Distinguishing anger in the brain using machine learning on EEG data

Bachelor's Project Thesis

Mihai Popescu, m.popescu.1@student.rug.nl
Supervisor: Dr. Marieke van Vugt

Abstract: Monastic debate is a form of meditation aimed at enhancing emotional regulation. It is an integral part of monastic training and serves as a complementary practice that promotes beneficial emotions and minimizes destructive ones in the process. This study attempted at attesting the validity of this form of meditation by examining, through the use of EEG, whether monks that are experienced in monastic debate have a lower tendency of becoming angry than inexperienced monks. The expectation is that compared to beginner monks, experienced monks would have benefited more from this type of training, enhancing their ability to regulate emotions and thus exhibit a lower number of occurrences of anger. Comparing moments of anger and non-anger, we found a significant difference in terms of oscillatory power in alpha, beta, and theta frequency bands across multiple electrodes that could potentially distinguish anger. In order to be able to differentiate between anger and non anger moments on a single-trial level, three support-vector machines with different kernels and a K-Nearest-Neighbour classifier with grid search optimization were built to predict occurrences of anger using two sets of feature vectors: i) discovered by our statistical analysis ii) retrieved from literature. The best accuracies obtained were 54.52% (SVM) and 80.61% (KNN). Experienced monks were found to have shorter and less frequent anger episodes (three anger moments that lasted four seconds on average), compared to inexperienced monks (13 anger moments that lasted 7.3 seconds on average), suggesting the fact that this form of training may be able to help regulate anger. The poor accuracy scores can be improved by collecting and making use of more data, and applying different validation techniques to reduce noise.

Keywords: *EEG, emotion, machine learning, debate, classification, SVM, KNN*

1 Introduction

More and more studies exploring emotions have emerged in the past decades, investigating their nature and role upon cognition. While defining emotions in an encompassing way can be a daunting task, emotions play a crucial role in the mental process of an individual. An example of this can be seen in one of the central components of the cognitive system; the working memory. Working memory has been found to be disrupted by unpleasant emotions, especially anger, in studies conducted on one's ability to perform working-memory intensive tasks (Perlstein et al., 2002). Several techniques to help regulate emotions have been suggested (Guendelman et al., 2017), though one that has seldom been mentioned is a form of analytical meditation called monastic debate.

Meditation practice has been shown to improve emotional regulation if performed on a daily basis (Tang et al., 2015). Although an abundance of literature exists on forms of meditation that revolve around concentrating on one object, there are many other forms of meditation derived from Buddhism that are left unexplored, each one having their own benefits.

Analytical meditation is a form of meditation that, at its roots, focuses on contemplating on experiences under the lens of reason and analysis. The starting point of the analysis for a practitioner is pondering the destructive effects of anger and assessing whether responding with anger to experiences yields constructive or destructive results.
Monastic debate is a form of analytical meditation that is complementary to other concentrative meditation practices (van Vugt et al., 2018). The purpose of this form of meditation is to harness knowledge about how subjective experiences are formed, minimizing destructive emotions and endorsing the beneficial ones in the process (van Vugt et al., 2018). The format of monastic debate is different in comparison to western forms of debate in terms of its physical manifestation, content, and structure, the debate showcasing certain participants stomping on the ground, clapping, and generally exhibiting more movement than usual. In place of convincing the other participant of one's standpoint, the goal of the debate is to collaboratively elucidate inconsistencies in logical lines of reasoning involving a plethora of topics that are part of the monastic curriculum. It involves a series of exchanges between participants in which at least two of the participants reply to a series of statements. The participants of the debate are assigned two possible roles: the defender who makes sure that every statement he replies to holds when compared to his previous statements, and the challenger who does not adhere to this constraint. The challenger attempts to trap the defender into a contradiction. The exchanges are of the form "What is the consequence of that?" (Tillemans, 2008).

Since the debates may be held in a social setting with possible interruptions from the general audience and teasing from the challenger, a defender may become rattled due to the circumstances and thus more susceptible to getting trapped into a contradiction. One crucial skill to becoming a successful debater is the ability to focus one's attention on the debate (van Vugt et al., 2018), effectively blocking out all external stimuli that might hinder one's own mental model. External stimuli, such as personal comments to the defender's address, might evoke an emotional response from the defender, interacting with their ability to focus on the debate. Challengers might evoke anger deliberately as a strategy to destabilize the defender. Defenders initially get destabilized, however overtime they learn not to respond to these attempts seriously and tend not to get destabilized as frequently.

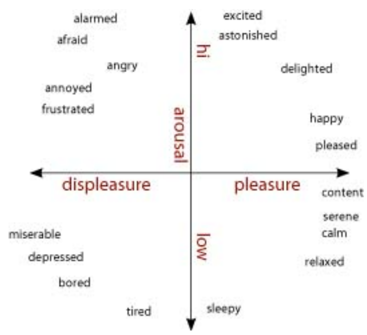

Figure 1.1: Circumplex model of emotions (Russel, 1980). Emotions are expressed by the (valence, arousal) coordinate system. Anger is high in arousal and displeasurable.

This study attempts at exploring whether the frequency and length of anger during debates changes as a function of experience. In order to answer this question, a description of how anger manifests itself in the brain in terms of localized brain area activation, as well as a means of automatically tracking emotions is sought after.

In the context of exploring emotion regulation strategies, the presence and intensity of emotional responses is of relevance to the study. One feasible means towards tracking the intensity of emotions is through the use of EEG since human emotions elicit physiological responses. Other techniques built for emotion recognition use visual, auditory, gesture, and posture-based features (Konar et. al, 2015). The appeal of using the physiological response of the body elicited by emotions needs to be emphasized, as the aforementioned techniques that Konar et al., 2015 present make use of external features that are the external expressions of emotions. These include facial features, bags of key words, auditory features such as pitch changes, or even posture. In an experiment setting, people may act differently in an attempt to suppress their emotions, potentially compromising the aforementioned visual, auditory, and posture-based techniques. EEG-based recognition methods are not susceptible to an individual masking their emotions, as they are able to capture the physiological responses that come bundled with emotions, counteracting the effects of external emotional suppression (Murugappan et al., 2010). As such, the use of EEG provides an additional layer of robustness.

With robustness comes complexity. Tracking emotions through EEG introduces new encumbrances that are not present in visual solutions. Measuring anger in the brain requires an objective definition of its manifestation in terms of brain area activation. Arriving at such a definition proves to be a challenging task. Although individual regions pertaining to one single emotion do not exist in the brain, there are certain brain areas that, when active at the same time, suggest the individual is experiencing an emotion (Kalat, 2019, p.355). From 'Figure 6.1', it is clear that not one brain region is single-handedly responsible for an emotion; rather, it is a combination of them.

Distinguishing anger through EEG requires a description that encompasses the areas in the brain that are active during anger moments. Discerning between anger and non-anger in the brain becomes a matter of extracting moments corresponding to both anger and non-anger during the debates and registering the active brain regions that pertain to both states. Thus, the study of these areas and their differences in terms of activation during anger moments is of interest to this study.

Although not one single definition of emotion encompasses all of its aspects, a psychologist by the name of James Russel created a model that has become the prevalent model used in emotional research. The circumplex model describes emotions in terms of two coordinates: valence and arousal (Lang et al., 1993). Valence of an emotion can be characterized from pleasant to unpleasant whilst arousal of an emotion relates to the degree of excitation which can vary from low to high (Russell, 1980). The dimensions may fluctuate independently of one another. According to the circumplex model, the combination of these two coordinates make up every human emotion. Anger is not unique in terms of its coordinates on the circumplex model, as, for instance, both anger and fear are high-arousal and unpleasant emotions. In our daily life we are able to distinguish between two such emotions with relative ease. However, this sort of characterization of emotions is useful in the emotion-recognition task when the features used could potentially pertain to two or more emotions of the same dimension. In such an event, the model highlights the need of an additional discerning feature that serves the role of breaking the ambiguity of the classification task for that specific emotion. For instance, using solely loudness, arguably an indicator of arousal, to detect astonishment would be insufficient since both astonishment and excitement are characterized by high arousal. Thus, at least one additional feature characteristic of astonishment would be required to break the ambiguity.

