Driving Towards Safety: Online PPG-based Drowsiness Detection with TCNs

Driving Towards Safety: Online PPG-based Drowsiness Detection with TCNs

Pierangelo Maria Rapa*†, Mattia Orlandi*, Andrea Amidei†, Alessio Burrello§, Roberto Rabbeni¶, Paolo Pavan†, Luca Benini*‡, Simone Benatti†*

* University of Bologna, Bologna, Italy — {pierangelomaria.rapa, mattia.orlandi, luca.benini}@unibo.it

† University of Modena and Reggio Emilia, Modena, Italy — {andrea.amidei, paolo.pavan, simone.benatti}@unimore.it

‡ ETH Zurich, Zurich, Switzerland — lbenini@iis.ee.ethz.ch

§ Polytechnic University of Turin, Turin, Italy — alessio.burrello@polito.it

¶ Maserati, Modena, Italy — roberto.rabbeni@maserati.com

***Abstract*—Increasing driver and driving safety is one of the most compelling needs of the automotive industry, both in terms of economic and social impact. Current approaches primarily focus on analyzing unsafe vehicle behavior, often overlooking the critical factor of the driver's physiological state. This paper introduces a novel solution leveraging temporal convolutional networks (TCNs) for unobtrusive driver drowsiness detection based on photoplethysmography (PPG). PPG data is collected seamlessly using sensors integrated into the steering wheel, providing a non-invasive assessment of the autonomic nervous system (ANS). We benchmarked our model on 16 subjects using a leave-one-subject-out (LOSO) cross-validation scheme, achieving an average accuracy of 77.03%. The model also shows good performance in avoiding false alarms when the driver is alert with a false positive ratio of just 8.21% and correctly detecting drowsiness with a low false negative ratio of 13.92%, improving the state-of-the-art for PPG based approaches. A quantized version of the model is deployed on a commercial ultra-low-power (ULP) system-on-a-chip (SoC), demonstrating real-world feasibility with an inference time of 4.8 ms and energy per inference of 117 µJ. This work represents a significant step towards unobtrusive, real-time physiological monitoring in driving environments, contributing to the ongoing efforts to improve driver's safety.**

I. Introduction and Related Works

Human-machine interaction (HMI) has recently garnered attention within the automotive sector, emerging as a compelling way to enhance the driving experience and bolster driver safety. Numerous research and commercial initiatives in academia and industry focus on addressing attention deficits and drowsiness, significant contributors to vehicular accidents. These solutions are increasingly available in high-end vehicles and encompass driving and driver monitoring devices, including lane assistants, braking aids, collision avoidance radar, and intermittent stop warnings. While these approaches analyze unsafe vehicle behavior, they often fall short of monitoring the primary causative factor: the driver's physiological state. Monitoring specific physiological parameters, such as fatigue and stress levels [1], offers crucial insights, enabling the implementation of proactive measures to avert hazardous situations rather than relying solely on reactive responses tied to vehicle behavior [1], [2].

Recent research underscores the efficacy of monitoring physiological signals, utilizing sensors integrated into the vehicle's cockpit or worn by the driver, for detecting fatigue or drowsiness states. This approach allows the car to inform the driver or initiate contingency actions promptly. Well-established methods in this domain involve biosignal analysis, encompassing electrocardiography (ECG) [3], photoplethysmography (PPG) [4], and electroencephalography (EEG) [5], alongside the examination of eye movements and blink rates via camera-based systems [6]. While camera-based techniques are integrated into selected high-end vehicles[1], challenges persist in terms of light condition variability and line-of-sight hindrances (e.g., glasses or sunglasses) [7].

From a clinical standpoint, direct measurement of sleep and wake states primarily relies on EEG due to its ability to detect autonomic nervous system (ANS) activity. However, direct skin and scalp contact is required, making EEG unsuitable for unobtrusive integration in a normal driving scenario.

