

Feature-Level Cross-Attentional PPG and Motion Signal Fusion for Heart Rate Estimation

Original

Feature-Level Cross-Attentional PPG and Motion Signal Fusion for Heart Rate Estimation / Kasnesis, P.; Toumanidis, L.; Burrello, A.; Chatzigeorgiou, C.; Patrikakis, C. Z.. - (2023). (Intervento presentato al convegno IEEE COMPSAC) [10.48550/arXiv.2210.11415].

Availability:

This version is available at: 11583/2980726 since: 2023-08-07T07:14:23Z

Publisher:

IEEE

Published

DOI:10.48550/arXiv.2210.11415

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

ACM postprint/Author's Accepted Manuscript

(Article begins on next page)

MULTI-HEAD CROSS-ATTENTIONAL PPG AND MOTION SIGNAL FUSION FOR HEART RATE ESTIMATION

Panagiotis Kasnesis^{§†}, Lazaros Toumanidis^{§†}, Alessio Burrello^{*},
Christos Chatzigeorgiou^{§†}, Charalampos Z. Patrikakis^{§†},

[§]ThinGenious PC, Maroussi, Greece

[†]Department of Electrical and Electronics Engineering, University of West Attica, Greece

^{*}Department of Electrical, Electronic and Information Engineering, University of Bologna, Italy

ABSTRACT

Nowadays, Heart Rate (HR) monitoring is a key feature of almost all wrist-worn devices exploiting photoplethysmography (PPG) sensors. However, arm movements affect the performance of PPG-based HR tracking. This issue is usually addressed by fusing the PPG signal with data produced by inertial measurement units. Thus, deep learning algorithms have been proposed, but they are considered too complex to deploy on wearable devices and lack the explainability of results. In this work, we present a new deep learning model, PULSE, which exploits temporal convolutions and multi-head cross-attention to improve sensor fusion's effectiveness and achieve a step towards explainability. We evaluate the performance of PULSE on three publicly available datasets, reducing the mean absolute error by 7.56% on the most extensive available dataset, PPG-DaLiA. Finally, we demonstrate the explainability of PULSE and the benefits of applying attention modules to PPG and motion data.

Index Terms— Deep Learning, Sensor Fusion, Heart Rate Monitoring, Attention, Photoplethysmography

1. INTRODUCTION

In recent years, wrist-worn devices (i.e., smartwatches) enable a 24h-monitoring of the subject's vital conditions thanks to miniaturized sensors, becoming increasingly popular in personalized health care and medical IoT applications [1]. One of the most important indices to monitor is Heart Rate (HR). Compared to first-generation monitoring devices, which exploit a simple 1-3 leads ECG connected through a chest strip, modern ones use photoplethysmographic (PPG) sensors, allowing HR monitoring to be integrated into the smartwatches [2]. However, a limitation of PPG-based HR monitoring is given by the presence of motion artifacts (MA). These are caused by variations of sensor position on the wrist or ambient light leaking in the gap between the sensor and the wrist. In literature, this problem has been first tackled utilizing filtering approaches. They use the correlation between acceleration data and the PPG signal to cancel the noise and

remove the MAs. Then, the HR is extrapolated from the cleaned signal [3, 4]. The critical limitation of these approaches is the high number of free hyper-parameters, which often limits their generalization.

Deep learning approaches have been proposed to improve generalization, bringing promising results on different public datasets [5, 6, 7, 8]. On the other hand, these models lack explainability, since acceleration and PPG data are fused in a *black box*. Until now, little attention has been posed to recent Transformers, given the usual high number of parameters required to train them. These models are based on the so-called *Attention Modules*, which correlate different tensors.

In this paper, we demonstrate that combining feature maps of convolutions with attention modules leads to improved accuracy in the PPG-based HR monitoring and allows to interpret the connection between acceleration and PPG data. The main contributions of this work are summarized as follows:

- We introduce a new state-of-the-art, yet lightweight (around 130M parameters), deep neural network to fuse PPG and motion signals for precise heart rate estimation. The model includes both temporal convolutional and multi-head cross-attention modules.
- We evaluate the effectiveness of the produced model on three publicly available datasets. On the largest one, PPG-DaLiA [5], we improve the mean absolute error (MAE) to 4.03 beats per minute (BPM), outperforming the best state-of-the-art model (a pure CNN) by 0.33 BPM.
- We demonstrate the explainability of the developed model and the benefits of applying attention modules to PPG and motion data by showcasing examples of attentional maps.

