ECE 590.03 Final Project Report

Yuchen Cao, Sophie Williams

Introduction and Background

Our brains are continuously processing and responding to an endless stream of visual stimuli, making it increasingly important to understand how our brains react to different types of imagery. Through this project, we aimed to investigate how recorded brain activity can be used to identify the subjective experience of individuals when they view different types of images. By analyzing these signals, we sought to uncover the neural mechanisms underlying aesthetic experiences and better understand how our brains process and respond to visual stimuli. Additionally, we hoped to learn more about the steps of running BCI experiments by conducting an experiment on our own.

We chose to use electroencephalography (EEG) to record brain activity as it is a commonly used non-invasive technique that measures the electrical activity of the brain through electrodes placed on the scalp. This method has been widely used to study the neural correlates of perception, cognition, and emotion. To collect and analyze EEG data, we utilized the BCI2000 data acquisition software, which is a widely-used platform for neuroscience research. This software provides many tutorials for setting up EEG experiments, collecting data, and analyzing results.

Methods

Experiment design

In order to investigate the brain's reactions to different types of imagery, we decided to conduct our own EEG experiment using the BCI2000 data acquisition software. Using readily available tutorials online, we were able to create our own stimulus presentation task. For simplicity sake, we decided to test the differences between brain responses to pleasant and unpleasant imagery with the use of the Geneva Affective PicturE Database (GAPED) which is a database of visual emotion stimuli. The database images are sorted into positive and negative categories. Positive categories included mainly images of human and animal babies as well as nature sceneries. The negative categories included images of spiders, snakes, and scenes of inhumane conditions for humans and animals (Fig 1).

Figure 1: Positive and Negative Image Stimuli

To create our stimulus presentation task, we selected ten images of each category and used a MATLAB script provided in the BCI2000 documentation to generate a parameters file that could be used by the software to sync stimulus presentation with EEG signal recording. Once the parameters file was created, we moved on to conducting our experiment by following a similar model to the steps outlined in the P300 BCI User Tutorial. We conducted two full trials, one for each partner. After placing the EEG cap on the subject, we created connections between the electrodes and scalp with the use of conductive gel and ensured all impedances were below 30kOhms to reduce noise in the signal (Fig 2).

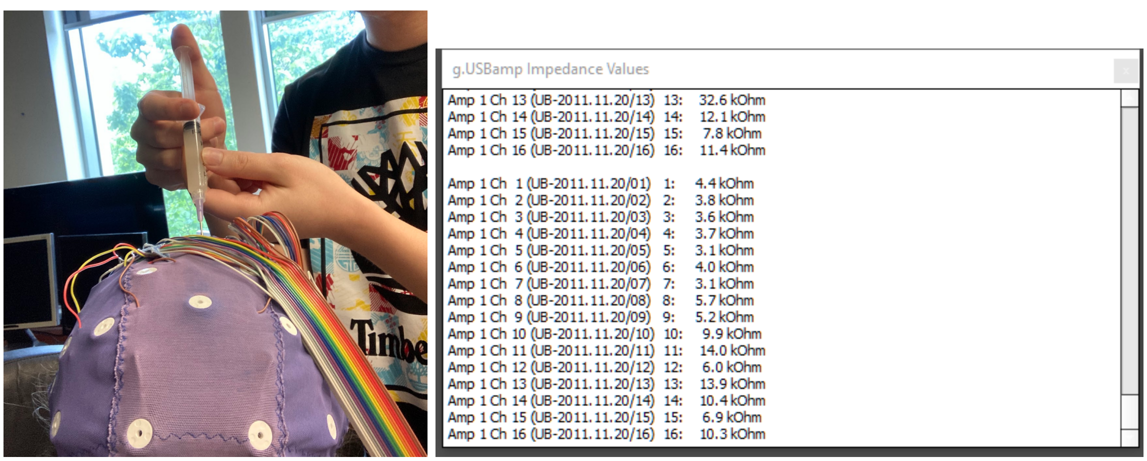

Figure 2: Electroconductive gel can improve noisy signals by decreasing the impedance between the electrode and scalp.