On the other hand, both ECG and PPG can be used for non-invasive assessment of ANS [8], and they have been analyzed in several studies to detect drivers' drowsiness. For instance, Babaeian et al. [3] and Fujiwara et al. [9] analyzed drowsiness using ECG electrodes attached to drivers, with accuracy reaching 80% and 92%, respectively. Despite their high accuracy in drowsiness detection, these studies are based on obtrusive acquisition systems, such as medical devices, which are impractical in a normal driving scenario, or smartwatches, which require a wireless connection and a data link into the car, which is not always possible or desirable for the driver.

Furthermore, many of these studies do not consider inter-subject variability, showing results that are not based on a solid cross-validation scheme and can vary enormously. For example, the work of Kundiger et al. [10] demonstrates a high accuracy of 92% when tuning parameters on a subject-dependent basis (i.e., intra-subject). However, the accuracy drops significantly to 73% when the model is applied inter-subject.

It is also important to note that these methods depend heavily on manually extracting features from PPG and ECG data, primarily those related to heart rate (HR) and heart rate variability (HRV). This leads to many manually adjusted parameters and considerable computational complexity.

This work targets PPG sensors seamlessly integrated into the steering wheel. These sensors offer a significant advantage over traditional methods such as ECG, as they do not require the driver to always have both hands on the steering wheel, thereby offering greater flexibility and comfort. To process the PPG data, we rely on a temporal convolutional network (TCN). The TCN model enables learning long-range features spanning significant intervals in time series, as required for a potentially slow and gradual process like drowsiness.

In our experiments, the drowsiness ground truth is obtained using the Karolinska sleepiness scale (KSS), a widely accepted measure [12]. Using leave-one-subject-out (LOSO) cross-validation

[1] https://www.lexus.com/models/GX/safety/pre-collision-system

Author	Bio-Signals	Sensor Type	Obtrusive	Testers	Drowsiness Ground-Truth	Cross-Val	Method	Accuracy
Lee, H. [11]	ECG and PPG	Wearable	Yes	6	External expert	Yes	CNN	ECG: 70% PPG: 64%
Kundinger, T. [10]	ECG and PPG	Elelctrodes and Smartwatch	Yes	15	External expert	Yes	KNN	ECG: 78% PPG: 73%
Fujiwara, K. [9]	ECG and EEG	Electrodes	Yes	17	EEG	No	MSPC	92%
Babaeian, M. [3]	ECG	Electrodes	Yes	24	External expert	No	KNN	80%
Our work	PPG	Probes on steering wheel	No	16	KSS	Yes	TCN	77%

TABLE I
SOA COMPARISON TABLE.

scheme, we also ensured the robustness and generalization of our model, which is tested in the inter-subject scenario.

The main contributions of this work are the following:

- We extend a published dataset [13] to enable more robust training and validation of our new data-driven approach.
- We present a PPG-based solution to detect driver's drowsiness using a TCN, reaching an average cross-validated accuracy on 16 subjects of 77.03%, which, to the best of our knowledge, is higher than the other PPG-based SoA solutions [11], [10].
- We deploy a quantized version of our model on a commercial ultra-low power (ULP) system-on-a-chip (SoC) to ensure that the full processing chain, alongside the unobtrusive acquisition, runs in real-time in a realistic driving environment. We evaluate the deployment in terms of inference speed and energy efficiency, achieving an inference time of just $4.8\,\mathrm{ms}$ and energy per inference of $117\,\mu\mathrm{J}$.

II. Material and Methods

A. Temporal convolutional networks

TCNs are a class of neural networks designed specifically to process time series [14], [15]. They are built upon 1D convolutions, which are applied along the temporal dimension of the input data, enabling the network to process sequences with variable length. In particular, TCNs have been proven to be more effective than recurrent neural networks (RNNs) for sequence modeling [15]. One of the key features of TCNs is *dilation*: by introducing a fixed step size d (i.e., the *dilation factor*) between filter inputs in the convolutional layers, the network increases its receptive field—and thus can capture long-term dependencies in the input sequence—with a reduced number of parameters. Another important aspect is *causality*, which ensures that each output $y(t_N)$ is based only on current data and on the previous history of the signal $x(t_I)$, with $t_I \leq t_N$.

