

Translating Emotions from EEG to Visual Arts

Piera Riccio^{1,2}[0000-0001-8602-8271], Francesco Galati³[0000-0001-6317-6298],
Maria A. Zuluaga³[0000-0002-1147-766X], Juan Carlos De
Martin⁴[0000-0002-7867-1926], and Stefano Nichele^{2,5}[0000-0003-4696-9872]

¹ ELLIS Unit Alicante Foundation, Spain

² Oslo Metropolitan University, Norway

³ EURECOM, France

⁴ Politecnico di Torino, Italy

⁵ Simula Metropolitan Centre for Digital Engineering, Norway
piera@ellisalicante.org; stenic@oslomet.no

Abstract. Exploring the potentialities of artificial intelligence (AI) in the world of arts is fundamental to understand and define how this technology is shaping our creativity. We propose a system that generates emotionally expressive paintings from EEG signals. The emotional information, encoded from the signals through a graph neural network, is inputted to a generative adversarial network (GAN), trained on a dataset of paintings. The design and experimental choices at the base of this work rely on the understanding that emotions are hard to define and formalize. Despite this, the proposed results witness an interaction between an AI system and a human, capable of producing an original and artistic re-interpretation of emotions. These results have a promising potential for AI technologies applied to visual arts.

Keywords: Art generation · Deep Learning · Computational Creativity · Affective Computing · Brain-Computer Interface.

1 Introduction

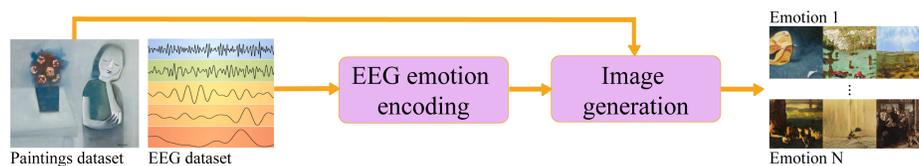


Fig. 1: General pipeline sketch of the proposed system. The emotions encoded in the EEG signals are transferred to a model devoted to the generation of expressive and original paintings. EEG image source: GOQii blog. Painting in the picture: *Room at Twilight* by Charles Blackman (1963).

The recent advances in artificial intelligence (AI) technologies enrich the debate concerning whether AI can make art. When approaching human-generated

artworks, we can speculate around the emotional sphere of the author or the ideas behind the work. Conversely, these thoughts are set aside when considering artworks produced by AI. We propose a human-machine interaction paradigm, in which a human agent provides the emotional signals that a machine could not produce otherwise, and the machine visualizes this information artistically. We aim at both the enhancement of machines' possibilities and the enhancement of human possibilities by their reciprocal interaction. In particular, the resulting emotional representation goes beyond the intentions and cultural background of the individual providing the signal. Therefore, such a representation would not be possible without this interaction.

We implement a brain-computer interface (BCI) based on a neural network pipeline that generates paintings from an inputted electroencephalographic (EEG) wave. The pipeline (in Figure 1) consists of two fundamental blocks: the first one encodes the emotional information in the EEG signals, while the second generates the paintings. For the training, two different datasets are employed. One of them is composed of recorded EEG waves, and the other is composed of paintings. Both the datasets have labels in the semantic space of emotions. We provide the code of this work in our repository⁶, and we encourage readers to try out some experiments.

In Section 2, we present the Background and Related Works, relevant for contextualizing our work. In Sections 3 and 4, we describe the needed datasets for this work and the implemented pipeline. In Sections 5 and 6, we provide an example experiment, with the evaluation of the obtained results. Finally, in Section 7, we discuss the artistic and cultural relevance of our results.

2 Background and Related Works

We propose an interdisciplinary work at the crossroads of two research areas: affective computing [27] and computational creativity [22]. Our framework generates paintings through a Generative Adversarial Network (GAN) conditioned on external data. GANs were originally introduced by Goodfellow et al. [9] and have had a remarkable impact on different fields, including art [39,7,29] and content generation [21]. StyleGAN2 [11] is considered among the state-of-the-art models to generate high-quality photorealistic images, and it is our reference architecture for this work. In the field of GAN conditioned with external data, the authors of [16] designed a model that generates paintings related to the composition epoch of an inputted piece of music. The music pieces and the paintings are related based on their historical epoch, which could be a simplistic and unrealistic correlation criterion.

To the best of our knowledge, no related work fulfills the translation of emotional states from EEG signals to paintings. In particular, we stress that emotions are vague and subjective phenomena that need a precise formalization to be processed by a machine [36]. Psychologists suggest mainly two formalization

⁶ <https://github.com/PieraRiccio/Emotions-fromEEG-toArt>

paradigms for emotions: a discrete [5] and a continuous one [32]. In our work, we rely on the discrete paradigm, with a finite set of basic emotions. EEG-based emotion recognition is an active area of research [42,40,41], as one of its main advantages is that EEGs represent inner phenomena that cannot be faked or controlled, differently from facial expressions, tone intonation, or words choice [15].

