

TS-MoCo: Time-Series Momentum Contrast for Self-Supervised Physiological Representation Learning

Philipp Hallgarten^{*1,2}, David Bethge³, Ozan Özdenizci^{4,5}, Tobias Grosse-Puppenthal¹, Enkelejda Kasneci²

¹Dr. Ing. h.c. F. Porsche AG, Stuttgart, Germany

²TU Munich, Germany

³LMU Munich, Germany

⁴Institute of Theoretical Computer Science, TU Graz, Austria

⁵TU Graz - SAL Dependable Embedded Systems Lab, Silicon Austria Labs, Graz, Austria

Abstract—Limited availability of labeled physiological data often prohibits the use of powerful supervised deep learning models in the biomedical machine intelligence domain. We approach this problem and propose a novel encoding framework that relies on self-supervised learning with momentum contrast to learn representations from multivariate time-series of various physiological domains without needing labels. Our model uses a transformer architecture that can be easily adapted to classification problems by optimizing a linear output classification layer. We experimentally evaluate our framework using two publicly available physiological datasets from different domains, i.e., human activity recognition from embedded inertial sensory and emotion recognition from electroencephalography. We show that our self-supervised learning approach can indeed learn discriminative features which can be exploited in downstream classification tasks. Our work enables the development of domain-agnostic intelligent systems that can effectively analyze multivariate time-series data from physiological domains.

Index Terms—self-supervised learning, physiological signal processing, EEG, emotion recognition, human activity recognition

I. INTRODUCTION

A major challenge in developing high-performing machine learning algorithms in the biomedical domain is the limited availability of labeled data, whereas state-of-the-art deep learning models require more labeled data to generalize better. This poses a significant problem if such models shall be leveraged to automatize tasks in this domain. Accordingly, in this paper, we focus on the following: “*Is it possible to learn physiological domain agnostic deep learning architectures?*” and “*Can we train such an architecture in a self-supervised manner without the need for labeled data?*”.

One promising avenue in this research domain focuses on deep representation learning [1–4]. These models learn to extract generalizable features, omitting the need for domain-specific hand-designed features. We particularly focus on self-supervised learning (SSL) frameworks, a machine learning paradigm where the model is trained without labeled data,




Fig. 1. Proposed *TS-MoCo* framework enables optimizing a physiological signal-agnostic encoder architecture for e.g., emotion recognition from EEG or human activity recognition from inertial sensory.

i.e., the supervision signal necessary for training is generated on-the-fly from the data itself. We investigate the architectural design space in depth and develop a novel SSL framework to learn generalized features across *multiple* physiological signals. Our proposed deep SSL framework, annotated as *TS-MoCo*, optimizes a general feature encoder agnostic to the physiological domain. We introduce a two-fold contrastive loss function for optimization based on time-series momentum contrast [5]. The feature encoder of *TS-MoCo* predicts physiological labels using only one dense layer that must be fitted on the labeled data on top of the learned embedding space. In contrast to prior research that focuses on single, specialized domains, we evaluate our framework in two physiological data domains, *i.e.*, emotion classification, and human activity recognition. Overall, we address the following questions:

- *How can we learn a generalizable feature encoder architecture and learning algorithm for diverse domains of physiological signal recordings?*
- *What is the technical design space for such a deep signal processing framework?*
- *How are the domains affecting the encoding architecture?*

* Corresponding Author: philipp.hallgarten@tum.de

Our framework thereby focuses on learning a signal domain agnostic feature encoder. We propose and explore self-supervised learning of general feature encoders as an alternative approach for medical deep learning researchers that enables use in downstream tasks with few labeled data.

II. RELATED WORK

A. Deep Learning for Physiological Signal Processing

In recent years various works explored deep learning based signal processing pipelines for physiological time-series data recordings. In [6], different autoencoder architectures for encoding physiological signals recorded through smartwatches are compared. [7] implements a transfer learning approach to train a deep learning model for human activity recognition from radar databases. Another set of works on physiological signal processing focuses on electroencephalography (EEG) signals. Cura et al. [8] use a deep convolutional neural network, to detect epilepsy from patients EEG recordings, and Geoffroy et al. [9] detect drowsiness from EEG and electrocardiograms (ECG) with a deep learning architecture.

