



HAL
open science

Transformer-Based Self-Supervised Learning for Emotion Recognition

Juan Vazquez-Rodriguez, Grégoire Lefebvre, Julien Cumin, James L. Crowley

► **To cite this version:**

Juan Vazquez-Rodriguez, Grégoire Lefebvre, Julien Cumin, James L. Crowley. Transformer-Based Self-Supervised Learning for Emotion Recognition. 26th International Conference on Pattern Recognition (ICPR 2022), Aug 2022, Montreal, Canada. hal-03634490v2

HAL Id: hal-03634490

<https://hal.science/hal-03634490v2>

Submitted on 1 Jun 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

TRANSFORMER-BASED SELF-SUPERVISED LEARNING FOR EMOTION RECOGNITION

Juan Vazquez-Rodriguez^{1,2}, Grégoire Lefebvre¹, Julien Cumin¹ and James L. Crowley²

¹Orange Labs, Grenoble, France

Email: {juan.vazquezrodriguez, gregoire.lefebvre, julien1.cumin}@orange.com

²Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG, Grenoble, France

Email: james.crowley@inria.fr

ABSTRACT

In order to exploit representations of time-series signals, such as physiological signals, it is essential that these representations capture relevant information from the whole signal. In this work, we propose to use a Transformer-based model to process electrocardiograms (ECG) for emotion recognition. Attention mechanisms of the Transformer can be used to build contextualized representations for a signal, giving more importance to relevant parts. These representations may then be processed with a fully-connected network to predict emotions.

To overcome the relatively small size of datasets with emotional labels, we employ self-supervised learning. We gathered several ECG datasets with no labels of emotion to pre-train our model, which we then fine-tuned for emotion recognition on the AMIGOS dataset. We show that our approach reaches state-of-the-art performances for emotion recognition using ECG signals on AMIGOS. More generally, our experiments show that transformers and pre-training are promising strategies for emotion recognition with physiological signals.

1 Introduction

When processing time-series signals with deep learning approaches, it is useful to be able to aggregate information from the whole signal, including long-range information, in a way that the most relevant parts are given more importance. One way of doing this is by employing an attention mechanism [1] that uses attention weights to limit processing to relevant contextual information, independent of distance.

Arguably, the Transformer [2] is one of the most successful attention-based approaches. Developed for Natural Language Processing (NLP), the Transformer uses attention mechanisms to interpret sequences of words, and is suitable for use in other tasks requiring interpretation of sequences, such as time series forecasting, [3], analysis of medical physiological signals [4, 5], and recognition of human activity from motion [6].

Physiological signal analysis can be seen as a form of time-series analysis and are thus amenable to processing with Transformers. Moreover, these signals can be used to predict emotions [7], and sensors for these types of signals can be incorporated into wearable devices, as a non-invasive means for monitoring the emotional reaction of users. Several works in this direction have emerged using signals like electrocardiograms (ECG) [8, 9], electroencephalograms (EEG) [10, 11], electrodermal activity (EDA) [12], and other types of physiological signals [13, 14].

Established approaches for deep learning with Convolutions and Recurrent networks require large datasets of labeled training data. However, providing ground truth emotion labels for physiological data is a difficult and expensive process, limiting the availability of data for training [15, 16, 17]. Pre-training models with self-supervised learning can help to overcome this lack of labeled training data. With such an approach, during pre-training the model learns general data representations using large volumes of unlabeled data. The model is then fine tuned for a specific task using labeled data. This approach has been successfully used in other domains including NLP [18, 19] and Computer Vision [20, 21]. It has also been successfully used in affective computing, in tasks like emotion recognition from

physiological signals [9, 22] and from speech [23], personality recognition [24], and facial expression recognition [25, 26, 27].

In this paper, we address the problem of predicting emotions from ECG signals. We are interested in obtaining contextualized representations from these signals using a Transformer-based architecture, and then using these representations to predict low/high levels of arousal and valence. We believe that the contextualized representations obtained with the Transformer should capture relevant information from the whole signal, which the performance of the downstream task of emotion recognition should benefit from. Our main contributions are: 1. We show that it is feasible to use a Transformer-based architecture for emotion prediction from ECG signals. 2. We show that using a self-supervised technique to pre-train the model is useful for ECG signals, achieving superior performance in emotion recognition than a fully-supervised approach. 3. We show that our pre-trained Transformer-based model reaches state-of-the-art performances on a dataset of the literature.

2 Related Work

Traditional techniques for emotion recognition from physiological signals include Gaussian naive Bayes, Support Vector Machines, k-Nearest Neighbours, and Random Forests. [16, 17, 28, 29, 30, 31]. These approaches typically use manually-selected time and frequency features derived from intuition and domain knowledge. Shukla et al. [12] show that commonly used features for arousal and valence prediction are not necessarily the most discriminant. This illustrates the difficulty of selecting good hand-crafted features.

To overcome this, researchers have increasingly used deep learning techniques to extract features from physiological signals for emotion recognition. A common approach, described by Santamaria et al. [8], is to use a 1D Convolutional Neural Network (CNN) to extract the features (also called representations), followed by a fully-connected network (FCN) used as classifier to predict emotions. As an alternative, Harper and Southern [32] use a Long Short-Term Memory (LSTM) network concurrently with a 1D-CNN. Siddharth et al. [33], first convert signals into an image using spectrograms [34], and then use a 2D-CNN for feature extraction, followed by an extreme learning machine [35] for classification.