In [33] and [2], the authors propose systems to create expressive self-portraits of people. In the case of [33], users manually choose the emotion they want to see in their portrait. In the case of [2], emotions are detected through facial expressions in a video. This system implies the possibility of faking emotions by pretending to smile or to frown. Both works are artistically limited: the representation of the emotions consists of applying predefined styles on existing pictures without any generative element.

The utilization of EEGs comes with several challenges [20]. Among them, the accuracy limitations in recognizing a wide range of emotions, the needed number of electrodes in the EEG recording devices, the scarcity of benchmark EEG databases for emotion recognition, and the non-linearity and instability of the signals. In the last decade, new feature extractors [4,24] and benchmark datasets have been introduced. However, it is challenging to achieve high classification accuracy on large sets of emotions [43] and the possibility of utilizing simpler off-the-shelf recording devices is being explored [13,28,26]. In the context of paintings creation from EEG waves, we mention [38] and [25]. In these cases, the generated images do not relate to semantically meaningful features of the EEG waves. The works presented in [6] and [30] propose paintings based on emotions identified from EEG waves but in a simplistic manner. In the case of [6], the images are composed with simple lines, colors, or fractals; in [30], they are based on bird swarms. In both works, we observe a low inter-painting variability, with the emotions being mainly represented through predefined colors or shapes and with poor generative elements.

In Figure 2, we provide an overview of the paintings resulting from these Related Works. We will refer to this image in later sections, to assess the visual quality of our results.

3 Preparation of Datasets

The training of our pipeline requires an explicit connection between emotions in EEG signals and visual representations. This connection relies on shared labels between two different datasets. One dataset is a collection of paintings, and their labels correspond to the emotions they evoke; the other is a collection of recorded EEG signals, with labels representing the emotion that people felt during the recording.

In the case of EEG datasets for emotion recognition, one of the most popular options is the SEED-IV [43]. The authors recorded the signals on 15 subjects with a device made of 62 channels, stimulating four distinct emotions (happiness, fear, sadness, neutral) through 72 short videos. Several researchers [18,19,17,35,44]

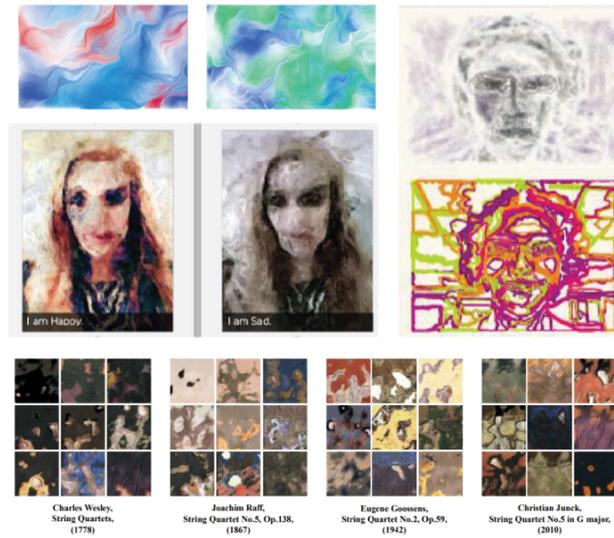


Fig. 2: Example pictures from some of the related works. Top right [30], medium right [33], left [2], bottom [16]. Images are taken from the relative papers, and from the website in the case of [30].

have reached high accuracy in recognizing emotions in this dataset. In this work, we perform a standard and reproducible experiment using the SEED-IV [43], but we also record signals with an off-the-shelf device (OpenBCI headband kit with eight dry-comb electrodes and Cyton board⁷). For the recordings, we follow the same procedure described by the authors of the SEED-IV dataset, and we extract differential entropy features [4] over five frequency bands (delta, theta, alpha, beta, and gamma).

In the dataset of paintings, we must have the same number of classes as the EEG dataset to allow matching physiological signals and artistic representations. In this work, we utilize the *WikiArt Emotions Dataset* [23], which contains thousands of paintings labeled with emotions according to a survey on several subjects. To create a matching between the four emotions in the SEED-IV, we select four different classes, namely *fear*, *sadness*, *anger* and *happiness*. The choice depends on the availability of samples in the dataset, and it does not extensively represent the spectrum of human emotions. We utilize a K-Means clustering technique to evaluate the color palettes of the paintings and to detect and delete stylistic outliers from the dataset. Some paintings appear in more than one class, but we keep them only in the least populated class in which they appear. To overcome class imbalance, we perform data augmentation in the classes containing fewer paintings, recursively cutting in four images bigger than 600x600 pixels until the corresponding class counts an adequate number of samples. We

⁷ <https://openbci.com/>

emphasize that this is a technical choice, and it implies an approximation in the correspondence between images and emotions, as different parts of the images may evoke different emotions. As a result of this augmentation technique, we obtain a dataset of around 2500 images, quite equally distributed among the four classes. These pre-processing steps are also available in our repository.