Frontal asymmetry and midline power are two criteria that have been used in order to differentiate between positive and negative emotions and between emotions that are similar in terms of valence with anger (Zhao et al., 2018; Balconi et al., 2009a; Harmon-Jones et al., 2001). Overall, a relative increase in the left hemisphere was observed to correlate with a positive emotional stimuli, whereas a greater right hemisphere activity was associated with negative emotions (Davidson et al., 2000; Balconi et al., 2009b; Poole et al., 2014). This criterion was used in order to differentiate in a broader sense between positive and negative emotions. Peering in to the more specific anger manifestation in the brain, there is evidence to suggest that analyzing the EEG activity at the frontal cortical level, as well as the activity in theta, alpha and beta bands can be used to detect anger. Higher levels of left frontal activity and lower right frontal activity were observed when individuals were insulted and felt angry (Zhao et al., 2018). Additionally, it is worth mentioning that in terms of arousal and valence, anger looks similar to anxiety, and so it is not self-evident that they are separable in the brain.

Since monastic debate is a form of meditation training aimed at reducing negative emotions, monks learn to control their emotions through prolonged practice. As such, this study attempted to uncover whether the effects of this practice are quantifiable in terms of occurrences of anger during debates. The hypothesis is that monks who are experienced in monastic debate would have a lower tendency of becoming angry in comparison to inexperienced monks. In particular, experienced monks may still get angry when provoked, however the duration of their anger would be lower compared to beginner monks. The hypothesis is supported by the fact that meditation enforces emotional regulation (Britta et al., 2011).

The appeal of automatically classifying anger can be seen in its applications both within and outside of the study. Methodologically, the ability to automatically classify anger from EEG allows us to explore whether EEG correlates of anger persist beyond the moments deduced in the manual analysis conducted by the researchers. Practically, having such a solution would reduce the time taken to investigate individual videos to establish occurrences of anger. Lastly, automatic classification would enable us to detect potential anger moments in new EEG segments, independent of the ones used in the study.

Throughout the literature, there is evidence to suggest that classifying human emotions based on EEG data using machine learning is possible (Murugappan et al., 2010),(Jalilifard et al., 2016) and (Yuan-Pin Lin et al., 2010). The prevailing machine learning techniques used in the majority of the literature include K-Nearest Neighbour classification (KNN) and support vector machines (SVM), yielding accuracy rates as high as 94% (Jalilifard et al., 2016) and 83.26% respectively. Classifying moments of anger from non-anger translates to a binary classification problem, a task suitable for both KNN and SVM classification methods.

To that end, both SVM and KNN classifiers were trained on EEG data recorded from monastic debates that corresponded to anger or non-anger moments. After the training phase, the classifiers predicted the occurrence of anger of monks engaged in monastic debate.

The research question of the study is: Is anger distinguishable in experienced and beginner monks? One natural methodological goal that follows from the research question is: what features in terms of brain area activation can be used in order to distinguish between anger and non-anger? Moreover, to test the degree of efficacy of monastic debate in terms of benefits seen in emotional regulation, we are interested in seeing whether there is a difference between experienced and inexperienced monks in terms of number of occurrences of anger, and especially whether experienced monks have shorter episodes of anger than beginner monks. The prediction for the research question is that anger will be distinguishable in both experienced and beginner monks using EEG, provided that sufficient accurate intervals pertaining to anger are used. Secondly, in order to establish brain areas to be used as features for the classification methods, different studies conducted on emotion recognition have been investigated (Davidson et al., 2000; Balconi, 2009b; Poole et al., 2014; Kalat, 2019, p. 357; et al., 2009).

Contralateral cortical regions seem to be asymmetrically responsible for processing affection (Poole et al., 2014). By inspecting a large number of photographs depicting emotions, Poole et al., (2014) found that a greater relative left-frontal activity corresponded with a negative approach-motivated state. Should this state be triggered by negative stimuli, such as an insult, there is a high likelihood that anger is occurring. Investigating frontal asymmetry prompts the investigation of the FP1, FP2, F3, and F4 electrodes which correspond to the frontal-cortical region. Petrantonakis, (2009) managed to obtain accuracy scores of 83.33% using an SVM classifier using the previously mentioned channels as features.

Higher temporal activation in the right hemisphere was found to correlate with negative emotions (Kalat, 2019, p. 357). The electrode corresponding to the right temporal lobe, T8, was also investigated as a potential descriptor of anger.

It is crucial to mention at this point that the aforementioned sites found to relate to anger were discerned in experimental contexts in which emotions were elicited through isolated means. By this, we mean that the pathway through which emotions were triggered consisted of solely either photographs, videos, words etc. This study however is far from being as controlled as the previously mentioned ones since we are dealing with a very naturalistic context. The medium through which emotions are elicited is not limited to one pathway as in the previous studies; rather monks are able to get triggered in any way through which they interact with one-another, be it visually, verbally, or otherwise. This poses a two-fold challenge.

Firstly, the features found in the literature may not immediately carry over in the context of monastic debate since they were found in a different context. Additionally, the range of human emotions in a natural environment spans a wider range than the six basic emotions used in studies mentioned (Petrantonakis, 2009). Because the participants of the previously mentioned study knew the possible classes of emotions before they would be subjected to them, a risk of mistaking recognition of emotions for experience of emotions exists.

Secondly, since the debates are emotional by nature, mistaking anger with seemingly similar emotions, such as annoyance or anxiety, is possible. In such an event, the features found in the brain to be discerning between anger and non-anger would potentially contradict the findings in the literature.

Based on the literature we predict that the brain regions corresponding with electrodes FP1, FP2, F3, F4, and T8 will be more active during anger moments than non-anger moments and thus constitute as anger-distinguishing features in the alpha (8-12Hz), beta (12.5-29.5Hz), and theta (4-7Hz) frequencies.

2 Methods

EEG recordings of monks engaged in monastic debate were taken and the different areas of activation in the brain belonging to the two states (anger/non-anger) were studied. The features found were used to train SVM and KNN classifiers in order to predict new occurrences of anger and non-anger in the context of monastic debate.

2.1 Design

The format of the debates consisted of one challenger and one defender engaged in monastic debate in front of a public audience that was not allowed to interfere during the debate, however it did happen on multiple occasions. Two possible classes of debates were taken into consideration as subject to our analysis: counting debates and logic debates. Counting debates involve participants recalling text, summaries, and definitions from the monastic literature. Logic debates involve reasoning about the text by starting from a set of premises and collaboratively figuring out what conclusions from the text are consistent and inconsistent (Dalai Lama, 2018 as cited in van Vugt, 2018). Logic debates were chosen because they were more likely to contain emotional responses.

2.1.1 Participants

The participants were monks between the ages of 20 and 39, ranging in experience. Monks who have practiced this form of debate for more than 4 years were considered experienced, otherwise, they were considered inexperienced. The final analysis included 10 inexperienced monks and 10 experienced monks.

2.1.2 Video recordings

In order to investigate the state of the brain during an anger moment, one must be able to determine when a person is experiencing anger. To accomplish this, 22 video recordings of the debates were taken and analyzed, each video corresponding to one debate. The average length of a debate was 15 minutes. Nine videos included experienced participants, whilst the other 13 videos included inexperienced participants. Three of the nine experienced videos were in a two defenders or two challengers format. All participants mentioned participated in at least one debate. In the inexperienced videos, every inexperienced participant had the chance to play both the defender and the challenger.

The debates were recorded using a video camera with audio recording capabilities. The audio of the camera had a sampling frequency of 48KHz and it was positioned at a side angle relative to the monks to allow for the investigation of both facial expressions. Syncing the videos with the EEG was done by either an initiation of a countdown at the beginning of the debate or by starting the video and EEG at the same time.

2.1.3 Video analysis

The analysis of the videos consisted of two researchers tasked with watching the videos independently of one another and recording the times during the videos corresponding, to their view, to an anger moment. Each video was watched twice, once per participant of the debate. The times found in common between the two researchers were established as the ground truths corresponding with

Figure 2.1: Typical format of a logic debate: The defender is seated and responds to the statements the challenger who is standing. EEG caps were used to record the electrical activity of different brain areas. The debate takes place in a Tibetan monastery. The typical debate lasted 15 minutes.

anger moments. The analysis resulted in having to remove 2 inexperienced debates due to lighting issues during the recordings, making it unable to see the participants' face. Four of the 11 remaining inexperienced videos did not have any common times between the two researchers, and so they had to be taken out of the final analysis.