B. Self-Supervised Learning for Physiological Domains

Self-supervised models have been previously explored for various physiological domains since physiological data monitoring is becoming ubiquitous, resulting in large amounts of unlabeled data. In [10], the authors present an SSL framework trained with a contrastive loss that learns representations for human activity recordings of inertial sensory data from multiple devices. The authors of [11] demonstrate that one can successfully improve emotion classification from electroencephalography (EEG) signals by training a generative adversarial network to synthesize data using SSL. Recently [1] proposed a vanilla SSL framework, *TS-TCC*, trained with a combination of temporal and contextual contrastive loss that achieved significant results across several domains.

C. Self-Supervised Learning with Momentum Contrast

In SSL with momentum contrast, two models of the same architecture are used: a student encoder and a teacher (or momentum) encoder. The student encoder is trained via loss backpropagation, while the teacher model parameters are set as an exponentially moving average of the student model parameters. Seminal work in [5] proposes *MoCo*, a framework that is trained with momentum contrast. The momentum encoder is hereby used to build a representation dictionary on-the-fly, which is then matched to the representation output by the student encoder network in a contrastive loss. *MoCo* achieves superior performance than supervised baselines on multiple vision tasks. In the *BYOL* framework by [12], both the momentum encoder and the student encoder predict representations where different input augmentations are applied. *BYOL* achieves comparable results to state-of-the-art baselines in vision tasks while neglecting the need for negative samples. Finally, [13] proposes *DINO* as an SSL framework trained with momentum contrast without negative samples using vision transformers [14] as a backbone and evaluated on vision tasks.

Momentum contrast reduces computational demand in comparison to conventional contrastive methods by neglecting negative samples, and exhibits a promising potential. However, to date, this SSL mechanism is not explored in depth for analyzing biomedical data. In our work, we tackle this research gap by proposing a framework based on momentum contrast and evaluating it on different physiological domains.

III. METHODOLOGY

A. Training Schema

We denote the signal with t discrete timesteps by $x = [x_t]_{t=0}^T$, and signals can be given with labels into C classes. The training procedure we used is twofold. Firstly, we use a self-supervised pre-training paradigm to learn a generalizable feature encoder. Secondly, we optimize a single linear classifier using the representations by this encoder to solve the downstream classification problem. During self-supervised pre-training, the supervision signal is created directly from the signal itself, omitting the need for laborious annotations. We use a combination of two supervision signals.

1) *Reconstructing the Future*: The first supervision signal is generated by splitting of the final K time-steps from the signal *i.e.* by splitting x into two parts $x_{\text{past}} = [x_t]_{t=0}^{T-K-1}$ and $x_{\text{future}} = [x_t]_{t=T-K}^T$. The former is encoded into a representation vector c by a feature encoder described through f_{ENC} with parameters θ , and the latter is used as target for a reconstruction head, that predicts the subsequent timesteps $t \in [T-K, T]$ of its input from a context vector. To compare the predicted subsequent time-steps \hat{x}_{future} with the ground truth x_{future} we use a mean squared error loss:

$$\mathcal{L}_{\text{Rec}} = \|x_{\text{future}} - \hat{x}_{\text{future}}\|_2^2. \quad (1)$$

2) *Momentum Contrast*: The second supervision signal is generated through momentum contrast [5]. Two identical architectures, teacher and student models, are used during training. The teacher model is provided with x_{past} as input and predicts an output, in our case, a representation vector:

$$c^{<T>} = f_{\text{ENC}}(x_{\text{past}}; \theta^{<T>}). \quad (2)$$

For the student model, the signal is first augmented with an augmentation function before computing the output:

$$c^{<S>} = f_{\text{ENC}}(A(x_{\text{past}}); \theta^{<S>}). \quad (3)$$

Training aims to guide the student towards computing the same representation as the teacher from an augmented version of the input. We compare these two output vectors using a cosine similarity loss metric:

$$\mathcal{L}_{\text{MC}} = 1 - \cos(c^{<S>}, c^{<T>}). \quad (4)$$

3) *Parameter Update*: After each iteration, student model parameters are updated by loss backpropagation

$$\mathcal{L}_{\text{SS}} = \mathcal{L}_{\text{Rec}} + \lambda \mathcal{L}_{\text{MC}}, \quad (5)$$

where the loss weight λ is set as a hyperparameter. The teacher model parameters $\theta^{<T>}$ are updated in order to state