One drawback of these CNN-based approaches is that they do not take context into account: after training, kernel weights of the CNN are static, no matter the input. For this reason, attention-based architectures such as the Transformer [2], capable of incorporating contextual information, have started to be used for emotion prediction. Transformers have been successfully used to recognize emotions with multimodal inputs composed of text, visual, audio and physiological signals [36, 37, 38, 39, 40]. In addition, Transformers have been used to process time-series in general [3, 41], and also to process uni-modal physiological signals in particular, with the aim of recognizing emotions. Arjun et al. [42] employ a variation of the Transformer, the Vision Transformer [43] to process EEG signals for emotion recognition, converting the EEG signals into images using continuous wavelet transform. Behinaein et al. [44] propose to detect stress from ECG signals, by using a 1D-CNN followed by a Transformer and a FCN as classifier.

Most of the approaches for measuring emotions, including those using multimodal physiological data, have relied on supervised learning, and thus are limited by the availability of labeled training data. Using self-supervised pre-training can improve performances of a model [45], as it allows to learn more general representations, thus avoiding overfitting in the downstream task. This is especially important for tasks with limited labeled data. Sarkar and Etemad [9] pre-train a 1D-CNN using a self-supervised task to learn representations from ECG signals. Their self-supervised task consists in first transforming the signal, with operations such as scaling or adding noise, and then using the network to predict which transformation has been applied. Ross et al. [22] learn representations from ECG signals using auto-encoders based on 1D-CNN. In both approaches, once the representations have been learned, they are used to predict emotions.

In contrast with the two previously mentioned approaches, we propose to take into account contextual information during pre-training by using a Transformer-based model. Such an approach has been used for pre-training Transformers from visual, speech and textual modalities [23, 46, 47, 48, 49]. Haresamudram et al. use this approach to pre-train a Transformer for human activity recognition using accelerometer and gyroscope data [6]. Zerveas et al. [50] develop a framework for multivariate time-series representation learning, by pre-training a Transformer-based architecture. However, none of these works deal with uni-modal physiological signals. In this work, we have extended this approach for use with ECG signals. Specifically, we investigate the effectiveness of pre-training a Transformer for ECG emotion recognition, which to the best of our knowledge has not been done before.

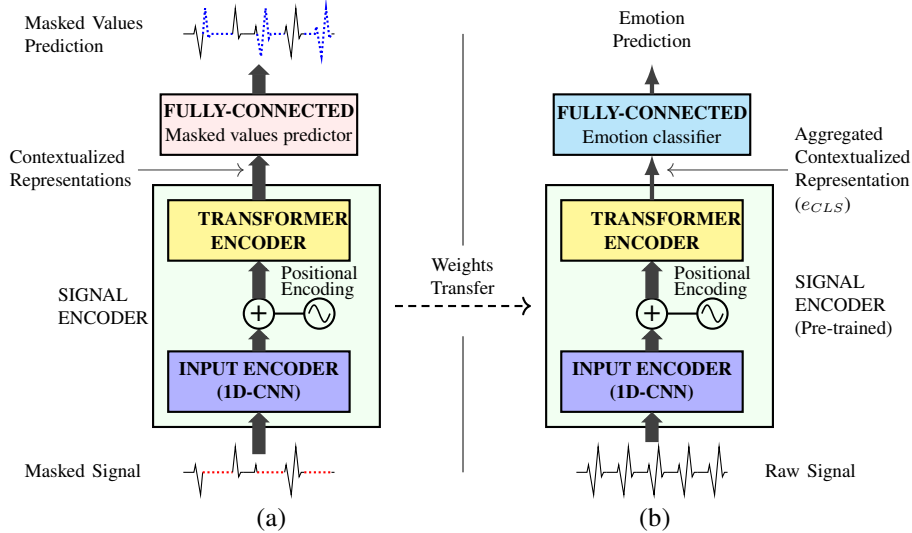


Figure 1: Our approach with self supervised learning based on a Transformer (a) and fine-tuning strategy for learning the final emotion predictor (b).

3 Our approach

Our framework for using deep learning for emotion recognition is based on the following two steps: first, we need to obtain contextualized representations from time-series signals using a deep model; then, we use those representations to perform the targeted downstream task. In this paper, the considered physiological time-series are raw ECG signals, and the downstream task is binary emotion recognition: predicting high/low levels of arousal, and high/low levels of valence.

For the first step (see Figure 1.a), we developed a signal encoder based on deep neural networks and attention, to obtain contextualized representations from ECG signals. The main component of the signal encoder is a Transformer [2]. This signal encoder is pre-trained with a self-supervised task, using unlabeled ECG data. For the second step (see Figure 1.b), we fine-tune the whole model (the signal encoder and the fully-connected classifier) for our downstream task of binary emotion recognition, using labeled ECG data.

In the following subsections, we describe in detail the different components of our approach.

3.1 Learning Contextualized Representations

At the heart of our signal encoder is a Transformer encoder [2], which we use to learn contextualized representations of ECG signals. In Transformers, contextual information is obtained through an attention mechanism, with the attention function considered as a mapping of a query vector along with a group of key-value vector pairs to an output. In the case of the Transformer encoder, each position in the output pays attention to all positions in the input. Several attention modules (also called *heads*) are used, creating various representation subspaces and improving the ability of the model to be attentive to different positions. The Transformer encoder is constructed by stacking several layers containing a multi-head attention module followed by a fully-connected network applied to each position, with residual connections. Since our implementation of the Transformer is almost identical to the one described in [2], we refer the readers to this paper for further details.

In Figure 2, we present our signal encoder, which we describe in the remainder of this subsection.

Input Encoder: to process an ECG signal with the Transformer, we first encode it into s feature vectors of dimension d_{model} that represent each one of the s values of the ECG signal. We use 1D Convolutional Neural Networks (1D-CNN) to perform this encoding, like in [6, 36, 51]. Thus, for a raw input signal $X = \{x_1, \dots, x_s\}$ where x_i is a single value, after encoding X with the input encoder we obtain features $F = \{f_1, \dots, f_s\}$ where $f_i \in \mathbb{R}^{d_{\text{model}}}$.

