Towards Self-Supervised Learning of ECG Signal Representation for the Classification of Acute Stress Types

Rajdeep K. Nath
VTT Technical Research Centre of Finland
Kuopio, Finland
rajdeep.nath@vtt.fi

Jaakko Tervonen
VTT Technical Research Centre of Finland
Espoo, Finland
jaakko.tervonen@vtt.fi

Johanna Närväinen
VTT Technical Research Centre of Finland
Kuopio, Finland
johanna.narvainen@vtt.fi

Kati Pettersson
VTT Technical Research Centre of Finland
Espoo, Finland
kati.pettersson@vtt.fi

Jani Mäntyjärvi
VTT Technical Research Centre of Finland
Oulu, Finland
jani.mantyjarvi@vtt.fi

ABSTRACT

We present a novel application of contrastive learning technique in learning the feature representation of ECG signal in a self-supervised manner for the classification of acute stress types. Acute stress types that occur for a very short period and are rapidly changing and alternating in nature are difficult to classify using conventional ECG features. This is because the change in conventional ECG features due to rapid and alternating acute stressors do not reflect instantaneously. We hypothesize that deep-learned features from ECG signals can better distinguish between the different stress types than conventional ECG features. Our proposed approach can generate distinct feature representations for the physical and mental stress task type using very short window lengths. Our results show that the deep-learned features perform better in terms of accuracy and F1 score in distinguishing between physical and mental stress task types. In the future, our proposed method can be used in a real world setting for understanding the dynamics of different stressors in a self-supervised fashion without the need for human labeling.

CCS CONCEPTS

• **Computing methodologies** → **Unsupervised learning**; **Neural networks**.

KEYWORDS

Acute stress, Electrocardiogram (ECG), Contrastive Learning, Smart Healthcare.

ACM Reference Format:

Rajdeep K. Nath, Jaakko Tervonen, Johanna Närväinen, Kati Pettersson, and Jani Mäntyjärvi. 2023. Towards Self-Supervised Learning of ECG Signal Representation for the Classification of Acute Stress Types. In *Proceedings of the Great Lakes Symposium on VLSI 2023 (GLSVLSI '23), June 5–7, 2023, Knoxville, TN, USA.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3583781.3590252

GLSVLSI '23, June 5–7, 2023, Knoxville, TN, USA

ACM ISBN 979-8-4007-0125-2/23/06.
https://doi.org/10.1145/3583781.3590252

1 INTRODUCTION

The reaction of the human body to changing and demanding environment elicits stress responses which are reflected through physiological and behavioral changes. To a certain extent, stress responses are necessary to maintain focus and deal with challenging situations. However, long-term continuous exposure to stressful situations is linked with the development of several chronic illnesses such as hypertension and anxiety disorder [2]. Furthermore, indirect effects of stress are related to poorer work performance and reduced quality of life which hampers the economy [18]. Although stress is an inevitable aspect of daily life, monitoring and managing stress can help mitigate its long-term detrimental effects. Acquiring and monitoring different types of data from a variety of end points can help us understand the stress dynamics that are harmful to individuals. Analyses of these data can enable us to build smart healthcare systems for real-time stress detection and management and minimize the chances of developing chronic stress-related illnesses later in life.

Physiological signals such as galvanic skin response (GSR) and electrocardiogram (ECG) are commonly used for stress detection as these can be easily acquired facilitating long-term continuous monitoring [14]. Although supervised machine learning techniques are most commonly used in the context of stress detection, unsupervised methods have been explored recently [6]. The stress detection models proposed in the literature are mostly aligned toward a binary classification of stressed and not-stressed states or towards a multi-class classification of stressed states based on some ground truth reference. There are only few works that have attempted to classify the different stress types, especially when the stress stimuli are rapid, alternating, and very short in duration, lasting for only about a minute. Distinguishing between these rapid and alternating stress stimuli is difficult because of the relatively slower response in sympathetic activation compared to the rapidly changing stress stimuli. Most of the existing studies that have attempted to distinguish between different stress types have used stress protocols with a recovery phase between the different stress tasks, which is unlikely in real-life scenarios.

