NeuroHarmony: Analyzing Engagement in Naturalistic Music Using EEG

Abstract:

Music deeply influences emotions and mental states. Traditional music recommendation systems rely on user preferences, tags, and listening history, but often miss the underlying neural responses that reveal how people truly connect with music. Recent advances in EEG (electroencephalography) make it possible to track brain activity in real time, offering a more personal way to understand musical engagement.

This research uses EEG data to predict listener engagement with naturalistic music and to identify musical features such as tempo, spectral patterns, and rhythmic structure that positively impact mood. Deep learning models are applied to link EEG responses with audio features, enabling the creation of curated song selections that align with desired mental states. The findings are integrated into NeuroHarmony, a web application that recommends music for calming, focus, or mood boosting, and can analyse user uploaded tracks to suggest their optimal use. This approach combines neuroscience, music signal processing, and artificial intelligence to advance both music therapy and personalised listening experiences.

Keywords:

Music, Machine Learning, Electroencephalography, Statistics, Feature Engineering, NLP

Introduction:

In a world where we can access music in mere instants, it's easier than ever to listen to the newest hits and trending genres. But most streaming platforms don't value the true power of music, as a tool to alter our state of mind and even change lives. This is the goal of NeuroHarmony, an alternative to commercial streaming platforms that aims to make the daily habit of music listening more intentional. By choosing your purpose for listening to music, you are able to reap far greater benefits for your mental health, productivity and mood.

EEG data is collected through electrodes placed on the scalps of test participants, and measures brain activity through detecting electrical signals produced by our brains' neurons. Through the detection of minute voltage fluctuations in the brain, an EEG machine detects levels brain activity. This is then presented as a graph showing brain waves either on computer or on paper. I chose to use EEG data for my project to understand the effect listening to music has on our mental state from a biological view point. EEG data acts as a clear, objective means of determining changes in brain activity and also shows changes in the frequency of brain waves (for example from delta frequency band to alpha frequency band) which parallel the changes in an individual's state of consciousness. Before creating an AI powered website that provides music recommendations based on individuals' personal needs, I needed to research the effects of music on our brains from a neurological point of view. By creating an AI model to find features in music that produce the intended effects and alter one's state of mind in the desired ways, I was able to understand how we can use music to better our lives. Having gained this knowledge, I hope NeuroHarmony acts as a valuable tool for people to integrate into their daily lives.

Aim:

The aim of my project is to provide a scientific approach to music therapy and enable people to reap the benefits of music therapy as a part of their daily routines. Through the creation of the website NeuroHarmony, I aim to transform the simple habit of streaming our favorite songs into a purposeful, intentional activity. I have been a fond lover of music for my entire life, and have played the piano since I was six. However, as my days have grown busier and life more stressful, I have transformed this hobby into a tool to allow me to reach my goals. I use instrumental, classical music and the soundtracks of my favorite movies to calm and regulate my mood when anxious or stressed; I use upbeat pop and edm music while on runs or on the way to examinations to help me feel more energised and alert; I use lofi, bineural beats and brown noise while studying or working on a mentally demanding task to boost focus. Inspired by the ways I use music in my own life, I created NeuroHarmony to curate and provide song recommendations based on listeners' desired purpose to listen to music: be it to calm, to boost their mood, or to be productive.

Objectives:

- Learning the basics of python
- Learning about the creation and working of key AI models
- EEG data statistical analysis
- Music signal analysis
- Learning relationships between EEG data and music data
- Developing a user friendly application for musical insights

Wider purpose:

Through my research, I hope to gain a deeper understanding of the effects music has on our brains. By analysing EEG data, I aim to draw meaningful conclusions about how we can better curate music choices to reap maximum benefits of improving our state of mind. With this knowledge I aim to create a web application called NeuroHarmony that allows users to harness the full power of music and use music to propel them towards achieving their fullest potentials. With a user-friendly application, and unique features like curated catalogues of music to 'boost mood', 'calm' or 'enhance focus', as well as 'Analyse my track' I hope that NeuroHarmony revolutionises the habit of listening to music for all users. NeuroHarmony is a tool people of varied age groups can use in their everyday lives to transform listening to music from a passive hobby to an intentional act that brings them closer to their goals.

Timeline and Plan:

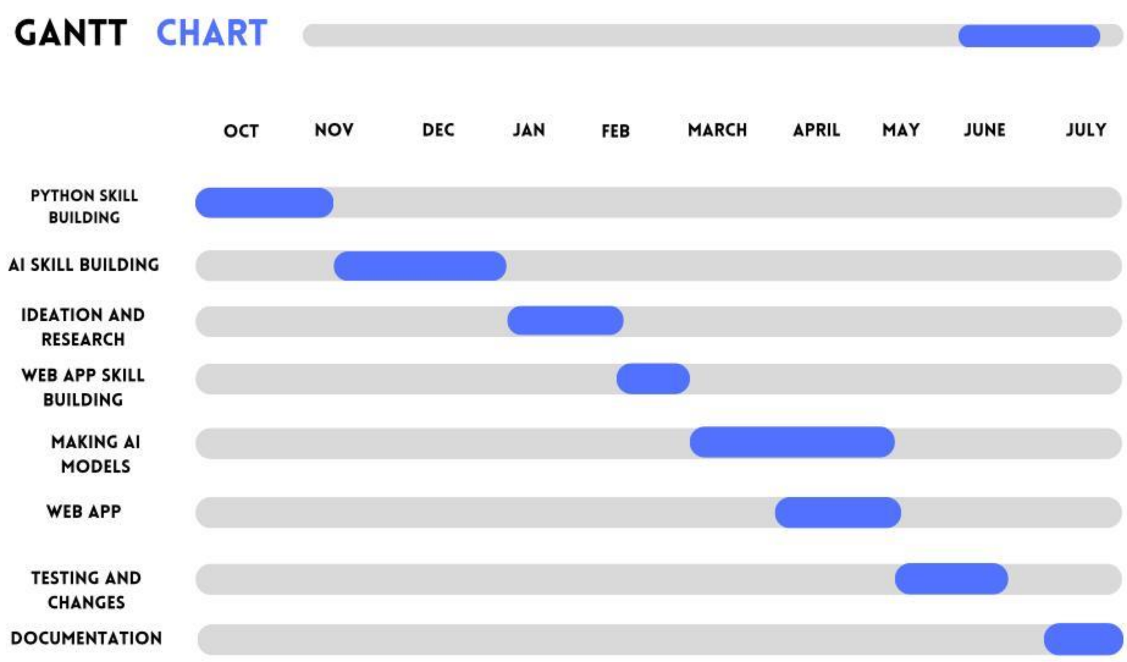

Total: 90 hours

<u>Python Skill Building:</u>
Before beginning the project, I completed a short course on the basics of Python to enhance my previous coding knowledge.

AI skill building:
I also learned about the principles of AI models, common models like regression models and convolutional neural networks to prepare for the later stages of my project.
Ideation and Research:
Being a fond lover of music, and having played the piano for over ten years, I decided to use these skills in my project synthesising two of my passions of music and neurobiology. After extensive research, I formulated the research procedure and found the dataset I would use. After brainstorming various potential useful applications of the research, I decided on a website.
Web App:
So, I spent a few weeks learning how to code a website using HTML as well as integrate Flask and Python to create a useful web application.
Making AI models:
By experimenting with multiple models like Random forest Classifier, regression models, 2-D CNN'S etc.I was able to create the required AI models used in the music classification and 'Analyse my track' features of the web application.
Testing and Changes:
Through trial and error I worked with my mentors to create AI models that fulfilled the needed purposes of the website features, and integrated those with the website's code.
Documentation:
Having kept a detailed log of all the tasks I accomplished relating to my project, while creating the final documentation I simply had to consolidate the information and organise it according to the given format.

Research and Related Work:

Recording brain activity while listening to music using wearable EEG devices combined with Bidirectional Long Short-Term Memory Networks by Jingyi Wang, et.al - This study uses EEG data recorded while participants listened to music to recognise varied emotional states

A study on the effects of music on our brains and emotions by Chloe Fong - This study aims at recognising the various roles music can serve towards enhancing our mental states, ranging from improving cognitive capabilities to reducing stress, fostering positive mood and even aiding in treatment of mental health disorders.

A study on the effect of different kinds of music on brainwaves by Shao Kuo Tai, et.al -
This study investigates the varied impacts different types of music can have on neural activity, it found that while listening to slow music can influence the power of alpha waves, listening to fast music can change the power of beta waves.

The Effect of Music on Human Brain; Frequency Domain and Time Series Analysis Using Electroencephalogram by Rab Nawaz, et.al- A study comparing EEG and mood changes from calming binaural beats versus participants' favourite music, measured over short and long periods.

The Positive Influence of Music on the Human Brain by Shiqi Zhang - A study examining how music affects cognition and memory, finding that listening to music improved students' memory, response times, and accuracy in a color–word matching task.

Science Behind This Project

My project is built on two main areas of science - how the brain reacts to music, and how music can be broken down into measurable parts.

How the brain reacts to music
When we hear music, our ears turn the sound waves into electrical signals, which travel to the brain. The temporal lobe does most of the "listening work," figuring out pitch, rhythm, and timbre.

To study this, I used EEG (electroencephalography). EEG works by detecting tiny voltage changes from brain cells firing. These patterns are grouped into "brainwaves":

- **Theta** - deep relaxation or daydreaming.
- **Alpha** - calm but alert.
- **Beta** - focused thinking and concentration.
- **Gamma** - high-level learning and processing.

Different types of music can increase certain brainwaves. For example, upbeat, fast music often boosts Beta and Gamma activity (focus and alertness), while slower, softer music can raise Alpha waves (relaxation).

How music is measured
I looked at things like:

- **Tempo** - how fast the beat is.

- **Loudness** - overall energy of the sound.
- **Brightness** - how sharp or mellow it sounds.
- **Timbre** - the tone colour that makes a violin sound different from a piano.

By linking these features to brainwave changes, I could explore the science of why certain songs make us feel or focus in specific ways.

Selection of Approach:

Datasets Explored:

Initially, I considered using the NMED-M dataset, which focuses on minimalistic music. However, due to its limited number of songs and lack of readily available audio, it was not suitable for my study. I then switched to the NMED-T dataset, which contains data from 20 participants listening to 10 commonly known songs, making it more viable for practical analysis.

Music Classification Models:

To build a music classification pipeline, I used a separate dataset comprising 10 music genres with 100 audio files each, where each file was 30 seconds long. Various models were evaluated:

- **Random Forest Classifier**: Achieved 72% accuracy on the GTZAN dataset.
- **XGBoost**: Improved performance with 85% accuracy.
- **Convolutional Neural Network (CNN)**: Achieved the highest performance with 92% accuracy, using Mel spectrograms as input features.

Rather than classifying songs purely by genre, the CNN model was refined to group songs into three affective categories based on musical features:

1. Calming songs
2. Songs for focus
3. Energy-boosting songs

This model was then used to classify new songs outside the dataset using extracted features via Librosa.