A moment of anger was common between the two researchers if any noted time of one researcher was within two seconds of another noted time of the other researcher. Should there be a difference of two seconds between two times, the second found at the midpoint between the two times is taken as an anger moment. Any second that is not labeled as an anger moment is labeled as a non-anger moment. Moreover, the duration of an anger moment was defined to be two seconds. The researchers noted down all of the times during which they found the participants to be angry, even if the times were in close succession of one another. This allowed the researchers to establish the average duration of an anger moment of both experienced and inexperienced monks.

The concordance score between the two researchers was computed to determine the matching rate of their labelings of anger and non-anger moments. The result of the concordance test showed an average matching rate of 84.2% concordance between the two.

The final list of anger moments used in the analysis was composed of a total of 62 anger moments found in common between the researchers, 27 of which belonged to experienced monks whilst 35 belonged to inexperienced monks.

2.2 Automatic classification based on audio

Since the method described above is susceptible to subjectivity, an automatic means of classification based on an objective criterion was developed. Since anger is typically associated with an increase in the loudness of the voice, an automatic anger classifier based on the loudness of the participants' voice was created. Detecting anger based on loudness is done by noting when a monks' voice significantly increases from its average volume. Detecting this change is done by isolating the monk's voice by removing background noise and skipping confounds, normalizing the data by subtracting the average from every sample point, and computing the z-score. Relative loudness should be taken into account between videos since different debates involve different people at different distances from the microphone. If the amplitude of a point is higher than a particular value of the z-score, the moment is marked as an anger moment. It is at these times that we are interested in investigating the activity of the brain.

A possible confound for anger detection are the claps generated by monks after posing questions. These claps are very loud and increase the relative loudness of the environment. In order to successfully isolate the voice of the monks, the claps must be taken into account and skipped. It was observed that each clap produced a hardware clip, a form of distortion that occurs in amplifiers when they are overdriven. Every clip corresponded with a clap and every clap corresponded with a higher relative loudness value than 0.9. As such, by setting a threshold of 0.9 relative loudness in the spectrum, the solution managed to detect 100% of the claps in every debate. There are rare instances where the voices of the participants generate distortions akin to a hardware clip, which are then detected fallaciously as claps. The solution was run on all of the videos and, out of every detected loud segment, 1.2% of them were voice distortions falsely labeled as claps. Nevertheless, since the solution managed to detect all claps present in the videos, and so skipping these intervals allowed for a better isolation of the monk's voice.

To keep the EEG and audio synchronized, the claps were kept in the original recording and the average sound was computed using every audio sample that is not within the intervals defined by the dots.

2.3 EEG recordings

EEG recordings were taken of the participants during the debates using an EEG headset (BrainProducts actiCAP) with 32 electrodes arranged in the international 10-20 system as reported in van Vugt et al., (2018). Both experienced and inexperi-

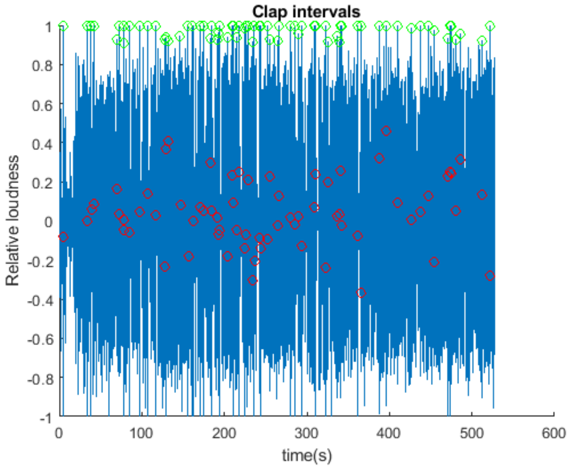

Figure 2.2: Clap intervals in the spectrum of one of the debates. Green dots correspond to the beginning of a clap, red dots designate the end of a clap. To isolate the voices of the monks, confounds such as claps must be skipped. Claps were defined by the interval corresponding with the width between a green dot followed by a red dot.

enced monks were recorded simultaneously using the setup described. Each electrode recorded activity at a sampling frequency of 500Hz. The layout of the electrodes can be seen in 'Figure 62.'.

2.3.1 Preprocessing

Due to the nature of the debate, EEG artifacts were bound to be generated during the recordings. Noting that the monks generally exhibit continuous movement throughout the debate, artifacts pertaining to muscle movement as well as to the electrode's leads and contacts are likely to be generated. Sudden movements and bounces as seen during the recordings by stomping the ground and balancing on one leg before clapping especially generate movement artifacts. Irrespective of monastic debate, EEG artifacts can be generated by other physiological factors as well, such as heartbeats, sweating, and eye blinks.

EEG artifacts are not a byproduct of brain activity, and so they should be removed as they are not relevant in the analysis. Artifact removal was done using independent component analysis (ICA) (Radüntz et al., 2015). ICA is able to remove several types of artifacts present in EEG data through linear decomposition (Jung et al., 2000). The channels were decomposed into a sum of spatially fixed components from which the relative projection strengths of every component at each sensor present on the scalp was used to determine the physiological origin of the components. There are well-known artefactual patterns that are expected when conducting EEG analysis, such as eye blinks and heartbeats. Once the physiological origin of a particular component has been traced to a non-brain-related activity, the signal can be corrected by excluding the troublesome component from the set of components that get projected back onto the scalp. The most common artifacts were produced by blinking, sweating, general movement, and clapping.

It is worth mentioning that measuring absolute oscillatory power at a site is wrong since skull thickness between participants may vary and volume conduction from the homologous site may not be the same. As such, a hierarchical regression model was used (Wheeler et al., 1993) in which the average power of all electrodes was entered. This removed the variance associated with whole-head power and power at the homologous site. Whatever residual power was left at the criterion electrode is then compared between low and high-anger conditions.

Sets of non-anger moments were taken per debate, ensuring that the number of non-anger moments equaled the number of anger moments. To reduce the odds of labeling an anger moment as a non-anger moment, random samples from the videos that were not labeled as anger on either of the researchers' lists were taken as non-anger. A buffer of 15 seconds from each anger moment was defined as an exclusion zone from which no non-anger moment could be defined. This arbitrary-length buffer was used to account for residual anger that might linger from an occurrence of anger.

The EEG data was divided into two-second intervals, meaning that an interval constituted 1000 samples across 32 channels of EEG data. The power at each channel was computed by centering the windows around values found at a distance of 50ms. This resulted in 41 samples of power spectra across 32 channels for two seconds of EEG.

2.4 Feature extraction

In the feature extraction phase, the raw EEG data was converted in order to conduct time-frequency analysis. This was done with the help of the Fieldtrip toolbox in Matlab. Applying $ft\_freqanalysis()$ on the original data for a given frequency interval yielded the power spectrum of all channels for that frequency. The spectra was calculated using the *mtmconvol()* function.

Spectral leakage was accounted for with the help of two discrete prolate spheroidal sequence (DPSS) tapers. The advantage of using multiple tapers is that they are able to detect non-stationary signals (van Vugt et., al, 2007), which is an appropriate task for the nature of signals generated by the brain (Kaplan et., al, 2005).

2.5 Channel selection

Anger moments were compared to non-anger moments by computing the oscillatory power at the channels under investigation. Previous studies have shown that an average relative increase in the following channels corresponded with anger: $FP1, FP2, F3, F4, T8$ at *alpha*, *beta*, and *theta* frequencies. A relative increase in activation of $FP1$ to $FP2$ during an anger moment is expected (Zhao et al., 2018).

As such, the oscillatory power of channels $FP1, FP2, F3, F4$, and $T8$ at α, β, and θ frequencies was used as the feature vectors for the classifiers.

2.6 Statistics

In order to determine whether average differences between experienced and beginner monks exist, a set of statistical tests were conducted.