Different stress types have different effects on the intensity of sympathetic activation and recovery. Research evidence indicates that changes in sympathetic activation during mental stress tasks are more profound compared to physical stress tasks [4, 9]. Therefore it is important to develop stress detection models that can capture the dynamics of such rapid, alternating, and short-term stress stimuli. If the different stress types in real-world dynamics can be distinguished, appropriate intervention and stress-type-specific management feedback can be initiated to help an individual calm down more efficiently. This will be an important step toward developing smart stress detection and management systems in the future.

In this paper, we propose a novel method for distinguishing between rapidly alternating physical and mental stress tasks with no recovery phases in between. We hypothesize that using deep-learned features from ECG signals will be able to better capture the differences between physical and mental stress tasks compared to conventional time-domain ECG features typically used for stress detection. The main contribution of our work is as follows:

- We propose a self-supervised ECG feature representation model based on a simple framework for contrastive learning (SimCLR).
- We propose a semi-supervised deep learning model (using 20% of training data) for the classification of ultra-short-term physical and mental stress tasks occurring rapidly and alternatively.
- We evaluate the performance of the deep-learned features as input to three benchmark machine learning algorithms for distinguishing between the rapid physical and mental stress task types.
- The performance of deep-learned features is compared with that of conventional ECG time-domain features for the classification of physical and mental stress.

2 RELATED WORK

Self-supervised approach for representation learning of ECG signal has gained popularity in recent years [8, 13]. SimCLR framework for contrastive learning has been predominantly used for computer vision tasks such as image recognition and object identification [1]. However, recently, these approaches are being applied to physiological signals for healthcare applications such as emotion recognition and stress detection [3, 12]. In [12], authors have used the SimCLR framework for classifying between baseline, stress, amusement, and meditation tasks using ECG signals. In [3], authors have used multiple physiological signals for various kinds of emotion detection validated on a different dataset. However, the problem of distinguishing between different ultra-short stress task types with no recovery period in between has not been investigated using these approaches. Instead, manual feature extraction methods have been used to extract features from ECG and electrooculogram (EOG) signals for distinguishing between different ultra-short stress task types with no recovery period in between [10]. Furthermore, the effect of ultra-short window lengths on model accuracy in distinguishing between cognitive load and rest has been analyzed [17]. It was observed that reducing window length significantly reduces the model accuracy and feature integration from multiple sensing modalities is required to efficiently distinguish between different stress types that are rapid and alternating in nature. In this work, we attempt to pursue a novel direction for distinguishing between rapid and alternating stress tasks using self-supervised learning.

3 OVERVIEW OF THE PROPOSED WORK

The overview of the proposed work is shown in Figure 1. The proposed work broadly consists of three parts:

- SimCLR framework for ECG feature representation and stress classification (Brown dotted box in Figure 1).
- Stress Classification using deep-learned features and benchmark machine learning algorithms (Blue dashed box in Figure 1).
- Conventional ECG feature extraction (Black solid box in Figure 1).

Before, we discuss the three parts of the proposed work in details, we will discuss the data collection and study protocol used in this study.

3.1 Data collection and signal processing

Data from 21 healthy young adults were used for this study. The study proposal was evaluated by the Ethics Committee in the Humanities and Social and Behavioural Sciences of the University of Helsinki. The whole study consisted of several tasks including the Maastricht acute stress test (MAST) [16]. For our analyses, only the MAST protocol was chosen because it has elements of alternating and short-term stress task types which is an important aspect of the study objective. It has been shown earlier that distinguishing between baseline and stressed state is easy, whereas distinguishing between the different stress task types is challenging [10]. Hence, for this analysis, our objective is only to distinguish between the different stress task types, and hence, the baseline measurement was excluded. MAST (Figure 2) consists of short and alternating physical stress tasks (cold pressor task) and mental stress (mental arithmetics) tasks. These stress tasks vary in duration (45 to 90 sec) with no recovery periods in between. During the cold pressor task, the participants were required to immerse their hands in cold water (temperature of 2° Celsius). During the mental math task, the participants were required to perform verbal subtractions fast and accurately under time pressure. More details on the experimental protocol and data collection are discussed in our previous work [10].