Ratings Model:

For understanding listener preference, I initially attempted a regression model to predict ratings but it performed poorly. The ratings were then binarized:

- **0–5** labeled as *Dislike*
- **6–10** labeled as *Like*

Using this binary classification, a Random Forest Classifier incorporating both musical features, EEG data, and user ratings achieved 71% accuracy.

Technical Implementation:

There are two parts to this project, the research with EEG Data to understand the effects of music and the Application.

Research:

1) Dataset - NMED T Dataset, an open dataset of electrophysiological and behavioral responses collected from 20 participants who heard a set of 10 commercially available musical works. Song stimuli span various genres and tempos.

The EEG for the participants was recorded using the HydroCel Geodesic Sensor Net (HCGSN) System which has a total of 128 channels:

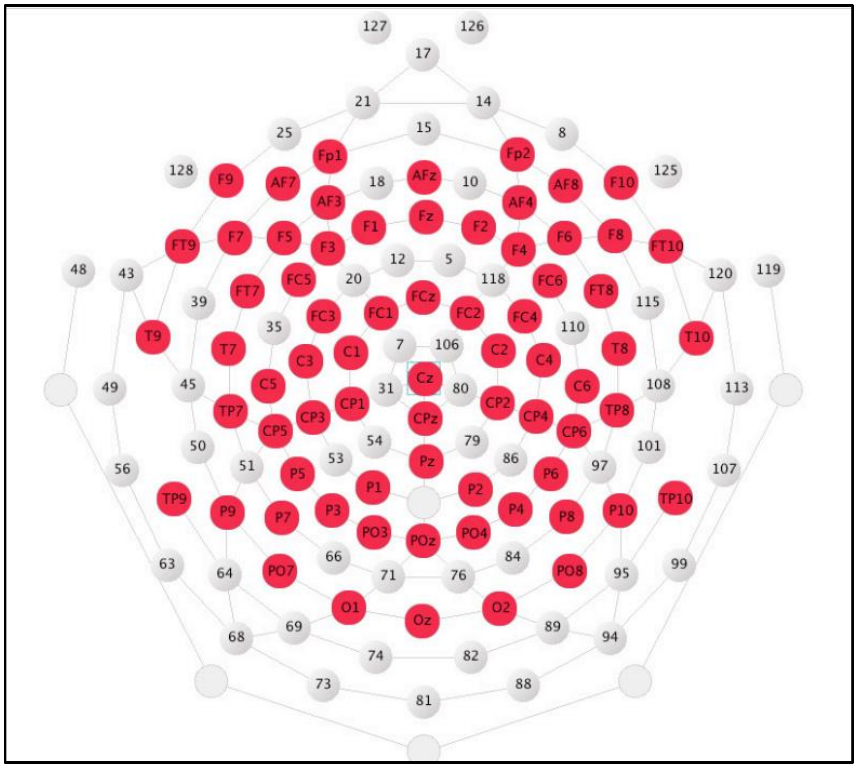

After the doctor's suggestion of using temporal channels only, I mapped this electrode configuration with the 10-10 system and used temporal channels

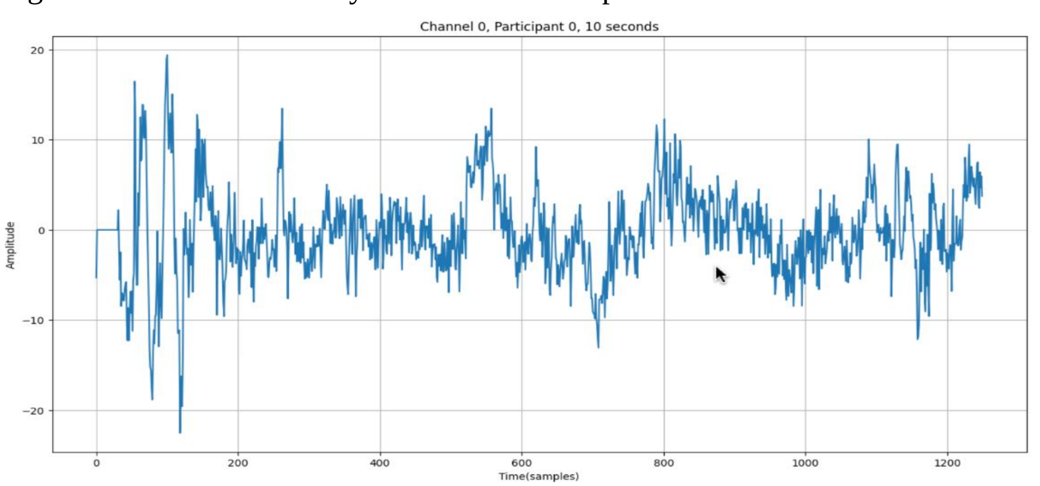

The EEG data is recorded for the duration of the entire song (for instance, song 1 has a duration of 4: 38, song 2 has a duration of 4: 12) , for all 10 songs, for each of the 20 participants.

There is also a mat file with enjoyment ratings which I then used to make a classifier.

2) Data Analysis:

Analyzing EEG frequency bands is useful as each band reflects specific cognitive and emotional states, such as relaxation, focus, or deep sleep. By examining these bands, I gained insights into how the brain responds to different stimuli like music, helping identify patterns of engagement, mood, and mental state.