2. BACKGROUND

Temporal Convolutional Networks: TCNs are 1D-Convolutional Neural Networks (CNNs), with the insertion of *dilation* in convolutional layers [9, 10]. The dilation is a fixed gap d inserted between input samples before being convolved with

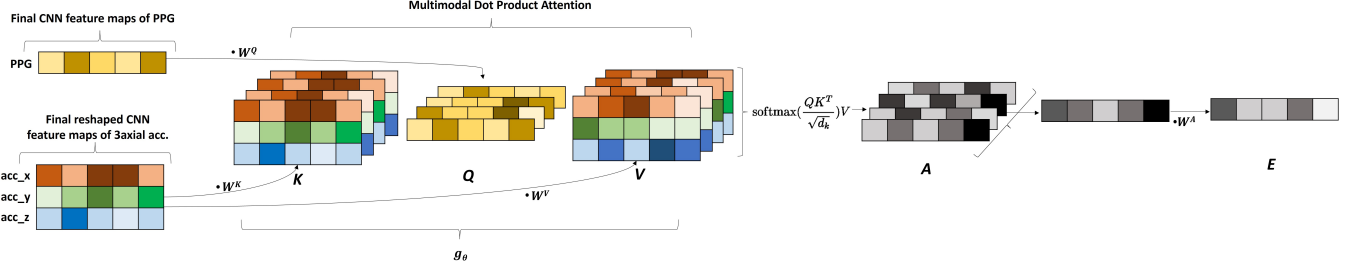


Fig. 1. Multi-head cross-attention module applied to PPG and 3axial accelerometer feature maps. The PPG embedding acts as a Query tensor, while the 3-axial accelerometer embeddings are Key and Value tensors.

the weights, thus increasing its temporal receptive field. A convolutional layer in a TCN is formulated as:

$$y_t^m = \sum_{i=0}^{K-1} \sum_{l=0}^{C_{in}-1} x_{t-di}^l \cdot W_i^{l,m} \quad (1)$$

where \mathbf{x} and \mathbf{y} are the input and output feature maps, t and m the output time-step and channel, respectively, \mathbf{W} the filter weights, C_{in} the number of input channels, d the dilation factor, and K the filter size.

Attention Module: An Attention Module is a neural network layer that correlates and emphasizes important parts of the input. The classical formulation is the multi-head self-attention (MHSA) [11]. MHSA takes a tensor of sequential data as input and correlates it with itself. First, it projects the sequence \mathbf{X} to three separate tensors, known as *queries* \mathbf{Q} , *keys* \mathbf{K} and *values* \mathbf{V} :

$$\mathbf{Q} = \mathbf{XW}_{\text{query}} \quad \mathbf{K} = \mathbf{XW}_{\text{key}} \quad \mathbf{V} = \mathbf{XW}_{\text{value}} \quad (2)$$

After, \mathbf{Q} , \mathbf{K} and \mathbf{V} are used to compute the *scaled dot-product attention* defined as

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) \doteq \mathbf{A} \doteq \text{SoftMax}_{\text{over keys}} \left(\frac{\mathbf{QK}^T}{\sqrt{d}} \right) \mathbf{V} \quad (3)$$

where \mathbf{A} is the *scaled dot-product attention*, and d the dimensionality of \mathbf{K} used as a scaling factor.

3. MATERIAL & METHODS

3.1. Network Architecture

We propose the use of Multi-Head Cross-Attentional (MHCA) module for sensor modality fusion (Fig. 1). MHCA in contrast to MHSA, which correlates the input tensor \mathbf{X} with itself, has as input two tensors \mathbf{X}_1 and \mathbf{X}_2 ; one of them represents the *key* \mathbf{K} and *value* \mathbf{V} tensors, while the other the *query* \mathbf{Q} tensor [12]. In our work, these tensors correspond to different modalities (i.e., PPG and 3axial accelerometer-based feature maps).

The proposed model architecture (selected after hyperparameter optimization) is called PULSE (Ppg and imU signal fuSion for heart rate Estimation) and is displayed in Fig.2. Three 1D convolutional blocks process the input tensor with

increasing channels (i.e., 32, 48 and 64). Each block comprises 3 consecutive dilated convolutions (similar to [7]). The selected dilation rate was equal to 2. Afterward, the produced feature maps are fed to the MHCA module (Fig 1). The PPG acts as a Query vector \mathbf{Q} , while the 3axial accelerometer embeddings as Key \mathbf{K} and Value \mathbf{V} vectors. Thus, for h heads, the cross-attention is computed, and, afterward, the h dot products (i.e., 4 for our case) are concatenated and transformed into \mathbf{E} using a dense layer again.