The ECG signal was acquired at a sampling rate of 1000 Hz using the NeurOne system (Bittium, Oulu, Finland) that employed a single lead ECG sensor between the left collarbone and right lower back. The sampling rate for ECG signal acquisition was high which is not usually the case in consumer-grade ECG sensors. Furthermore, a high sampling rate is not suitable for long-term signal acquisition and monitoring. Hence, we down-sampled the ECG signal to 250 Hz before further analyses. The down-sampled ECG signals were then filtered using a bandpass Butterworth filter in the frequency range of 2-30 Hz. Other than filtering, no other preprocessing is performed on the ECG signal. Subsequently, the ECG signals are segmented into 8-second lengths with an overlap of 4 sec for each stress task type (Figure 2). It is to be noted that, no overlapping signal segment was extracted in between the stress

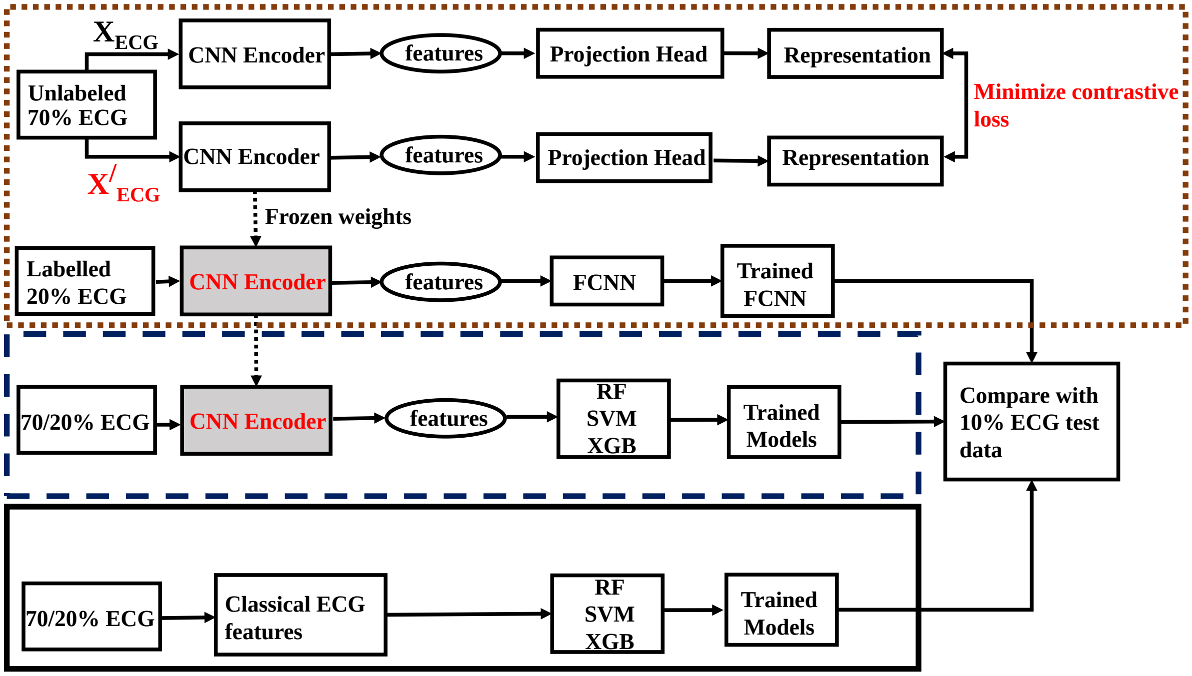

Figure 1: Overview of the proposed work. The brown dotted box represents the SimCLR framework for self-supervised learning of ECG features (upstream task) and the supervised training of the fully connected neural network (FCNN) during the downstream task. The blue dashed box represents training the benchmark machine learning algorithms with deep-learned features and the black solid box represents training the benchmark machine learning algorithms with conventional time-domain ECG features.