Firstly, in order to reach our methodological goal of defining unique features pertaining to anger in terms of brain-area activation, moments of anger were compared with moments of non-anger and the brain areas associated with the electrodes that were found to be different in the two states were reported. In order to determine whether a significant difference exists between the anger conditions at the channel level, a linear-mixed effects model is proposed. One advantage of using a linear-mixed effects model is that it is less susceptible to individual differences present between participants (Baayen et al., 2008). The model was implemented using the open source R programming language with the package *lme4*. The proposed model is used to predict interbrain synchrony at 32 channels between the two conditions. Because the samples used in the analysis originated from multiple debates with different debaters, every sample was assigned the debate from which it originated, represented by a number. This number was taken as a random effect in the model to account for the possibility of other variables that may differ between participants interfering in the analysis. It essentially accounts for the uncontrollable variability found between participants. It is this variability that may influence the oscillatory power measured during the experiment in a random way, and so the individual participants that may differ in other ways between one another are treated as random effects in the model. The fixed effects in the model were the channels, taking anger as a factor. Fixed effects are the items that are controlled in the experiment, where a prediction is made based on some systematic change that is deliberate. Here, we expect the oscillatory power of the brain areas associated with the electrodes to be different between anger and non-anger moments. As such, the channels where the oscillatory power is measured are taken as the fixed effects, and anger is the factor that we predict would change the values found at the channels in a systematic way.

Secondly, the average length of an anger episode was computed by segmenting occurrences of anger into episodes of anger. Episodes of anger were created by chaining occurrences of anger that were found to be occurring consecutively. Bearing in mind that an occurrence of anger lasted two seconds, an episode's duration was defined by the difference between the last and first occurrence of anger that were registered consecutively. This way, we were able to determine whether, on average, the duration of anger of an experienced monk was lower than that of a beginner's.

All anger moments found in common between the researchers across debates were used in the analysis. Every sample that was labeled as anger was concatenated in a vector to be then used in a linear-mixed effects model. Because 35 and 27 occurrences of anger were found per beginners and experienced monks respectively, with one occurrence of anger being two seconds, the total amount of anger used for beginners and experienced were 70 and 54 seconds respectively. In terms of power spectra samples, since 41 samples were taken per occurrence of anger on 32 channels, 45920 samples were used for beginners, whilst 35424 samples were used for experienced.

The same number of non-anger moments were sampled from all debates. The final values used were found by randomly sampling values that were labeled as non-anger, with the precaution that the samples did not correspond to an anger moment that either of the researchers found individually.

Since 32 different statistical analyses were made, the likelihood of making type-I errors increased (Ranganathan et al., 2016). To adjust for this, a Bonferroni correction was made that set the p-value threshold as follows:

$$p - threshold = 0.05/nTests = 0.05/32 = 0.0015625$$

Any outliers found in the data were removed. An outlier was defined as any value that was five standard deviations from the mean.

2.7 Classification

One goal of the study was to explore the possibility of automatically classifying anger moments from EEG data. To meet that end, a multitude of classification techniques have been considered, taking into account the nature of the data. The chosen classifier for the job would have to have certain favorable characteristics: be a supervised learning algorithm, be able to do binary classification, show good performance when dealing with high-dimensional data, and be able to adequately generalize from a small number of samples. A support-vector machine and a K-Nearest-Neighbour classifier with grid search hyperparameter optimization were chosen and created using the open source programming language *Python 3.7* coupled with the open source machine learning library *Scikit-learn*. Support vector machines have shown good performance when classifying high dimensional data on a small number of samples due to their convex optimization property (Richhariya, 2018). KNN classification have shown good results throughout the literature of EEG classification and present an intuitive and simple concept to implement. A grid search hyperparameter optimization was implemented on KNN to determine the best K-nearest neighbors. To avoid overfitting 10-folds cross-validation was used.

A support-vector machine using three different kernels were built: polynomial (degree = 2), gaussian radial basis, and sigmoid. The regularization parameter of the SVM (C-parameter) was kept at the default value to reduce overfitting. The gamma parameter was set to $1/nFeatures$. The data were normalized within-participants before conducting the classification process. A K-fold cross-validation technique was used as a means to reduce overfitting and to attest whether the classifier is generalizable on unseen data. The classifier was trained and tested using two sets of features, and the performances of each set was compared. The features used were: i) the features found significant by the linear-mixed effects model and ii) the features used in the literature (FP1, FP2, F3, F4, T8) x (α, β, θ).

The classifiers' accuracy, precision, recall, and F1 score were reported in order to verify whether the solutions are able to generalize and accurately classify new data. The best performance obtained on training and testing sets were reported. Testing sets show a more realistic performance of the algorithms since they constitute data that have not yet been seen by the classifiers. The scores were computed by averaging all individual scores obtained from the folds.

The notion of having generalizable algorithms in the context of machine learning entails a present awareness of the bias-variance trade-off that characterizes underfitting and overfitting. The notion of generalizability underwent an array of attempts towards a formal definition with little success (Sharma et al., 2014). A high precision leads to low bias but high variance, while a low precision leads to high bias, but low variance (Sharma et al., 2014). A proper choice of precision balances the amount of bias and variance. The F1 score can be used as a measure of a classifier's performance that takes into account its precision, as well as its recall rate. The F1 score represents the harmonic average of precision and recall, its value ranging from 0 to 1. A perfect classifier with perfect recall and precision would yield an F1 score of 1.

All intersecting anger moments found in the manual inspection were used, and an equal number of non-anger moments were sampled in a similar fashion as practiced with the Linear-mixed-effects model.

The best classifier was used to predict the number of anger moments between experienced and inexperienced monks and to explore whether experienced monks had a lower number of anger moments compared to inexperienced monks.

2.8 Code

Every script used during the study is available at: https://github.com/uberVelocity/EEG-classification

3 Results

In order to answer the research question, the anger moments extracted from the videos and EEG were statistically compared. Then, SVM and KNN classifiers were run in order to determine whether anger moments can be distinguished on a single-trial level.

3.1 Automatic anger classification based on audio

The performance of the audio classifier was computed for every debate used in the analysis. The average score of every metric can be seen in 'Table 3.1'.

Derivation	Average
Accuracy	71%
Misclassification rate	29%
True positive rate	42 %
False positive rate	29 %
True negative rate	71 %
Precision	1%
Prevalence	0.7%

Table 3.1: Performance of automatic audio anger classification: derivations obtained by averaging all confusion matrices from each debate the classifier was run on.

The loudness-based classifier manages to correctly classify non-anger moments more often than not as can be seen by the true negative rate. Due to the low number of anger moments present in the data relative to non-anger moments, as shown by prevalence, the classifier shows a very poor precision rate as it falsely labels non-anger loud moments as anger.

3.2 Linear-Mixed-Effects

The following section describes the average difference in terms of oscillatory power of monks during anger and non-anger moments.

As expected, a difference in terms of brain area activation was observed between experienced and inexperienced monks. Channels FP1, FP2, F4, and T8 conformed to our expectation and were found to be higher during anger moments in experienced monks. Unexpectedly, F7, CP6, and P4 were found to have higher activation levels in the anger condition for experienced monks. Contrary to the prediction, no difference in activation between the two conditions could be seen for F3 ($\beta = -0.3521$, df=52916, p= 0.766).

Channel	β	t-value	p
FP1	15.982	4.896	<.001
FP2	13.815	3.783	< .001
F7	2.309	3.701	< .001
F4	13.701	0.517	< .001
T8	2.756	7.016	< .001
CP6	8.308	17.748	< .001
P4	2.489	10.107	< .001

Table 3.2: Channels distinguishing between anger and non-anger in experienced monks.

For inexperienced monks, the test reported the presence of a difference at all 32 channels investigated in the β frequency ($p < .001$).

Topographical plots in 'Figure 6.4' show two-second moments of activation of brain areas during anger and non-anger for alpha, beta, and theta frequencies. 'Figure 6.5' showcases the topographical plots for inexperienced monks. Experienced monks follow a centralized pattern of activation during non-anger moments. The focus of the activation shifts in the frontal cortex during an anger moment, an effect more prominent in the alpha and theta frequencies. No visible pattern can be seen for inexperienced monks within conditions. Notable is the predominance of the beta frequency in the anger condition relative to other frequencies, as well as the lower difference in brain area activity between the anger and non-anger moments for experienced monks relative to inexperienced monks.