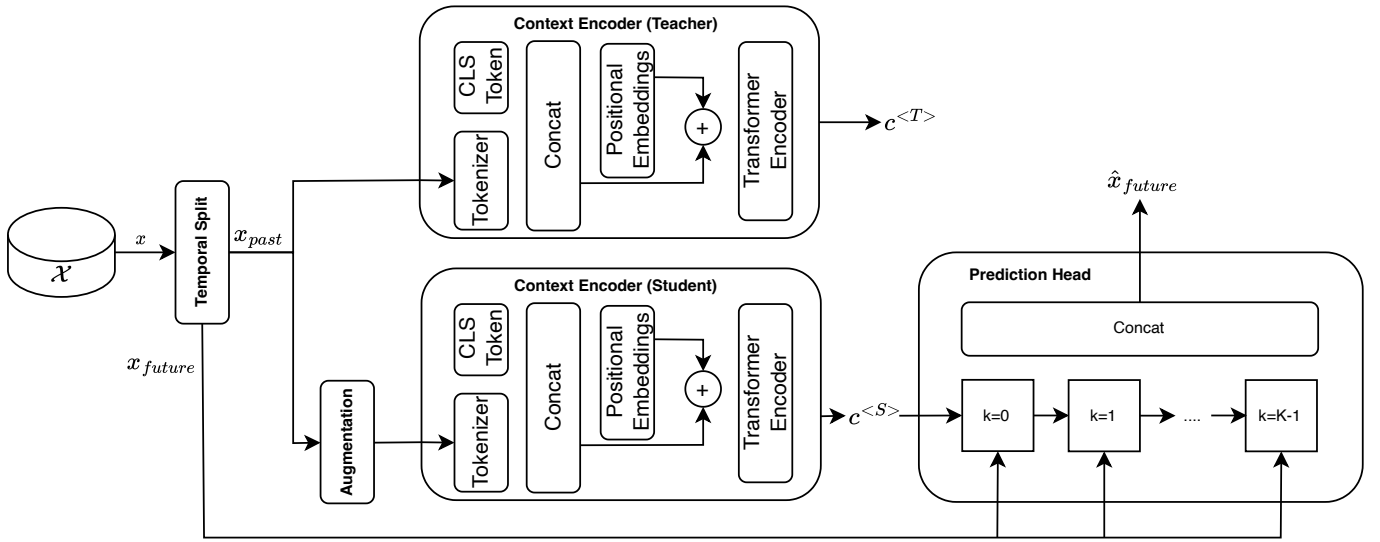


Fig. 2. Architecture of the *TS-MoCo* framework consisting of a student and teacher context encoder, and a GRU-based reconstruction head.

an exponentially moving average (EMA) of the student model parameters $\theta^{<S>}$, as utilized in [5]:

$$\theta^{<T>} = \tau\theta^{<T>} + (1 - \tau)\theta^{<S>}, \quad (6)$$

with τ being the momentum weight hyperparameter.

4) *Linear Evaluation*: Following self-supervised pre-training, we employ a linear evaluation scheme to evaluate the learned feature encoder. We first use the pre-trained frozen encoder to compute representations of non-augmented input signals x . Then a single linear layer that predicts class logits from these representations is utilized according to

$$\hat{y}_i = \log q_i = f_{\text{CLA}}(c_i^{<S>} : \theta^{<C>}), \quad (7)$$

where a conventional cross-entropy loss is used.

B. *TS-MoCo* Framework Architecture

1) *Augmentation Function*: Whether or not training the proposed framework is successful heavily depends on the used augmentation function as it defines the encoding and reconstruction task. If the augmentation is too weak, the tasks become trivial and can be solved by the models without learning the underlying concepts of the data. On the other hand, if the augmentation is too strong, training the model may fail due to limited capabilities of the architecture. Further, we explicitly refrain from domain-specific augmentation functions to allow a fast and easy adaption of our framework to various domains. By employing a *window-wise temporal masking*, we fulfill aforementioned constraints to the augmentation function. Hereby, a certain percentage p of the signal is overwritten with a masking token *i.e.* 0. Especially for high-frequency signals, masking singular timesteps scattered along the time-axis would state an easy task, therefore we assure that all masked timesteps span a continuous window.