CLS token: given that our downstream task is a classification task, we need to obtain a single representation of the whole processed signal at the output of our signal encoder. Similar to what is done in BERT [19], we append a special classification token (CLS) at the start of the feature sequence F , resulting in the sequence $F' = \{CLS, f_1, \dots, f_s\}$. We use a trainable vector of dimension d_{model} as CLS token. At the output of the Transformer, we obtain an embedding of




Figure 2: Our Transformer-based signal encoder that produces contextualized representations. The aggregated representation e_{CLS} is used for classification.

the CLS token (e_{CLS}), along with the rest of the representations of the signal (see Figure 2 and Equation 2). Through the attention mechanisms of the Transformer, e_{CLS} is capable of aggregating information from the entire input signal and its contextualized representations. For this reason, at classification time, e_{CLS} can be used as input for the classifier network.

Positional Encoding: positional information of each input is required so that the Transformer can take into account the actual ordering of time-steps in the input sequence. As in [2], we use fixed sinusoidal positional embeddings. We sum the positional embeddings with the features F' :

$$Z = \{CLS + pe_0, f_1 + pe_1, \dots, f_s + pe_s\}, \quad (1)$$

where $pe_i \in \mathbb{R}^{d_{\text{model}}}$ is the positional embedding for time-step i . We then apply layer normalization [52] to Z . Please refer to [2] for details on how to obtain the positional embeddings.

Transformer Encoder: we obtain contextualized representations E using a Transformer encoder with h heads and l layers on the sequence Z :

$$E = \{e_{CLS}, e_1, \dots, e_s\} = \text{Transformer}_{h,l}(Z). \quad (2)$$

We then use the representations E for emotion recognition, as is described in Section 3.3

3.2 Pre-training Task

To pre-train our signal encoder, we employ a self-supervised approach inspired in BERT [19]. We mask random segments of a certain length by replacing them with zeros, and then we train our model to predict the masked values, as shown in Figure 1a. Labeled data is not needed for this step.

Similar to [51], a proportion p of points is randomly selected from the input signal as starting points for masked segments, and then for each starting point the subsequent M points are masked. The masked segments may overlap.

To predict masked points, we use a fully-connected network (FCN) on top of the signal encoder, as shown in Figure 1a. We only predict values of masked inputs, as opposed to reconstructing the whole signal. We use the mean square error between predicted and real values as the reconstruction loss \mathcal{L}_r during pre-training:

$$\mathcal{L}_r = \frac{1}{N_m} \sum_{j=1}^{N_m} (\hat{x}_j - x_{p(j)})^2, \quad (3)$$

where N_m is the number of masked values, \hat{x}_j is the prediction corresponding to the j^{th} masked value, and $x_{p(j)}$ is the original input value selected to be the j^{th} masked value, whose position is $p(j)$ in the input signal.

3.3 Fine-tuning

We fine-tune our model to perform binary emotion prediction, as shown in Figure 1b. This step is supervised, using labeled data. To make the prediction, a FCN is added on top of the signal encoder, using e_{CLS} as input. We initialize the signal encoder with the weights obtained after pre-training, while the FCN is randomly initialized. We then fine-tune all the parameters of the model, including the pre-trained weights. For this task, we minimize the binary cross-entropy loss \mathcal{L}_{ft} :

$$\mathcal{L}_{ft} = -w_p y \log[\sigma(out)] - (1 - y) \log[1 - \sigma(out)] \quad (4)$$

Table 1: Comparison of different strategies of our approach on AMIGOS dataset

		Arousal Acc.	Arousal F1	Valence Acc.	Valence F1
Aggregation Method	Max-Pooling 1	$0.85\pm 6.6e^{-3}$	$0.84\pm 6.4e^{-3}$	$0.78\pm 6.5e^{-3}$	$0.78\pm 6.6e^{-3}$
	Max-Pooling 2	$0.86\pm 7.4e^{-3}$	$0.84\pm 7.3e^{-3}$	$0.8\pm 6.3e^{-3}$	$0.8\pm 5.9e^{-3}$
	Average-Pooling 1	$0.87\pm 8.3e^{-3}$	$0.87\pm 7.3e^{-3}$	$0.82\pm 6.2e^{-3}$	$0.82\pm 6.7e^{-3}$
	Average-Pooling 2	$0.88\pm 4.4e^{-3}$	$0.87\pm 4.6e^{-3}$	$0.83\pm 6.4e^{-3}$	$0.83\pm 6.6e^{-3}$
	Last Representation	$0.85\pm 1.3e^{-2}$	$0.84\pm 1.2e^{-2}$	$0.8\pm 7.6e^{-3}$	$0.8\pm 8.0e^{-3}$
Segment Length	40 seconds	$0.86\pm 1.2e^{-2}$	$0.85\pm 1.1e^{-2}$	$0.82\pm 1.0e^{-2}$	$0.81\pm 9.9e^{-3}$
	20 seconds	$0.87\pm 5.6e^{-3}$	$0.86\pm 6.4e^{-3}$	$0.82\pm 7.8e^{-3}$	$0.82\pm 8.1e^{-3}$
Our Best Approach	CLS with 10s segment	$0.88\pm 5.4e^{-3}$	$0.87\pm 5.4e^{-3}$	$0.83\pm 7.8e^{-3}$	$0.83\pm 7.4e^{-3}$

Table 2: No Pre-training vs pre-trained model

Pre-train	Arousal Acc.	Arousal F1	Valence Acc.	Valence F1
No	$0.85\pm 5.6e^{-3}$	$0.84\pm 5.8e^{-3}$	$0.8\pm 6.5e^{-3}$	$0.8\pm 6.4e^{-3}$
Yes	$0.88\pm 5.4e^{-3}$	$0.87\pm 5.4e^{-3}$	$0.83\pm 7.8e^{-3}$	$0.83\pm 7.4e^{-3}$

where y is an indicator variable with value 1 if the class of the ground truth is positive and 0 if it is negative, out is the output of the classifier, σ is the sigmoid function, and w_p is the ratio of negative to positive training samples, used to compensate unbalances that may be present in the dataset.