The final model tested the interaction between experience and anger in order to establish for which channels and frequencies combinations a different anger effect between experienced and inexperienced monks is showcased. The results can be seen in 'Table 6.6'.

3.3 Classification

In order do be able to automatically detect the occurrence of anger in new debates, as well as register the possible residual anger after an occurrence of anger, automatic classification methods are sought after. This section presents the results obtained after building SVM and KNN classification methods. Results for both training and testing sets were reported.

3.3.1 Support-Vector Machines

Two support vector machines were trained; one was trained on the features found significant by the Linear-mixed effects model, while the other was trained on the features presented in the literature. Each test was conducted using a different kernel as specified in 'section 2.7'.

The feacture vector found to perform the best was the one containing the channels found significant by the Linear-mixed-effects model of inexperienced monks (all beta channels):

Kernel	Accuracy	Precision	Recall	F1
RBF	77.89%	81.01%	72.87%	0.76
Sigmoid	53.89%	53.62%	53.20%	0.53
Poly	63.33%	77.39%	37.73%	0.50

Table 3.3: Scores obtained on training sets by SVM using different kernels trained on anger moments using feature vector found by the linear mixed effects model of inexperienced monks (all beta channels).

Comparing the results of each kernel, it seems that the SVM that implemented the RBF kernel performed the best on the training sets, whilst the sigmoid kernel performed at random odds. The polynomial kernel showed above random odds accuracy, however it returned few results with a recall rate of 37.73%.

Kernel	Accuracy	Precision	Recall	F1
RBF	54.52%	54.91%	34.02%	0.42
Sigmoid	53.60%	53.62%	53.20%	0.53
Poly	51.68%	50.54%	18.50%	0.27

Table 3.4: Scores obtained on testing sets by SVM using different kernels trained on anger moments using feature vector found by the linear mixed effects model of inexperienced monks (all beta channels).

On the testing sets, the classifiers did not outperform a random classifier as can be seen in the above table.

By running the classifiers on the features found in the literature, the performance obtained was at random odds for both training and testing.

3.3.2 K-Nearest-Neighbour Classification

Two KNN classifiers were trained on two different sets of feature vectors. This constituted the ones found in the literature, and the ones found significant by the Linear-mixed effects models.

The best feature vector in terms of the classifiers' scores was the one containing the features found significant by the inexperienced LME model (all beta channels).

Accuracy	Precision	Recall	F1	K
98.24%	96.63%	99.97%	0.98	7

Table 3.5: Scores obtained on training sets by KNN with the help of grid search k-optimization. The feature vector used consisted of channels found significant in the LME model for inexperienced monks.

Accuracy	Precision	Recall	F1	K
80.61%	87.75%	65.95%	0.75	7

Table 3.6: Scores obtained on testing sets by KNN the same KNN classifier as referred to in 'Table 3.5'. The feature vector used consisted of channels found significant in the LME model for inexperienced monks.

Overall, the KNN classifier out-performed the support vector machines, as it was able to obtain a 80.61% accuracy rating on the testing sets with an F1 score of 0.75.

3.4 Episodes of anger

The average length of episodes of anger was of particular interest to the study to determine the efficacy of monastic debate as an emotional regulation technique. Overall, by computing the overall average length of all episodes of anger, a higher average duration of anger episodes was observed for inexperienced monks (4.2 seconds, n = 35) compared to experienced monks (2.1 seconds, n = 27). However, looking only at consecutive episodes of anger (occurences of anger that were longer than two seconds), inexperienced monks showed a longer average length more episodes (7.3 seconds, n = 13) compared to experienced monks (4.0 seconds, n = 3). The longest episode of anger observed from inexperienced monks was 21 seconds, while the longest for experienced monks was four seconds.

3.5 Anger prediction

The best KNN classifier was used in order to explore whether experienced monks have a lower tendency of becoming angry than inexperienced monks. The classifier ran on four experienced and four inexperienced debates using all beta channels. On average, all debates had approximately 15 minutes, so the number of samples is comparable between experienced and inexperienced monks. The average duration of anger predicted was 6.9% of the time for experienced monks, whilst inexperienced monks were predicted by the classifier to be angry for 7.7% of the time during debates.

4 Discussion

This study attempted to determine whether anger can be distinguished in experienced and inexperienced monks. To reach this end, unique features pertaining to anger using EEG analysis were defined with the help of a Linear-mixed effects model and two types of classifiers were created. The analysis of localized brain features showcased mixed results for experienced and inexperienced monks. On the one hand, the features found to correlate with anger in the literature were a subset of the features found by the Linear-Mixed effects model in experienced monks, indicating that anger could be differentiated using channels FP1, FP2, F3, F4, T8, CP6, and F4. On the other hand, the features found by the LME in the inexperienced group were disjointed from our prediction, and were showcasing predominant activation at the β frequency. These findings suggest an array of possibilities.

As predominance of the beta frequency is indicative of anxiety, stress, and a heightened focus state (Kropotov, 2016), its occurrence in inexperienced monks might indicate a mismatch between the targeted emotion and the detected one. Anxiety and fright might be a product of unfamiliarity with the task and would be reinforced by the added context of an experiment. Higher activation of the frontal and temporal lobes in the left hemisphere suggests activity at the behavioral activation system (BAS) whilst symmetrical higher activation in the right temporal lobe suggests activity in the behavioral inhibition system(BIS)(Gray, 1970). BAS activity relates to happiness or anger, whilst BIS activity relates to heightened attention, arousal, and inhibitory action which could be characteristic of fear and anxiety. As T8 relates to the temporal lobe of the right hemisphere, an increased activation would suggest the presence of fear.

While these findings are interesting, there are some critical points that need to be discussed in terms of potential issues with our results. First, such a finding might highlight the conceivable problem that a different notion of anger is understood between the two researchers, and as such moments of anger can be mixed with other similar emotions in the context of the circumplex model, like fear or annoyance. During experienced debates, the researchers noticed the fact that most of the moments of anger that the experienced monks showcased resembled short outbursts of anger usually lasting no longer than two seconds. These short outbursts were immediately followed by laughter, making it difficult to agree upon whether the participants actually felt angry in that moment. Converesely, anger in beginner monks could be seen much clearer compared to experienced monks, however overall monks seem to be able to switch between emotions very rapidly, making the classification process much more difficult, but reinforcing the need for a physiological mean of measuring emotions. In order to improve validity and reinforce the findings of the study, certain improvements to the methodology should be made. These improvements pertain mostly to the way occurrences of anger are identified.

Firstly, people from within a culture are better able at recognizing emotions from that culture than people from outside of it (Gendron et al., 2014). Since perceptions of emotions from faces are not culturally universal, one suggestion would be to have Tibetan people watch the videos and have them report the occurrences of anger, as they are able to understand Tibetan and would be able to better recognize emotional responses.

Secondly, a means towards objectively labeling anger based on some external measurable feature would eliminate the subjectivity when choosing anger moments, an aspect which is currently intrinsic to the methodology. One might argue that using external methods of capturing anger moments loses the value of using physiological means. However true, the inclusion of external measures should not exclude the EEG. External measures could potentially complement the findings of the people conducting the manual anger detection task. The hypothetical methods could individually classify moments of anger and be tied in an ensemble learning context in which each solution votes on the answer. Majority voting could then be used with different weights assigned to each contributor to arrive at a conclusion. The EEG signal associated with the moments labeled as anger would then be analyzed in the same manner as in the unimodal approach.

The audio classification technique constituted an attempt at constructing a system that automated the process of classifying anger through the use of amplitude. Techniques involving other prosodic features were considered, such as formant analysis and pitch change, however voice intensity is a good indicator of arousal. Using only one feature dimension, however, is inherently unable of truthfully capturing any emotion in the context of the two-dimensional definition of an emotion in the circumplex model. Additional features, such as gestures and facial expressions are proposed to be used in tandem with amplitude as an improvement to the current classifier (Konar et al., 2015).