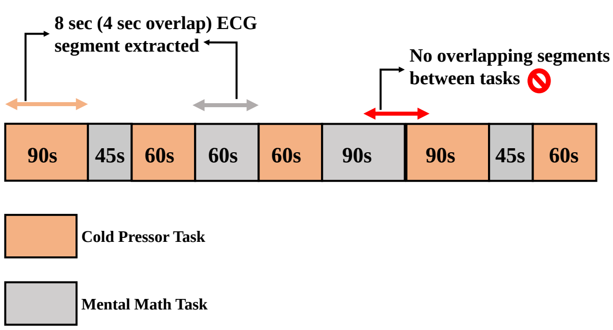

Figure 2: Overview of the MAST stress induction protocol

task types to avoid information leakage between stress task types. For subsequent tasks, the whole dataset consisting of 8 sec ECG signal segments are randomly divided into three parts (70% and 20% for training, and 10% for testing).

3.2 SimCLR framework for ECG feature representation and stress classification

The SimCLR framework consists of two tasks, (i) upstream task and (ii) downstream task. The upstream task consists of two modules, (i) CNN-based encoder networks and (ii) Projection heads. The architecture of our CNN encoder (Table 1) is inspired by the CNN encoder presented in [13] with slight modifications. In the upstream task, the original ECG signal segment (X_{ECG}), and the augmented version of the ECG signal segment (X'_{ECG}) are fed as input in parallel to two CNN encoders. The CNN encoders generate a higher-level representation, known as the feature space of the ECG signal segments. The augmented version of the ECG signal is generated by adding random Gaussian noise, scaling, and introducing random DC shift to the original ECG signal segment. The feature space is then converted into a more expressive representation using a projection head that is attached to the output of the CNN encoder. The contrastive learning algorithm is then applied to the representations of X_{ECG} and X'_{ECG} to maximize the similarity between the representations generated by X_{ECG} and X'_{ECG}. X_{ECG} and X'_{ECG} form a positive pair while both X_{ECG} and X'_{ECG} form a negative pair with all other ECG signal segments. When an augmented representation of an ECG signal segment forms a positive pair with the original ECG signal segment, the contrastive learning algorithm tries to learn the underlying features that make the original ECG signal segment and its augmented counterpart similar to each other. For the upstream task, 70% of unlabeled ECG signal segments are used, and once trained, the trained CNN encoder is used for the downstream task.

In the downstream task, the weights of the trained CNN encoder are kept frozen throughout the training of the fully connected neural network (FCNN) (Table 2). During the downstream task, the FCNN for stress-type classification tasks is trained using only 20%

Table 1: Architecture details of CNN encoder and projection head used in the upstream task.

Module Name	Layer Details	No. of Parameters
Encoder	Conv1D (filters = 32)	1056
	Conv1D (filters = 32)	32800
	MaxPool	0
	Conv1D (filters = 64)	32832
	Conv1D (filters = 64)	65600
	MaxPool	0
	Conv1D (filters = 128)	65664
	Conv1D (filters = 128)	131200
	GlobalMaxPool	0
	Flatten	0
	Dense (N = 128)	16512
Projection Head	Dense (N = 256)	33024
	Dense (N = 128)	32896

of the labeled ECG segment. The upstream and downstream tasks are implemented using Keras with the Tensorflow backend.

Table 2: Architecture details of fully connected neural network for stress task type classification.