In our work, we employ a modified version of TEMPONet, a TCN model designed for gesture classification and hand kinematics regression from surface electromyography (sEMG) signals [16], [17]. The feature extraction part consists of three main blocks, each composed of:

- two convolutional layers (each followed by batch normalization and ReLU non-linearity) with filter size 3×1, full padding and variable dilation;
- one convolutional layer with filter size 5×1, variable padding and stride $s = 4$, followed by an average pooling with kernel 2×1, batch normalization and ReLU non-linearity.

As in [16] and [17], the three blocks are characterized by dilation factor $d = 2, 4, 8$, respectively. The convolutional layers in the three blocks produce 8, 16, and 32 output channels, respectively. In [16] and [17], the stride varies between the three blocks ($s = 1, 2, 4$); conversely, we apply the same stride $s = 4$ in all the blocks: this choice ensures that the receptive field is large enough to cover the whole $2\,\mathrm{min}$-long (2400 samples at $20\,\mathrm{Hz}$) window in input. Lastly, the classifier part consists of two fully-connected (FC) layers: the first, comprising 64 neurons, is followed by batch normalization, ReLU non-linearity, and a dropout layer with probability $p = 0.5$ for regularization purposes; the second, comprising 1 neuron, is followed by a Sigmoid non-linearity and outputs the probability of the subject being drowsy in the time window considered.

In addition to the PPG signal, we include as a second input the driving time $\mathrm{DT}(t)$, which provides the network with a simple but important information about the driver's state, taking into account fatigue-induced drowsiness. We modeled $\mathrm{DT}(t)$ as an exponential function:

$$\mathrm{DT}(t) = 1 - e^{-t/\tau} \tag{1}$$

where the time constant τ is heuristically set to $2\,\mathrm{h}$. Compared to a simple linear time feature, the usage of Eq. 1 delivers a double advantage: it guarantees that the network inputs are in a well-defined range, as $\mathrm{DT}(t) \in [0, 1[\ \forall t \geq 0$; it weighs more the latest stages of the driving—in particular, $\tau = 2\,\mathrm{h}$ ensures that Eq. 1 saturates at $\sim 9\,\mathrm{h}$, which is the limit recommended by EU regulations[2]. The model's architecture is summarized in Fig. 1.

B. Dataset

We evaluate our model on a PPG-based dataset for driver drowsiness detection presented in [13]. Initially comprising 14 subjects, the dataset has been extended with 7 additional subjects. All the recordings take place in the Maserati Innovation Lab in Modena (IT) using the static driving simulator DIL (Driver-In-the-Loop). The simulator room comprises a fully functional vehicle cockpit, a projector, and an audio system to ensure an immersive driving experience. The simulated scenario involves a two-lane road with light traffic to keep the driver engaged. The dataset features two PPG channels acquired at $1\,\mathrm{ksps}$ from PPG sensors integrated into the steering wheel cover and the driver's perceived level of drowsiness through the KSS questionnaire [18]. The KSS is a 9-point scale that matches verbal sentences to the psychophysical status experienced, where 1 is the lowest and 9 is the highest level. Drivers were to report their drowsiness level every $5\,\mathrm{min}$ using a tablet positioned in the simulator's cockpit.

The acquisition protocol was designed to induce drowsiness thanks to the night-time recording sessions and the simulation room maintained completely dark and sound-isolated. The drivers were introduced to the KSS and instructed how to report the score using the tablet, they were also asked to keep their hands on the steering wheel in a normal *"10 and 2 o'clock"* position to

[2]https://transport.ec.europa.eu/transport-modes/road/social-provisions/driving-time-and-rest-periods_en