Figure 3 provides a high-level view of the images in this dataset. The associated color palettes for each class are estimated through the same clustering technique utilized for the outliers detection. In *happiness*, paintings have bright and gaudy colors, while *fear* presents a bleak and dark color palette. In *anger*, we observe a prevalence of red and warm shades, which leave space for blue and cold shades in *sadness*.



Fig. 3: Grids showing the images generated in the dataset, divided according to their labels. From left to right: *anger*, *happiness*, *sadness*, *fear*.

4 Pipeline

We provide a detailed overview of the proposed pipeline in Figure 4. EEG signals are highly variable among different subjects. Models trained on EEG signals are commonly trained in a subject-dependent fashion (trained and tested on the same individual). In this work, we have performed this kind of training, starting from the state-of-the-art model on the SEED-IV dataset. This model, *Regularized Graph Neural Network* (RGNN) [44], processes the information on the topological disposition of the electrodes on the scalp. We perform its training separately from the rest of the pipeline. In Figure 4, the purpose of the Encoder E (RGNN) and its loss \mathcal{L}_E is to extrapolate relevant features from the EEG signals $\mathbf{e} \in \mathbb{R}^{62 \times 5}$, where 62 are the space channels and 5 the frequency bands. The extracted features $\mathbf{v} \in \mathbb{R}^{50}$ are used as input to StyleGAN2 [11], which concatenates them with random latent vectors $\mathbf{z} \in \mathbb{R}^{512}$. The insertion of \mathbf{v} allows to train StyleGAN2 conditionally and generate paintings with relevant characteristics when representing different emotions. To make the conditioning process more effective, we should input the predicted emotion from the EEG signals (one-hot encoded vectors). However, we aim to represent the complexity and the richness of the human emotional sphere. To get closer to this effect, we utilize latent vectors with a higher dimension (50 entries). Such vectors, although summarizing the emotion in the EEG signals, also preserve their uniqueness.

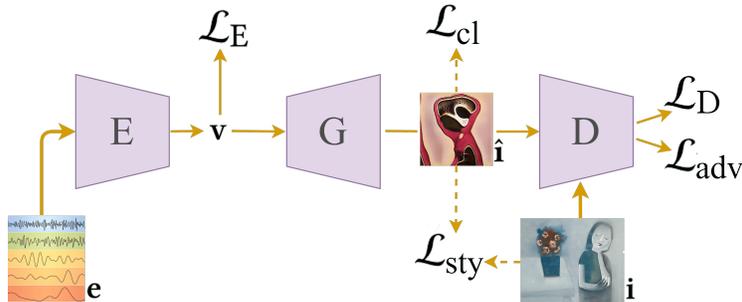


Fig. 4: Detailed visual sketch representing the different elements and losses in the pipeline, when trained on a subject of the SEED-IV dataset.

Being constituted by a GAN model, the training of this pipeline is adversarial [9]. This means that the generation block G aims to fool the discriminator D , another neural network in charge to distinguish between real paintings and generated paintings $\hat{\mathbf{i}}/\mathbf{i} \in \mathbb{R}^{3 \times W \times H}$ (see losses \mathcal{L}_D and \mathcal{L}_{adv}). The utilized model, StyleGAN2, is generally trained on datasets containing tens of thousands of images. Given the need for human effort to rigorously match emotional labels to paintings, it is unlikely to access datasets of this dimension. To address this problem, we refer to the work in [12], in which the authors consider implementing an adaptive discriminator augmentation to create a model (StyleGAN2ADA) that can be trained successfully also on smaller datasets. The remaining elements of Figure 4 are described in the following section.

4.1 Extra losses

To shorten the time required by the training process, we provide the possibility of utilizing two extra losses (introduced in [16]): \mathcal{L}_{cl} and \mathcal{L}_{sty} , visible in Figure 4. \mathcal{L}_{cl} is obtained through an auxiliary classifier that identifies the emotion conveyed by the produced paintings and checks whether it is equal to the same emotion as the label of the inputted EEG wave. It is based on AlexNet architecture [14], pre-trained on ImageNet database [3] and retrained on our dataset (using a proportion 80:20), via transfer learning [8]. \mathcal{L}_{sty} , instead, provides a comparison between the style of a generated image with the style of a real image with the same label as the inputted EEG wave: the more similar they are, the lower this loss. For this purpose, we utilize a pre-trained VGG-19 [34]. Overall, the three losses can be combined into one generator loss:

$$\mathcal{L}_G = \mathcal{L}_{adv} + \lambda_1 \mathcal{L}_{cl} + \lambda_2 \mathcal{L}_{sty} \quad (1)$$

where λ_1 and λ_2 are scalar hyper-parameter factors.

