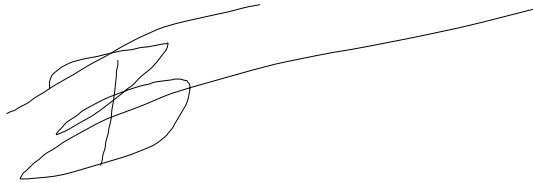


Declaro, bajo mi responsabilidad, que el Proyecto presentado con el título
“An Optimized and Explainable Pre-Trained Transformer Model for Accurate Stress
Detection using ECG Signals”

en la ETS de Ingeniería - ICAI de la Universidad Pontificia Comillas en el
curso académico 2023/24 es de mi autoría, original e inédito y

no ha sido presentado con anterioridad a otros efectos.

El Proyecto no es plagio de otro, ni total ni parcialmente y la información que ha sido
tomada de otros documentos está debidamente referenciada.



Fdo.: Francisco Barragán Castro

Fecha: 11/ 07/ 2024

Autorizada la entrega del proyecto

EL DIRECTOR DEL PROYECTO

Fdo.: Berta Ruíz González Fecha: 16/07/2024

Álvaro López López Fecha: 16/07/2024

An Optimized and Explainable Pre-Trained Transformer Model for Accurate Stress Detection using ECG Signals

Francisco Barragán Castro
fbarraganc@alu.comillas.edu

Berta Ruiz González
brgonzalez@comillas.edu

Álvaro López López
allopez@comillas.edu

Abstract

This paper presents an enhanced and interpretable pre-trained transformer model designed for precise stress detection using 1-lead electrocardiogram (ECG) signals. Our model combines self-supervised learning with transformer architecture to efficiently process raw ECG data while ensuring explainability. By utilizing self-supervised learning, we effectively manage unlabeled data and achieve leading accuracy in stress detection. We prioritize model interpretability with saliency maps, ensuring compliance with medical standards and fostering trust in AI-driven healthcare solutions. This methodology establishes a new standard for ECG-based stress detection on the WESAD dataset, merging advanced AI techniques with practical medical applications.

Keywords: Transformer Models, Electrocardiogram (ECG), Stress Detection, Self-Supervised Learning, AI in Healthcare, Medical AI Interpretability, WESAD Dataset

1 Introduction

Stress-related health issues are becoming increasingly prevalent worldwide and seriously impact people’s mental health and quality of life [1]. Advancements in wearable technology present a promising alternative for stress analysis, with sensors continuously monitoring physiological signals such as electrocardiograms (ECGs). ECGs reflect the heart’s electrical activity and can provide insights into an individual’s stress response.

Recent technological advancements have paved the way for innovative applications across various sec-

tors. Notably, the synergy between deep learning and sensor technology holds significant promise for healthcare and wellness monitoring. Wireless measurements of critical health parameters like ECGs could soon become commonplace[2, 3], revolutionizing traditional industries and fostering innovation by opening the door for continuous non-intrusive measurement.

1.1 Related Work

The literature indicates that transformer-based models are at the forefront of this technological revolution, particularly for analyzing ECG signals[4, 5, 6]. Transformers, characterized by their ability to handle sequential data and self-attention capabilities, have demonstrated superior performance compared to other approaches like Convolutional Neural Networks (CNNs)[7, 8]. Their inherent ability to capture long-term dependencies in data makes them ideal for ECG signal analysis, where the temporal relationship between different parts of the signal is crucial. This superiority in handling ECG data, combined with their versatility and scalability, underscores their potential as key tools in developing emotion-sensitive models [5].

1.2 Methodology

Our approach integrates self-supervised learning with transformer architecture to effectively process raw ECG data while ensuring explainability. By leveraging self-supervised learning, our model autonomously handles unlabeled data, significantly improving its ability to learn useful representations from the data itself. The transformer architecture, known for its prowess in capturing long-term dependencies in sequential data, further enhances the

model’s accuracy.