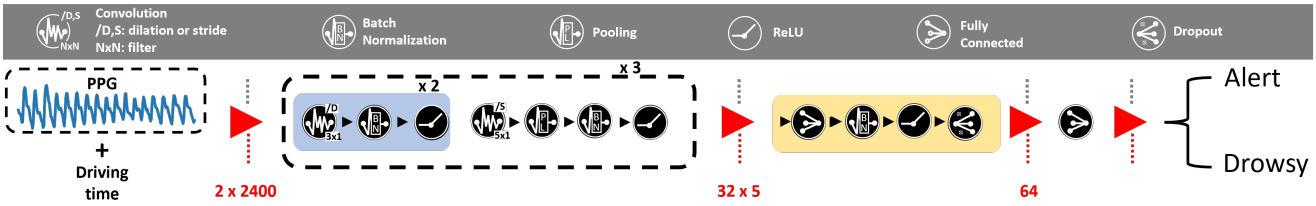

Fig. 1. TEMPONet [16], [17] architecture adapted for drowsiness detection.

ensure proper sensors coverage. The recording time was not fixed, but the driving sessions were monitored from the simulator's control room, and they were stopped when the driver fell asleep, or the driving style became very dangerous, i.e., collisions with other vehicles or objects. The dataset includes more than 22 h of recording with ~ 83 min average duration per driving session.

C. Data preparation

Starting from the original dataset, we focus the analysis only on the PPG channel positioned on the left part of the steering wheel; since the left hand was not used to prompt KSS scores, the recorded signal shows a better quality. Moreover, we apply a zero-phase band pass filter between 0.5 Hz and 10 Hz, and we downsample the PPG signal to 20 Hz. Then, the signal is divided into overlapped windows of 2 min with a step size of 30 s.

We employed KSS scores as the ground truth labels for our classification model. The KSS was binarized into two classes, denoted as *Alert* and *Drowsy*. The *Alert* class includes scores ranging from 1 to 6, while the *Drowsy* class ranges from 7 to 9 on the KSS. Subsequently, we upsampled the binarized KSS scores to match the frequency of PPG measurements. Ultimately, we designated the last KSS score as the class label within each window.

To ensure homogeneity in the dataset, we evaluated the class distribution for each subject and excluded five subjects from the analysis. Notably, three subjects exclusively exhibited the *Alert* class, while the remaining two displayed a significant imbalance, favoring the *Drowsy* class. Consequently, all experimental procedures were conducted on a subset of 16 subjects.

D. GAP9 SoC

To test the feasibility of our model in a real driving application, we deployed our model on GAP9[3], a commercial ULP SoC, featuring a 9-core RISC-V compute cluster, an AI accelerator and a single-core RISC-V controller, showing state-of-the-art (SoA) performance in terms of energy efficiency[4]. GAP9 doesn't have a data cache but adopts a hierarchical memory model comprising a fast-access L1 memory (128 kB) shared by all cores belonging to the cluster; a retentive L2 RAM (1.6 MB), accessible by all the cores and direct memory access (DMA) controllers; a non-volatile L3 on-chip flash memory (2 MB). The peripheral subsystem has a dedicated DMA component (μDMA), allowing the autonomous control of data transfers between peripherals and L2 memory without keeping the control core busy. We used *GAPflow*, a set of tools that allow to port a pre-trained neural network algorithm on GAP9. *GAPflow* is composed of two main tools: *NNTool*, a Python package that takes a high-level network description and translates all the layers and parameters to a model description that can be fed to the second tool, the *AutoTiler*, that optimizes data memory movements running on the GAP platform and generates code optimized for a multi-core cluster.

[3]https://greenwaves-technologies.com/gap9_processor/
[4]https://mlcommons.org/en/inference-tiny-10/

III. Experimental Results

A. Experimental setup

We evaluate our model utilizing a LOSO cross-validation scheme with 16 folds. Each fold contains 13 subjects for training, 2 subjects for validation, and 1 subject for the test.

Using *NNtool*, we run two optimization steps on the original network: (i) we reduce the number of operations and the memory overhead by fusing different basic operations (e.g., convolution followed by pooling and ReLU can be fused in a single "ConvPoolRelu" operation), (ii) we quantize the network to reduce memory usage and enable integer-only arithmetic using a post-training 8-bit quantization. The intermediate representation of the graph generated after these two steps is fed to the *AutoTiler*, which generates the actual C code using optimized software library primitives. This step is crucial, given that there is no data cache in GAP9. In order to be as efficient as possible, data must be stored in the L1 memory with a limited capacity (128 kB). The AutoTiler finds the optimal way to split the data and move them into the L1 memory when needed.