5 Experiments and results

Different experiments can be carried out with this pipeline, by utilizing different EEG signals (e.g. self-recorded), by changing the resolution of the generated images, or by adding the extra losses in the training. Our experiments are performed with a Tesla V100-SXM2-16GB GPU, associated with a 25GB RAM. In the following section, we provide a detailed overview of the highest resolution experiment that we have carried out. However, other experiments have led to these main findings:

- The implemented pipeline is adaptable to different computational needs. The visual results get better when the resolution settings are higher, but they are also meaningful at lower resolutions.
- The RGNN architecture can be trained on signals recorded with other devices. In particular, we have recorded signals utilizing an off-the-shelf device with eight electrodes. This characteristic gives the possibility of utilizing this pipeline in practical contexts without having professionals handling the recording device.
- The utilization of extra losses does not provide a visual enhancement of the obtained images, but it allows the model to converge faster. For this reason, we suggest utilizing extra losses when there are time constraints on the training period.
- Transfer learning can be a useful technique under some conditions. For example, if the reader wishes to train the pipeline on a smaller dataset of paintings, we suggest using as a starting point of the training the weights learned in an experiment of the SEED-IV.

In our repository, we provide different experiments with short explanations and checkpoints of the weights.

Ideally, the output paintings visually express the emotion encoded in the inputted EEG wave. However, there are no objective metrics that define such property in a painting. The evaluation of the results is, therefore, divided into different stages:

- Evaluation of the Frechet-Inception Distance (FID) metric [10] over iterations. This metric estimates the statistical similarity between features extracted through the Inception v3 model [37] on both the generated paintings and the dataset. These features are approximated with two Gaussian distributions, and the Frechet distance among them is computed as reported in Equation 2 (μ_1 and μ_2 being the means of the two distributions, σ_1 and σ_2 their variances).

$$d^2 = \|\mu_1 - \mu_2\|^2 + Tr(\sigma_1 + \sigma_2 - 2 * \sqrt{(\sigma_1 * \sigma_2)}) \quad (2)$$

- Qualitative evaluation of the conditioning: we generate grids of paintings for every class and we estimate their color palettes, comparing them among each other and with the palettes estimated on the dataset of paintings (Figure 3).

- Detailed analysis regarding some of the generated paintings, by looking at some hand-picked results in terms of details, shapes and colors.
- Statistical analysis of the results of an online survey to which several users have taken part.

5.1 Example experiment

In this example experiment, we generate images at 512x512 pixels, without using extra losses. The RGNN is trained on subject 15 of the SEED-IV dataset. In Figure 5, we provide an overview of the FID metric. As expected, its value diminishes with the progression of the training, showing that the generated images become statistically closer to the images in the dataset. The value of this metric converges below 50.

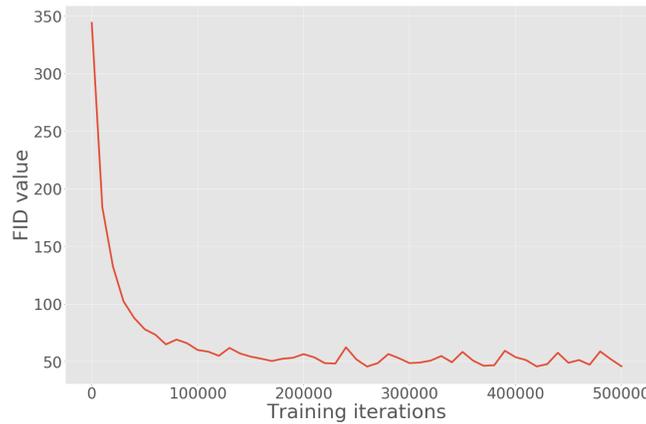


Fig. 5: FID metric over iterations for the experiment on Subject 15 of the SEED-IV dataset at resolution 512x512 pixels.

In Figure 6, we provide a high-level overview of the generated images. Comparing these grids with the ones of the dataset, shown in Figure 3, we observe that the pipeline provided stylistically similar images in every class. In addition, we emphasize that the implemented pipeline produces paintings that also show expressive content and shapes. We provide some details in Figure 7; such images show high potential for the artistic possibilities and applications of this pipeline, which qualitatively outperforms the results obtained in [16] and in the other related works shown in Figure 2. Compared to [30,33,2], our paintings stand out for two main reasons: their content is generated ad-hoc on inputted EEG signals, and we have not pre-defined stylistic norms to express emotions. Our results are heterogeneous and original, providing a re-interpretation of human emotions that we have not observed in any previous work in the technical or artistic literature.