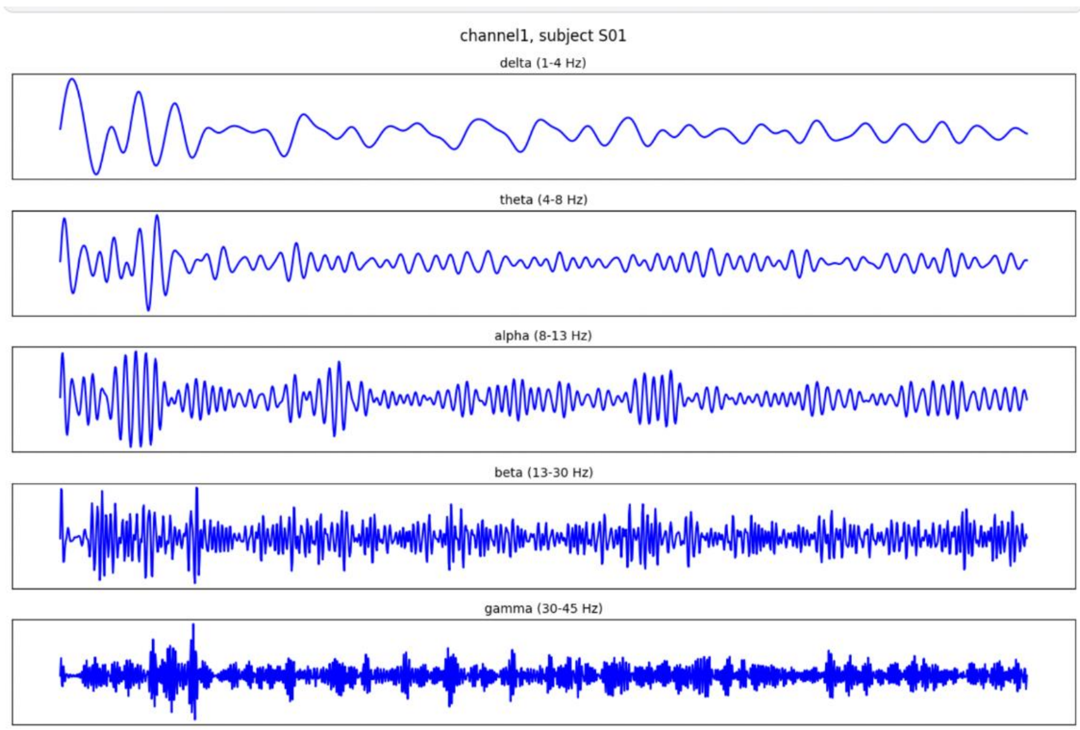

Delta Waves (0.5–4 Hz): Linked to deep sleep and healing, delta waves are the brain's slowest rhythms.

Theta Waves (4–8 Hz): Associated with deep relaxation and creativity, often seen in meditation and REM sleep.

Alpha Waves (8–13 Hz): Represent calm focus and relaxed alertness, common during light meditation or restful wakefulness.

Beta Waves (13–30 Hz): Indicate active thinking and concentration, typical during problem-solving and mental tasks.

Gamma Waves (30–100 Hz): Reflect peak cognitive performance and rapid information processing in demanding tasks.
I started by extracting audio features from each song using the Librosa library in Python. These included:

- Tempo – beats per minute
- Spectral Centroid – brightness or sharpness of the sound
- Spectral Flux – how quickly the spectrum is changing (captures rhythmic variation)
- RMS Energy – loudness or overall energy of the signal
- Zero-Crossing Rate – related to noisiness or percussive nature
- MFCCs – a compact representation of timbre

Then, I extracted EEG features from temporal channels, since they are known to be most involved in auditory processing. Instead of using the raw signals, I transformed them into features like:

- Band Power in the Theta, Alpha, Beta, and Gamma ranges

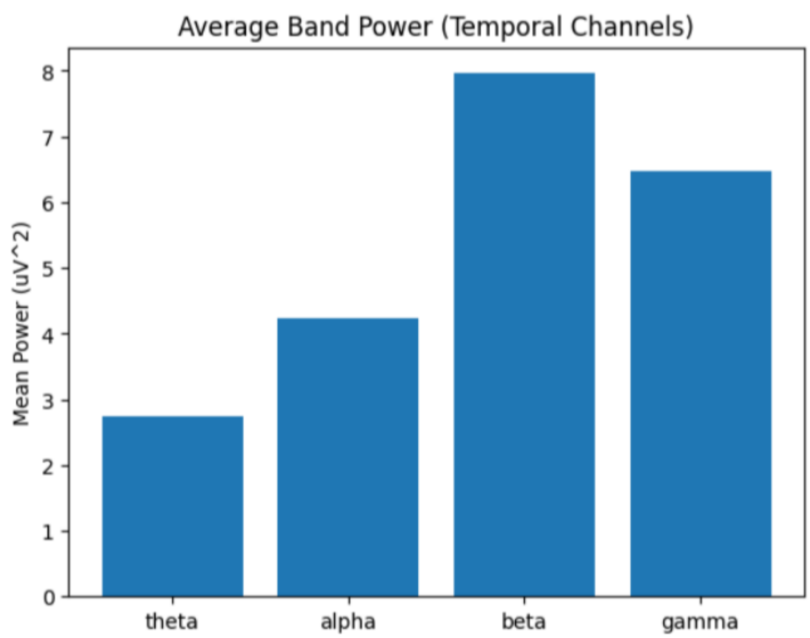

Using Spearman correlation analysis and librosa, I found consistent patterns in how the brain responded to certain music features which were then used in making the model:

- RMS Energy showed a positive correlation with Alpha power.

- MFCC means were positively correlated with Alpha, Beta, and Gamma activity.

- Spectral Flux had negative correlations with Beta and Gamma activity.

- Spectral Centroid and ZCR were positively correlated with Beta and Gamma power.

