

Hybrid Handcrafted and Learnable Audio Representation for Analysis of Speech Under Cognitive and Physical Load

Gasser Elbanna^{1,2*}, Alice Biryukov^{1,2}, Neil Scheidwasser-Clow¹, Lara Orlandic¹, Pablo Mainar², Mikolaj Kegler³, Pierre Beckmann⁴, Milos Cernak²

¹École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

²Logitech Europe S.A., Lausanne, Switzerland

³Imperial College London, London, United Kingdom

⁴Université de Lausanne, Lausanne, Switzerland

gasser.elbanna@epfl.ch, milos.cernak@ieee.org

Abstract

As a neurophysiological response to threat or adverse conditions, stress can affect cognition, emotion and behaviour with potentially detrimental effects on health in the case of sustained exposure. Since the affective content of speech is inherently modulated by an individual's physical and mental state, a substantial body of research has been devoted to the study of paralinguistic correlates of stress-inducing task load. Historically, voice stress analysis has been conducted using conventional digital signal processing (DSP) techniques. Despite the development of modern methods based on deep neural networks (DNNs), accurately detecting stress in speech remains difficult due to the wide variety of stressors and considerable variability in individual stress perception. To that end, we introduce a set of five datasets for task load detection in speech. The voice recordings were collected as either cognitive or physical stress was induced in the cohort of volunteers, with a cumulative number of more than a hundred speakers. We used the datasets to design and evaluate a novel self-supervised audio representation that leverages the effectiveness of handcrafted features (DSP-based) and the complexity of data-driven DNN representations. Notably, the proposed approach outperformed both extensive handcrafted feature sets and novel DNN-based audio representation learning approaches.

Index Terms: computational paralinguistics, voice stress analysis, audio representation learning, DSP features, deep learning

1. Introduction

The emergence of machine learning methods for the paralinguistic analysis of speech signals has been of considerable influence for healthcare-related applications. This includes both the detection of long-term traits such as speech pathology [1] and shorter-term mental states such as emotion, effort, and acute stress [2–4]. For the latter, inter-individual differences in stress perception constitutes, among others, one of the major limitations of voice stress analysis (VSA) studies [5].

Among the main sources of acute stress [4], effort-induced stress can be elicited by increased physical or cognitive demands to complete the task (i.e., physical and cognitive load, respectively). Both behavioural states have been the subject of numerous paralinguistic studies [6–10]. On the one hand, physical load is known to create notable differences in speech production, especially due to competition between breathing processes for speech generation and increased oxygen consumption

*GE performed this work as an intern at Logitech.

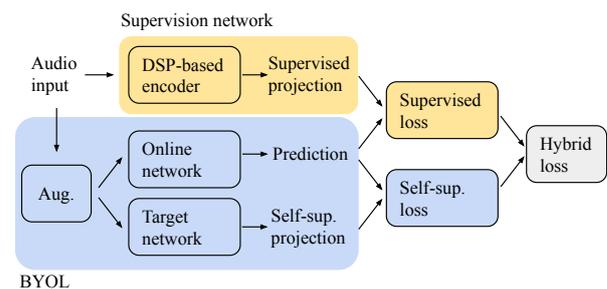


Figure 1: Hybrid audio representation learning: combining the self-supervised BYOL framework (blue) with supervision from DSP features (yellow).

to meet metabolic needs [6]. Regarding voice-related low-level descriptors (LLDs), several studies have reported increases in mean fundamental frequency after intensive physical exercise [7–9]. Similar effects on speech production have been observed in scenarios eliciting cognitive load [7, 10]. As a matter of fact, paralinguistic studies of stress have historically relied on handcrafted acoustic feature sets based on digital signal processing (DSP). Such feature sets could then be used for classification of cognitive or physical load levels using various classical machine learning approaches, including GMM supervector systems [11], *i*-vector frameworks [12], and ensemble learning [13].