B. Experimental evaluation

In Tab. II, we show the accuracy and the F1 score for each subject before and after the *INT8* quantization. Overall, we reached an average accuracy and F1 score of 77.30% and 77.03%, respectively, with a negligible drop when quantization is used.

We extend our analysis by arranging the original reported KSS scores into three groups—*alert* (1–4), *hypovigilant* (5–6) and *drowsy* (7–9)—and evaluating the binary predictions of our model for each of them: in particular, for the *alert* and *hypovigilant* states we expect the model to predict 0 (i.e., "not sleepy"), while for the *drowsy state* we expect it to predict 1 (i.e., "sleepy"). We compute the accuracy of the model's predictions against the three groups, and we obtain a 91.42% of accuracy for *alert*, a 68.63% of accuracy for *hypovigilant* and a 83.48% of accuracy for *drowsy*. Note that the reported accuracies correspond to the true negative ratio (TNR) for the *alert* and *hypolvigilant* groups and to the true positive ratio (TPR) for the *drowsy* group. Furthermore, we obtain a false positive ratio (FPR) of 8.21% for *alert*, a FPR of 32.43% for *hypolvigilant* and a false negative ratio (FNR) of 13.92% for *drowsy*.

There is a clear indication that our model has a drop in performance in the *hypovigilant* state. Fig. 2 shows, in addition to the predictions of the network and the binary target classes, the three regions corresponding to the three groups: the top plot illustrates a successful prediction by our model, while the bottom plot depicts a prediction with low accuracy, marked by a notable misclassification in the *hypovigilant* state. The top plot of Fig. 3

		S01	S02	S03	S05	S06	S07	S08	S09	S10	S11	S13	S15	S16	S17	S18	S19	Avg ± Std
F1 Score	FP32	82.05%	68.10%	99.07%	83.21%	62.50%	87.50%	88.65%	70.97%	71.60%	58.33%	93.83%	63.58%	99.46%	95.88%	52.56%	59.46%	77.30% ± 15.26%
	INT8	81.58%	67.86%	99.38%	83.21%	63.56%	87.50%	88.65%	73.42%	81.82%	57.31%	91.57%	59.52%	97.85%	93.94%	55.21%	59.21%	77.60% ± 14.83%
Accuracy	FP32	86.14%	63.22%	98.61%	80.34%	59.46%	86.00%	85.00%	64.57%	76.77%	60.23%	93.01%	63.79%	99.17%	94.94%	55.15%	66.10%	77.03% ± 14.75%
	INT8	86.14%	62.81%	99.07%	80.34%	61.26%	86.00%	85.00%	66.93%	83.84%	58.52%	90.21%	60.92%	96.69%	92.41%	55.76%	64.97%	76.93% ± 14.40%

TABLE II
RESULTS OF THE LOSO CROSS-VALIDATION FOR BOTH ORIGINAL AND QUANTIZED MODELS.

further demonstrates this phenomenon, revealing heightened variability and lower median accuracy in the *hypovigilant* state.

Going deeper into the analysis, we further divide the *hypovigilant* and *drowsy* groups into their original KSS scores (i.e., 5–9), and we evaluate the model's predictions: for scores 5 and 6, we expect the model to predict a 0, whereas for scores between 7 and 9, a 1. As shown in the bottom plot in Fig. 3, we observe a significant variability in the accuracy on score 6, indicating challenges in accurately classifying scores on the border between *hypovigilant* and *drowsy* groups. Conversely, the accuracy on KSS scores far from that border is substantially higher: for instance, the accuracy on score 7 (near the border) is 76.47% while the accuracy on score 9 (far from the border) is 98.33%. We attribute *hypovigilant*-related misclassifications to the inherent self-reported nature of the KSS and to the ambiguity of the verbal sentences that describe the driver's status.