4 Experimental Setup

In this section, we describe the experimental choices taken to evaluate our approach for a downstream task of binary emotion recognition (high/low levels of arousal and valence), on ECG signals. We present the datasets used, the pre-processes employed, and the parametrization of our two steps of pre-training and fine-tuning.

4.1 Datasets

For pre-training, we only require datasets that contain ECG signals, regardless of why they were actually collected or which labeling they have, if any. The datasets that we use in our experiments are: ASCERTAIN [16], DREAMER [53], PsPM-FR [54], PsPM-HRM5 [55], PsPM-RRM1-2 [56], and PsPM-VIS [57]. We also employ the AMIGOS dataset [17], taking care of not using the same data for pre-training and evaluating our model, as this dataset is also used for the downstream task. To gather as much data as possible, we use all the ECG channels available in the datasets. For ASCERTAIN, we discard some signals according to the quality evaluation provided in the dataset: if a signal has a quality level of 3 or worse in the provided scale, it is discarded. In total, there are around 230 hours of ECG data for pre-training.

To fine-tune our model to predict emotions, we use the AMIGOS dataset [17]. In this dataset, 40 subjects watched videos specially selected to evoke an emotion. After watching each video, a self-assessment of their emotional state is conducted. In this assessment, subjects rated their levels of arousal and valence on a scale of 1 to 9. Of the 40 subjects, 37 watched a total of 20 videos, while the other 3 subjects watched only 16 videos. During each trial, ECG data were recorded on both left and right arms. We use data only from the left arm to fine-tune our model. AMIGOS includes a pre-processed version of the data, that was down-sampled to 128Hz and filtered with a low-pass filter with 60Hz cut-off frequency. We use these pre-processed data for our experiments, including the pre-training phase. The ECG data that we use for fine-tuning amounts to around 65 hours of recordings.

4.2 Signal Pre-processing

We first filter signals with an 8th order Butterworth band-pass filter, having a low-cut-off frequency of 0.8Hz and a high-cut-off frequency of 50Hz. We then down-sample the signals to 128 Hz, except for AMIGOS which already has that sampling rate. Signals are normalized so they have zero-mean and unit-variance, for each subject independently. Signals are finally divided into 10-second segments (we also report results for segments of 20 seconds and 40 seconds).

4.3 Pre-training

As stated previously, we use ASCERTAIN, DREAMER, PsPM-FR, PsPM-RRM1-2, PsPM-VIS, and AMIGOS for pre-training. Since we also use AMIGOS for fine-tuning, we need to avoid using the same segments both for pre-training and for evaluating the model. To do this, we pre-train two models, one using half of the data from AMIGOS, and the second using the other half. When testing our model with certain segments from AMIGOS, we fine-tune the model that was pre-trained with the half of AMIGOS that do not contain those segments. More details are given in Section 4.4. In total, both of our models are pre-trained with 83401 10-second segments.

We select a proportion of $p = 0.0325$ points from each input segment to be the starting point of a masked span of length $M = 20$, resulting in around 47% of the input values masked.

The input encoder is built with 3 layers of 1D-CNN with ReLU activation function. We use layer normalization [52] on the first layer, and at the output of the encoder. Kernel sizes are (65, 33, 17), the numbers of channels are (64, 128, 256) and the stride for all layers is 1. This results in a receptive field of 113 input values or 0.88s. We selected this receptive field size because it is comparable with the typical interval between peaks on an ECG signal, which is between 0.6s and 1s, including when experiencing emotions [58].

The Transformer in our signal encoder has a model dimension $d_{\text{model}} = 256$, 2 layers and 2 attention heads, with its FCN size of $d_{\text{model}} \cdot 4 = 1024$. The FCN used to predict the masked values consists of a single linear layer of size $d_{\text{model}}/2 = 128$ followed by a ReLU activation function. An additional linear layer is used to project the output vector to a single value, which corresponds to the predicted value of a masked point.

We pre-train the two models for 500 epochs, warming up the learning rate over the first 30 epochs up to a value of 0.001 and using linear decay after that. We employ Adam optimization, with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and L_2 weight decay of 0.005. We use dropout of 0.1 at the end of the input encoder, after the positional encoding, and inside the Transformer.

We tuned the number of layers and heads in the Transformer, the learning rate, and the warm-up duration using the Ray Tune framework [59] with BOHB optimization [60].

4.4 Fine-Tuning

We fine-tune our model (both the signal encoder and FCN classifier) for emotion recognition with the AMIGOS dataset, using each of the 10-second segments as a sample. As labels, we use the emotional self-assessments given in the dataset. Since these assessments provide values of arousal and valence on a scale 1 to 9, we use the average arousal and the average valence as threshold value to determine a low or a high level.

We use 10-fold cross-validation to evaluate our approach. Recall that we pre-train two signal encoders. After dividing AMIGOS into 10 folds, we use folds 1 to 5 to pre-train one signal encoder (SE_1), and folds 6 to 10 to pre-train the second one (SE_2) (and all data from the other datasets, for both). Then, when we fine-tune the models to be tested with folds 1 to 5, we use the weights from SE_2 to initialize the signal encoder parameters. In a similar fashion, we use SE_1 as initialization point of the signal encoder when we fine-tune the models to be tested with folds 6 to 10. This method allows us to pre-train, fine-tune and test our model in a more efficient way than pre-training 10 different models, one for each fold, while retaining complete separations between training and testing data.