Alternatively, several deep learning frameworks have been proposed for cognitive and physical load detection [14, 15]. Yet, training such models directly on task load corpora remains a challenging issue due to their small sizes, as larger networks tend to overfit due to their higher model capacity. To overcome this issue, several studies have shown that pre-training models on large corpora of speech or audio can produce state-of-the-art performance in various data-constrained speech processing and paralinguistic tasks such as speech emotion recognition [16], speaker or language identification [17–19].

With that in mind, in this paper, we focus on studying the detection of cognitive and physical load from voice recording using data-driven audio representation models. We introduce four novel datasets related to cognitive load with varying task complexities, languages, and protocols, and one dataset related to physical load. We used the datasets to evaluate a novel hybrid training protocol for general-purpose speech representation (Figure 1). The proposed approach aims at taking advantage of both self-supervised representation learning and non-trainable DSP-based feature sets. However, unlike existing ensemble approaches, which simply combine outputs from different mod-

els, our method incorporates the two models during the training phase to learn a new and unique representation better than the sum of its parts. We employed the proposed model to detect cognitive and physical stress from voice recordings in the above-outlined datasets. We showed that this *hybrid* speech representation generalizes across different datasets and produces better performance than existing DSP-based, as well as DNN-based representation learning methods. Importantly, the proposed hybrid model yielded better results than the ensemble of the methods used for its pre-training.

2. Materials & Methods

2.1. Tasks and Datasets

The speech datasets proposed in this study vary in language, duration, protocol, and more importantly task complexity. A brief description of the datasets is addressed in Table 1. Cognitive load was induced in the speakers through multitasking, based on the idea that cognitive load increases with the number of tasks to be performed simultaneously [20, 21]. On the other hand, physical load was induced by recording speech samples before and after the speaker performed an intensive aerobic physical exercise (VO_2 max running test). In each dataset, the recorded speech samples were labelled in accordance with the experimental condition (i.e., with load / without load).

Cognitive Load 1 (Cog1) - This dataset contains 21 speakers (3 female, 18 male) who participated in their native language (English or French). The experiment comprises six blocks. The first block consisted of a scripted speech task, in which participants read aloud at most three sentences from SIWIS, a multilingual speech and text database [22]. The other five blocks included the same scripted speech task, but this time participants were instructed to count spoken syllables mixed in additive Gaussian white noise (AGWN) while reading aloud. Syllables were selected from the Articulation Index dataset [23]. Each block contained five trials lasting 10 seconds each, producing approximately five minutes of speech material per speaker.

Cognitive Load 2 (Cog2) - 20 Chinese Mandarin speakers (10 female, 10 male) were recruited to perform a similar task to the Cog1 task. The experiment included 40 trials per participant: 20 of them included only reading scripted text from randomly selected articles across various media sources, while the other 20 trials comprised an additional working memory task, as participants had to listen to and memorise three numeric digits while reading. Each trial lasted 15 seconds, giving 10 min of speech material per participant.

Cognitive Load 3 (Cog3) - This task was identical to Cog2, except for the working memory task, as participants were asked to memorize alphanumeric digits (e.g., “A5” instead of “5” in Cog2). The dataset comprised 40 15-second trials from 20 speakers, producing 10 min of speech data per speaker.

Cognitive Load 4 (Cog4) - This dataset was collected from 22 subjects (7 female, 15 male). Audio recordings were collected as participants performed an experiment in which cognitive load was induced based on multitasking, with an experimental design inspired by [24]. The experiment was carried out in three blocks, a baseline block lasting four minutes, followed by two multitasking blocks lasting two minutes each. In the baseline block, participants were shown an image on a screen and were asked to describe it in as much detail as possible in their native language. The images shown on the screen were presented in colour and included cityscapes, landscapes, and famous paintings. Every 30 seconds, a new image appeared, and the par-

ticipants began a new recording. In the multitasking blocks, a high level of cognitive load was applied by asking participants to perform four tasks simultaneously, in addition to describing the image appearing on the screen. The four tasks consisted of an auditory response task, a visual vigilance task, an arithmetic task, and a memory task. The participants were instructed to pay equal attention to all tasks and perform them simultaneously to the best of their ability, while at the same time describing the image on the screen in as much detail as possible. For these blocks, images were selected from commonly used images to study various mental and cognitive impairments through speech (e.g., Cookie-theft picture [25]). As in the baseline block, the presented image changed every 30 seconds. Each participant produced a total of eight minutes of speech data.