To test that the EEG signals were being read correctly, we attempted to create muscle artifacts through eye movement and muscle clenching. The distortion of the signal in time with these actions indicated that the electrodes were an indication that electrodes were taking correct measurements (Fig 3).

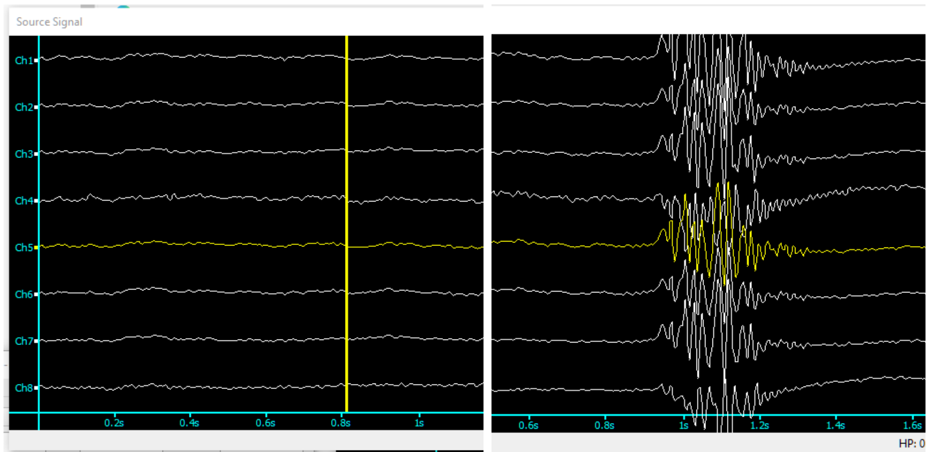

Figure 3: A difference between calm brain signals (left) and brain signals with muscle artifacts (right) indicates that the EEG data is being correctly reported.

Once the functionality of the system was confirmed, the stimulus presentation task was run. The images were presented in a random order for two seconds at a time with a brief pause between images. A 256 Hz sample rate was used resulting in 512 data points per channel for each stimuli. Each subject conducted five trials with small breaks in between and collected approximately 400 responses per subject (383 for Yuchen and 391 for Sophie). Data was collected from the first 16 electrodes labeled on the map in Figure 4 and exported to a file for further analysis.

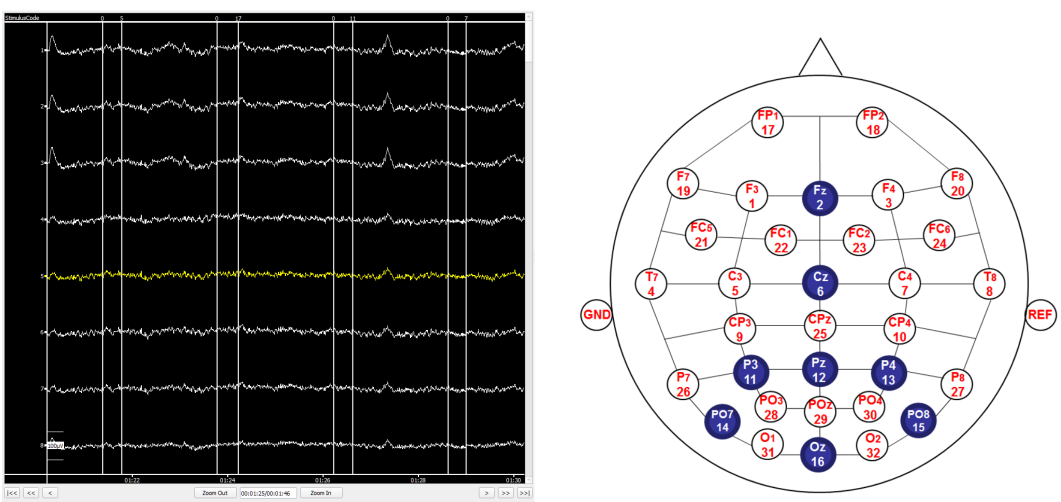

Figure 4: The output of each electrode during a trial of the stimulus presentation task (left) and map of electrode placements (right). The displayed outputs correspond to electrodes 1 through 8 but data was collected from electrodes 9-16 as well.