To ensure the model’s interpretability, we incorporate saliency maps. These tools provide visual insights into the model’s decision-making process, aligning with medical standards and building trust in AI-driven healthcare applications. Our model sets a new benchmark for ECG-based stress detection, demonstrating superior performance on the WE-SAD dataset [9] compared to existing methods.

The rest of the paper is organized as follows: In Section 2, we delve into the methodology employed in our study, detailing the data collection and implementation specifics. In Section 3, we present and analyze the results of the proposed method, showcasing its performance compared to existing models. In Section 4, we explore the interpretability of these results, offering insights and explanations for the observed patterns and trends through saliency maps. Finally, in Section 5, we draw our conclusions, highlighting the implications of our findings and suggesting potential avenues for future research.

2 Proposed Method

We apply a two stage training process to build an encoder-decoder architecture. The first pretraining stage consists uses self-supervised learning to build a robust encoder that can capture the most vital parts of the ECG signal in a latent space. On the second phase, we leverage the pretrained weights of this encoder to fine-tune them whilst training a decoder to a stress detection model.

2.1 Self-supervised pretraining

The first stage employs a self-supervised learning model to autonomously identify patterns and features within ECG data, establishing a robust foundational representation. This foundation is critical because it allows the model to learn from a large amount of unlabeled data, enhancing its capability to generalize and perform well in subsequent stress classification tasks.

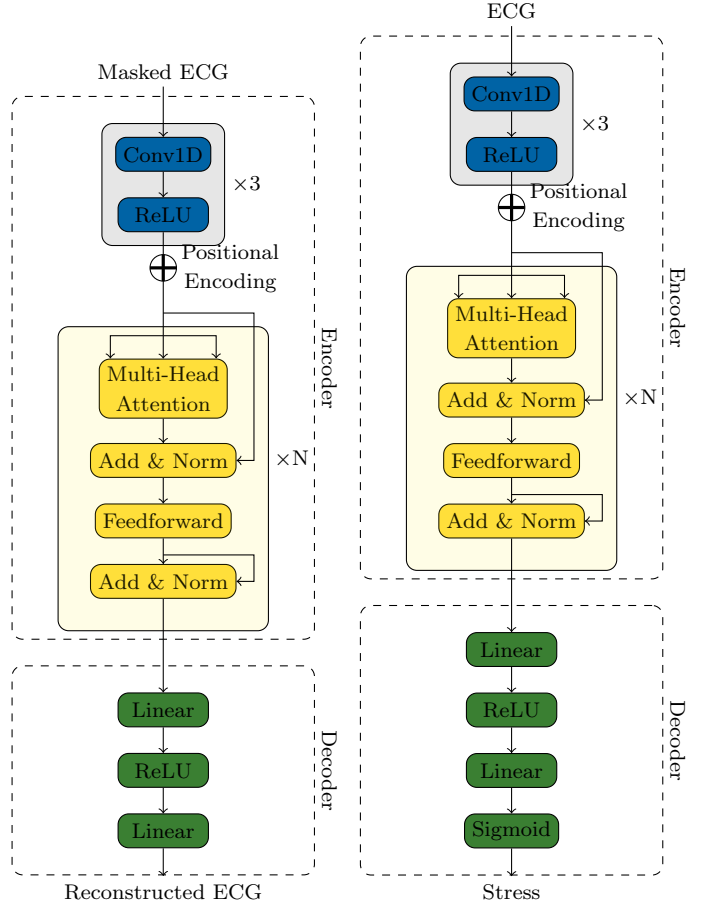


Figure 1: Architecture of the models

2.1.1 Masked Reconstruction Task

We chose a task-based approach, specifically masked reconstruction, over contrastive learning [10] due to its effectiveness in learning representations from time-series data[11][12][13]. Masked reconstruction requires the model to predict occluded parts of the ECG signal, guided by a binary mask. This task forces the model to understand the underlying structure of the ECG waveform, which is essential for accurate stress detection.