Physical Load (Phys) - 28 participants (13 female, 15 male) were recruited to perform a VO_2 max running test [26]. Before and after the exercise, participants were instructed to speak spontaneously in their native language (English or French) for three minutes about personal topics, e.g., their current work, their favourite physical activities, or their daily routine. Subsequently, 10 utterances were extracted from each speaker, five for each condition (before and after exercise), where each utterance lasted between 5 and 15 seconds.

2.2. Audio Representation Models

All cognitive and physical load speech datasets were used as downstream tasks to benchmark different deep learning-based audio representations against well-established handcrafted acoustic feature sets. Importantly, the data-driven models were not fine-tuned directly on the task-specific data, but only pre-trained on independent datasets.

2.2.1. Handcrafted models

openSMILE (OS) - openSMILE [27] is an open-source feature extraction toolkit that generates handcrafted low-level descriptors (LLD) of audio inputs. This toolkit comprises three feature sets obtained by computing functionals on LLD contours. The most prominent of them is ComParE (OS/ComParE) [2], which contains 6373 static features. Another commonly used feature set is eGeMAPS (OS/eGeMAPS) [28] which includes 88 features. Both approaches are purely DSP-based and require no pre-training on auxiliary datasets.

2.2.2. Data-driven models

VGGish - VGGish [29] is a commonly used audio feature extractor derived from the convolutional neural network VGG-16 [30]. This model was pre-trained in a supervised fashion to classify audio events from the Youtube-8M dataset [31].

YAMNet - YAMNet [32] is a DNN pre-trained to classify events from AudioSet, a large collection of audio samples from YouTube videos [33]. Its feature extractor employs the MobileNetV1 convolutional architecture [34].

TRILL, layer 19 - TRILL [17] is a model pre-trained in a self-supervised manner on speech samples from AudioSet. The CNN model based on ResNetish [29] was originally designed as a universal speech representation. It uses triplet-loss learning to create an embedding space where temporally adjacent samples are mapped close together. Here, we use the features generated from layer 19, as found by [17] to perform best on non-semantic speech tasks.

BYOL-A - Unlike contrastive learning frameworks, Bootstrap Your Own Latent for Audio (BYOL-A) [18] generates audio

Table 1: Cognitive/physical load speech datasets summary. ENG: English, FRA: French, CMN: Mandarin Chinese.

Dataset	Language	Utterances	Speakers	Duration (h)	Speech data acquisition protocol
Cog1	ENG/FRA	586	21	1.75	Count audio stimuli (syllables) while reading scripted text
Cog2	CMN	800	20	3.33	Memorise three numeric digits while reading scripted text
Cog3	CMN	800	20	3.33	Memorise three alphanumeric digits while reading scripted text
Cog4	9 languages	349	22	2.93	Describe an image while performing four cognitive tasks
Phys	ENG/FRA	280	28	0.70	Spontaneous speech before/after VO ₂ max running test

representations using two augmented views of a single audio sample, inspired by the success of BYOL [35] for image representation. To obtain audio representations, the log-mel spectrogram (LMS) of an input audio sample is first fed to a data enhancement module, producing two randomly augmented copies of the input LMS (Figure 1). Subsequently, the augmented spectrograms are respectively passed to an online and a target network. Both networks share a similar architecture, with an encoder and a projector module. However, the online network comprises an additional module called predictor to avoid having collapsed representations [35]. The aim of this training paradigm is for the online network to predict the output generated by the target network by minimising a mean squared error (MSE) loss function. BYOL-A has achieved competitive performance in various speech processing tasks, e.g., language and speaker identification.

