossi. 2020. A deep learning approach for mood recognition from wearable data. In *IEEE International Symposium on Medical Measurements and Applications*. IEEE, 1–5.

[54] Matthew P Walker and Els van Der Helm. 2009. Overnight therapy? the role of sleep in emotional brain processing. *Psychological bulletin*, 135, 5, 731.

[55] Junxiang Wang, Guangji Bai, Wei Cheng, Zhengzhang Chen, Liang Zhao, and Haifeng Chen. 2023. Prompt-based domain discrimination for multi-source time series domain adaptation. *arXiv preprint arXiv:2312.12276*.

[56] Claire Whiting, Nellia Bellaert, Christen Deveney, and Wan-Ling Tseng. 2023. Associations between sleep quality and irritability: testing the mediating role of emotion regulation. *Personality and Individual Differences*, 213, 112322.

[57] Christian D Wiesner, Julika Pulst, Fanny Krause, Marike Elsner, Lioba Baving, Anya Pedersen, Alexander Prehn-Kristensen, and Robert Göder. 2015. The effect of selective rem-sleep deprivation on the consolidation and affective evaluation of emotional memories. *Neurobiology of Learning and Memory*, 122, 131–141.

[58] Mark Lawrence Wong, Esther Yuet Ying Lau, Jacky Ho Yin Wan, Shu Fai Cheung, C Harry Hui, and Doris Shui Ying Mok. 2013. The interplay between sleep and mood in predicting academic functioning, physical health and psychological health: a longitudinal study. *Journal of Psychosomatic Research*, 74, 4, 271–277.

[59] Hao Xue and Flora D Salim. 2023. Promptcast: a new prompt-based learning paradigm for time series forecasting. *IEEE Transactions on Knowledge and Data Engineering*.

[60] Alexandros Zenonos, Aftab Khan, Georgios Kalogridis, Stefanos Vatsikas, Tim Lewis, and Mahesh Sooriyabandara. 2016. Healthyoffice: mood recognition at work using smartphones and wearable sensors. In *IEEE International Conference on Pervasive Computing and Communication Workshops*. IEEE, 1–6.

[61] Nan Zhang et al. 2023. Psychosocial factors of insomnia in depression: a network approach. *BMC Psychiatry*, 23, 1, 949.

[62] Lili Zhu, Petros Spachos, and Stefano Gregori. 2023. Electrodermal activity for emotion recognition using cnn and bi-gru model. In *IEEE International Conference on Communications*. IEEE, 5533–5538.

[63] Zack Zhu, Hector F Satizabal, Ulf Blanke, Andres Perez-Uribe, and Gerhard Tröster. 2015. Naturalistic recognition of activities and mood using wearable electronics. *IEEE Transactions on Affective Computing*, 7, 3, 272–285.

[64] Martin Zihlmann, Dmytro Perekrestenko, and Michael Tschannen. 2017. Convolutional recurrent neural networks for electrocardiogram classification. In *Computing in Cardiology*. IEEE, 1–4.