Notably, we obtain a low FPR of 8.21% in the *alert* group, which indicates that our model is able to effectively reduce the number of false alarms when the driver is certainly awake but also a low FNR of 13.92% in the *drowsy* group, which is the most dangerous state and, consequently, the state in which it is crucial that our model performs correctly. Even though this value leads SoA for PPG, the approach can be integrated with others, e.g., based on driving events, to further reduce the FNR.

C. Deployment on GAP9

We evaluate the performance of the model in terms of energy consumption, time per inference, and memory occupation. The profiling was conducted setting GAP9 cluster frequency to 240 MHz and the $V_{\text{dd core}}$ to 0.65 V. We use a power profiler from Nordic Semiconductor to measure the current draw and a GPIO to synchronize the measurement with the code execution. The final deployed model has 24.07k parameters and the final memory occupancy is 92.67 kB. Given the small memory footprint, we promote most of the tensors and weights of the network into the fast-access L1 memory to minimize memory movements.

The network requires 1.51 MOPs to be executed, resulting in a time per inference of 4.83 ms, totally compatible with the *online* constraint of a new prediction every 30 s, and an energy consumption of only $\sim 117\,\mu$J.

IV. CONCLUSIONS

This paper introduces a novel solution for unobtrusive driver drowsiness detection. Leveraging TCNs and photoplethysmography PPG data seamlessly collected by sensors integrated into the steering wheel, the proposed model achieves a SoA average cross-validated accuracy of 77.60% across 16 subjects. Our model is able to effectively reduce the number of false alarms when the driver is clearly awake, as evidenced by our low FPR of 8.21% in the *alert* group. We also obtain a low FNR of 13.92% in the *drowsy* group, which is the most hazardous state and, therefore, the state in which it is critical that our model performs correctly. The deployment of a quantized version of the model

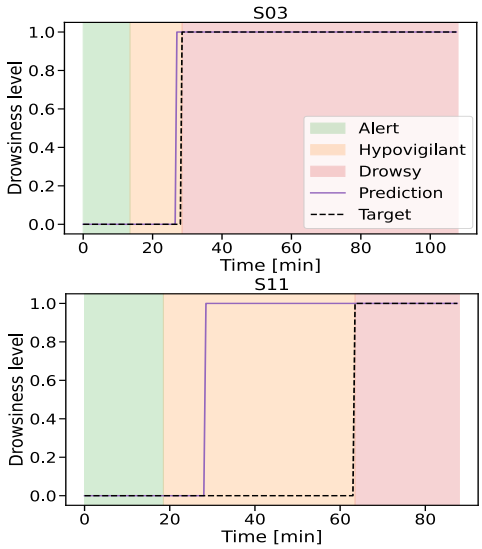

Fig. 2. Comparison of model's predictions output (in purple) and the ground truth target (in black) for two subjects, with the colored regions indicating the three driver's states: above, a subject that shows high accuracy; below, a subject with low accuracy and misclassified predictions in *hypovigilant* state.

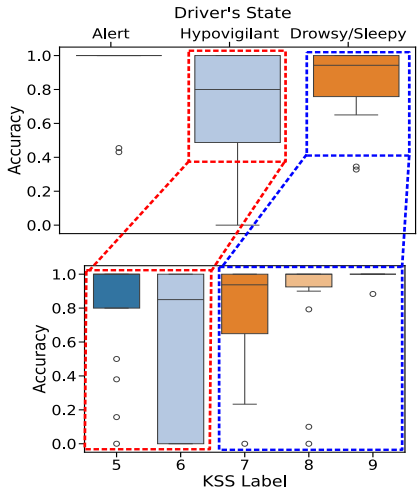

Fig. 3. Above, the accuracy boxplot on the three driver's states; below, the close-up view of the accuracy boxplot on the original KSS scores for *hypovigilant* and *drowsy* states.

on a commercial ULP SoC demonstrates real-world feasibility, with an inference time of 4.8 ms and energy consumption of 117 μJ. Future work will concentrate on integrating the optimized TCN model into online drowsiness detection systems, creating a comprehensive end-to-end application. Moreover, leveraging the computational capabilities of the GAP9 processor opens avenues for system scalability, potentially exploring the integration of the second PPG channel and other types of sensors.