BYOL-A Extensions - Following the same implementation of BYOL-A¹, we proposed in a previous work BYOL-S² [16], BYOL for speech, which was pre-trained on a speech subset of AudioSet. In addition, we modified the default CNN encoder architecture used in BYOL-A (a 5-layer CNN) using a lightweight version of the Convolutional vision Transformer (CvT)³ [36] to leverage the best of CNNs and Transformers. All proposed encoders produce an embedding of size 2048. The proposed models and associated methods are openly available online⁴.

2.2.3. Hybrid models

Considering the effectiveness of handcrafted features for various paralinguistic challenges [2, 3], we derived a new training protocol for BYOL-S by combining hand-made and data-driven features in a hybrid format. More specifically, we augmented the BYOL framework with a supervision network which generates DSP-based features extracted from the ComParE feature set in openSMILE, as shown in Figure 1. Therefore, in addition to the self-supervised BYOL loss (\mathcal{L}_{ss}), we computed a supervised loss \mathcal{L}_{sup} between the outputs of the supervised and online networks, set as a MSE. The final loss function was defined as a weighted sum of the self-supervised and supervised losses:

$$\mathcal{L}_{\text{hybrid}} = \alpha_{ss}\mathcal{L}_{ss} + \alpha_{sup}\mathcal{L}_{sup} \quad (1)$$

where α_{ss} and α_{sup} weigh the parts of the self-supervised and supervised losses, respectively.

Given the size of the ComParE feature set (6373), we projected the BYOL-S model outputs, from both online and target networks, to produce 6373-dimensional embeddings to compute the hybrid loss ($\mathcal{L}_{\text{hybrid}}$). However, the output size of the pre-trained encoder module remained unchanged (2048).

¹<https://github.com/nttcs/ab/byol-a>

²<https://github.com/Neclow/SERAB>

³<https://github.com/lucidrains/vit-pytorch>

⁴<https://github.com/GasserElbanna/serab-byols>

2.3. Evaluation Pipeline

All models were evaluated on the five cognitive and physical load tasks mentioned above. Audio samples were processed to produce their representation vectors per model. Each downstream dataset was divided into training and test sets with ratios of 70% and 30%, respectively. To ensure a fair and robust evaluation, the partitions were constructed to be speaker-independent (i.e., all utterances from a given speaker belonged to either the training or the test set) and to have a similar gender distribution. Each partition was also standardised (to zero mean and unit variance) to alleviate inter-speaker variability. Subsequently, classification was performed by training and validating linear support-vector machines (SVMs) using five-fold cross-validation. A grid search was applied to optimise the penalty hyperparameter, with values ranging from 10^{-5} to 10^5 . Using such a simple classifier with linear kernel allows us to place more emphasis on the predictive power of the extracted features. Lastly, model performance was evaluated by computing the unweighted average recall on the test set to compensate for any imbalanced data distributions. All procedures in the evaluation pipeline were implemented using `scikit-learn` [37].

3. Results

Table 2 presents the performance of each model in all five tasks by reporting the unweighted average recall (UAR, in percentages) on an unseen and speaker-independent test set. Importantly, using the ComParE handcrafted features exhibited better performance than most deep learning-based models, including recent self-supervised models for speech and audio representation (TRILL, BYOL-A and BYOL-S). Only BYOL-S/CvT achieved better performance, although only slightly better (1.1% average improvement). This result highlights the effectiveness of DSP-based features in paralinguistic tasks.

On the other hand, hybrid representation models, which aim to take advantage of both data-driven and handcrafted features, yielded significantly better performance. The two proposed hybrid models, one using BYOL-A’s CNN encoder [18] and one with CvT encoding [36], consistently outperformed their corresponding BYOL-S models. This result suggests that adding the DSP-based supervision to the self-supervised representation learning framework helped to improve its generalization capacity in cognitive/physical load detection tasks.