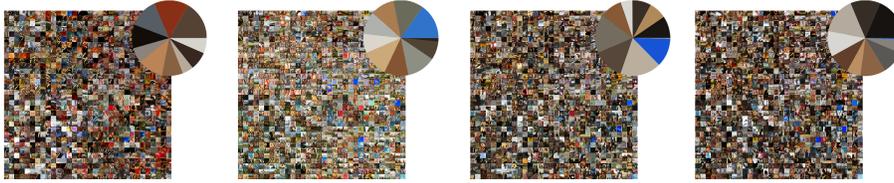


Fig. 6: Grids showing the images generated in the training with the SEED-IV dataset, divided according to their labels. From left to right: *anger*, *happiness*, *sadness*, *fear*.

6 Results assessment: online survey

Given the purely subjective nature of the phenomena underlying this research, we push the evaluation a step further, asking people to rate the emotions they perceive from the paintings. We stress that some paintings do not necessarily fall into the idea of the emotion from which they were generated: they may either evoke something different or evoke more emotions at once. However, we assume that if we wanted to diminish this effect, we could have utilized the emotional prediction of the EEG signals as input to the StyleGAN2 instead of utilizing latent vectors of higher dimensions, as explained in Section 3. Given this design choice, we have organized an online survey allowing people to rate emotions on a scale from 0 to 5 for each painting. In this way, people can select more emotions together, specifying their intensity. Users are also allowed to put all 0s in case none of the proposed emotions is somewhat related to the image. The survey⁸ proposes 96 different paintings, picked from different experiments. The following analysis considers 10285 answers given by 532 different users that have navigated through this website. In this survey, we have not collected any information regarding the demographics or background knowledge of the participants. We mark answers with 0s on every emotion as non-valid. Although these answers may refer to paintings that evoke different emotions rather than the four listed, we delete non-valid answers from our analysis, as considering them would require further measures that fall out of the scope of this work. Only 7% of the answers are non-valid, suggesting that users perceive emotional content in the proposed paintings.

In Figure 8, we provide an overview of the accuracy of users in recognizing the emotions in the paintings. In Figure 8a, the percentages are computed on each vote separately. A vote is considered right when the emotion that has received the highest mark is the same as the emotion of the corresponding EEG signal. For *happiness*, *fear* and *sadness* we have a rather satisfying sensitivity (around 58%), that is much lower for *anger* (around 22%). However, *anger* is the least perceived emotion across all the votes and, in fact, only rather small percentages of other votes are for this class (with high specificity). On the contrary, the emotions of *sadness* and *fear* are perceived more commonly across these votes. Figure 8

⁸ <https://pierariccio.pythonanywhere.com/>



(a) Some of the generated paintings in the *happiness* class.



(b) Some of the generated paintings in the *anger* class.

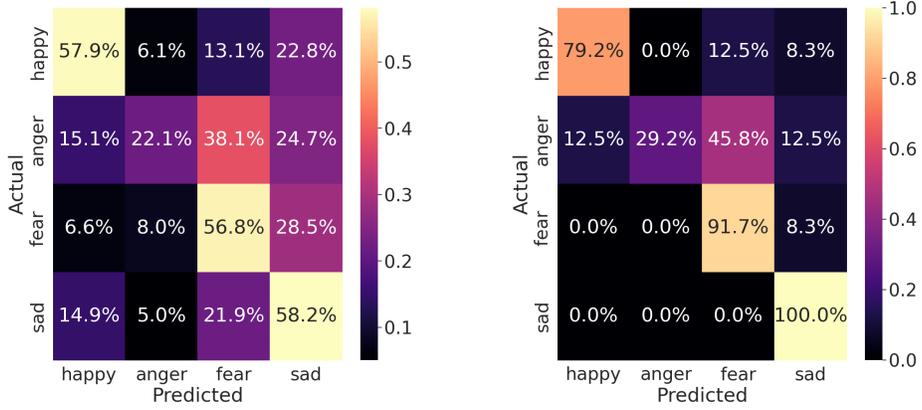


(c) Some of the generated paintings in the *sadness* class.



(d) Some of the generated paintings in the *fear* class.

Fig. 7: Details of hand-picked results from the different classes.



(a) Confusion matrix considering all the votes separately. (b) Confusion matrix grouping the votes on the same paintings.

Fig. 8: Confusion matrices regarding the answers obtained from the online survey.

generally shows a partial overlapping between *fear* and *anger*. Such overlapping seems to be asymmetric: *fear* is confused with *anger*, but the opposite does not seem to happen so often. Interestingly, for all classes, the majority of paintings that are perceived as belonging to emotion do belong to the perceived emotion. In Figure 8b, the votes are grouped and summed according to the painting they refer to. In this way, we show how many paintings in each class are, on average, correctly classified by users. It is outstanding that all the paintings (100%) belonging to the class *sadness* are perceived as such, followed by *fear* with a percentage of 91.7% and *happiness* with almost 79.2%. In the case of *anger*, only 29.2% of the paintings are classified correctly. Although this emotion is the least perceived by users, the totality (100%) of the paintings that are classified as *anger* do belong to this class.