Our masking strategy involves fixed-width masks of 39 samples, randomly generated with a 1.66% likelihood and capable of overlapping. This ensures that the model learns to reconstruct significant portions of the ECG signal, capturing essential features relevant to stress detection, very similar to [13].

2.1.2 Loss Function

We opted for Mean Square Error (MSE) loss, focusing on the masked sections of the ECG signal. This choice directs the model to learn efficiently from the occluded data, emphasizing the reconstruction of the most challenging parts of the signal. This approach improves the model’s ability to handle real-world data where noise and occlusions are common.[13][12]

2.1.3 Encoder Architecture

Our encoder comprises a CNN followed by a Transformer. The CNN is designed to highlight specific heartbeat patterns, essential for detecting subtle stress-related changes in the ECG signal. The Transformer captures global dependencies within the sequence, making it well-suited for long-term temporal analysis inherent in ECG data. This combination leverages the strengths of both architectures, providing a detailed and comprehensive understanding of the ECG signals.

Layer	Configuration
1	Conv1D ($d = 64, k = 123, s = 1, p = 61$) + LN + ReLU + dropout = 0.1
2	Conv1D ($d = 128, k = 65, s = 1, p = 32$) + LN + ReLU + dropout = 0.1
3	Conv1D ($d = 256, k = 33, s = 1, p = 16$) + LN + ReLU + dropout = 0.1

Table 1: Convolutional embedding configurations

In order for the CNN to work as intended, the perceptive field must be of similar length to one heartbeat [5], therefore the CNNs kernel, stride, and padding should be calculated accordingly. Given our sampling rate of 250 Hz, the parameters shown in Table 1, create a perceptive field of 0.8 seconds (75bpm), slightly below the resting heart rate of a healthy adult (65bpm)[13]. This design choice aims to enhance performance for stressed heart rates, typically higher due to tachycardia. The Transformer’s fixed positional encoding and multi-head attention mechanisms further ensure accurate temporal modeling of the ECG data. In the transformer encoder we denote the number of Transformer blocks

as N , the hidden size as d_{model} , the number of self-attention heads as h , and the dimension of the feed-forward layer as d_{ff} . We present results for two different Transformer model sizes shown in Table 2.

Model	N	d_{model}	d_{ff}	h
$MODEL_{Large}$	2	256	512	2
$MODEL_{Little}$	1	256	512	2

Table 2: Transformer Model Configurations

In both cases, the Transformer blocks incorporate a dropout rate of 0.1, Batch Normalization, and utilize the GeLU as an activation function. We will refer to the model using two layers as $MODEL_{Large}$, and when applying the single layer as $MODEL_{Little}$.

2.1.4 Reconstruction decoder

The decoder consists of a fully connected network of two layers, where the dimensionality of the encoder’s output is gradually reduced to a univariate ECG. The first layer reduces the dimensionality from 256 to 128 and the second layer from 128 to 1. Between the two, a ReLU activation function is placed to introduce some non-linearity.

2.1.5 Dataset

We used the ALSEDAS dataset [14] for pre-training, containing 48,000 ECG samples filtered and noise-reduced. The large size and diversity of this dataset allow the model to learn robust representations. Downsampling to 250 Hz prevents artifacts and ensures the data is manageable for real-time processing.

2.2 Fine-tuning for Stress Detection

In the second stage, we fine-tune the pre-trained model specifically for stress detection. Reusing the pre-trained CNN and Transformer encoders capitalizes on the robust representations learned in the first stage. This transfer learning approach significantly reduces the amount of labeled data needed for training the stress classifier, enhancing efficiency and effectiveness [10].

2.2.1 Loss function

We selected Binary Cross-Entropy with Logits Loss for the stress detection task. This loss function is particularly suitable for binary classification problems, such as stress vs. no-stress, and handles imbalanced datasets effectively, ensuring accurate predictions across different stress levels.