As the proposed hybrid loss consists of a weighted sum of its self-supervised and supervised counterparts (Eq. 1), we reported model performance on the five datasets for different $\alpha_{ss} : \alpha_{sup}$ ratios. Here, we selected the **Hybrid BYOL-S/CvT** variant, which yielded the best overall performance. The results shown in Table 3 indicate a steady increase in model performance with increasing weight assigned to the supervised loss (α_{sup}). The 1:2 loss ratio ($\alpha_{ss} : \alpha_{sup}$) yielded the best model for detecting cognitive and physical load in the considered tasks. Increasing the weight of supervised loss further led to a consistent decrease in performance across all tasks.

Table 2: Test UAR (%) on cognitive and physical load datasets. {} indicates the concatenation of feature vectors. OS: openSMILE [27]. The best scores are shown in bold. The models are sorted according to the UAR averaged across the five tasks.

Model	Cog1	Cog2	Cog3	Cog4	Phys	Average
VGGish [29]	64.0	69.6	90.0	77.1	55.0	71.1
TRILL, layer 19 [17]	58.8	72.9	93.3	74.0	62.5	72.3
OS/eGeMAPS [28]	62.9	65.4	93.8	72.9	67.5	72.5
BYOL-A [18]	68.0	70.4	90.0	81.3	61.3	74.2
YAMNet [32]	58.1	71.7	99.2	83.3	62.5	75.0
BYOL-S [16]	65.3	72.9	91.7	80.2	66.3	75.3
OS/ComParE [27]	78.7	74.6	100.0	79.2	52.5	77.0
BYOL-S/CvT [16]	79.2	80.0	100.0	74.0	57.5	78.1
Hybrid & ensemble models:						
{BYOL-S [16]; OS/ComParE [27]}	76.2	77.9	100.0	83.3	56.3	78.7
{BYOL-S/CvT [16]; OS/ComParE [27]}	81.2	77.1	100.0	79.2	57.5	79.0
Hybrid BYOL-S ($\alpha_{ss} = \alpha_{sup} = 1$)	66.5	75.4	90.8	93.8	71.3	79.5
Hybrid BYOL-S/CvT ($\alpha_{ss} = \alpha_{sup} = 1$)	78.4	74.6	99.2	83.3	66.3	80.3
Hybrid BYOL-S/CvT ($\alpha_{ss} = 1, \alpha_{sup} = 2$)	83.8	72.5	98.3	87.5	70.0	82.4

Table 3: Impact of hybrid loss weighting (Eq. 1) in Hybrid BYOL-S/CvT pre-training on test UAR (%).

$\alpha_{ss} : \alpha_{sup}$ (Eq. 1)	Cog1	Cog2	Cog3	Cog4	Phys	Average
4 : 1	73.5	71.3	97.5	87.5	60.0	77.9
2 : 1	72.2	72.5	99.2	92.7	55.0	78.3
4 : 3	76.7	70.0	98.3	88.5	62.5	79.2
1 : 1	78.4	74.6	99.2	83.3	66.3	80.3
3 : 4	80.2	75.0	99.2	86.5	65.0	81.2
1 : 2	83.8	72.5	98.3	87.5	70.0	82.4
1 : 4	81.4	72.5	97.9	86.5	63.8	80.4

Finally, we further validated the proposed hybrid framework by examining the results obtained by concatenating ComParE features with either BYOL-S or BYOL-S/CvT embeddings (8421-D feature vector). In particular, the best hybrid BYOL-S/CvT model yielded a 3.4% higher UAR than the concatenation of the two representations used in its pre-training.

4. Discussion & Conclusions

In this paper, we introduced four novel datasets for large-scale detection of cognitive and physical load from speech. We used the datasets to evaluate existing DPS-based and DNN-based speech representation learning methods, and to develop a novel hybrid (supervised / self-supervised) approach.