Example for Bonobo - Forest Fires for all participants for three frequency bands:

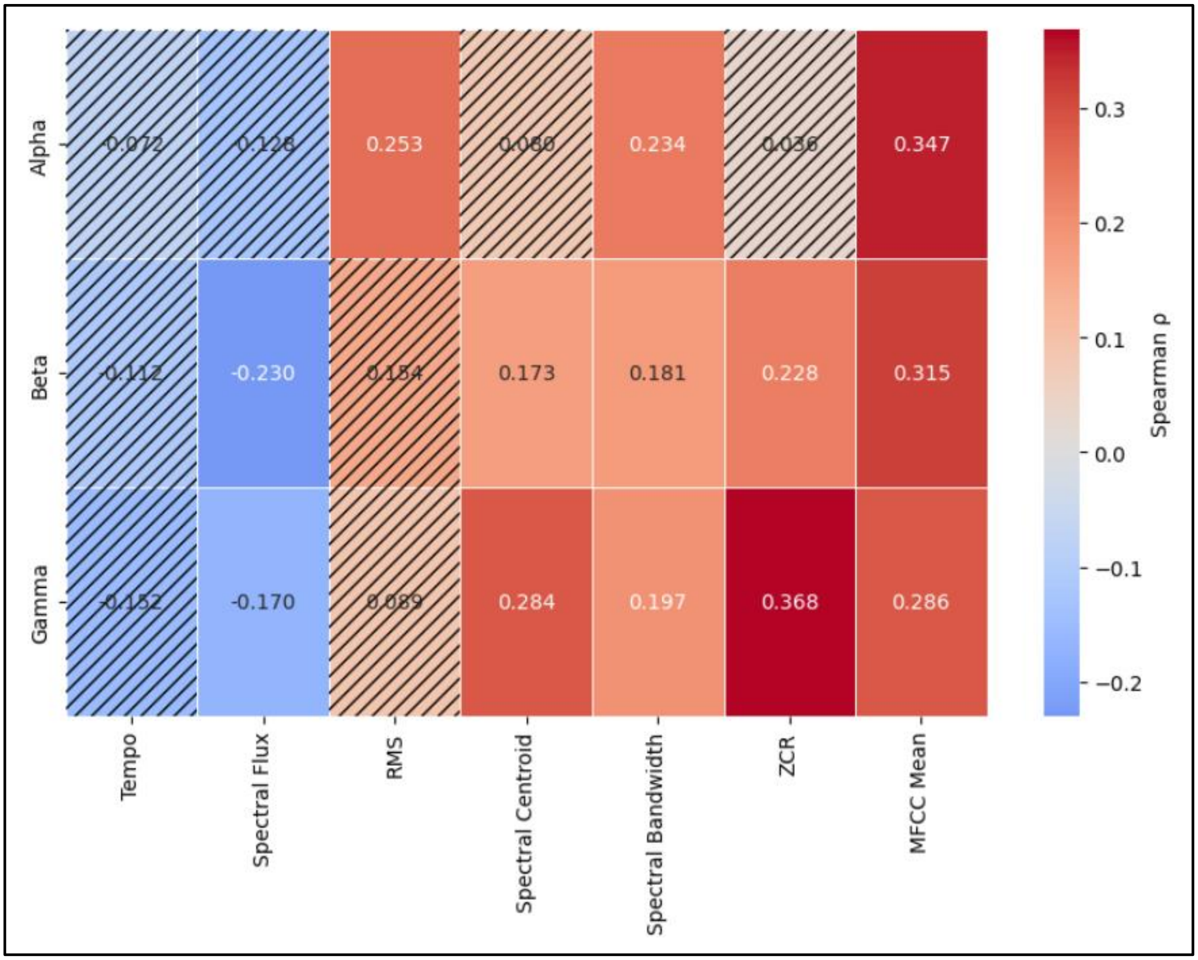

Applied Discrete Wavelet Transform (DWT) to decompose EEG signals into distinct frequency subbands while preserving how these frequencies changed over time. Excluded delta wave subband based on doctor's advice, as it is primarily linked to deep sleep and unrelated to music listening or cognitive engagement.

Focused on alpha, beta, and theta subbands, which are more relevant to attention, relaxation, and music processing. Extracted key features from each subband over short, consecutive time windows:

- Mean – average signal amplitude.
- Standard deviation – variation in amplitude.
- Energy – overall signal strength.
- Entropy – signal complexity and unpredictability.