2.2.2 Encoder fine-tuning

A critical aspect of the downstream task stage is the flexibility in choosing to either fine-tune the entire encoder, only the last layer, or leave it frozen. Fine-tuning the entire encoder was chosen, where all layers of the encoder are updated based on the new stress detection task. This approach allows the model to adjust all learned features to better suit the specific nuances of stress-related ECG patterns. Fine-tuning the full encoder can lead to higher accuracy because the model can refine its understanding of the data at all levels. However, this method is computationally intensive and requires more data to prevent overfitting. Fine-tuning the entire encoder is particularly beneficial when there is sufficient labeled data available to support the training process.

2.2.3 Stress Decoder Architecture

The stress decoder transforms the encoded data into a stress prediction. It consists of two linear layers with batch normalization and ReLU activation functions. The first layer reduces the dimensionality from 256 to 128, and the second layer further reduces it to 1, resulting in a single stress prediction for the entire sequence. The use of batch normalization helps stabilize training and prevents overfitting by normalizing the activations, while the ReLU activation introduces non-linearity to capture complex relationships in the data.

2.2.4 Dataset

We used the WESAD dataset[9], used for classifier training, provides comprehensive physiological and motion data with accurate stress labels. Segmenting the data into 8-second chunks simulates real-time sensor readings, aligning with our goal of real-time stress detection. Downsampling to 250 Hz ensures

computational efficiency without losing critical signal details.

3 Experiments

In our experiments, we aimed to evaluate the impact of various model configurations on the performance of stress detection on the WESAD dataset. We focused on two main aspects: the size of the transformer encoder and the length of the sampled time window (8 seconds vs. 4 seconds). This resulted in two distinct configurations. One called $MODEL_{Large}$ with a context window of 8 seconds and two transformer layers, and a second one $MODEL_{Small}$ with a 4-second window and a single transformer layer.

Hyperparameter	$MODEL_{little}$	$MODEL_{large}$
Sequence Length (sec)	4	8
Accumulated Gradient	8	8
Batch Size	32	32
Dropout Rate	0.6	0.4
Learning Rate	5.947×10^{-4}	1.585×10^{-5}
Epochs	100	100

Table 3: Optimal Hyperparameter Configurations for $MODEL_{Little}$ and $MODEL_{Large}$


We implemented these models using PyTorch Lightning and trained them on NVIDIA GeForce RTX 4090 GPUs.

In terms of validation, we employ a hold-out scheme, dividing each dataset into three subsets: 60% for training, 20% for validation, and 20% for testing

3.1 Pretraining phase

For the pre-training phase, we utilized the Adam optimizer with a fixed learning rate of 0.001 and gradient clipping prevents overfitting and stabilizes the learning process. Training consisted of 100 epochs with a batch size of 64, whether processing 8-second or 4-second ECG signals.

The presence of low losses(2), in both cases, serves as a strong indicator that the self-supervised task



(a) Losses of the model $MODEL_{little}$



(b) Losses of the model $MODEL_{large}$

Figure 2: Losses of the self-supervised task

has been carried out effectively, implying that the encoder has successfully captured a high-quality representation of an ECG signal.

3.2 Stress detection

Hyperparameter sweeps and Bayesian optimization were employed to fine-tune the model, focusing on learning rate and weight decay. This meticulous tuning ensures the model performs optimally, balancing accuracy and computational efficiency. Obtaining the parameters shown in Table 3.

To evaluate the model’s performance, we used both accuracy and F1 metrics. While accuracy is useful, especially for unbalanced datasets, the F1-score is a more appropriate measure. This is because the F1 score offers a balanced evaluation, taking into account both majority and minority class performance.

The results for stress prediction on the WESAD database, as shown in Table 4, highlight the exceptional performance of our models in terms of accuracy and F1-score. Specifically, our $MODEL_{large}$ achieves an accuracy of 0.996 and an F1-score of 0.992, while our $MODEL_{little}$ achieves an accuracy of 0.992 and an F1-score of 0.983.