ACKNOWLEDGMENT

The work was supported by EU project TRISTAN (g.a 101095947). A special thanks to Eng. Di Mare Giancarlo and Eng. Gianfranco Bordonaro for the Maserati DIL simulator setup.

REFERENCES

[1] M. Q. Khan *et al.*, "A comprehensive survey of driving monitoring and assistance systems," *Sensors*, vol. 19, no. 11, 2019.

[2] M. Doudou *et al.*, "Driver drowsiness measurement technologies: current research, market solutions, and challenges," *International Journal of Intelligent Transportation Systems Research*, pp. 1–23, 2019.

[3] M. Babaeian *et al.*, "Driver drowsiness detection algorithms using electrocardiogram data analysis," in *2019 IEEE 9th Annual Computing and Communication Workshop and Conference (CCWC)*. Las Vegas, NV, USA: IEEE, Jan. 2019, pp. 0001–0006.

[4] A. Amidei *et al.*, "Validating photoplethysmography (PPG) data for driver drowsiness detection," in *2021 IEEE International Workshop on Metrology for Automotive (MetroAutomotive)*, Jul. 2021, pp. 147–151.

[5] J. Cui *et al.*, "EEG-based cross-subject driver drowsiness recognition with an interpretable convolutional neural network," *IEEE Transactions on Neural Networks and Learning Systems*, 2022.

[6] M. H. Baccour *et al.*, "Camera-based eye blink detection algorithm for assessing driver drowsiness," in *2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2019, pp. 987–993.

[7] A. Dasgupta *et al.*, "A vision-based system for monitoring the loss of attention in automotive drivers," *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 4, pp. 1825–1838, Dec. 2013.

[8] B. Liu *et al.*, "The assessment of autonomic nervous system activity based on photoplethysmography in healthy young men," *Frontiers in Physiology*, vol. 12, p. 733264, 2021.

[9] K. Fujiwara *et al.*, "Heart rate variability-based driver drowsiness detection and its validation with EEG," *IEEE Transactions on Biomedical Engineering*, vol. 66, no. 6, pp. 1769–1778, Jun. 2019, conference Name: IEEE Transactions on Biomedical Engineering.

[10] T. Kundinger *et al.*, "Assessment of the potential of wrist-worn wearable sensors for driver drowsiness detection," *Sensors*, vol. 20, no. 4, p. 1029, Jan. 2020, number: 4 Publisher: Multidisciplinary Digital Publishing Institute.

[11] H. Lee *et al.*, "Using wearable ECG/PPG sensors for driver drowsiness detection based on distinguishable pattern of recurrence plots," *Electronics*, vol. 8, no. 2, p. 192, Feb. 2019, number: 2 Publisher: Multidisciplinary Digital Publishing Institute.

[12] S. Soares *et al.*, "Analyzing driver drowsiness: from causes to effects," *Sustainability*, vol. 12, no. 5, p. 1971, 2020.

[13] A. Amidei *et al.*, "ANGELS-smart steering wheel for driver safety," in *2023 9th International Workshop on Advances in Sensors and Interfaces (IWASI)*. IEEE, 2023, pp. 15–20.

[14] C. Lea *et al.*, "Temporal convolutional networks for action segmentation and detection," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 1003–1012.

[15] S. Bai *et al.*, "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling," 2018.

[16] M. Zanghieri *et al.*, "Robust real-time embedded EMG recognition framework using temporal convolutional networks on a multicore IoT processor," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 14, no. 2, pp. 244–256, 2020.

[17] M. Zanghieri *et al.*, "sEMG-based regression of hand kinematics with temporal convolutional networks on a low-power edge microcontroller," in *2021 IEEE International Conference on Omni-Layer Intelligent Systems (COINS)*, 2021, pp. 1–6.

[18] T. Åkerstedt *et al.*, "Subjective and objective sleepiness in the active individual," *International journal of neuroscience*, vol. 52, no. 1-2, pp. 29–37, 1990.