Due to the nature of the debate and its context, practical issues arise when conducting analysis with the goal of recognizing emotions. Firstly, because monks tend to clap during the debates, audio recognition techniques requiring voice isolation undergo extensive pre-processing to eliminate said claps. Secondly, since the monks are speaking Tibetan, it is not possible to use language analysis or prosodic features to infer the emotional state of an individual.

Arriving to the research question, based on the average times measured in the analysis, it seems that our hypothesis was correct in predicting that experienced monks have a lower tendency of becoming angry compared to beginner monks. Experienced monks were found to have a lower number of occurrences and, on average, have shorter episodes of anger altogether. The finding potentially suggests that this type of training underwent by the monks improves emotional regulation.

One of the bigger motivations of using machine learning was to be able to predict occurrences of anger in unseen data. Using the classifier with the highest overall accuracy of approximately 80%, experienced monks were found to have, on average, lower-length episodes of anger compared to inexperienced monks.

All in all, the findings suggest the presence of a difference at the level of brain area activation between the two groups during anger and non-anger moments. The features found from our LME model partially correspond with the ones found in present literature, giving light to the possibility of distinguishing anger in experienced in inexperienced monks. Implementing the suggestions mentioned above would reinforce the validity of the data and of the analysis as a whole.

For the future of the project, multiple systems should be used to determine occurrences of anger in the participants. Such systems could potentially include convolutional neural networks for facial recognition, hidden Markov models for voice analysis, or simply a friendly Tibetan willing to help researchers identify moments of anger during the debates. Additionally, because the duration of a manifestation of anger in the brain could occur faster than the sampled anger interval of two seconds, a smaller time margin of anger could be used in an attempt to better pinpoint the occurrence of anger and reduce the noise of the data. Finally, it never hurts to have more data to use not just in the analysis to distinguish characteristical brain areas pertaining to anger, but also to train and test the classifiers.

5 References

Baayen, Harald & Davidson, Douglas & Bates, Douglas. (2008). Mixed-effects modeling with crossed random effects for subjects and items. Journal of Memory and Language. 59. 390-412. 10.1016/j.jml.2007.12.005.

Balconi, M., & Pozzoli, U. (2009a). Arousal effect on emotional face comprehension: Frequency band changes in different time intervals. Physiology & Behavior, 97, 455-462.

Balconi, Michela & Mazza, Guido. (2009b). Lateralisation effect in comprehension of emotional facial expression: A comparison between EEG alpha band power and behavioural inhibition (BIS) and activation (BAS) systems. Laterality. 15. 361-84. 10.1080/13576500902886056.

P. Ekman and W. V. Friesen, “Pictures of facial affect,” in Human Interaction Laboratory. San Francisco, CA: Univ. California Medical Center, 1976.

Fitts PM, Posner MI. Human Performance. Brooks/Cole Pub. Co; Belmont, CA: 1967

Gendron, M., Roberson, D., van der Vyver, J. M., & Barrett, L. F. (2014). Perceptions of emotion from facial expressions are not culturally universal: Evidence from a remote culture. Emotion, 14(2), 251-262.

Guendelman S, Medeiros S, Rampes H. Mindfulness and Emotion: Insights from Neurobiological, Psychological, and Clinical Studies. Front Psychol. 2017;8:220. Published 2017 Mar 6. doi:10.3389/fpsyg.2017.00220.

Harmon-Jones, Eddie & Sigelman, Jonathan. (2001). State anger and prefrontal brain activity: Evidence that insult-related relative left-prefrontal activation is associated with experienced anger and aggression. Journal of personality and social psychology. 80. 797-803. 10.1037//0022-3514.80.5.797.

Henriques, Jeffrey & J. Davidson, Richard. (2000). Decreased responsiveness to reward in depression. Cognition & Emotion. 14. 711-724. 10.1080/02699930050117684.

Holzel, Britta & Lazar, Sara & Gard, Tim & Schuman-Olivier, Zev & Vago, David & Ott, Ulrich. (2011). How Does Mindfulness Meditation Work? Proposing Mechanisms of Action From a Conceptual and Neural Perspective. Perspectives on Psychological Science. 6. 537-559. 10.1177/1745691611419671. http://dx.doi.org/10.1037/a0036052 http://dx.doi.org/10.1109/TSMC.1983.6313160

Lang, P., Greenwald, M., Bradley, M. and Hamm, A. (1993). Looking at pictures: Affective, facial, visceral, and behavioral reactions. Psychophysiology, 30(3), pp.261-273.

Jalilifard, Amir & Pizzolato, Ednaldo & Islam, Md. Kafiul. (2016). Emotion Classification Using Single-Channel Scalp-EEG Recording. Conference proceedings: ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference. 2016. 10.1109/EMBC.2016.7590833.

Jung T-P, Makeig S, Humphries C , Lee TW, McKeown MJ, Iragui V, and Sejnowski TJ, "Removing Electroencephalographic Artifacts by Blind Source Separation," Psychophysiology, 37:163-78, 2000

Kalat, J. W. (2019) Biological Psychology, 13th Edition. Boston, MA: Cengage Learning

Kaplan, Alexander & Fingelkurts, Alexander & Andrew & Fingelkurts, Alexander & Andrew & Borisov, Sergey & Darkhovsky, Boris. (2005). Nonstationary nature of the brain activity as revealed by EEG/MEG: Methodological, practical and conceptual challenges. Signal Processing. 85. 2190-2212. 10.1016/j.sigpro.2005.07.010.

Konar, Amit & Chakraborty, Aruna. (2015). Emotion Recognition: A Pattern Analysis Approach. 10.1002/9781118910566.

Kropotov, J.D.. (2016). Functional Neuromarkers for Psychiatry: Applications for Diagnosis and Treatment.

Lazarou, Ioulietta & Nikolopoulos, Spiros & , Panagiotis & Kompatsiaris, Ioannis & Tsolaki, Magdalini. (2018). EEG-Based Brain–Computer Interfaces for Communication and Rehabilitation of People with Motor Impairment: A Novel Approach of the 21 Century. Frontiers in Human Neuroscience. 12. 10.3389/fnhum.2018.00014.

Lin, Yuan-Pin & Wang, Chi-Hong & Wu, Tien-Lin & Jeng, Shyh-Kang & Chen, J. (2009). EEG-based emotion recognition in music listening: A comparison of schemes for multiclass support vector machine. 489-492. 10.1109/ICASSP.2009.4959627.

M, Murugappan & Ramachandran, Nagarajan & Sazali, Yaacob. (2010). Classification of human emotion from EEG using discrete wavelet transform. J. Biomedical Science and Engineering. 334054. 390-396. 10.4236/jbise.2010.34054.

Perlstein WM, Elbert T, Stenger VA (2002) Dissociation in human prefrontal cortex of affective influences on working memory-related activity. Proc Natl Acad Sci U S A 99:1736-1741.

Petrantonakis, Panagiotis & Hadjileontiadis, Leontios. (2009). Emotion Recognition From EEG Using Higher Order Crossings. IEEE transactions on information technology in biomedicine : a publication of the IEEE Engineering in Medicine and Biology Society. 14. 186-97. 10.1109/TITB.2009.2034649.

Phan, K. Luan & Wager, Tor & Taylor, Stephan & Liberzon, Israel. (2002). Functional Neuroanatomy of Emotion: A Meta-Analysis of Emotion Activation Studies in PET and fMRI. NeuroImage. 16. 331-48. 10.1006/nimg.2002.1087.

Poole, Bryan & Gable, Philip. (2014). Affective motivational direction drives asymmetric frontal hemisphere activation. Experimental brain research. 232. 10.1007/s00221-014-3902-4.

Radüntz, Thea & Scouten, Jon & Hochmuth, O & Meffert, B. (2015). EEG Artifact Elimination by Extraction of ICA-Component Features using Image Processing Algorithms. Journal of Neuroscience Methods. 243. 84-93. 10.1016/j.jneumeth.2015.01.030.

Ranganathan, Priya & S Pramesh, C & Buyse, Marc. (2016). Common pitfalls in statistical analysis: The perils of multiple testing. Perspectives in Clinical Research. 7. 106. 10.4103/2229-3485.179436.

Rasmussen, J. (1983). Skills, rules, and knowledge; signals, signs, and symbols, and other distinctions in human performance models. IEEE Transactions on Systems, Man, & Cybernetics, SMC-13(3), 257-266.