Model	Method	Time window(s)	Acc.	F1
[15]	Deep ECGNet	10	0.908	0.857
[16]	QDA	40	0.857	-
[17]	Image	-	0.925	-
[9]	LDA	60	0.854	0.813
[4]	TF	30	0.911	0.833
Ours $MODEL_{little}$	SSL + TF	4	0.992	0.983
Ours $MODEL_{large}$	SSL + TF	8	0.996	0.992

Table 4: ECG stress-based detection on WESAD database. Explanation QDA = Quadratic Discriminant Analysis, LDA = Linear Discriminant Analysis, TF = Transformer, SSL = Self-supervised learning

An interesting finding is that reducing the model complexity and input sequence length does not significantly affect the performance. This suggests that a simpler model can be effectively used for stress detection without compromising accuracy. The reduced input sequence length allows for quicker stress detection in real-time applications, and the model’s parameter count is significantly decreased from 3,993,476 to 2,807,940.

Table 4 also compares our models with several other approaches applied to the WESAD database for stress versus non-stress detection. Our models consistently outperform the alternatives, achieving state-of-the-art results in both accuracy and F1-score. Notably, our method demonstrates superior performance even with a shorter time window, which is critical for real-time stress detection. The pre-training phase in our approach plays a crucial role in enhancing the model’s ability to detect stress effectively, setting it apart from other transformer-based methods.

4 Model Explainability

While achieving high accuracy and F1 scores is critical for the stress detection model, interpretability is equally important, particularly in the healthcare sector. Understanding the model’s reasoning helps clinicians trust its decisions and enables identifica-

tion and correction of any biases or errors. This section focuses on the techniques used to enhance the interpretability of our model, specifically saliency maps.

4.1 Saliency maps

Saliency maps are used to visualize the decision-making process of the stress detection model. They highlight which parts of the input data were most important for the model when making a prediction. This technique involves backpropagation from the output layer to the input layer, calculating the gradient of the output concerning the input. The magnitude of these gradients indicates the importance of each feature, which is then visually represented as a saliency map.[18] In our context, bright red bars on the map indicate strong evidence supporting stress, while blue bars indicate evidence against stress.

For instance, in a true negative case, the map might focus on a short T wave as evidence against stress, while in a true positive case, it highlights a faint T wave typical of stress. This detailed view helps us understand the temporal dependencies the model relies on.

4.2 Artificial ECGs

To further understand and validate our model, we created synthetic ECGs by altering specific waveform features, as shown in Figure 3. This process allows us to observe the model’s responses to controlled changes, ensuring they align with medical understanding. We used mathematical models to generate synthetic ECG waves[19], adjusting parameters to simulate different stress indicators [20][21][22].

- RR Interval: By shortening the RR interval, simulating an increased heart rate, the predicted probability of stress increased significantly. This aligns with the clinical expectation that a higher heart rate can indicate stress.
- PR Interval: Shortening the PR interval increased the stress probability, demonstrating the model’s sensitivity to changes in atri-




Figure 3: Artificial ECG waveform compared with a real signal



Figure 4: Artificial ECG waveform compared with real signal predictions

oventricular conduction time, a known stress marker.

- T Wave Amplitude: Doubling the T wave amplitude and increasing its duration led to a noticeable increase in the stress prediction probability, reflecting the model’s ability to recognize stress-related changes in ventricular repolarization.
- QT Interval: Shortening the QT interval, representing a faster depolarization-repolarization cycle, also increased the predicted stress probability, in line with medical knowledge.

These synthetic ECGs confirmed that our model’s predictions are consistent with established medical indicators of stress, thereby enhancing trust in its reliability. Some examples of the changes in the predictions can be seen in the Figures 4 5 6 and 7.