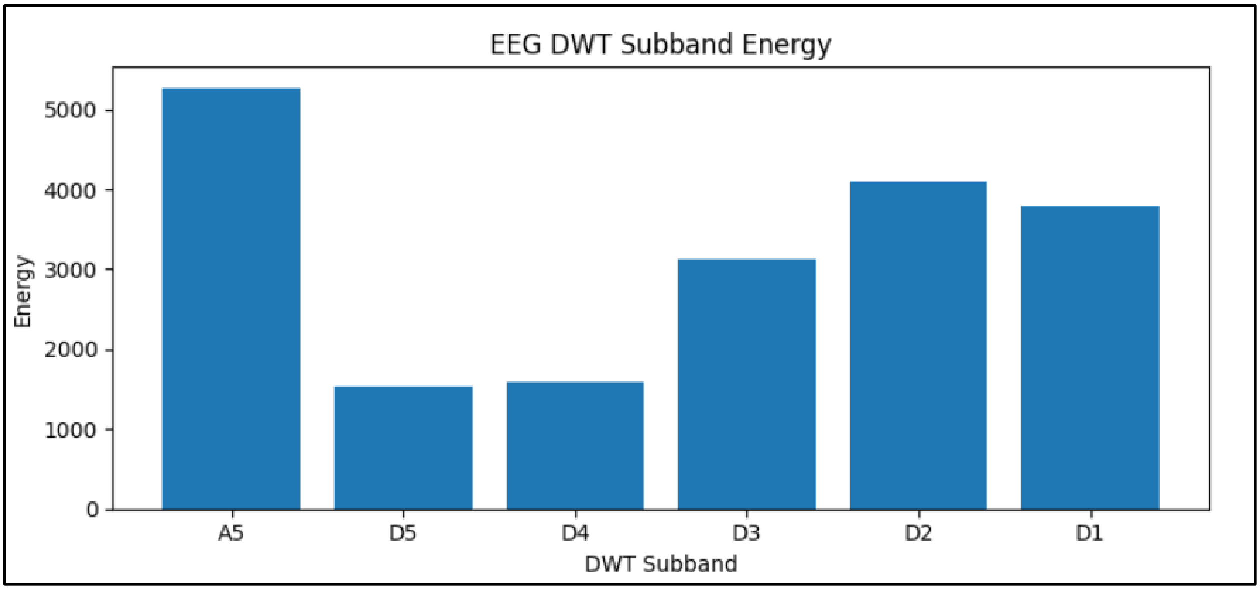

Once I had both sets of features:

- Matched EEG responses from each participant with the corresponding audio features for each window.
- Averaged the EEG band powers and wavelet stats across temporal channels in each window

Each participant's brain responds differently, and with only 20 people, it's hard to make strong generalizations. But I did see repeating trends, which supports the idea that specific musical features can trigger distinct EEG patterns, especially in the temporal areas of the brain.

Ratings Model:

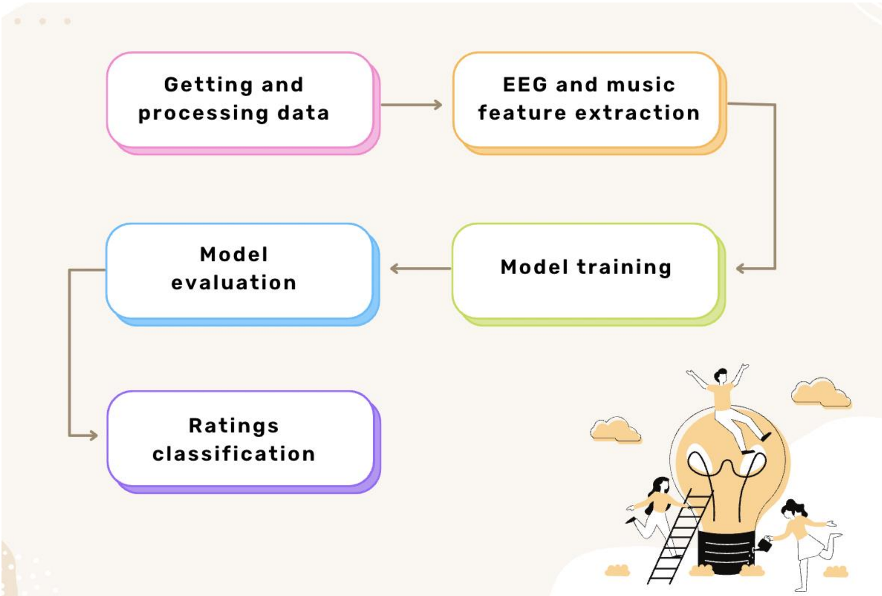

EEG Feature Extraction

The EEG data was recorded using a 128-channel HydroCel Geodesic Sensor Net. After talking to a neurologist, I focused only on the temporal channels because they are most involved in auditory processing. Based on a paper that mapped the sensor net to the 10–10 system.

I extracted two types of features from the EEG:

1. **Frequency Band Power**
   I looked at power in four frequency bands for each 2-second segment of EEG:
   - Theta (4–8 Hz)
   - Alpha (8–13 Hz)
   - Beta (13–30 Hz)
   - Gamma (30–100 Hz)

     I didn't include delta because it's more related to deep sleep and not useful here.

     I found out that beta and gamma were the most active bands during music listening across participants.

2. **Wavelet Features (Using DWT)**
   To capture both time and frequency details, I applied a 5-level Discrete Wavelet Transform using the Daubechies 4 (db4) wavelet. From each sub-band, I extracted:
   - Mean
   - Standard deviation
   - Energy

- Entropy

Music Feature Extraction

I also extracted features directly from the audio files using Librosa. I did this for the same 2-second windows as the EEG. The features I used were:

- Spectral centroid (how bright the sound is)
- Spectral flux (how much the sound changes)
- Zero-crossing rate
- RMS energy
- MFCCs (which represent the overall shape of the sound spectrum)

These features helped me understand the characteristics of the song at each moment.

Labels:

The dataset includes enjoyment ratings from each participant for every song. To simplify things, I turned the ratings into binary labels:

- If a participant rated a song 5, I labeled it as “liked”
- If it was below or equal to 5, I labeled it as “disliked”

This way, I could train a classifier on EEG and audio features to predict this binary label.

Training the Model

I tried three different models:

- Logistic Regression
- Random Forest
- SVM

I trained models using:

- Only EEG features
- Only audio features
- Combined EEG + audio features

The Random Forest model with combined features gave the best results.

Application:

Web application:

Goal of the website:

The website aims to allow users to listen to curated catalogues of songs that fulfill one of the three key purposes of music - to boost mood, to calm, and to enhance focus. The goal is to allow users of varied age groups to reap the full benefits of music in their daily lives.