Richhariya, Bharat & Tanveer, M. (2018). EEG signal classification using universum support vector machine. Expert Systems with Applications. 106. 10.1016/j.eswa.2018.03.053.

Russell, James (1980). "A circumplex model of affect". Journal of Personality and Social Psychology. 39 (6): 1161–1178. doi:10.1037/h0077714

Shehieb, Wessam & Alansari, Sara & Jadallah, Nada. (2017). EEG-based communication system for patients with locked-in syndrome using fuzzy logic. 1-5. 10.1109/BME-iCON.2017.8229168.

Siuly, Siuly & Li, Yan & Zhang, Yanchun. (2016). EEG Signal Analysis and Classification: Techniques and Applications. 10.1007/978-3-319-47653-7.

Tang Y., Hölzel B. K., Posner M. I. (2015). The neuroscience of mindfulness meditation. Nat. Rev. Neurosci. 16, 213–225. 10.1038/nrn3916.

Tatum, William & AM, Husain & Benbadis, Selim & Kaplan, Peter. (2007). Handbook of EEG interpretation.

Tillemans, T.J.F. Argumentation (2008) 22: 1. https://doi.org/10.1007/s10503-007-9072-9

van Vugt, M. K., Pollock, J., Johnson, B., Gyatso, K., Norbu, N., Lodroe, T., … Fresco, D. M. (2018, November 15). Inter-brain synchronization in the practice of Tibetan monastic debate. Retrieved from osf.io/qej3y.

van Vugt, M. K., Sederberg, P. B., & Kahana, M. J. (2007). Comparison of spectral analysis methods for characterizing brain oscillations. Journal of neuroscience methods, 162(1-2), 49–63. doi:10.1016/j.jneumeth.2006.12.004

WM, Elbert T, Stenger VA (2002) Dissociation in human prefrontal cortex of affective influences on working memory-related activity. Proc Natl Acad Sci U S A 99:1736-1741.

Zhao G, Zhang Y and Ge Y (2018) Frontal EEG Asymmetry and Middle Line Power Difference in Discrete Emotions. Front. Behav. Neurosci. 12:225. doi: 10.3389/fnbeh.2018.00225

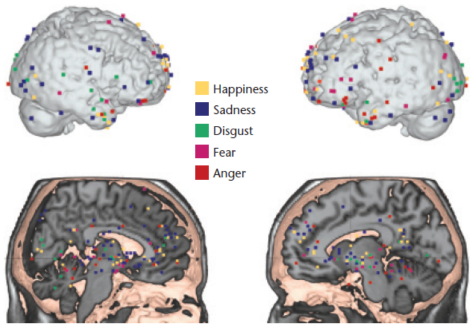

Figure 6.1: Brain areas active during a range of emotions (Phan et al., 2002 as cited in Kalat, 2019). No specific brain region is uniquely associated with one particular emotion. The figure is a result of an fMRI meta-analysis.

6 Appendix

The appendix contains referenced images too big to fit in-text. It also contains extra images which proved to be interesting or that offered different means of visualizing the data.

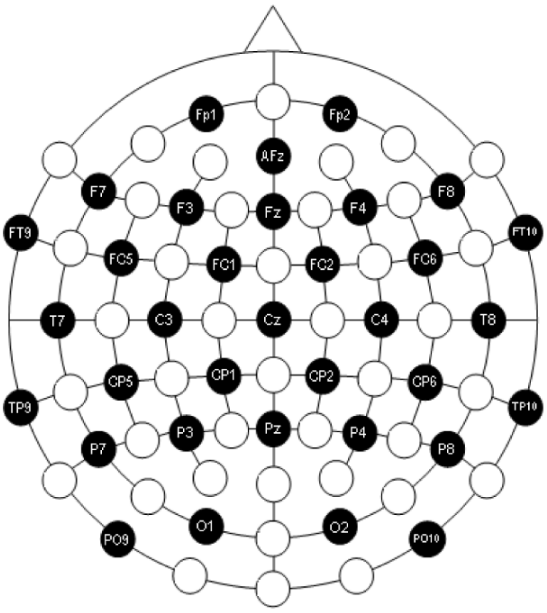

Figure 6.2: Electrode layout showcasing all 32 channels used in the analysis.

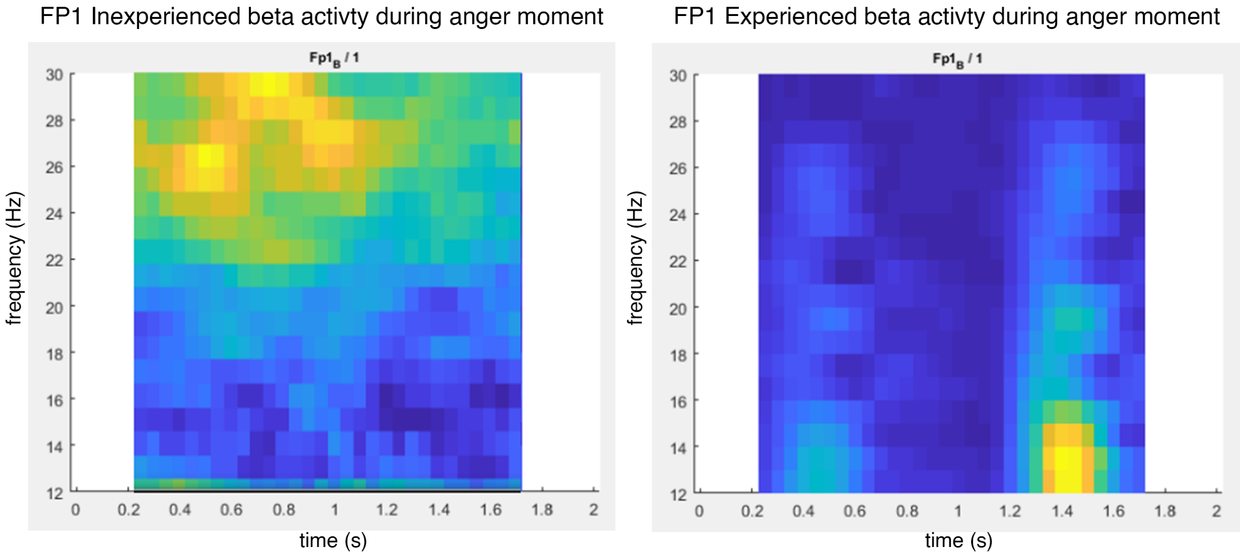

Figure 6.3: Single FP1 channel comparison between experienced and inexperienced monks during anger moments at beta frequency. Activity is more pronounced in the inexperienced group during anger moments noted by the researchers in the beta frequency.