2) *Feature Encoder*: The used feature encoder consists of a tokenizer, the addition of positional embeddings, and a transformer encoder. As a first step, the signal used as input to the encoder is mapped into an embedding space by a single linear layer named tokenizer

$$[h_i^{(1)}]_{t=0}^{T-K-1} = \theta_{\text{tokenizer}} * [x_i]_{t=0}^{T-K-1}. \quad (8)$$

Next, like in [15], a classification token ($< CLS >$) is prepended to the signal. The value of this token is initialized randomly and learned during training.

$$[h_i^{(2)}] = [< CLS >, [h_i^1]_{t=0}^{T-K-1}] \quad (9)$$

In order to allow the feature encoder to exploit positional information of the signal values, we add positional embeddings to the tokenized signals, *i.e.* we add a unique vector to the signal values of each timestep. We create the positional embedding matrix \mathbf{P} from sinusoidal according to previous work [16]. We also introduce a hyperparameter $\alpha \in \{0, 1\}$, allowing us to enable/disable positional embeddings:

$$[h_i^{(3)}] = [h_i^{(2)}] + \alpha P_i. \quad (10)$$

Finally, the latent signal is input to a transformer encoder [16] of depth d , outputting a contextual embedding vector for each input token $[h_i^{(d+3)}]$. We use the output embedding for the classification token as context encoding c .

3) *Reconstruction Head*: We use Gated Recurrent Units (GRU) for the reconstruction head. The context vector output by the student feature encoder is thereby used as initial hidden state. Further, we apply teacher-forcing for predicting the future timesteps, *i.e.* we pass the ground truth values for a timestep besides a hidden state to the GRU cell for predicting the values of the next timestep. Our overall architecture is illustrated in Figure 2.

TABLE I

COMPARISON WITH BASELINE MODELS. WE COMPARE OUR SELF-SUPERVISED LEARNING FRAMEWORK WITH A SUPERVISED TRAINED BASELINE MODEL WITH THE SAME ENCODING ARCHITECTURE AND A RANDOM CLASSIFIER ON THE TEST SET.

Dataset	Task	Signal	# classes / channels	Supervised	TS-TCC	TS-MoCo (ours)	Random Classifier
SEED	Emotion Classification	EEG	3 / 62	0.79	0.42	0.43	0.33
UCIHAR	Activity Recognition	Inertial Sensory	6 / 9	0.89	0.90	0.52	0.16

TABLE II

ABLATION STUDY. *TS-MoCo* TEST SET ACCURACIES FOR VARIOUS SETTINGS OF RELEVANT HYPERPARAMETERS. “—” INDICATES THE SAME VALUE AS IN THE FIRST ROW.

Hyperparameter Settings					Classification Accuracies	
κ	λ	K	p_M	α	SEED	UCIHAR
0.9	1.0	6	0.5	0	0.43	0.36
0.99	—	—	—	—	0.42	0.47
0.7	—	—	—	—	0.42	0.34
—	100	—	—	—	0.41	0.32
—	0.01	—	—	—	0.42	0.52
—	—	24	—	—	0.39	0.37
—	—	2	—	—	0.41	0.37
—	—	—	0.75	—	0.41	0.33
—	—	—	0.25	—	0.42	0.40
—	—	—	—	1	0.43	0.38

IV. EXPERIMENTAL STUDY

A. Datasets

1) *SEED*: The SJTU Emotion EEG Dataset (*SEED*) [17], [18] is a dataset of EEG recordings from 15 subjects during viewing of emotion-eliciting videos. Signals were recorded from 62 channels at 1000 Hz and later downsampled to 200 Hz. Following recent work [4], [19], [20], we applied a 4–40 Hz bandpass filter and segmented the signals into 2 s non-overlapping windows. EEG recording signals are labeled to be of either negative, neutral, or positive emotions.

2) *UCIHAR*: This dataset [21] for human activity recognition comprises recordings from 30 subjects during six activities of daily living: *walking*, *walking upstairs*, *walking downstairs*, *sitting*, *standing*, *laying*, measured from waist-mounted smartphone accelerometer and gyroscope. Class imbalance is handled via stochastic undersampling [19].

B. Baseline Comparison Models

1) *Random Classifier*: We use this as a baseline that is expected to make predictions based on the relative frequency of each class in the training data. We employ a stratified strategy, *i.e.* predicted class label is sampled from a multinomial distribution with empirical priors.

2) *Supervised Trained Models*: To demonstrate an upper bound of performance, we also compare our framework to a fully-supervised pipeline. The architecture of this model is similar to that of the self-supervised trained model during evaluation *i.e.*, it consists of a transformer-based feature encoder followed by a single dense layer for classification.