Layer Details	No. of Parameters	Activation
Dense (N = 128)	16512	Relu
Dense (N = 64)	8256	Relu
Dense (N = 32)	2080	Relu
Dense (N = 2)	66	Softmax

3.3 Stress Classification using deep-learned features and benchmark machine learning algorithms

To evaluate the efficacy of deep-learned features from the trained CNN encoder, they are used as an input to three benchmark machine learning algorithms, random forest (RF) classifier, support vector machine (SVM), and XGBoost (XGB). The training of these benchmark machine learning algorithms is performed with both 70% and 20% of labeled ECG signal segments.

3.4 Conventional ECG feature extraction

The conventional time-domain features from ECG signals are extracted from the 8-second ECG signal segment for training the three benchmark machine learning algorithms to distinguish between the stress types tasks. The extracted time-domain features from ECG include heart rate, interbeat interval (IBI), respiratory frequency, and 18 heart rate variability (HRV) time-domain features. These features are extracted using the NeuroKit2 library in python [7]. A total of 21 features are used as conventional ECG features. It is to be noted that, several HRV features cannot be extracted since the window length used is too short [15]. Finally, the performance of the trained models using all three approaches is evaluated using 10% of the ECG segment (test data).

4 RESULTS

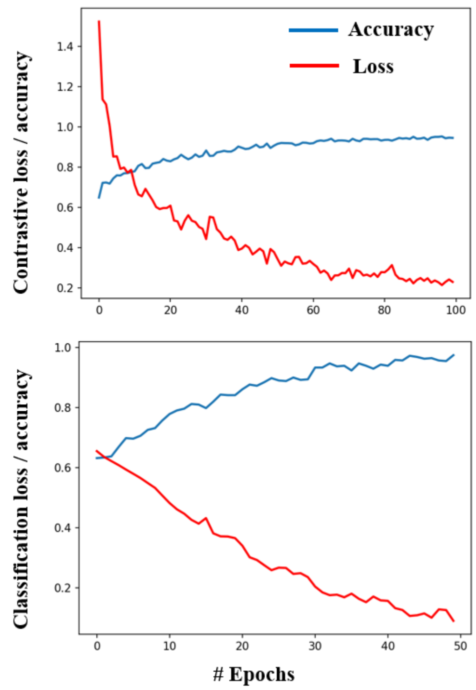

Figure 3: Training performance of the upstream task (top) and classification task (bottom).

The upstream contrastive learning task was trained for 100 epochs with a batch size of 64. Adam optimizer was used to optimize the contrastive loss function. During the upstream tasks, there were a total of 411,584 trainable parameters, of which, the CNN encoder had 345,664 trainable parameters. The contrastive loss and accuracy continued to improve until 90 epochs (Figure 3), after which both loss and accuracy start to saturate.

Figure 4 visualizes the average feature activation values generated by the CNN encoder for the whole dataset. It can be observed that the feature activations during both the physical and mental stress task have similar patterns. However, some of the feature activations during the physical stress tasks have higher amplitude compared to the activations during the mental math task. These apparent differences in the amplitude of feature activations could serve as an important indicator of the difference between the two stress task types.

The FCNN for the downstream task was also trained similarly with the Adam optimizer and categorical cross entropy as the loss function. The network was trained for 50 epochs with a batch size of 64. The FCNN was trained for only 50 epochs because of the relatively low number of trainable parameters (26,914). The classification loss and accuracy both showed an almost linear improvement trend (Figure 3) with the number of epochs and saturated by the end of 50 epochs. Subsequently, the test data set (10% of ECG segments) is fed into the trained CNN encoder to generate the feature space representation of the test data and finally as an input to the trained FCNN model for the classification of stress types.