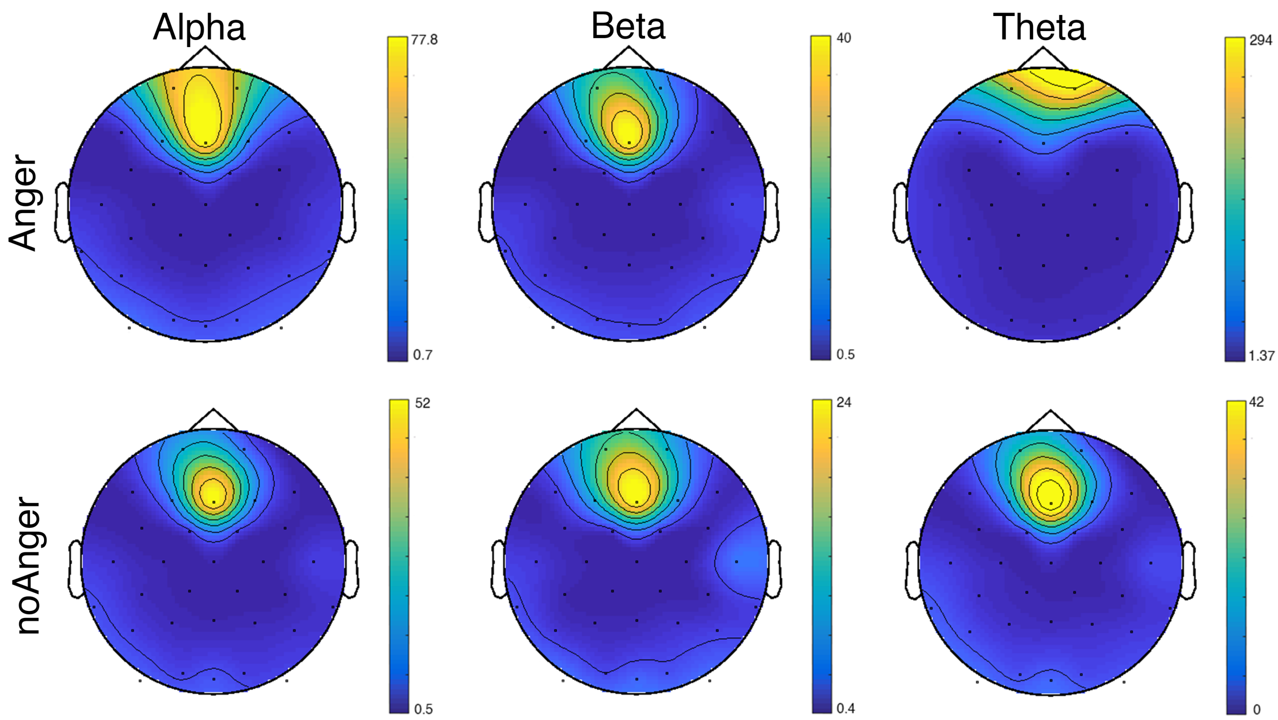

Figure 6.4: Topographical plots of experienced participants showing brain area activation in both conditions for alpha, beta, and theta frequencies. The plots were created using anger and non-anger moments sampled from the pool used in the analysis. Brighter colors are indicative of higher relative activation. A higher activation can be observed in the anger condition in the frontal cortex in the alpha and theta frequencies.

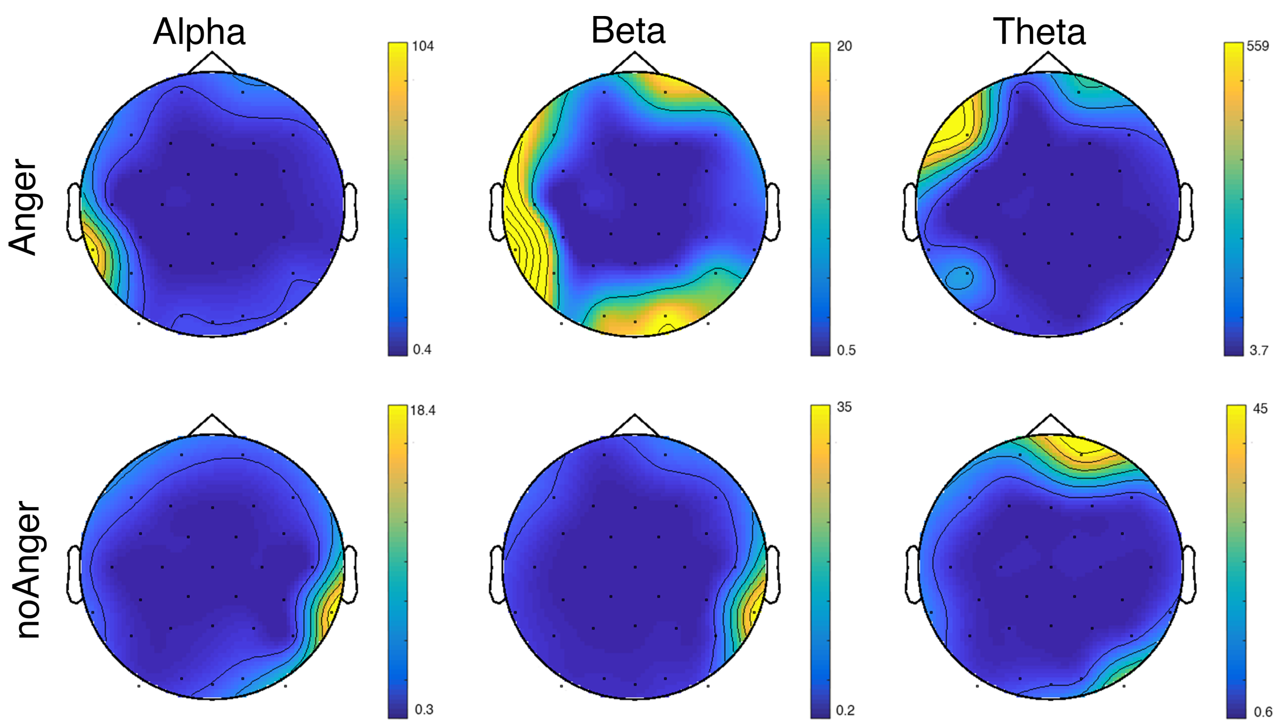

Figure 6.5: Topographical plots of inexperienced participants. The areas of activation are less structured compared to experienced monks. Beta frequency is most prominent in multiple brain regions.

Channel	Frequency	β	t-value	p
FP1	α	18.688	5.612	< 0.001
	θ	62.361	4.165	< 0.001
FP2	α	21.552	4.193	< 0.001
	θ	88.793	3.820	< 0.001
F7	β	-3.691	-4.380	< 0.001
F3	β	-3.024	-3.503	< 0.001
Fz	α	8.111	4.545	< 0.001
	β	5.654	5.492	< 0.001
F4	θ	12.510	5.497	< 0.001
F8	α	-6.507	-3.919	< 0.001
	β	-6.194	-6.850	< 0.001
FC5	α	-5.038	-3.450	< 0.001
	β	-5.162	-5.844	< 0.001
	θ	-13.999	-5.056	< 0.001
FC1	β	-5.590	-11.239	< 0.001
	θ	-7.334	-3.813	< 0.001
FC6	α	-4.691	-3.343	< 0.001
	β	-2.809	-4.343	< 0.001
T7	α	-6.976	-4.338	< 0.001
	β	-7.604	-8.231	< 0.001
	θ	-18.405	-5.486	< 0.001
C3	β	-2.548	-5.061	< 0.001
	θ	-11.696	-4.627	< 0.001
Cz	β	-1.871	-3.614	< 0.001
	θ	-10.253	-4.284	< 0.001
C4	β	-2.588	-4.711	< 0.001
	θ	-12.576	-4.744	< 0.001
T8	α	-4.730	-3.451	< 0.001
	β	-5.605	-8.009	< 0.001
	θ	-9.390	-3.630	< 0.001
TP9	α	-13.274	-7.206	< 0.001
	β	-7.165	-7.893	< 0.001
	θ	-19.804	-5.076	< 0.001
CP5	α	-5.859	-4.305	< 0.001
	β	-3.797	-7.684	< 0.001
	θ	-16.325	-6.802	< 0.001
CP1	α	-135.947	-4.739	< 0.001
	θ	-434.499	-7.087	< 0.001
CP2	β	-2.358	-4.305	< 0.001
	θ	-11.659	-4.125	< 0.001
CP6	θ	16.466	6.311	< 0.001
TP10	α	-26.136	-3.506	< 0.001
	β	-27.585	-7.717	< 0.001
	θ	-70.663	-4.545	< 0.001
P7	α	-5.755	-5.55	< 0.001
	β	-4.020	-8.550	< 0.001
	θ	-19.359	-7.877	< 0.001
P3	β	-2.548	-4.579	< 0.001
	θ	-16.684	-5.584	< 0.001
Pz	α	-315.825	-3.880	< 0.001
P4	α	1.852	3.760	< 0.001
	θ	8.043	7.744	< 0.001
P8	α	-38.613	-3.989	< 0.001
	β	-24.104	-6.822	< 0.001
	θ	-150.883	-5.067	< 0.001
PO9	β	-7.012	-5.290	< 0.001
O1	β	-4.253	-6.379	< 0.001
	θ	-10.477	-3.334	< 0.001
Oz	β	-41.583	-7.454	< 0.001
O2	α	-62.659	-4.562	< 0.001
	θ	-145.452	-5.435	< 0.001
PO10	θ	-70.524	-4.978	< 0.001

Table 6.6: Distinguishing channels and frequency combinations for the interaction between experience and anger.