Figure 5: Effect of the RR interval on the model’s predictions

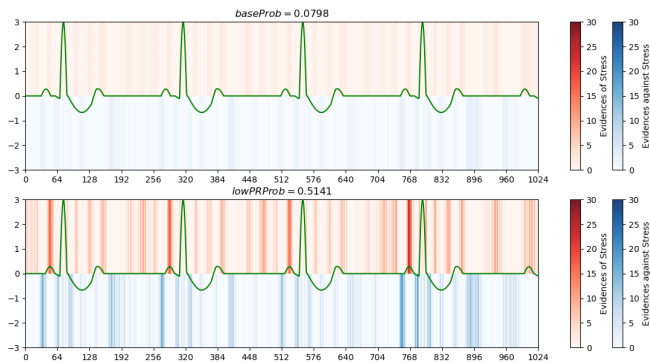


Figure 6: Effect of the PR interval on the model’s predictions




Figure 7: Effect of the T-wave amplitude on the model’s predictions

5 Conclusions and Future Work

5.1 Conclusions

The successful intersection of medical technology and artificial intelligence was a hallmark of this project. By drawing upon the latest advancements in Electrocardiogram (ECG) technology and integrating them with cutting-edge Deep Learning techniques, our endeavor has significantly elevated the potential of stress detection using raw ECG data over previous works. This synergy opens up promising avenues for the integration of these sensors for multiple applications, such as the automotive field.

In terms of methodology, our research underscored the power of a two-stage process, which placed substantial emphasis on self-supervised learning. This pivotal decision gifted our model with an enhanced capability to recognize and decode intricate patterns from ECG signals over traditional HRV analysis, thus amplifying its accuracy and reliability manifold.

Further, by harnessing state-of-the-art computational paradigms like Transformers and Self-Supervised Learning, we not only positioned the project at the forefront of innovation but also served as a testament to the transformative power of contemporary AI techniques when wielded in medical contexts. Such an adoption not only reflects our commitment to the best in technology but also casts a spotlight on the possibility of real-time and non-invasive stress detection.

Parallely, our deep dive into interpretability studies, particularly through tools like saliency maps, has demystified the AI decision-making process, reinforcing the imperative of transparency in AI-driven outcomes which is crucial in the medical field.

5.2 Future Work

Even though we are delighted with our results, we recognize that there is always room for improvement and expansion in such a dynamic field.

One immediate area of potential lies in the expansion of our dataset. The WESAD dataset used is

quite limited in size and might be a limiting factor for the model’s real-life adoption. By integrating a broader and more varied spectrum of ECG signals, we can potentially enable our model to achieve even greater levels of accuracy and generalization. This diversity in data can be instrumental in mirroring the true complexity and variability seen in real-world stress patterns.

A natural progression of this project would be its deployment in real-world scenarios. Given the potential applications in diverse fields like vehicular safety, continuous healthcare monitoring, and even workplace stress management, our model can be a game-changer. Experimenting with its deployment in these arenas can yield practical insights and pave the way for broader adoption.

Linking with the previous point, in a real deployment of the model, the recorded signals are likely to have some sensor distortion which needs to be calibrated for, varying for each sensor model. Our two-stage training model will ensure this retraining is minimal and more focused on the latter stages of the model, without a complete retraining, for more efficient and cost-effective adaptation to different deployment environments.

Another area that warrants further exploration is the architectural nuances of the model. One aspect that could significantly lighten the model for IoT applications would be reducing the CNN encoder’s parameters, as they largely contribute to the overall model’s size. There are several ways to do this, but it is important for the architecture to maintain the receptive field of approximately 60bpm.

One avenue would be reducing the sampling rate to 125Hz, which would almost halve the number of parameters used. And, as shown while explaining the methodology, theoretically, it should be possible to lower the sampling rate up to 125Hz without significantly reducing the delineation, but we would rather not push the model so far in this study.