It is important to note that the supervised, trained model task is much easier to accomplish. However, self-supervised

trained models can offer additional advantages of being suitable in scenarios where labeled data is limited and expensive. Furthermore, the self-supervised model learns task-agnostic embeddings, which allows the feature encoder to be reused for other tasks arbitrarily. We highlight these factors for consideration while interpreting our results.

3) *Self-Supervised Trained Models*: We compare *TS-MoCo* against another self-supervised learning framework for physiological data, *TS-TCC* as introduced in Section II-B. In contrast to *TS-MoCo*, *TS-TCC* uses a conventional contrastive loss instead of momentum contrast.

C. Parameter Optimization

We pre-trained our models for 100 epochs and then performed 100 (*UCIHAR*) / 150 (*SEED*) epochs of linear evaluation with early stopping after 20 epochs based on the validation loss. Baseline models were similarly trained for 100 epochs. We used a 60-20-20 split for the train, validation, and test.

V. RESULTS

A. Comparison with Baseline Models

Table I presents a comparison of the self-supervised *TS-MoCo* framework with a supervised trained baseline model with same encoding architecture (*Supervised*), the self-supervised *TS-TCC* framework, and a random classifier based on the strategy introduced in Section IV-B1. Note that the supervised model demonstrates an upper bound of performance to the problems. Our results clearly show that *TS-MoCo* achieves significantly above chance-level accuracies on both datasets. Classification accuracies of the self-supervised models are comparable on the *SEED* dataset, however, the simplified training pipeline of *TS-MoCo* appears to introduce a classification accuracy trade-off on the *UCIHAR* dataset.

B. Ablation Studies

1) *Model Components*: To evaluate the influence of the hyperparameters κ , λ , K , p_M , and α , we report the results of an ablation study in Table II. We observe different influences for different domains. For the *SEED* dataset, varying the hyperparameters barely results in different classification performance, as the best configuration achieves only +0.04 better accuracy than the worst configuration.

For *UCIHAR*, we observe results being highly dependent on the choice of the hyperparameters and can result in low accuracies *e.g.*, if the value of λ is increased.

2) *Randomly Initialized Encoder*: To evaluate the feasibility of the pre-training phase of *TS-MoCo*, we compare it against a model with same encoding architecture, but the encoder of this setup is randomly initialized and frozen *i.e.*, only a classifier is fitted to the outputs of a *randomly initialized encoder*. By comparing the classification results on the *UCI-HAR* dataset, we observe that the classification accuracy of the *random encoder* falls short of that achieved through *TS-MoCo* by 0.06 (0.52 vs 0.46), *i.e.* pre-training of *TS-MoCo* does help to slightly improve classification accuracy.

VI. DISCUSSION

We present *TS-MoCo*¹, the first transformer encoder-based self-supervised learning framework with momentum contrast for physiological signal recording domain datasets. We performed an experimental pilot study using a linear evaluation scheme to demonstrate the representational capacity of our self-supervised encoding architecture.

We observed that the performance of *TS-MoCo* does not reach to supervised trained baseline models. However importantly, representations learned by *TS-MoCo* offer the additional benefit of being task agnostic and do not require labeled data to be trained. Further, we note that simplifying the training schema with momentum contrast can limit classification accuracies in certain domains compared to conventional contrastive learning. We also observed that a strong influence of used hyperparameters can occur depending on the data domain.

Although the supervised baselines performed better than *TS-MoCo*, our model introduces a valuable encoding mechanism for physiological signal domains where no labeled data is available at all when supervised learning is not possible. Especially in the medical domain, labeling data is often cumbersome and dissemination of such information often leads to concerns regarding data privacy [22].

In future work, we aim to analyze the amount of labeled data necessary to train a classification layer based on the *TS-MoCo* encoded representations. We expect this amount to be significantly lower than needed for the supervised trained models. This would favor the use of our pipeline for domains where only very little labeled data is available.

ACKNOWLEDGEMENT

This research was partly funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) in TRR161 (Quantitative methods for visual computing, Project ID 251654672) in Project C06.

REFERENCES

[1] E. Eldele, M. Ragab, Z. Chen, M. Wu, C. K. Kwok, X. Li, and C. Guan, "Time-series representation learning via temporal and contextual contrasting," in *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, 2021, pp. 2352–2359.