Creating features offered:

To create curated catalogues of songs, I developed an AI model that classifies popular music into three groups depending on which of the key purposes of music it fulfills. By first researching about the key features of a song that can be isolated as the input data for the model, I was able to determine that values like the chroma, zero crossing rate, tempo, etc of the song could be used as inputs in the music classification model. By applying this model to popular music of varied genres I was able to create the website feature 'Stream Now' which showcases music based on the purpose it fulfills.

To create the 'Analyse My Track' feature that classifies a music file uploaded by users into one of the three groups, I used a google gemini API. Through trial and error, I was able to create a prompt that accurately determined the purpose a given song served and gave a short response on the particular features of the song enable it to serve this purpose.

Designing the website:

I used a colour scheme of shades of blue and white to create a relaxing backdrop for the website features. I created a simple, minimalist user interface to enhance ease of navigation and usage. To improve user experience, the website also allows users to choose an 'avatar' relating to popular music genres after they log in to their account.

The website was developed using Flask as the backend framework, handling routing, user sessions, and integration with Firebase services. The music library includes 6 songs in the Focus category, 7 songs in Calm, and 7 songs in Boost.

Pyrebase was used to implement secure email authentication, enabling users to register, log in, and manage their accounts seamlessly. The Firebase Admin panel provides additional management tools, including viewing each user's last played track, monitoring playlists they have created, and tracking their selected avatar.

Curated song recommendations:

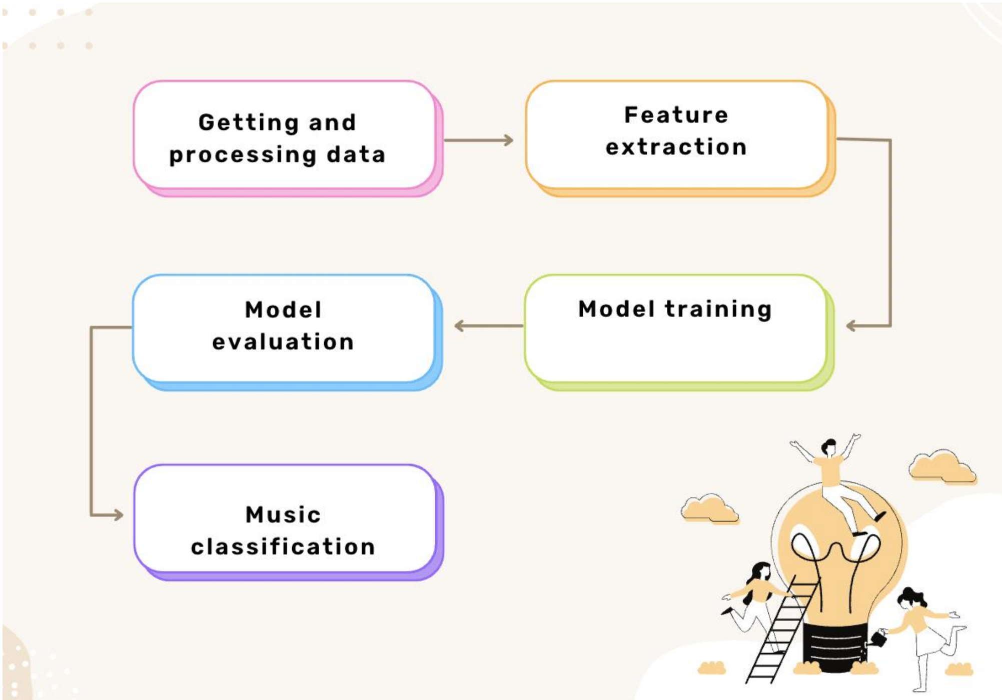

By developing an AI model that accurately classifies songs into one of the three categories (music to calm, music to boost mood, and music to focus), I could create curated lists of popular songs for users to listen to depending on their needs. I used the GTZAN dataset.