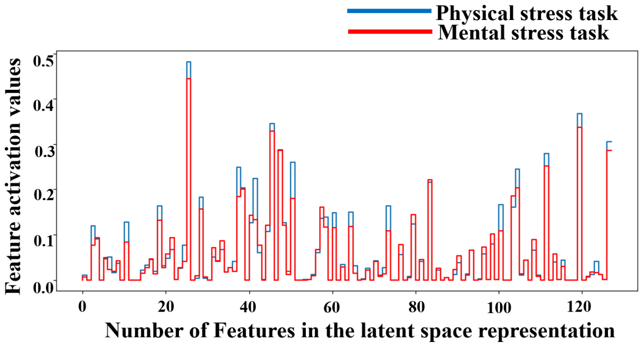

Figure 4: Average feature activation values generated by the trained CNN encoder for the entire dataset.

Table 3: Performance of the FCNN model in classifying between the physical and mental stress tasks.

	Precision	Recall	F1-Score	Support
Physical task	0.77	0.80	0.78	177
Mental task	0.64	0.61	0.62	107
Macro avg	0.71	0.70	0.70	284
Weighted avg	0.72	0.73	0.72	284
Accuracy = 73%				

Table 4: Confusion matrix for the FCNN model in classifying between the physical and mental stress tasks.

		Predicted		
		Physical task	Mental task	Total
Actual	Physical task	141	36	177
	Mental task	42	65	107
	Total	183	101	N

The overall accuracy achieved by the FCNN model in classifying between the physical and mental stress task-type is 73% (Table 3). In specific task-types, FCNN has a higher F1-score (0.78) in detecting physical stress tasks compared to mental stress tasks (0.62). The fact that there are fewer ECG signal segments representing the mental stress task compared to the physical stress task is one of the likely factors behind the higher F1-score of the model in detecting the physical stress tasks. Out of 177 instances of physical stress task-type ECG segment, the FCNN model was able to correctly classify 141 instances (Table 4). Whereas, out of 107 instances of mental stress task-type ECG segment, only 65 instances were correctly detected by the FCNN model.

4.1 Comparison with benchmark machine learning algorithms and conventional ECG features

The benchmark algorithms; random forest classifier, support vector machine, and XGBoost are trained using similar hyperparameters for both deep-learned features and classical HRV features. For the random forest classifier, the number of estimators was 200 and the maximum depth was 20. For the support vector machine, a linear kernel was used with the value of regularization parameter C fixed at 500. For the XGBoost classifier, the number of estimators was 200 with a maximum depth of 30 and a learning rate of 0.00001. The values of hyperparameters were selected to be roughly optimized based on the classical ECG training features.

Random forest classifier achieved an accuracy of 55% and 62% on the test set when trained with 20% and 70% of data using the conventional ECG features respectively (Table 5). However, the random forest classifier was able to achieve an accuracy of 68% when trained with only 20% of the data using deep-learned features. However, when the random forest classifier was trained with 70% of data using deep-learned features, the accuracy increased to 77% and the F1-score was 0.75. A similar trend was also observed with the support vector machine and XGBoost classifier as both of these classifiers performed better when trained with deep-learned features as compared to conventional ECG features. For example, SVM trained with 20% of data using deep-learned features achieved 17.9% higher performance in terms of accuracy compared to SVM trained with conventional ECG features (Table 5). Similarly, the XGBoost classifier when trained with 20% of data using deep-learned features achieved about a 20% increase in accuracy compared to training with conventional ECG features (Table 5).