As for the transformer, it undoubtedly leads to great results even when using just one layer. But some of its hyperparameters, such as the number of heads, and the d_{model} dimensionality (which comes from

the CNN), could also be studied.

Lastly, to facilitate training, it might be a good idea to explore other optimizer methods rather than ADAM, in particular, SAM [23] and its variants seem like a good fit because of the nature of the task and the training instabilities found, where performance was very dependent on hyperparameter initialization.

References

- [1] L. E. Søvold, J. A. Naslund, A. A. Kousoulis, S. Saxena, M. W. Qoronfleh, C. Grobler, and L. Münter. “Prioritizing the Mental Health and Well-Being of Healthcare Workers: An Urgent Global Public Health Priority”. In: *Frontiers in Public Health* 9 (2021). ISSN: 2296-2565. DOI: [10.3389/fpubh.2021.679397](https://doi.org/10.3389/fpubh.2021.679397). URL: <https://www.frontiersin.org/journals/public-health/articles/10.3389/fpubh.2021.679397>.
- [2] U. Ha, S. Madani, and F. Adib. “WiStress: Contactless stress monitoring using wireless signals”. In: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5.3 (2021), pp. 1–37.
- [3] J. Chen, D. Zhang, Z. Wu, F. Zhou, Q. Sun, and Y. Chen. “Contactless electrocardiogram monitoring with millimeter wave radar”. In: *IEEE Transactions on Mobile Computing* (2022).
- [4] B. Behinaein, A. Bhatti, D. Rodenburg, P. Hungler, and A. Etemad. “A Transformer Architecture for Stress Detection from ECG”. In: *ISWC ’21* (2021). DOI: [10.1145/3460421.3480427](https://doi.org/10.1145/3460421.3480427).
- [5] J. Vazquez-Rodriguez, G. Lefebvre, J. Cumin, and J. L. Crowley. “Transformer-based self-supervised learning for emotion recognition”. In: *2022 26th International Conference on Pattern Recognition (ICPR)*. IEEE, 2022, pp. 2605–2612. DOI: [10.1109/ICPR56361.2022.9956027](https://doi.org/10.1109/ICPR56361.2022.9956027).
- [6] G. Yan, S. Liang, Y. Zhang, and F. Liu. “Fusing transformer model with temporal features