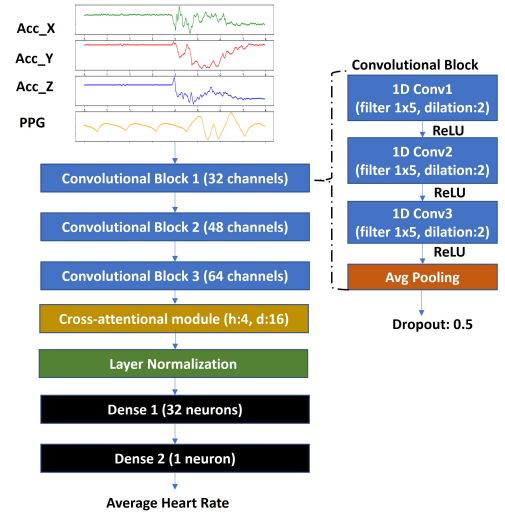


Fig. 2. PULSE network architecture

The proposed MHCA-based fusion module operates on human actions, i.e., high-level features extracted by the temporal convolutions applied to short-term variations of the 3axial acceleration and the arterial translucency [19]. Such an action maybe an arm forth movement (walking activity) or steering the wheel (driving activity). Thus, MHCA takes as input a set of action embeddings $e = \{e_1^1, e_2^1, \dots, e_1^2, e_2^2, \dots, e_n^m\}$ extracted from signal values included in a time window (i.e., 32 actions for each modality are produced in an 8-second window), where e_n^m is the n -th action embedding measured by the m -th sensor modality; MHCA is responsible for discovering the relationships between the action embeddings measured by the 3axial accelerometer conditioned by those measured by the PPG sensor, while for the case of MHSA they would be

Table 1. Per subject MAE performance of PULSE on the PPG-DaLiA dataset compared to state-of-the-art algorithms.

Model	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10	S11	S12	S13	S14	S15	Mean
Deep learning models																
[13]	7.73	6.74	4.03	5.90	18.51	12.88	3.91	10.87	8.79	4.03	9.22	9.35	4.29	4.37	4.17	7.65
[14]	5.46	5.01	3.74	6.48	12.68	10.52	3.31	8.07	7.91	3.29	7.05	6.76	3.84	4.85	3.57	6.02
[7]	4.29	3.62	2.44	5.73	10.33	5.26	2.00	7.09	8.60	3.09	4.99	6.25	1.92	3.02	3.55	4.81
[7]*	3.78	3.36	2.33	4.84	9.95	4.38	2.20	5.88	7.59	2.74	4.55	5.20	2.14	2.99	3.47	4.36
Ours	4.75	3.31	2.22	5.25	7.43	4.22	2.28	8.93	6.95	2.93	3.98	6.57	1.70	3.22	2.88	4.44
Ours*	3.78	3.04	2.20	4.41	6.95	3.71	2.39	8.17	6.19	2.60	3.85	5.22	1.98	3.13	2.79	4.03
Mathematical models																
[15]	33.1	27.8	18.5	28.8	12.6	8.7	20.65	21.8	22.3	12.6	21.1	22.8	27.7	12.1	16.4	20.5
[16]	8.86	9.67	6.40	14.1	24.06	11.34	6.31	11.25	16.04	6.17	15.15	12.03	8.50	7.76	8.29	11.06
[17]	4.50	4.50	3.20	6.00	5.00	3.40	2.80	6.30	8.00	2.90	5.10	4.70	3.10	5.00	4.10	4.57
[18]	5.40	4.30	3.00	8.00	2.20	2.80	3.30	8.50	12.60	3.60	3.60	6.10	3.00	5.50	3.70	5.04

also conditioned by themselves (increasing the complexity to quadratic [20]). After applying MHCA the **E** tensor is produced, which is normalized and passed through 2 dense layers, with the last one producing the estimated heart rate.