As further analysis of these results, in Figure 9, we provide a ranking of 10 paintings for each class, based on the percentage of users that have voted the images as belonging to that class (the misplaced paintings are highlighted in red). The majority of the paintings in this ranking are in the correct class and with generally high percentages of votes. The first and the ninth paintings in *fear*, as well as the sixth one in *sadness*, actually belong to class *anger*. These three examples reiterate that, when the conveyed feeling is perceived as negative, it is easier to associate it with sadness or fear, especially when the painting lacks strong red shades. In addition, the sixth painting in *anger* presents strong red tones, but it originally belongs to class *happiness*. The voting percentages for this painting are: 19.4% in class *happiness*, 34.7% in class *anger*, 36.7% in *fear* and 9.2% in *sadness*. The wide distribution of these votes suggests a perceived ambiguity.

	1	2	3	4	5	6	7	8	9	10
Happiness	 88.7%	 87.1%	 82.3%	 75.5%	 74.6%	 74.3%	 64.7%	 63.7%	 63.3%	 61.9%
Anger	 44.4%	 40.6%	 39.6%	 39.1%	 36.5%	 34.7%	 34.4%	 31.8%	 30.9%	 30.8%
Fear	 72.1%	 67.6%	 67.0%	 67.0%	 66.7%	 66.7%	 66.3%	 64.1%	 61.9%	 61.5%
Sadness	 79.6%	 77.9%	 77.2%	 73.8%	 71.4%	 65.7%	 63.9%	 62.9%	 60.0%	 60.0%

Fig. 9: Ranking of paintings in each class, according to the percentage of users that have voted them. The paintings highlighted in red belong to a different class.

Overall, the results of this survey are satisfying and relevant for the scope of this research. Considering that the accuracy of the EEG encoder model was not 100%, some of the EEG signals are misclassified before being inputted into the generative model. As a natural consequence, some paintings convey a different emotion than the intended one. Despite this, this survey shows a general alignment between the emotion in the original EEG signal and the emotion that is perceived by users.

7 Discussion

The existence of an effective and recognizable emotional content in these paintings is a significant result, which can stimulate further reflections spanning a broad spectrum of research disciplines, which depart from the purely technical level.

From an aesthetic point of view, the proposed paintings are generated by a GAN trained on WikiArt Emotions, a dataset mainly containing works from Western artists. As a future work idea, we suggest merging the utilized paintings with other sources. In particular, it would be interesting to experiment with the utilization of a broader dataset, generating images devoid of artistic biases intrinsic in the utilized dataset.

Despite this cultural limitation, the present work can be placed in the context of critical disabilities studies from two perspectives. First, this pipeline is a tool that can be used for art therapy, detaching emotional and artistic expression from the body and technical skills. Furthermore, the shapes of the characters

in these paintings are far from the aesthetic canons that have often stigmatized imperfections within works of art. In this sense, it is interesting to question the role of AI in establishing new artistic canons, or reinforcing those incorporated into the art world starting from the modern art era.

From an anthropological and philosophical perspective, the paintings we propose in this paper represent individual emotions, but they do not correlate with the specific intentions of the people providing the EEG signals. In this sense, it is interesting to speculate on one of the most common questions in classical philosophy concerning the existence of rationality external to the human, which follows a logic above individuals. This work proposes a pipeline that attempts to simulate an artistic expression external to the human one, and it can be food for thought on the complexity of the interaction between AI and humans in our cultural system.

For a further deepening on these topics, we refer readers to our paper [31], in which we have discussed our results from a philosophical perspective, focusing on fairness, inclusion, and aesthetics.

8 Conclusions

In this work, we have designed a brain-computer interface (BCI) that generates paintings according to the emotions detected in EEG signals. We have based our design choices on the rationale that, often, a single word (or label) is not enough to describe a feeling. Our pipeline consists of two different components, and the training needs a dataset made of paintings and one made of EEG signals, both labeled with emotions. The signals are processed by an encoder model that produces latent vectors inputted to another model generating the paintings. We have studied one of the most popular EEG datasets for emotion recognition (SEED-IV) and trained one of the state-of-the-art deep learning models (RGNN) on this dataset, also providing experimental training and results on recordings that we took with an OpenBCI device. We have pre-processed and analyzed a dataset made of paintings with emotional labels (WikiArt Emotions), performing outliers detection and elimination, and data augmentation to fight class imbalance. We have successfully integrated StyleGAN2ADA into our pipeline. Finally, we have provided a quantitative and qualitative evaluation of our results. To assess the expressive properties of these paintings, we have set up an online survey, which has confirmed the success of our experiments.

9 Acknowledgments

This work has been partially funded by the DeepCA project (Norwegian Research Council, Young Research Talent program grant agreement 286558) and the FELT project [1] (Norwegian Artistic Research Programme).