- for ECG heartbeat classification”. In: *2019 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. IEEE. 2019, pp. 898–905.
- [7] P. Sarkar and A. Etemad. “Self-supervised ECG representation learning for emotion recognition”. In: *IEEE Transactions on Affective Computing* 13.3 (2020), pp. 1541–1554.
- [8] U. R. Acharya, S. L. Oh, Y. Hagiwara, J. H. Tan, M. Adam, A. Gertych, and R. San Tan. “A deep convolutional neural network model to classify heartbeats”. In: *Computers in biology and medicine* 89 (2017), pp. 389–396.
- [9] P. Schmidt, A. Reiss, R. Duerichen, C. Marberger, and K. Van Laerhoven. “Introducing WESAD, a Multimodal Dataset for Wearable Stress and Affect Detection”. In: *ICMI’18*, (2018). DOI: [10.1145/3242969.3242985](https://doi.org/10.1145/3242969.3242985).
- [10] S. Rabbani and N. Khan. “Contrastive self-supervised learning for stress detection from ecg data”. In: *Bioengineering* 9.8 (2022), p. 374.
- [11] G. Zerveas, S. Jayaraman, D. Patel, A. Bhamidipaty, and C. Eickhoff. “A Transformer-based Framework for Multivariate Time Series Representation Learning”. In: *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2021), pp. 2114–2124. DOI: [10.1145/3447548.3467401](https://doi.org/10.1145/3447548.3467401).
- [12] H. Choi and P. Kang. *Multi-Task Self-Supervised Time-Series Representation Learning*. 2023. arXiv: [2303.01034](https://arxiv.org/abs/2303.01034) [cs.LG]. URL: <https://arxiv.org/abs/2303.01034>.
- [13] J. Vazquez-Rodriguez, G. Lefebvre, J. Cumin, and J. L. Crowley. *Transformer-Based Self-Supervised Learning for Emotion Recognition*. 2022. arXiv: [2204.05103](https://arxiv.org/abs/2204.05103) [q-bio.NC]. URL: <https://arxiv.org/abs/2204.05103>.
- [14] J. Zheng, J. Zhang, S. Danioko, H. Yao, H. Guo, and C. Rakovski. “A 12-lead electrocardiogram database for arrhythmia research covering more than 10,000 patients”. In: *Scientific Data* (2020). DOI: [10.1038/s41597-020-0386-x](https://doi.org/10.1038/s41597-020-0386-x).
- [15] P. Prajod and E. André. “On the Generalizability of ECG-based Stress Detection Models”. In: *2022 21st IEEE International Conference on Machine Learning and Applications (ICMLA)*. IEEE. 2022, pp. 549–554.
- [16] P. Bota, C. Wang, A. Fred, and H. Silva. “Emotion assessment using feature fusion and decision fusion classification based on physiological data: Are we there yet?” In: *Sensors (Switzerland)* 20 (2020), pp. 1–17. DOI: [10.3390/s20174723](https://doi.org/10.3390/s20174723).
- [17] S. Ishaque, N. Khan, and S. Krishnan. “Detecting stress through 2D ECG images using pretrained models, transfer learning and model compression techniques”. In: *Machine Learning with Applications* 10 (2022), p. 100395. DOI: [10.1016/j.mlwa.2022.100395](https://doi.org/10.1016/j.mlwa.2022.100395).
- [18] P. S. Parvatharaju, R. Doddaiiah, T. Hartvigsen, and E. A. Rundensteiner. “Learning Saliency Maps to Explain Deep Time Series Classifiers”. In: *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. CIKM ’21. Virtual Event, Queensland, Australia: Association for Computing Machinery, 2021, pp. 1406–1415. ISBN: 9781450384469. DOI: [10.1145/3459637.3482446](https://doi.org/10.1145/3459637.3482446). URL: <https://doi.org/10.1145/3459637.3482446>.
- [19] L. Michaeli, J. Šaliga, I. Andras, and P. Dolinsky. “A new simple ECG signal model”. In: Sept. 2019.
- [20] J. Barry, A. P. Selwyn, E. G. Nabel, M. B. Rocco, K. Mead, S. Campbell, and G. Rebecca. “Frequency of ST-segment Depression Produced by Mental Stress in Stable Angina Pectoris from Coronary Artery Disease”. In: *The American Journal of Cardiology* 61.13 (1988), pp. 989–993. ISSN: 0002-9149. DOI: [10.1016/0002-9149\(88\)90112-9](https://doi.org/10.1016/0002-9149(88)90112-9).
- [21] A. Chaudhuri, M. Ray, K. Arpan, and Ghosh. “Effect of Perceived Stress on Electrocardiographic Parameters of Female Health Professionals: A Cross Sectional Observational Study”. In: (Oct. 2019).

- [22] A. Bhide, R. Durgaprasad, L. Kasala, V. Velam, and N. Hulikal. “Electrocardiographic Changes during Acute Mental Stress”. In: *Int J Med Sci Public Health* 5.5 (2016), p. 835. ISSN: 2320-4664. DOI: [10.5455/ijmsph.2016.19082015137](https://doi.org/10.5455/ijmsph.2016.19082015137). (Visited on 08/23/2023).
- [23] P. Foret, A. Kleiner, H. Mobahi, and B. Neyshabur. *Sharpness-Aware Minimization for Efficiently Improving Generalization*. 2021. arXiv: [2010.01412](https://arxiv.org/abs/2010.01412) [cs.LG]. URL: <https://arxiv.org/abs/2010.01412>.