In spite of the effectiveness of data-driven features, handcrafted acoustic feature sets maintained competitive performance, a result consistent with outcomes of previous paralinguistic challenges [3]. To leverage their efficacy, we proposed a novel hybrid approach obtained by combining data-driven features from a BYOL-derived model (BYOL-S/CvT) with handcrafted features extracted from openSMILE [27] during pre-training. This hybrid training protocol outperformed all data-driven and handcrafted feature sets presented herein, including fusions of different individual representations. Unlike model ensemble methods, in the inference stage, the proposed hybrid model involves only a single encoder (i.e., a single forward-pass), thus making it lighter & faster than conventional approaches, which usually require processing of the input samples through several models. Following the experimentation with hybrid loss weighting, we found that shifting the loss emphasis to the fixed handcrafted features improved the overall model

performance. However, increasing this emphasis beyond the optimal loss weighting ratio ($\alpha_{ss} = 1$ & $\alpha_{sup} = 2$, Eq. 1) led to a poorer performance, which illustrates the synergistic interaction between the two parts of the hybrid framework.

Importantly, the obtained representation differs from a simple concatenation of BYOL-S/CvT and openSMILE features. Indeed, the hybrid training protocol allowed the model to learn a more robust speech representation for detecting cognitive and physical load. The inclusion of an auxiliary loss to update the online network from projected DSP-based features likely helps improve pre-training stability, a known issue in BYOL-derived frameworks [35] (usually circumvented by exponential averaging of weight updates). Furthermore, the use of fixed DSP-based features in the representation learning, likely reduces any potential overfitting to the pre-training dataset. As a matter of fact, such a hybrid model also produced highly competitive results within the HEAR benchmark, outperforming several state-of-the-art approaches [38].

The proposed hybrid representation learning approach could be easily incorporated into any existing data-driven audio representation learning methods. Moreover, both components of the proposed framework could be substituted with other (potentially larger/complex) models. In particular, the fixed DSP-based feature extractor could be replaced with another pre-trained DNN-based representation. In a similar vein to self-distillation approaches [39], the two learnable representations could be iteratively updated in the pre-training process to produce a potentially more robust representation. Alternatively, the DSP-based feature extractor could also generate a specific family of acoustic features (e.g., prosodic features) which could be beneficial for building more specialized audio representations.