The FCN classifier used to predict emotions has two hidden layers of sizes [1024, 512] with ReLU activation functions, and an output layer that projects the output to a single value. We fine-tune one model to predict arousal and another to predict valence. For each task, we fine-tune our model for 100 epochs using Adam optimization, with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and L_2 weight decay of 0.00001. We start with a learning rate of 0.0001, and decrease it every 45 epochs by a factor of 0.65. We keep using a dropout of 0.1 at the end of the input encoder, after the positional encoding, and inside the Transformer. We use dropout of 0.3 in the FCN classifier.

We used Ray Tune with BOHB, as we did on pre-training, to tune the learning rate, the learning rate schedule, the shape and dropout of the FCN classifier, and the L_2 weight decay.

5 Results

In our results, we use as metrics the mean accuracy and mean F1-score between positive and negative classes. We report the mean and confidence intervals of the metrics across our 10 folds of cross-validation. The confidence intervals are calculated using a t-distribution with 9 degrees of freedom, for a two-sided 95% confidence.

Table 3: Comparison of different methods on AMIGOS dataset

	Model	Subj. Ind.	Input Seg. Size	Arousal Acc.	Arousal F1	Valence Acc.	Valence F1
Various experiment protocols	Gaussian Naive Bayes [17]	Yes	20 seconds	-	0.551	-	0.545
	1D-CNN [8]	No	200 peaks	0.81	0.76	0.71	0.68
	2D-CNN [33]	Yes	Not segmented	0.83	0.76	0.82	0.80
	1D-CNN with LSTM [32]	Yes	Not segmented	-	-	0.81	0.80
	Convolutional autoencoder [22]	No	10 seconds	0.85	0.89	-	-
Our protocol	Pre-trained CNN [9]	No	10 seconds	$0.85 \pm 5.4e^{-3}$	$0.84 \pm 5.3e^{-3}$	$0.77 \pm 5.5e^{-3}$	$0.77 \pm 5.1e^{-3}$
	Pre-trained Transformer (ours)	No	10 seconds	$0.88 \pm 5.4e^{-3}$	$0.87 \pm 5.4e^{-3}$	$0.83 \pm 7.8e^{-3}$	$0.83 \pm 7.4e^{-3}$

5.1 Comparing Aggregation Methods and Segment Lengths

We report in Table 1 the performances of our approach for different strategical choices. Firstly, we compare different aggregation approaches to combine the contextualized representations at the output of the signal encoder, given to the FCN classifier. Secondly, we compare performances for different segment lengths used to divide the input signals.

Aggregation Method: we compared 4 strategies for aggregating representations, to be given as input to the FCN: max-pooling, average-pooling, using only the last representation e_s , and using only the embedding of the CLS token e_{CLS} (we call this strategy CLS). Max-pooling 1 and average-pooling 1 are the result of max-pooling and average-pooling across all representations, to obtain a single representation of size $d_{model} = 256$. Max-pooling 2 was optimized on the validation set: representations are reduced to a size of 64, divided into two groups, then max-pooling was applied on each group and the results concatenated to obtain a single representation of size 128. Average-pooling 2 was optimized on the validation set: representations are divided into 4 groups, average-pooling is applied on each group and the results concatenated to obtain a single representation of size 1024.

We see in Table 1 that the best results were obtained with average-pooling strategies and with CLS, with accuracies up to 0.88 for arousal, for example. In the following experiments, we will thus use CLS as our aggregation method. Indeed, although results are practically identical for CLS and average-pooling 2 (e.g. $0.88 \pm 5.4e^{-3}$ compared to $0.88 \pm 4.4e^{-3}$ accuracies for arousal), CLS has the advantage of being a commonly-used strategy for Transformers, which does not require any kind of tuning on validation data, contrary to average-pooling 2.

Segment length: we compare 3 different segment lengths for dividing ECG signals into input instances: 10, 20, and 40 second segments. We can see in Table 1 that shorter segments lead to better results on average, both for arousal and valence. For example for arousal, 10-second segments lead to an accuracy of $0.88 \pm 5.4e^{-3}$, compared to $0.87 \pm 5.6e^{-3}$ for 20-second segments, and $0.86 \pm 1.2e^{-2}$ for 40-second segments.

Two explanations emerge for this observation: firstly, since emotions are relatively volatile states, longer segmentation might cover fluctuating emotional states, thus making it harder to characterize emotion; secondly, longer segments should require more complex models (i.e. bigger Transformer and FCN), which are harder to train with the relatively restricted amount of labeled data in AMIGOS. Moreover, shorter segments are faster to process, allowing a high number of training epochs and smaller learning rates. In the following experiments, we will thus use 10-second segments.

5.2 Effectiveness of Pre-training

To demonstrate the effectiveness of our pre-training approach, we tested our architecture by fine-tuning our model on AMIGOS with all parameters randomly initialized, instead of using a pre-trained signal encoder (thus skipping step (a) of our process in Figure 1). As reported in Table 2, the pre-trained model is on average significantly better than the model with no pre-training, for both accuracy and F1-score. For example, for arousal, the pre-trained model reaches an average accuracy of $0.88 \pm 5.4e^{-3}$, compared to $0.85 \pm 5.6e^{-3}$ for the model with no pre-training. These results illustrate the benefits of pre-training Transformers for our task. Moreover, during our experiments, we observed that the model with no pre-training had a tendency to overfit quickly, which was not the case for the pre-trained model. Pre-training the model on many different datasets should increase its robustness to overfitting when fine-tuning on a specific dataset.

5.3 Comparisons With Other Approaches

We report in Table 3 various state-of-the-art results for emotion recognition from ECG signals on the AMIGOS dataset. The first section of the table contains results from works which all use different experiment protocols, such as different segment sizes, different separations of data into training and test sets, subject dependent and independent evaluations, etc. These results are therefore not directly comparable with one another, nor are they directly comparable with ours. Nevertheless, we report them to showcase the variety of state-of-the-art approaches published for this task, and give a relative idea of achieved performances on AMIGOS.