3.2. Datasets

PPG-DaLiA: The first dataset used for evaluation is PPG-DaLiA, a public dataset including physiological and motion data for PPG-based heart rate estimation. The dataset was collected using a wrist-worn placed on the subjects' non-dominant wrist to produce the one-channel PPG (sampling rate 64Hz) and the 3axial accelerometer (sampling rate 32Hz) signals while they wear a smart belt on their chest for the ground truth ECG labels. Fifteen subjects aged 21–55 years, are included in the dataset. It is worth noticing that the subjects performed a wide range of activities under close to real-life conditions, which are included as extra labels for human activity recognition [21]. The dataset includes 64,697 samples after segmentation (8-second window with 6-second overlap).

IEEE.Training & IEEE.Test: These datasets are recorded using a wrist-worn device to produce a two-channel PPG sensor and a three-axis accelerometer. The sampling rate of all signals is 125 Hz. In *IEEE.Training*, twelve subjects participated (aged 18–35 years), performing a running activity at varying speeds. In *IEEE.Test*, 8 different subjects (aged 19–58 years) participated while performing arm rehabilitation exercises or intensive arm movements (such as boxing). In addition, the datasets provide the average heart rate on 8-second sliding windows (window shift: 2 seconds) extracted from the raw ECG signal to be used as ground truth. The datasets include 1,768 samples (*IEEE.Training*) and 1,328 samples *IEEE.Test*.

4. RESULTS

We used Python 3.8 and the PyTorch framework to design and train the neural network. We validated all models following the cross-validation protocol proposed in [13, 7], denoted as Leave-One-Session-Out (LOSO) cross-validation, where the

15 subjects are divided in four data folds; 3 are used as training set, while the remaining one is subdivided to form the test set (1 subject), and the validation set. This approach leads to 15 training iterations ensuring the generalizability of the produced model to unseen subject data. Both the datasets are downsapled to 32 Hz, using input windows of 4×256 (PPG-DaLiA) and 5×250 (IEEE datasets). The windows are then normalized using per-channel z-score. We selected Adam [22] as network optimizer, having the following hyperparameters: learning rate equal to 0.0005 ($\beta_1:0.9$, $\beta_2:0.999$, $\epsilon_1:1e-08$). Moreover, the batch size is set to 256 and the number of epochs to 500, with a patience of 150. The validation model that had the lowest error is used test set. Finally, we followed the same post-processing method described in [7], where the output values are clipped in case the prediction is more or less 10% of the averaged 10 last estimated values.

4.1. Results on PPG-DaLiA

Table 1 presents the MAE results produced by PULSE compared with state-of-the-art (SOTA) deep learning-based and mathematical models, with * symbol denoting the use of post-processing. Our approach achieved an average MAE equal to 4.44 BPM, reduced to 4.03 BPM when using post-processing. For most subjects, PULSE surpassed the SOTA performance of both classical and deep learning models. Furthermore, the worse MAE obtained across patients is 8.17 BPM, which still allows for a good assessment of the health conditions.

We also performed an error analysis (MAE) w.r.t. the activity classes. As expected and shown in Fig. 3, the static activities (e.g., working in an office) are more robust than the motion activities (e.g., walking) which produce higher MAE. The lowest error is achieved when the subjects performed the sitting activity and the highest when they are ascending/descending stairs.

Table 2 discusses different options for input data and attention modality, switching between PPG only and PPG + acceleration as input, and varying between all the possible combinations in the attention module. As can be observed, the PULSE architecture is not only the most effective one, but also the most lightweight in terms of parameters. Using only




Fig. 3. MAE results per activity on PPG-DaLiA. the PPG as input and the MHSA module, lead to the same amount of parameters but a higher MAE of 4.52 BPM. Modifying the MHCA module and assigning the PPG to the K and V vectors, instead, increase both the MAE to 4.44 BPM and the number of parameters to 197.36k. We want also to underline that the self-attention-based network having as input PPG and accelerometer signals is the larger architecture and it achieves the best MAE on many subjects, but does not generalize across some others (e.g., subject 5), causing a worse overall average MAE. Our intuition is that using the MHCA module allows for a lower overfitting of the dataset, guiding the network to exploit correlations similar to the ones of SOTA classical algorithms.

Table 2. Ablation study on input sensors & attention types.