5. References

- [1] J. C. Vázquez-Correa, T. Arias-Vergara, J. R. Orozco-Arroyave, B. Eskofier, J. Klucken, and E. Nöth, "Multimodal Assessment of Parkinson's Disease: A Deep Learning Approach," *IEEE J. Biomed.*, vol. 23, no. 4, pp. 1618–1630, 2018.
- [2] B. Schuller, S. Steidl, A. Batliner, A. Vinciarelli, K. Scherer, F. Ringeval, M. Chetouani, F. Weninger, F. Eyben, E. Marchi *et al.*, "The INTERSPEECH 2013 Computational Paralinguistics Challenge: Social Signals, Conflict, Emotion, Autism," in *Proc. Interspeech 2013*, 2013, pp. 148–152.
- [3] B. Schuller, S. Steidl, A. Batliner, J. Epps, F. Eyben, F. Ringeval, E. Marchi, and Y. Zhang, "The INTERSPEECH 2014 Computational Paralinguistics Challenge: Cognitive & Physical Load," in *Proc. Interspeech 2014*, 2014, pp. 427–431.
- [4] M. Van Puyvelde, X. Neyt, F. McGlone, and N. Pattyn, "Voice stress analysis: A new framework for voice and effort in human performance," *Front. Psychol.*, vol. 9, 2018.
- [5] C. L. Giddens, K. W. Barron, J. Byrd-Craven, K. F. Clark, and A. S. Winter, "Vocal Indices of Stress: A Review," *J. Voice*, vol. 27, no. 3, p. 390.e21–390.e29, 2013.
- [6] S. E. Baker, J. Hipp, and H. Alessio, "Ventilation and Speech Characteristics During Submaximal Aerobic Exercise," *J. Speech Lang. Hear. Res.*, vol. 51, no. 5, pp. 1203–1214, 2008.
- [7] B. Johannes, P. Wittels, R. Enne, G. Eisinger, C. A. Castro, J. L. Thomas, A. B. Adler, and R. Gerzer, "Non-linear function model of voice pitch dependency on physical and mental load," *Eur. J. Appl. Physiol.*, vol. 101, no. 3, pp. 267–276, 2007.
- [8] J. Trouvain and K. P. Truong, "Prosodic characteristics of read speech before and after treadmill running," in *Proc. Interspeech 2015*, 2015, pp. 3700–3704.
- [9] K. W. Godin and J. H. Hansen, "Analysis and Perception of Speech Under Physical Task Stress," in *Proc. Interspeech 2008*, 2008, pp. 1674–1677.
- [10] L. J. Rothkrantz, P. Wiggers, J.-W. A. v. Wees, and R. J. v. Vark, "Voice stress analysis," in *Proc. TSD 2004*, 2004, pp. 449–456.
- [11] J. M. K. Kua, V. Sethu, P. Le, and E. Ambikairajah, "The UNSW Submission to INTERSPEECH 2014 ComParE Cognitive Load Challenge," in *Proc. Interspeech 2014*, 2014, pp. 746–750.
- [12] M. Van Segbroeck, R. Travadi, C. Vaz, J. Kim, M. P. Black, A. Potamianos, and S. S. Narayanan, "Classification of Cognitive Load from Speech using an i-vector Framework," in *Proc. Interspeech 2014*, 2014, pp. 751–755.
- [13] H. Jing, T.-Y. Hu, H.-S. Lee, W.-C. Chen, C.-C. Lee, Y. Tsao, and H.-M. Wang, "Ensemble of Machine Learning Algorithms for Cognitive and Physical Speaker Load Detection," in *Proc. Interspeech 2014*, 2014, pp. 447–451.
- [14] G. Gosztolya, T. Grósz, R. Busa-Fekete, and L. Tóth, "Detecting the intensity of cognitive and physical load using AdaBoost and Deep Rectifier Neural Networks," in *Proc. Interspeech 2014*, 2014, pp. 452–456.
- [15] A. Gallardo Antolín and J. M. Montero Martínez, "A saliency-based attention LSTM model for cognitive load classification from speech," in *Proc. Interspeech 2019*, 2019, pp. 216–220.
- [16] N. Scheidwasser-Clow, M. Kegler, P. Beckmann, and M. Cernak, "SERAB: A multi-lingual benchmark for speech emotion recognition," in *ICASSP*, 2022, pp. 7697–7701.
- [17] J. Shor, A. Jansen, R. Maor, O. Lang, O. Tuval, F. C. Quitry, M. Tagliasacchi, I. Shavitt, D. Emanuel, and Y. Haviv, "Towards Learning a Universal Non-Semantic Representation of Speech," in *Proc. Interspeech 2020*, 2020, pp. 140–144.
- [18] D. Niizumi, D. Takeuchi, Y. Ohishi, N. Harada, and K. Kashino, "BYOL for Audio: Self-Supervised Learning for General-Purpose Audio Representation," in *IJCNN*, 2021, pp. 1–8.
- [19] J. Shor, A. Jansen, W. Han, D. Park, and Y. Zhang, "Universal paralinguistic speech representations using self-supervised conformers," in *ICASSP*, 2022, pp. 3169–3173.