To compare our approach with another state-of-the-art approach as fairly as possible, it is required that both use exactly the same experiment protocol. For this, we fully retrained and tested the pre-trained CNN approach proposed by Sarkar and Etemad [9], with the experiment protocol we presented. To this end, we use the implementation provided by the authors¹. To ensure fair comparisons, the exact same data was used to pre-train, fine-tune, and test both our approach and also Sarkar and Etemad’s approach, for each fold of cross-validation.

We see in Table 3 that our approach achieves better performance on average than Sarkar and Etemad’s approach with the same experiment protocol, for both arousal and valence. For example, our approach achieves an F1-score of $0.83 \pm 7.4e^{-3}$ for valence, compared to $0.77 \pm 5.1e^{-3}$ for the pre-trained CNN. These results are statistically significant with $p < 0.01$ following a t-test.

This final set of results shows that our approach, and more generally self-supervised Transformer-based approaches, can be successfully applied to obtain contextualized representations from ECG signals for emotion recognition tasks.

6 Conclusions and Perspectives

In this paper, we investigate the use of transformers for recognizing arousal and valence from ECG signals. This approach used self-supervised learning for pre-training from unlabeled data, followed by fine-tuning with labeled data. Our experiments indicate that the model builds robust features for predicting arousal and valence on the AMIGOS dataset, and provides very promising results in comparison to recent state-of-the-art methods. This work showcases that self-supervision and attention-based models such as Transformers can be successfully used for research in affective computing.

Multiple perspectives emerge from our work. New pre-training tasks can be investigated: other methods such as contrastive loss or triplet loss might be more efficient with regards to the specificities of ECG signals, compared to masked points prediction which we used in this work. Extending our work to other input modalities (EEG, GSR, and even non-physiological inputs such as ambient sensors) and, in general, to process multimodal situations could prove useful for improving performances of emotion recognition. Finally, larger scale experiments, with new datasets captured in varied situations, will allow for a better understanding of the behaviour of our approach.

Acknowledgements: This work has been partially supported by the MIAI Multidisciplinary AI Institute at the Univ. Grenoble Alpes: (MIAI@Grenoble Alpes - ANR-19-P3IA-0003).

References

- [1] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [2] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All you Need. *Advances in Neural Information Processing Systems*, 30, 2017.
- [3] Shiyang Li, Xiaoyong Jin, Yao Xuan, Xiyong Zhou, Wenhui Chen, Yu-Xiang Wang, and Xifeng Yan. Enhancing the Locality and Breaking the Memory Bottleneck of Transformer on Time Series Forecasting. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [4] Genshen Yan, Shen Liang, Yanchun Zhang, and Fan Liu. Fusing Transformer Model with Temporal Features for ECG Heartbeat Classification. In *2019 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 898–905, November 2019.

¹<https://code.engineering.queensu.ca/pritam/SSL-ECG>

- [5] David Ahmedt-Aristizabal, Mohammad Ali Armin, Simon Denman, Clinton Fookes, and Lars Petersson. Attention Networks for Multi-Task Signal Analysis. In *2020 42nd Annual International Conference of the IEEE Engineering in Medicine Biology Society (EMBC)*, pages 184–187, July 2020.
- [6] Harish Haresamudram, Apoorva Beedu, Varun Agrawal, Patrick L. Grady, Irfan Essa, Judy Hoffman, and Thomas Plötz. Masked reconstruction based self-supervision for human activity recognition. In *Proceedings of the 2020 International Symposium on Wearable Computers*, pages 45–49, New York, NY, USA, September 2020. Association for Computing Machinery.
- [7] Lin Shu, Jinyan Xie, Mingyue Yang, Ziyi Li, Zhenqi Li, Dan Liao, Xiangmin Xu, and Xinyi Yang. A Review of Emotion Recognition Using Physiological Signals. *Sensors*, 18(7):2074, July 2018.
- [8] L. Santamaria-Granados, M. Munoz-Organero, G. Ramirez-González, E. Abdulhay, and N. Arunkumar. Using Deep Convolutional Neural Network for Emotion Detection on a Physiological Signals Dataset (AMIGOS). *IEEE Access*, 7:57–67, 2019.
- [9] P. Sarkar and A. Etemad. Self-Supervised Learning for ECG-Based Emotion Recognition. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3217–3221, May 2020.
- [10] T. Song, W. Zheng, P. Song, and Z. Cui. EEG Emotion Recognition Using Dynamical Graph Convolutional Neural Networks. *IEEE Transactions on Affective Computing*, 11(3):532–541, July 2020.
- [11] H. Becker, J. Fleureau, P. Guillotel, F. Wendling, I. Merlet, and L. Albera. Emotion Recognition Based on High-Resolution EEG Recordings and Reconstructed Brain Sources. *IEEE Transactions on Affective Computing*, 11(2):244–257, April 2020.
- [12] J. Shukla, M. Barreda-Angeles, J. Oliver, G. C. Nandi, and D. Puig. Feature Extraction and Selection for Emotion Recognition from Electrodermal Activity. *IEEE Transactions on Affective Computing*, pages 1–1, 2019.
- [13] Shuhao Chen, Ke Jiang, Haoji Hu, Haoze Kuang, Jianyi Yang, Jikui Luo, Xinhua Chen, and Yubo Li. Emotion Recognition Based on Skin Potential Signals with a Portable Wireless Device. *Sensors*, 21(3):1018, January 2021.
- [14] Mangesh Ramaji Kose, Mitul Kumar Ahirwal, and Anil Kumar. A new approach for emotions recognition through EOG and EMG signals. *Signal, Image and Video Processing*, 15(8):1863–1871, November 2021.
- [15] Sander Koelstra, Christian Muhl, Mohammad Soleymani, Jong-Seok Lee, Ashkan Yazdani, Touradj Ebrahimi, Thierry Pun, Anton Nijholt, and Ioannis Patras. DEAP: A Database for Emotion Analysis ;Using Physiological Signals. *IEEE Transactions on Affective Computing*, 3(1):18–31, January 2012.
- [16] Ramanathan Subramanian, Julia Wache, Mojtaba Khomami Abadi, Radu L. Vieriu, Stefan Winkler, and Nicu Sebe. ASCERTAIN: Emotion and Personality Recognition Using Commercial Sensors. *IEEE Transactions on Affective Computing*, 9(2):147–160, April 2018.
- [17] J. A. Miranda Correa, M. K. Abadi, N. Sebe, and I. Patras. AMIGOS: A Dataset for Affect, Personality and Mood Research on Individuals and Groups. *IEEE Transactions on Affective Computing*, pages 1–1, 2018.
- [18] Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep Contextualized Word Representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana, June 2018. Association for Computational Linguistics.
- [19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.
- [20] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A Simple Framework for Contrastive Learning of Visual Representations. In *Proceedings of the 37th International Conference on Machine Learning*, pages 1597–1607. PMLR, November 2020.
- [21] Bichen Wu, Chenfeng Xu, Xiaoliang Dai, Alvin Wan, Peizhao Zhang, Zhicheng Yan, Masayoshi Tomizuka, Joseph Gonzalez, Kurt Keutzer, and Peter Vajda. Visual Transformers: Token-based Image Representation and Processing for Computer Vision. *arXiv:2006.03677 [cs, eess]*, November 2020.
- [22] Kyle Ross, Paul Hungler, and Ali Etemad. Unsupervised multi-modal representation learning for affective computing with multi-corpus wearable data. *Journal of Ambient Intelligence and Humanized Computing*, October 2021.