Constructing Models

The second task was to use the collected data in order to create a classifier to detect whether the recipient was viewing a pleasant or unpleasant image stimuli. Because the collected data was sequence data, we used sequence neural network models like RNN, GRU, LSTM to build the classifier, and train and test the models on each recipient.

RNN[1]

Recurrent Neural Networks (RNNs) are a class of artificial neural networks specifically designed for processing sequential data. They are characterized by their ability to maintain an internal state or "memory" of past inputs, enabling them to capture temporal dependencies within the data. RNNs have been widely used in various applications such as natural language processing, speech recognition, and time-series forecasting. However, despite their effectiveness in modeling sequential data, RNNs can suffer from vanishing and exploding gradient problems, limiting their capacity to learn long-range dependencies.

GRU[2]

Gated Recurrent Units (GRUs) are a type of recurrent neural network architecture that addresses some of the limitations of traditional RNNs. GRUs incorporate gating mechanisms that control the flow of information within the network. These mechanisms allow the model to selectively remember or forget information, improving its ability to capture long-range dependencies and mitigating the vanishing gradient problem.

LSTM[3]

Long Short-Term Memory (LSTM) networks are a popular type of recurrent neural network designed to address the shortcomings of traditional RNNs. LSTMs feature a unique cell structure that incorporates input, output, and forget gates, which allow the model to selectively store and retrieve information over long time periods. This design effectively mitigates the vanishing gradient problem and enables the network to learn long-range dependencies more efficiently. Due to their robust performance, LSTMs have been widely employed in time-series prediction.

BiLSTM[4]

Bidirectional Long Short-Term Memory (BiLSTM) networks are an extension of the LSTM architecture that allows for the efficient processing of sequential data in both forward and backward directions. By maintaining two separate LSTM layers—one processing the input sequence in its original order and the other in reverse—BiLSTMs are able to capture complex dependencies and contextual information from both past and future time steps. This bidirectional approach has been demonstrated to improve the performance of various sequence-to-sequence tasks, including machine translation, named entity recognition, and sentiment analysis, by providing a more comprehensive understanding of the input data.

Model Design

In this project, our aim was to assess the ability of deep learning models, using a basic set of parameters and minimal parameter selection. We utilized a sequence model to process the input and mapped all output to a linear layer to generate a single output indicating whether the text was pleasant or unpleasant. For optimization, we selected Adam, and trained the model for 10 epochs.

To test the performance of our model, we split the data for each recipient into a training set, which comprised 80% of the data, and a test set, which comprised 20% of the data. This allowed us to evaluate the model's ability to generalize to new, unseen data.

Results

Figure 5: Average wave data of Sophie & Yuchen when watching the stimuli

In order to gain a sense of the data, the average responses to the pleasant and unpleasant stimuli were displayed for each individual participant. Comparing the average responses, it is clear that there is much difference in each person's response. Yuchen’s responses to different stimuli types vary significantly more than Sophie’s.

For Sophie (Figure 6), the average pleasant and unpleasant image responses are very similar. Only channel 8 (electrode T8) shows perceptible differences. This made classification a bit more inaccurate for the machine learning models. For the data of Yuchen (Figure 7), the responses are perceptibly different for most of the channels. This significant variation between test subjects may be attributed to the variation in the way their brains process information.