- [20] K. Ryu and R. Myung, "Evaluation of mental workload with a combined measure based on physiological indices during a dual task of tracking and mental arithmetic," *Int. J. Ind. Ergon.*, vol. 35, no. 11, pp. 991–1009, 2005.
- [21] T. T.-J. Chong, M. A. Apps, K. Giehl, S. Hall, C. H. Clifton, and M. Husain, "Computational modelling reveals distinct patterns of cognitive and physical motivation in elite athletes," *Sci. Rep.*, vol. 8, no. 1, pp. 1–11, 2018.
- [22] J.-P. Goldman, P.-E. Honnet, R. Clark, P. N. Garner, M. Ivanova, A. Lazaridis, H. Liang, T. Macedo, B. Pfister, M. S. Ribeiro *et al.*, "The SIWIS database: a multilingual speech database with acted emphasis," in *Proc. Interspeech 2016*, 2016, pp. 1532–1535.
- [23] J. Wright, "Articulation index LDC2005S22," <https://catalog.ldc.upenn.edu/LDC2005S22>, Linguistic Data Consortium, 2005.
- [24] A. Holm, K. Lukander, J. Korpela, M. Sallinen, and K. Müller, "Estimating brain load from the EEG," *Sci. World J.*, vol. 9, pp. 639–651, 2009.
- [25] H. Goodglass, E. Kaplan, and B. Barresi, *Boston Diagnostic Aphasia Examination (3rd ed.)*. Philadelphia, PA: Lippincott Williams & Wilkins, 2001.
- [26] L. Leger and R. Boucher, "An indirect continuous running multi-stage field test: the Université de Montréal Track Test," *Canadian J. of Appl. Sport Sciences*, vol. 5, pp. 77–84, 1980.
- [27] F. Eyben, M. Wöllmer, and B. Schuller, "openSMILE - The Munich Versatile and Fast Open-Source Audio Feature Extractor," in *Proc. ACM Multimedia (MM)*, 2010, pp. 1459–1462.
- [28] F. Eyben, K. R. Scherer, B. W. Schuller, J. Sundberg, E. André, C. Busso, L. Y. Devillers, J. Epps, P. Laukka, S. S. Narayanan *et al.*, "The Geneva Minimalistic Acoustic Parameter Set (GeMAPS) for Voice Research and Affective Computing," *IEEE Trans. Affect. Comput.*, vol. 7, no. 2, pp. 190–202, 2015.
- [29] S. Hershey, S. Chaudhuri, D. P. Ellis, J. F. Gemmeke, A. Jansen, R. C. Moore, M. Plakal, D. Platt, R. A. Saurous, B. Seybold *et al.*, "CNN architectures for large-scale audio classification," in *ICASSP*, 2017, pp. 131–135.
- [30] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," in *ICLR*, 2015.
- [31] S. Abu-El-Haija, N. Kothari, J. Lee, A. Natsev, G. Toderici, B. Varadarajan, and S. Vijayanarasimhan, "Youtube-8M: A large-scale video classification benchmark," *arXiv preprint arXiv:1609.08675*, 2016.
- [32] M. Plakal and D. Ellis, "YAMNet," <https://github.com/tensorflow/models/tree/master/research/audioset/yamnet>, 2020.
- [33] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, "Audio Set: An ontology and human-labeled dataset for audio events," in *ICASSP*, 2017, pp. 776–780.
- [34] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "MobileNets: Efficient convolutional neural networks for mobile vision applications," *arXiv preprint arXiv:1704.04861*, 2017.
- [35] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. Richemond, E. Buchatskaya, C. Doersch, B. A. Pires, Z. D. Guo, M. G. Azar *et al.*, "Bootstrap your own latent - a new approach to self-supervised learning," in *NeurIPS*, 2020, pp. 21271–21284.
- [36] H. Wu, B. Xiao, N. Codella, M. Liu, X. Dai, L. Yuan, and L. Zhang, "CvT: Introducing convolutions to vision transformers," in *ICCV*, 2021, pp. 22–31.
- [37] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg *et al.*, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, 2011.
- [38] G. Elbanna, N. Scheidwasser-Clow, M. Kegler, P. Beckmann, K. El Hajal, and M. Cernak, "Byol-S: Learning self-supervised speech representations by bootstrapping," *arXiv preprint arXiv:2206.12038*, 2022.
- [39] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin, "Emerging properties in self-supervised vision transformers," in *ICCV*, 2021, pp. 9650–9660.