- [23] Manon Macary, Marie Tahon, Yannick Estève, and Anthony Rousseau. On the Use of Self-Supervised Pre-Trained Acoustic and Linguistic Features for Continuous Speech Emotion Recognition. In *2021 IEEE Spoken Language Technology Workshop (SLT)*, pages 373–380, January 2021.
- [24] Siyang Song, Shashank Jaiswal, Enrique Sanchez, Georgios Tzimiropoulos, Linlin Shen, and Michel Valstar. Self-supervised Learning of Person-specific Facial Dynamics for Automatic Personality Recognition. *IEEE Transactions on Affective Computing*, pages 1–1, 2021.
- [25] Olivia Wiles, A. Sophia Koepke, and Andrew Zisserman. Self-supervised learning of a facial attribute embedding from video. *arXiv:1808.06882 [cs]*, August 2018.
- [26] Yong Li, Jiabei Zeng, Shiguang Shan, and Xilin Chen. Self-Supervised Representation Learning From Videos for Facial Action Unit Detection. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10916–10925, Long Beach, CA, USA, June 2019. IEEE.
- [27] Shuvendu Roy and Ali Etemad. Self-supervised Contrastive Learning of Multi-view Facial Expressions. In *Proceedings of the 2021 International Conference on Multimodal Interaction*, pages 253–257, Montréal QC Canada, October 2021. ACM.
- [28] Martin Gjoreski, Blagoj Mitrevski, Mitja Luštrek, and Matjaž Gams. An Inter-domain Study for Arousal Recognition from Physiological Signals. *Informatica*, 42(1), March 2018.
- [29] Değer Ayata, Yusuf Yaslan, and Mustafa E. Kamasak. Emotion Recognition from Multimodal Physiological Signals for Emotion Aware Healthcare Systems. *Journal of Medical and Biological Engineering*, 40(2):149–157, April 2020.
- [30] Y. Hsu, J. Wang, W. Chiang, and C. Hung. Automatic ECG-Based Emotion Recognition in Music Listening. *IEEE Transactions on Affective Computing*, 11(1):85–99, January 2020.
- [31] Lin Shu, Yang Yu, Wenzhuo Chen, Haoqiang Hua, Qin Li, Jianxiu Jin, and Xiangmin Xu. Wearable Emotion Recognition Using Heart Rate Data from a Smart Bracelet. *Sensors*, 20(3):718, January 2020.
- [32] R. Harper and J. Southern. A Bayesian Deep Learning Framework for End-To-End Prediction of Emotion from Heartbeat. *IEEE Transactions on Affective Computing*, pages 1–1, 2020.
- [33] S. Siddharth, T. Jung, and T. J. Sejnowski. Utilizing Deep Learning Towards Multi-modal Bio-sensing and Vision-based Affective Computing. *IEEE Transactions on Affective Computing*, pages 1–1, 2019.
- [34] Sean A. Fulop and Kelly Fitz. Algorithms for computing the time-corrected instantaneous frequency (reassigned) spectrogram, with applications. *The Journal of the Acoustical Society of America*, 119(1):360–371, January 2006.
- [35] Guang-Bin Huang, Qin-Yu Zhu, and Chee-Kheong Siew. Extreme learning machine: Theory and applications. *Neurocomputing*, 70(1-3):489–501, December 2006.
- [36] Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J. Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. Multimodal Transformer for Unaligned Multimodal Language Sequences. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6558–6569, Florence, Italy, July 2019. Association for Computational Linguistics.
- [37] Z. Wu, X. Zhang, T. Zhi-Xuan, J. Zaki, and D. C. Ong. Attending to Emotional Narratives. In *2019 8th International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 648–654, September 2019.
- [38] Jian Huang, Jianhua Tao, Bin Liu, Zheng Lian, and Mingyue Niu. Multimodal Transformer Fusion for Continuous Emotion Recognition. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3507–3511, May 2020.
- [39] Cong Cai, Yu He, Licai Sun, Zheng Lian, Bin Liu, Jianhua Tao, Mingyu Xu, and Kexin Wang. Multimodal Sentiment Analysis based on Recurrent Neural Network and Multimodal Attention. In *Proceedings of the 2nd on Multimodal Sentiment Analysis Challenge*, pages 61–67, Virtual Event China, October 2021. ACM.
- [40] Woan-Shiuan Chien, Huang-Cheng Chou, and Chi-Cun Lee. Self-assessed Emotion Classification from Acoustic and Physiological Features within Small-group Conversation. In *Companion Publication of the 2021 International Conference on Multimodal Interaction*, pages 230–239, Montreal QC Canada, October 2021. ACM.
- [41] Neo Wu, Bradley Green, Xue Ben, and Shawn O’Banion. Deep Transformer Models for Time Series Forecasting: The Influenza Prevalence Case. *arXiv:2001.08317 [cs, stat]*, January 2020.
- [42] Arjun Arjun, Aniket Singh Rajpoot, and Mahesh Raveendranatha Panicker. Introducing Attention Mechanism for EEG Signals: Emotion Recognition with Vision Transformers. In *2021 43rd Annual International Conference of the IEEE Engineering in Medicine Biology Society (EMBC)*, pages 5723–5726, November 2021.