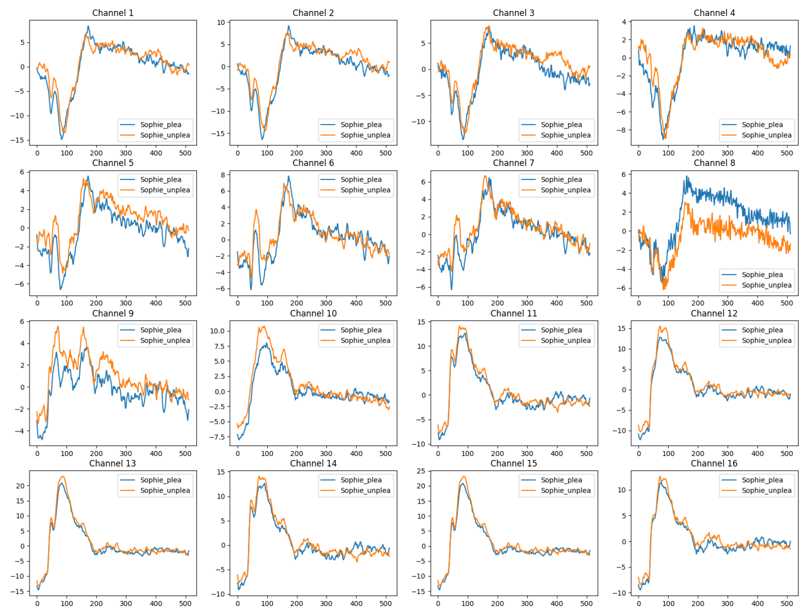

Figure 6: Average EEG response of Sophie (Pleasant vs Unpleasant)

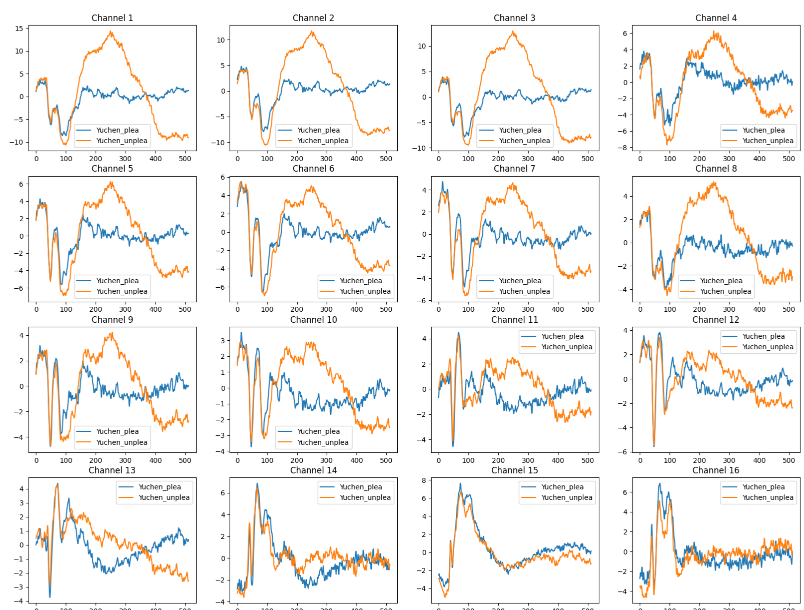

Figure 7: Average EEG response of Yuchen (Pleasant vs Unpleasant)

Classification Results

Using the data in conjunction with the models outlined in the Methods section resulted in the classifications results shown in Table 1. For Sophie's data, RNN model reached the highest accuracy of 62% along with F1 score being 0.62. For Yuchen's data, LSTM played well, achieved 74% accuracy along with 0.74 F1 score. The higher accuracy of Yuchen's classifier compared with Sophie's may be connected to the more easily perceptible differences in Yuchen's responses to the different stimuli categories.

Other differences in models can be seen in the conduction matices of Figure 8. For Sophie's data, pleasant stimuli are harder to differentiate, and the models performed well on the unpleasant data. For Yuchen's data, models generally miscategorised the unpleasant stimuli.

Subject	Model	Accuracy	Recall	Precision	F1-score
Sophie	**RNN**	**62.03%**	**0.62**	**0.62**	**0.62**
	GRU	54.43%	0.54	0.54	0.54
	LSTM	51.9%	0.51	0.51	0.51
	BiLSTM	54.43%	0.54	0.54	0.54
Yuchen	RNN	64.94%	0.64	0.64	0.64
	GRU	71.43%	0.72	0.70	0.71
	LSTM	**74.03%**	**0.74**	**0.74**	**0.74**
	BiLSTM	70.13%	0.71	0.69	0.70