[2] M. Han, O. Özdenizci, T. Koike-Akino, Y. Wang, and D. Erdoğmuş, "Universal physiological representation learning with soft-disentangled rateless autoencoders," *IEEE Journal of Biomedical and Health Informatics*, vol. 25, no. 8, pp. 2928–2937, 2021.

[3] O. Özdenizci and D. Erdoğmuş, "Stochastic mutual information gradient estimation for dimensionality reduction networks," *Information Sciences*, vol. 570, pp. 298–305, 2021.

[4] D. Bethge, P. Hallgarten, T. Grosse-Puppenthal, M. Kari, L. L. Chuang, O. Özdenizci, and A. Schmidt, "EEG2Vec: Learning affective EEG representations via variational autoencoders," in *IEEE International Conference on Systems, Man, and Cybernetics*, 2022, pp. 3150–3157.

[5] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 9729–9738.

[6] M. Panagiotou, A. Zlatintsi, P. Filntisis, A. Roumeliotis, N. Efthymiou, and P. Maragos, "A comparative study of autoencoder architectures for mental health analysis using wearable sensors data," in *30th European Signal Processing Conference*, 2022, pp. 1258–1262.

[7] J. Fix, I. Hinojosa, C. Ren, G. Manfredi, and T. Letertre, "Transfer learning for human activity classification in multiple radar setups," in *30th European Signal Processing Conference*, 2022, pp. 1576–1580.

[8] O. K. Cura, M. A. Ozdemir, and A. Akan, "Epileptic EEG classification using synchrosqueezing transform with machine and deep learning techniques," in *28th European Signal Processing Conference*, 2021, pp. 1210–1214.

[9] G. Geoffroy, L. Chaari, J.-Y. Tourneret, and H. Wendt, "Drowsiness detection using joint EEG-ECG data with deep learning," in *29th European Signal Processing Conference*, 2021, pp. 955–959.

[10] Y. Jain, C. I. Tang, C. Min, F. Kawsar, and A. Mathur, "ColloSSL: Collaborative self-supervised learning for human activity recognition," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 6, no. 1, pp. 1–28, 2022.

[11] Z. Zhang, S.-h. Zhong, and Y. Liu, "GANSER: A self-supervised data augmentation framework for EEG-based emotion recognition," *IEEE Transactions on Affective Computing*, 2022.

[12] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. Richemond, E. Buchatskaya, C. Doersch, B. Avila Pires, Z. Guo, M. Gheshlaghi Azar *et al.*, "Bootstrap your own latent—a new approach to self-supervised learning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 21 271–21 284, 2020.

[13] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin, "Emerging properties in self-supervised vision transformers," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9650–9660.

[14] A. Dosovitskiy *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," in *International Conference on Learning Representations*, 2021.

[15] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.

[16] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in Neural Information Processing Systems*, vol. 30, 2017.

[17] W.-L. Zheng and B.-L. Lu, "Investigating critical frequency bands and channels for EEG-based emotion recognition with deep neural networks," *IEEE Transactions on Autonomous Mental Development*, 2015.

[18] R.-N. Duan, J.-Y. Zhu, and B.-L. Lu, "Differential entropy feature for EEG-based emotion classification," in *6th International IEEE/EMBS Conference on Neural Engineering*. IEEE, 2013, pp. 81–84.

[19] D. Bethge, P. Hallgarten, T. Grosse-Puppenthal, M. Kari, R. Mikut, A. Schmidt, and O. Özdenizci, "Domain-invariant representation learning from EEG with private encoders," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2022, pp. 1236–1240.

[20] D. Bethge, P. Hallgarten, O. Özdenizci, R. Mikut, A. Schmidt, and T. Grosse-Puppenthal, "Exploiting multiple EEG data domains with adversarial learning," in *44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society*, 2022, pp. 3154–3158.

[21] D. Anguita, A. Ghio, L. Oneto, X. Parra Perez, and J. L. Reyes Ortiz, "A public domain dataset for human activity recognition using smartphones," in *Proceedings of the 21th International European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*, 2013, pp. 437–442.

[22] M. Benčević, M. Habijan, I. Galić, and A. Pizurica, "Self-supervised learning as a means to reduce the need for labeled data in medical image analysis," in *30th European Signal Processing Conference*, 2022, pp. 1328–1332.

¹<https://github.com/philipp77/TS-MoCo>