5 DISCUSSIONS AND FUTURE WORK

In this work, we propose a self-supervised approach for learning ECG signal representation for the classification of physical and mental stress states, induced using MAST. These two stressors in MAST are rapid and alternating in nature with no recovery phase in between. Our results showed that deep-learned features are more efficient in classifying stress task-types than conventional features extracted from ECG signals. Our results are in line with previous work that has explored the deep-learning-based approach in analyzing ECG signals in applications such as stress detection and detection of paroxysmal atrial fibrillation (PAF) [5, 11]. In [11], researchers have found that CNN-based neural networks were competitive in learning key features from raw ECG signals for PAF patient screening. Similarly, in [5] deep-learning-based approach for feature representation from ECG signal was competitive in detecting stressful states. Our approach of using a self-supervised method for ECG feature representation in a novel application of classifying between rapid and alternating short-term stress types is likely to be useful in a real-world scenario where labeling of stress types is challenging. Furthermore, our work shows that deep-learned features can prove to be superior in situations where conventional ECG features might not be sufficient, such as in capturing the rapidly changing stress dynamics. Overall the results are promising and call for future research with more rigorous exploration and evaluation.

The final set of feature representations generated by the CNN encoder resulted in 128 features which is significantly higher than the 21 classical features extracted from the ECG signal. However, it is to be noted that, it is not possible to extract a higher number of features from the ECG signal in ultra-short window length [15]. Moreover, it has been observed that increasing the number of

Table 5: Performance comparison on test data with benchmark machine learning algorithms using deep-learned features and classical features.

Classifiers	Deep Learned Features		Conventional ECG Features	
	Training (70%)	Training (20%)	Training (70%)	Training (20%)
Random Forest	Acc = 77%, F1-Score = 0.75	Acc = 68%, F1-Score = 0.64	Acc = 62%, F1-Score = 0.56	Acc = 55%, F1-Score = 0.51
Support Vector Machine	Acc = 65%, F1-Score = 0.61	Acc = 67%, F1-Score = 0.61	Acc = 58%, F1-Score = 0.42	Acc = 55%, F1-Score = 0.44
XGBoost	Acc = 70%, F1-Score = 0.69	Acc = 64%, F1-Score = 0.60	Acc = 60%, F1-Score = 0.57	Acc = 51%, F1-Score = 0.50
FCNN	NA	Acc = 73%, F1-Score = 0.70	NA	NA

features does not necessarily result in higher performance [10] in stress detection. Furthermore, the deep-learning-based approach automatically diminishes feature values that are not important and hence is a desirable attribute for deep-learning-based feature learning approach.

A possible limitation of the work is that since we have randomly split the ECG signal segments into 70% for feature learning, 20% for classification, and 10% for testing, it is possible that ECG signal segments from the same subject were present in all or at least two of the subsets. Hence, there is a possibility of information leakage on a subject level. However, for this work, the objective was to verify if such an approach is feasible in the context of distinguishing between stress types that are rapid and alternating in nature. Hence, a more rigorous training and testing approach will be implemented in the future extension of this work. Moreover, the comparison with the conventional technique is also performed with the same data split method to ensure a fair comparison. In addition, while segmenting the ECG signal into 8 seconds length, we made sure not to include overlapping segments between stress tasks. This step was crucial in enforcing that there was no information leakage between the stress task types.

In the future, the proposed method for self-supervised ECG signal representation for stress type classification can be extended to other signals that can also be monitored continuously in everyday life such as galvanic skin response (GSR) and photoplethysmogram (PPG). Further, multi-modal feature representation using the proposed method can enhance the performance of stress type classification. In the long-term, such self-supervised models are suitable for deployment in the real-world as they can potentially benefit in learning from the vast amount of data collected in different contexts. Moreover, our proposed SimCLR based stress detection framework is suitable for adapting in a client-edge-cloud architecture with potential for improvements in data privacy as all the data are encoded.

ACKNOWLEDGMENTS

The work was funded by the Academy of Finland under Grant-Nos.: 334092, 313401, 351282 and VTT.

REFERENCES

[1] 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.

[2] George P Chrousos. 2009. Stress and disorders of the stress system. *Nature reviews endocrinology* 5, 7 (2009), 374–381.

[3] Vipula Dissanayake, Sachith Seneviratne, Rajib Rana, Elliott Wen, Tharindu Kaluarachchi, and Suranga Nanayakkara. 2022. Sigrep: Toward robust wearable emotion recognition with contrastive representation learning. *IEEE Access* 10 (2022), 18105–18120.