References

1. Bergaust, K., Nichele, S.: FeLT-the futures of living technologies pp. 90–97 (2019). <https://doi.org/10.14236/ewic/POM19.14>
2. Colton, S., Valstar, M.F., Pantic, M.: Emotionally aware automated portrait painting. In: Proceedings of the 3rd International Conference on Digital Interactive Media in Entertainment and Arts. pp. 304–311. DIMEA '08, Association for Computing Machinery, New York, NY, USA (2008). <https://doi.org/10.1145/1413634.1413690>
3. Deng, J., Dong, W., Socher, R., Li, L., Kai Li, Li Fei-Fei: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. pp. 248–255 (2009). <https://doi.org/10.1109/CVPR.2009.5206848>
4. Duan, R., Zhu, J., Lu, B.: Differential entropy feature for eeg-based emotion classification. In: 2013 6th International IEEE/EMBS Conference on Neural Engineering (NER). pp. 81–84 (2013). <https://doi.org/10.1109/NER.2013.6695876>
5. Ekman, P.: An argument for basic emotions. *Cognition and Emotion* **6**(3-4), 169–200 (1992). <https://doi.org/10.1080/02699939208411068>
6. Ekster, G.: The cognichrome (2018), <http://www.cognichrome.com/>. Last accessed 22 Feb 2022
7. Elgammal, A., Liu, B., Elhoseiny, M., Mazzone, M.: Can: Creative adversarial networks generating “art” by learning about styles and deviating from style norms. In: Proceedings of the 8th International Conference on Computational Creativity, ICCO 2017. Georgia Institute of Technology (2017)
8. Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press (2016)
9. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A.C., Bengio, Y.: Generative adversarial nets. In: NIPS. pp. 2672–2680 (2014)
10. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: GANs trained by a two time-scale update rule converge to a local nash equilibrium. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. pp. 6629–6640. NIPS'17, Curran Associates Inc., Red Hook, NY, USA (2017)
11. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8107–8116 (2020). <https://doi.org/10.1109/CVPR42600.2020.00813>
12. Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., Aila, T.: Training generative adversarial networks with limited data. In: NeurIPS 2020 (2020)
13. Katsigiannis, S., Ramzan, N.: 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 (2018). <https://doi.org/10.1109/JBHI.2017.2688239>
14. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Advances in Neural Information Processing Systems. vol. 25. Curran Associates, Inc. (2012)
15. Larradet, F., Niewiadomski, R., Barresi, G., Caldwell, D.G., Mattos, L.S.: Toward emotion recognition from physiological signals in the wild: Approaching the methodological issues in real-life data collection. *Frontiers in Psychology* **11**, 1111 (2020). <https://doi.org/10.3389/fpsyg.2020.01111>

16. Lee, C.C., Lin, W.Y., Shih, Y.T., Kuo, P.Y.P., Su, L.: Crossing You in Style: Cross-Modal Style Transfer from Music to Visual Arts, pp. 3219–3227. Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3394171.3413624>
17. Li, H., Jin, Y.M., Zheng, W.L., Lu, B.L.: Cross-subject emotion recognition using deep adaptation networks. In: Cheng, L., Leung, A.C.S., Ozawa, S. (eds.) *Neural Information Processing*. pp. 403–413. Springer International Publishing, Cham (2018). https://doi.org/10.1007/978-3-030-04221-9_36
18. Li, Y., Wang, L., Zheng, W., Zong, Y., Qi, L., Cui, Z., Zhang, T., Song, T.: A novel bi-hemispheric discrepancy model for eeg emotion recognition. *IEEE Transactions on Cognitive and Developmental Systems* pp. 1–1 (2020). <https://doi.org/10.1109/TCDS.2020.2999337>
19. Li, Y., Zheng, W., Zong, Y., Cui, Z., Zhang, T., Zhou, X.: A bi-hemisphere domain adversarial neural network model for eeg emotion recognition. *IEEE Transactions on Affective Computing* pp. 1–1 (2018). <https://doi.org/10.1109/TAFFC.2018.2885474>
20. Liu, Y., Sourina, O., Nguyen, M.K.: Real-time eeg-based emotion recognition and its applications. In: Gavrilova, M.L., Tan, C.J.K., Sourin, A., Sourina, O. (eds.) *Transactions on Computational Science XII: Special Issue on Cyberworlds*. pp. 256–277. Springer Berlin Heidelberg, Berlin, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22336-5_13
21. Mao, X., Li, Q.: *Generative Adversarial Networks for Image Generation*. Springer Nature (2021)
22. McCormack, J., d’Inverno, M.: On the future of computers and creativity. In: *AISB 2014 Symposium on Computational Creativity*, London. Citeseer (2014)
23. Mohammad, S., Kiritchenko, S.: WikiArt emotions: An annotated dataset of emotions evoked by art. In: *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. European Language Resources Association (ELRA), Miyazaki, Japan (2018)
24. Nawaz, R., Cheah, K.H., Nisar, H., Yap, V.V.: Comparison of different feature extraction methods for eeg-based emotion recognition. *Biocybernetics and Biomedical Engineering* **40**(3), 910–926 (2020). <https://doi.org/10.1016/j.bbe.2020.04.005>
25. NeuroSky: Eeg art - neurosky (2015), <http://neurosky.com/2015/11/beautiful-brainwaves-creating-eeg-art/>. Last accessed 22 Feb 2022
26. Pham, T.D., Tran, D.: Emotion recognition using the emotiv epoc device. In: Huang, T., Zeng, Z., Li, C., Leung, C.S. (eds.) *Neural Information Processing*. pp. 394–399. Springer Berlin Heidelberg, Berlin, Heidelberg (2012). https://doi.org/10.1007/978-3-642-34500-5_47
27. Picard, R.W.: *Affective computing*. MIT Press (1997)
28. Plass-Oude Bos, D.: Eeg-based emotion recognition. *The Influence of Visual and Auditory Stimuli* (2006)
29. Praramadhan, A.A., Saputra, G.E.: Cycle generative adversarial networks algorithm with style transfer for image generation. *arXiv preprint arXiv:2101.03921* (2021)
30. Quark, R.: *The art of feeling* (2017), <https://www.randomquark.com/work/the-art-of-feeling>. Last accessed 22 Feb 2022
31. Riccio, P., Bergaust, K., Christensen-Scheel, B., Zuluaga, M.A., De Martin, J.C., Nichele, S.: AI-based artistic representation of emotions from EEG signals: a discussion on fairness, inclusion, and aesthetics. *arXiv preprint* (2022). <https://doi.org/10.48550/arXiv.2202.03246>