Modality	Attention	MAE	#KParams
PPG	MHSA	4.52	131.82
PPG, Acc	MHSA	5.11	230,13
PPG (Q) , Acc (K,V)	MHCA	4.03	131.82
PPG (K,V) , Acc (Q)	MHCA	4.44	197.36

4.2. Results on IEEE datasets

Table 3 presents the obtained results of the developed PULSE model in comparison with the DeepPPG [13] for the cases of the *IEEE Train* and *IEEE Test* datasets. In the former dataset, DeepPPG achieved lower MAE, but we should mention that this model is a huge ensemble of networks, consisting of 7 deep convolutional models, each one of them having 8.5M parameters, whilst our has around 130K parameters. So, a good trade-off between performance and complexity is demonstrated using PULSE. On the other hand, PULSE obtained the best performance when applied to *IEEE Test* dataset. It is worth noticing that apart from having two PPG channels as input instead of one, no modification and hyper-parameter tuning are applied to PULSE (tuned on PPG-DaLiA), which indicates its generalizability.

4.3. Model explainability

Apart from the low obtained MAE, one of the benefits of using an attention-based layer is that of visualizing the attentional maps to achieve explainability. Fig. 4 presents few

Table 3. MAE performance of PULSE on the IEEE datasets

Model	IEEE Train	IEEE Test
DeepPPG [13]	4.00	16.51
Ours	5.03	16.54
Ours*	4.42	13.53

examples and the corresponding pre-processed input signals. While there has been no defined explainability metric for wearable-based heart rate estimation, it is useful to check in which accelerometer axis the PPG signals pay more attention (shown in dark orange). The attentional map is divided, using white horizontal lines, into three 64x64 attentional maps, i.e., one for each accelerometer axis (x, y, and z) combined with the PPG. Thus, in Fig. 4a. we observe that for the displayed “lunch” example, the model pays more attention to the y-axis (even though it has the least oscillations), since y-axis is usually the more stressed while eating. Fig. 4b. displays that for the selected “walking” example, the most important values belong to the x and z-axis, which share a similar pattern, given that their values are associated to the oscillations of the superior arms. Finally, Fig. 4c (transition) shows that, as expected, the model’s output is not based on a particular axis (input signals have many oscillations).

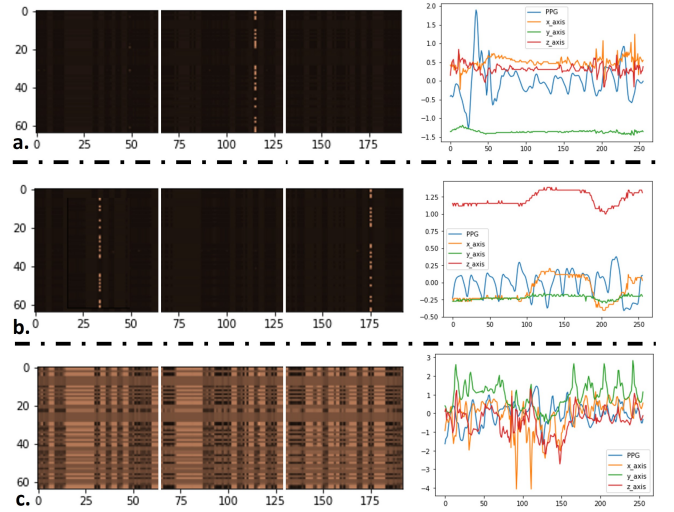


Fig. 4. Example of the produced attentional maps. The activities performed are: a. lunch, b. walking, and c. transition.

5. CONCLUSION

In this work, we propose a deep neural network based on a cross-attention module to fuse physiological and motion signals for heart rate detection. The method achieved 4.03 BPM as MAE on the PPG Dalia dataset, outperforming state-of-the-art algorithms, and generalize over multiple datasets, achieving state-of-the-art results on a benchmarked dataset. Further, we show how the attention maps can be exploited for explainability of the model.