- [43] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *arXiv:2010.11929 [cs]*, October 2020.
- [44] Behnam Behinaein, Anubhav Bhatti, Dirk Rodenburg, Paul Hungler, and Ali Etemad. A Transformer Architecture for Stress Detection from ECG. In *2021 International Symposium on Wearable Computers*, pages 132–134, Virtual USA, September 2021. ACM.
- [45] Dumitru Erhan, Aaron Courville, Yoshua Bengio, and Pascal Vincent. Why Does Unsupervised Pre-training Help Deep Learning? In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 201–208. JMLR Workshop and Conference Proceedings, March 2010.
- [46] Aparna Khare, Srinivas Parthasarathy, and Shiva Sundaram. Multi-Modal Embeddings Using Multi-Task Learning for Emotion Recognition. In *Interspeech 2020*, pages 384–388. ISCA, October 2020.
- [47] Wasifur Rahman, Md Kamrul Hasan, Sangwu Lee, AmirAli Bagher Zadeh, Chengfeng Mao, Louis-Philippe Morency, and Ehsan Hoque. Integrating Multimodal Information in Large Pretrained Transformers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2359–2369, Online, July 2020. Association for Computational Linguistics.
- [48] S. Siriwardhana, T. Kaluarachchi, M. Billingham, and S. Nanayakkara. Multimodal Emotion Recognition With Transformer-Based Self Supervised Feature Fusion. *IEEE Access*, 8:176274–176285, 2020.
- [49] Aparna Khare, Srinivas Parthasarathy, and Shiva Sundaram. Self-Supervised Learning with Cross-Modal Transformers for Emotion Recognition. In *2021 IEEE Spoken Language Technology Workshop (SLT)*, pages 381–388, January 2021.
- [50] George Zerveas, Srideepika Jayaraman, Dhaval Patel, Anuradha Bhamidipaty, and Carsten Eickhoff. A Transformer-based Framework for Multivariate Time Series Representation Learning. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 2114–2124, Virtual Event Singapore, August 2021. ACM.
- [51] Alexei Baeovski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. Wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations. *Advances in Neural Information Processing Systems*, 33:12449–12460, 2020.
- [52] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer Normalization. *arXiv:1607.06450 [cs, stat]*, July 2016.
- [53] Stamos Katsigiannis and Naeem Ramzan. DREAMER: A Database for Emotion Recognition Through EEG and ECG Signals From Wireless Low-cost Off-the-Shelf Devices. *IEEE Journal of Biomedical and Health Informatics*, 22(1):98–107, January 2018.
- [54] Athina Tzovara, Dominik R. Bach, Giuseppe Castegnetti, Samuel Gerster, Nicolas Hofer, Saurabh Khemka, Christoph W. Korn, Philipp C. Paulus, Boris B. Quednow, and Matthias Staib. PsPM-FR: SCR, ECG and respiration measurements in a delay fear conditioning task with visual CS and electrical US., August 2018.
- [55] Philipp C. Paulus, Giuseppe Castegnetti, and Dominik R. Bach. PsPM-HRM5: SCR, ECG and respiration measurements in response to positive/negative IAPS pictures, and neutral/aversive sounds, June 2020.
- [56] Dominik R. Bach, Samuel Gerster, Athina Tzovara, and Giuseppe Castegnetti. PsPM-RRM1-2: SCR, ECG, respiration and eye tracker measurements in response to electric stimulation or visual targets, September 2019.
- [57] Yanfang Xia, Filip Melinščak, and Dominik R. Bach. PsPM-VIS: SCR, ECG, respiration and eyetracker measurements in a delay fear conditioning task with visual CS and electrical US, July 2020.
- [58] Yan Wu, Ruolei Gu, Qiwei Yang, and Yue-jia Luo. How Do Amusement, Anger and Fear Influence Heart Rate and Heart Rate Variability? *Frontiers in Neuroscience*, 13:1131, 2019.
- [59] Richard Liaw, Eric Liang, Robert Nishihara, Philipp Moritz, Joseph E. Gonzalez, and Ion Stoica. Tune: A Research Platform for Distributed Model Selection and Training. *arXiv:1807.05118 [cs, stat]*, July 2018.
- [60] Stefan Falkner, Aaron Klein, and Frank Hutter. BOHB: Robust and Efficient Hyperparameter Optimization at Scale. In *Proceedings of the 35th International Conference on Machine Learning*, pages 1437–1446. PMLR, July 2018.