[4] Suh-Yeon Dong, Miran Lee, Heesu Park, and Inchan Youn. 2018. Stress resilience measurement with heart-rate variability during mental and physical stress. In *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 5290–5293.

[5] Bosun Hwang, Jiwoo You, Thomas Vaessen, Inez Myin-Germeys, Cheolsoo Park, and Byoung-Tak Zhang. 2018. Deep ECGNet: An optimal deep learning framework for monitoring mental stress using ultra short-term ECG signals. *TELEMEDICINE and e-HEALTH* 24, 10 (2018), 753–772.

[6] Talha Iqbal, Adnan Elahi, William Wijns, and Atif Shahzad. 2022. Exploring unsupervised machine learning classification methods for physiological stress detection. *Frontiers in Medical Technology* 4 (2022).

[7] 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 (feb 2021), 1689–1696. https://doi.org/10.3758/s13428-020-01516-y

[8] Temesgen Mehari and Nils Strodthoff. 2022. Self-supervised representation learning from 12-lead ECG data. *Computers in biology and medicine* 141 (2022), 105114.

[9] Ákos Móra, Zsolt Komka, József Végh, István Farkas, Gyöngyi Szilágyi Kocsisné, Edit Bosnyák, Márta Szmodis, Roland Ligetvári, Éva Csöndör, Gábor Almási, et al. 2022. Comparison of the cardiovascular effects of extreme psychological and physical stress tests in male soccer players. *International Journal of Environmental Research and Public Health* 19, 2 (2022), 715.

[10] Kati Pettersson, Jaakko Tervonen, Johanna Närväinen, Pentti Henttonen, Ilmari Määttänen, and Jani Mäntyjärvi. 2020. Selecting Feature Sets and Comparing Classification Methods for Cognitive State Estimation. In *2020 IEEE 20th International Conference on Bioinformatics and Bioengineering (BIBE)*. IEEE, 683–690.

[11] Bahareh Pourbabaee, Mehrsan Javan Roshtkhari, and Khashayar Khorasani. 2017. Deep convolutional neural networks and learning ECG features for screening paroxysmal atrial fibrillation patients. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 48, 12 (2017), 2095–2104.

[12] Suha Rabbani and Naimul Khan. 2022. Contrastive Self-Supervised Learning for Stress Detection from ECG Data. *Bioengineering* 9, 8 (2022), 374.

[13] Pritam Sarkar and Ali Etemad. 2020. Self-supervised ECG representation learning for emotion recognition. *IEEE Transactions on Affective Computing* 13, 3 (2020), 1541–1554.

[14] Sibi C Sethuraman, Pranav Kompally, Saraju P Mohanty, and Uma Choppali. 2021. MyWear: a novel smart garment for automatic continuous vital monitoring. *IEEE Transactions on Consumer Electronics* 67, 3 (2021), 214–222.

[15] Fred Shaffer and Jay P Ginsberg. 2017. An overview of heart rate variability metrics and norms. *Frontiers in public health* (2017), 258.

[16] Tom Smeets, Sandra Cornelisse, Conny WEM Quaedflieg, Thomas Meyer, Marko Jelicic, and Harald Merckelbach. 2012. Introducing the Maastricht Acute Stress Test (MAST): a quick and non-invasive approach to elicit robust autonomic and glucocorticoid stress responses. *Psychoneuroendocrinology* 37, 12 (2012), 1998–2008.

[17] Jaakko Tervonen, Kati Pettersson, and Jani Mäntyjärvi. 2021. Ultra-short window length and feature importance analysis for cognitive load detection from wearable sensors. *Electronics* 10, 5 (2021), 613.

[18] Peter Warr and Karina Nielsen. 2018. Wellbeing and work performance. *Handbook of well-being. Salt Lake City, UT: DEF Publishers* (2018).