6. REFERENCES

- [1] A. Yeole and D. R. Kalbande, "Use of internet of things (iot) in healthcare: A survey," in *Proc. of the ACM Symposium on Women in Research*, 2016, pp. 71–76.
- [2] Z. Zhang, Z. Pi, and B. Liu, "Troika: A general framework for heart rate monitoring using wrist-type photoplethysmographic signals during intensive physical exercise," *IEEE Transactions on biomedical engineering*, vol. 62, no. 2, pp. 522–531, 2014.
- [3] Z. Zhang, "Photoplethysmography-based heart rate monitoring in physical activities via joint sparse spectrum reconstruction," *IEEE transactions on biomedical engineering*, vol. 62, no. 8, pp. 1902–1910, 2015.
- [4] N. Huang and N. Selvaraj, "Robust ppg-based ambulatory heart rate tracking algorithm," in *IEEE EMBC*. IEEE, 2020, pp. 5929–5934.
- [5] A. Reiss, I. Indlekofer, P. Schmidt, and K. Van Laerhoven, "Deep ppg: Large-scale heart rate estimation with convolutional neural networks," *Sensors*, vol. 19, no. 14, pp. 3079, 2019.
- [6] D. Biswas, L. Everson, M. Liu, M. Panwar, B. Verhoef, S. Patki, C. H. Kim, A. Acharyya, C. V. Hoof, et al., "Cornet: Deep learning framework for ppg-based heart rate estimation and biometric identification in ambulant environment," *IEEE transactions on biomedical circuits and systems*, vol. 13, no. 2, pp. 282–291, 2019.
- [7] A. Burrello, D. Jahier Pagliari, M. Risso, S. Benatti, E. Macii, L. Benini, and M. Poncino, "Q-ppg: Energy-efficient ppg-based heart rate monitoring on wearable devices," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 15, pp. 1196–1209, 2021.
- [8] A. Burrello, D. J. Pagliari, P. M. Rapa, M. Semilia, M. Risso, T. Polonelli, M. Poncino, L. Benini, and S. Benatti, "Embedding temporal convolutional networks for energy-efficient ppg-based heart rate monitoring," *ACM Transactions on Computing for Healthcare*, vol. 3, no. 2, pp. 1–25, 2022.
- [9] S. Bai, J. Z. Kolter, and V. Koltun, "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling," *arXiv:1803.01271*, 2018.
- [10] C. Lea, R. Vidal, A. Reiter, and G. D. Hager, "Temporal convolutional networks: A unified approach to action segmentation," in *European Conference on Computer Vision*. Springer, 2016, pp. 47–54.
- [11] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *arXiv preprint arXiv:1706.03762*, 2017.
- [12] Y. Li, A. W. Yu, T. Meng, B. Caine, J. Ngiam, D. Peng, J. Shen, B.-X. Wu, Y. Lu, D. Zhou, Q. V. Le, A. L. Yuille, and M. Tan, "Deepfusion: Lidar-camera deep fusion for multi-modal 3d object detection," *ArXiv*, vol. abs/2203.08195, 2022.
- [13] A. Reiss, I. Indlekofer, P. Schmidt, and K. V. Laerhoven, "Deep ppg: Large-scale heart rate estimation with convolutional neural networks," *Sensors*, vol. 19, 2019.
- [14] S. b. Song, J. W. Nam, and J. H. Kim, "Nas-ppg: Ppg-based heart rate estimation using neural architecture search," *IEEE Sensors Journal*, vol. 21, pp. 14941–14949, 2021.
- [15] T. Schäck, M. Muma, and A. M. Zoubir, "Computationally efficient heart rate estimation during physical exercise using photoplethysmographic signals," in *EU-SIPCO*, pp. 2478–2481, 2017.
- [16] S. M. A. Salehizadeh, D. Dao, J. B. Bolkhovskiy, C. H. Cho, Y. Mendelson, and K. H. Chon, "A novel time-varying spectral filtering algorithm for reconstruction of motion artifact corrupted heart rate signals during intense physical activities using a wearable photoplethysmogram sensor," *Sensors*, vol. 16, 2016.
- [17] N. Huang and N. Selvaraj, "Robust ppg-based ambulatory heart rate tracking algorithm," in *EMBC*, pp. 5929–5934, 2020.
- [18] M. Zhou and N. Selvaraj, "Heart rate monitoring using sparse spectral curve tracing," in *EMBC*, pp. 5347–5352, 2020.
- [19] P. Kasnesis, C. Chatzigeorgiou, C. Z. Patrikakis, and M. Rangoussi, "Modality-wise relational reasoning for one-shot sensor-based activity recognition," *Pattern Recognit. Lett.*, vol. 146, pp. 90–99, 2021.
- [20] C.-F. Chen, Q. Fan, and R. Panda, "Crossvit: Cross-attention multi-scale vision transformer for image classification," *2021 IEEE/CVF ICCV*, pp. 347–356, 2021.
- [21] D. Y. Hwang, P. C. Ng, Y. Yu, Y. Wang, P. Spachos, D. Hatzinakos, and K. N. Plataniotis, "Hierarchical deep learning model with inertial and physiological sensors fusion for wearable-based human activity recognition," in *IEEE ICASSP*, pp. 21–25, 2022.
- [22] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2015.