Table 1: Classification results regarding models and recipients

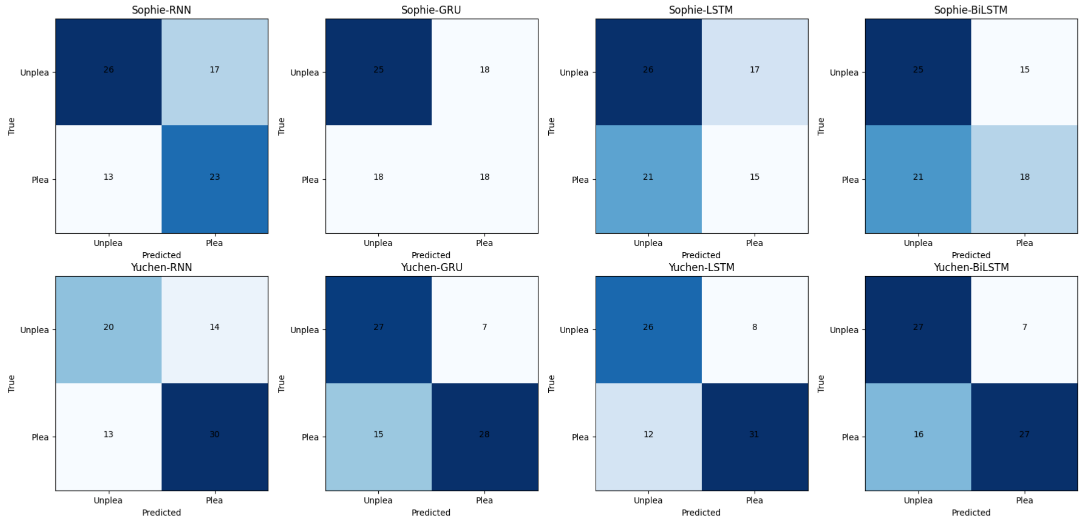

Figure 8: Confusion Matrix for all models

Discussion

An analysis of the EEG data shows that responses to positive and negative visual stimuli depends greatly on the individual participant. For Sophie, there was very little difference in the brain response to pleasant and unpleasant imagery, but Yuchen's responses displayed more variation. This allowed for a more accurate classification model for Yuchen's data.

What remained consistent, however, was the general shape of the output waveform; for each image, a peak in the data would occur shortly after the initial presentation of the stimulus. This provides evidence to confirm the idea of the P300 potential which refers to a spike in activity approximately 300ms following the presentation of the target stimulus. Additionally, the ability of the classifier to surpass the baseline accuracy of 50% (probability of picking correctly from two options) for both participants indicates that pleasant and negative imagery is processed differently in the brain.

Although this experiment provided an excellent introduction to the steps involved in conducting an EEG experiment, there are many directions for future work. Because response varied so significantly between two subjects, a greater number of test subjects is necessary. Additionally, pleasant and unpleasant imagery is very subjective. What one participant finds pleasing may not necessarily appeal to another participant. In the future, it may be beneficial to have a greater number and variety of images to ensure positive and negative responses are actually elicited from the participant. Furthermore, fatigue was a large issue in test subjects. By having shorter trials and frequent breaks, fatigue from image viewing could be greatly reduced and this may elicit more useful responses from participants.

References

1. Rumelhart D E, Hinton G E, Williams R J. Learning internal representations by error propagation[R]. California Univ San Diego La Jolla Inst for Cognitive Science, 1985.
2. Chung J, Gulcehre C, Cho K H, et al. Empirical evaluation of gated recurrent neural networks on sequence modeling[J]. arXiv preprint arXiv:1412.3555, 2014.
3. Hochreiter S, Schmidhuber J. Long short-term memory[J]. Neural computation, 1997, 9(8): 1735-1780.
4. Huang Z, Xu W, Yu K. Bidirectional LSTM-CRF models for sequence tagging[J]. arXiv preprint arXiv:1508.01991, 2015.
5. Dan-Glauser, E.S., Scherer, K.R. The Geneva affective picture database (GAPED): a new 730-picture database focusing on valence and normative significance. *Behav Res* 43, 468–477 (2011). https://doi.org/10.3758/s13428-011-0064-1