32. Russell, J.A.: A circumplex model of affect. *Journal of personality and social psychology* **39**(6), 1161—1178 (1980). <https://doi.org/10.1037/h0077714>
33. Salevati, S., DiPaola, S.: A creative artificial intelligence system to investigate user experience, affect, emotion and creativity. In: *Proceedings of the Conference on Electronic Visualisation and the Arts*. pp. 140—147. EVA '15, BCS Learning & Development Ltd., Swindon, GBR (2015). <https://doi.org/10.14236/ewic/eva2015.13>
34. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014)
35. Song, T., Zheng, W., Song, P., Cui, Z.: Eeg emotion recognition using dynamical graph convolutional neural networks. *IEEE Transactions on Affective Computing* **11**(3), 532–541 (2020). <https://doi.org/10.1109/TAFFC.2018.2817622>
36. Strongman, K.T.: *The psychology of emotion: From everyday life to theory*. Wiley Chichester (2003)
37. Xiaoling Xia, Cui Xu, Bing Nan: Inception-v3 for flower classification. In: *2017 2nd International Conference on Image, Vision and Computing (ICIVC)*. pp. 783–787 (2017). <https://doi.org/10.1109/ICIVC.2017.7984661>
38. Xiong, J.: *Mind art* (2014), <https://www.behance.net/gallery/22054167/Mind-Art->. Last accessed 22 Feb 2022
39. Yi, R., Liu, Y.J., Lai, Y.K., Rosin, P.L.: APDrawingGAN: Generating artistic portrait drawings from face photos with hierarchical GANs. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 10735–10744 (2019). <https://doi.org/10.1109/CVPR.2019.01100>
40. Zhang, J., Chen, P., Nichele, S., Yazidi, A.: Emotion recognition using time-frequency analysis of eeg signals and machine learning. In: *2019 IEEE Symposium Series on Computational Intelligence (SSCI)*. pp. 404–409. IEEE (2019). <https://doi.org/10.1109/SSCI44817.2019.9003057>
41. Zhang, J., Li, J., Nichele, S.: Instantaneous mental workload recognition using wavelet-packet decomposition and semi-supervised learning. In: *2019 IEEE Symposium Series on Computational Intelligence (SSCI)*. pp. 410–416. Institute of Electrical and Electronics Engineers (IEEE) (2020). <https://doi.org/10.1109/SSCI44817.2019.9002997>
42. Zhang, J., Yin, Z., Chen, P., Nichele, S.: Emotion recognition using multi-modal data and machine learning techniques: A tutorial and review. *Information Fusion* **59**, 103–126 (2020). <https://doi.org/10.1016/j.inffus.2020.01.011>
43. Zheng, W.L., Liu, W., Lu, Y., Lu, B.L., Cichocki, A.: Emotionmeter: A multimodal framework for recognizing human emotions. *IEEE Transactions on Cybernetics* **49**(3), 1110–1122 (2019). <https://doi.org/10.1109/TCYB.2018.2797176>
44. Zhong, P., Wang, D., Miao, C.: EEG-based emotion recognition using regularized graph neural networks. *IEEE Transactions on Affective Computing* pp. 1–1 (2020). <https://doi.org/10.1109/TAFFC.2020.2994159>