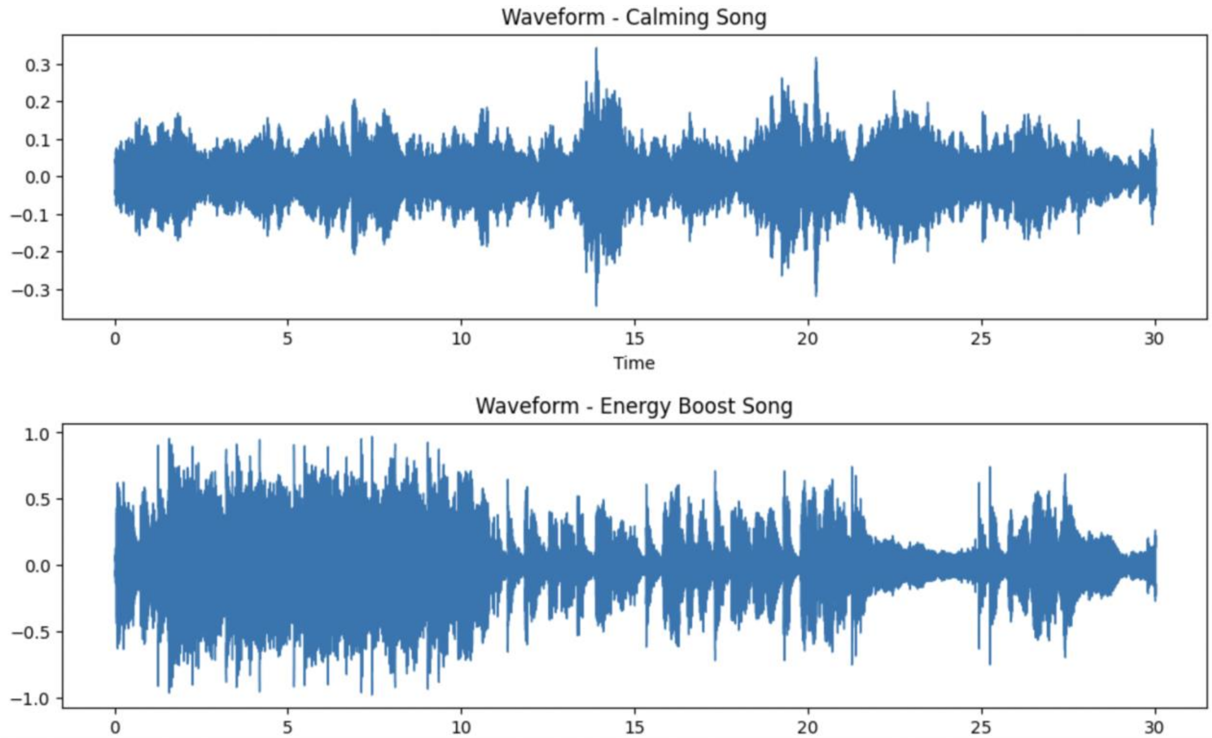

Waveform - Calming Song
0.3
0.2
0.1
0.0
−0.1
−0.2
−0.3
0
5
10
15
20
25
30
Time
Waveform - Energy Boost Song
1.0
0.5
0.0
−0.5
−1.0
0
5
10
15
20
25
30

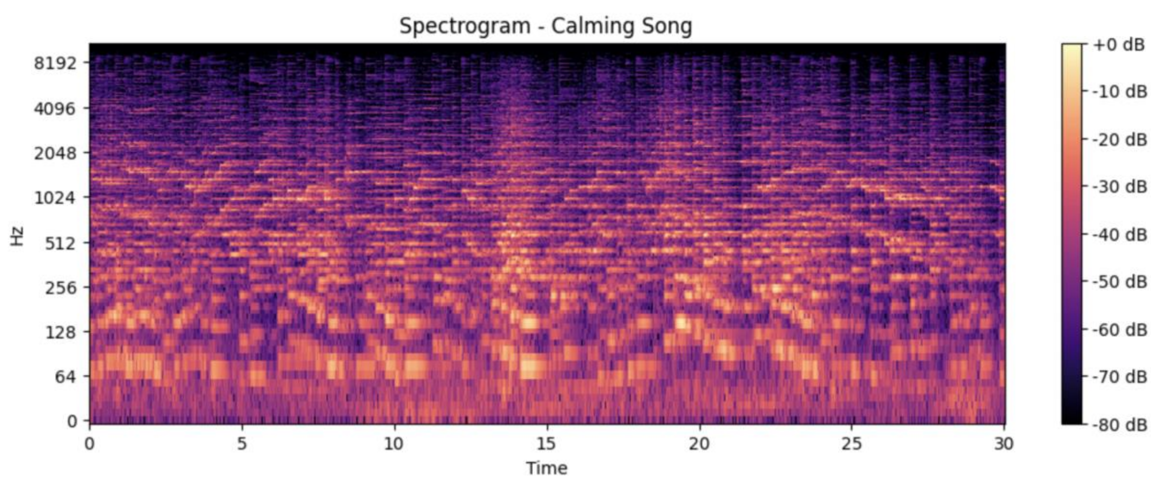

Spectrogram - Calming Song
8192
4096
2048
1024
512
256
128
64
0
Hz
0
5
10
15
20
25
30
Time
+0 dB
-10 dB
-20 dB
-30 dB
-40 dB
-50 dB
-60 dB
-70 dB
-80 dB

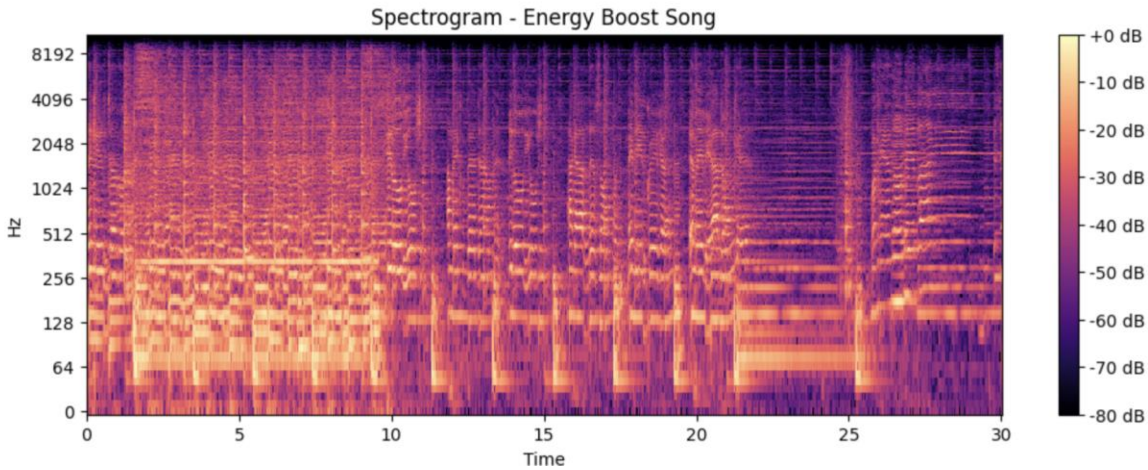

CNN using Mel Spectrograms ended up giving the most accuracy (92%) in classifying the songs into three categories (calm, focus, energy boost) which was then used to classify new songs to put in the application.

Analyse my track feature:

- By using a Google Gemini API, I built a feature that classifies an uploaded song by the users of the NeuroHarmony web application into one of the three key purposes music serves - calming, energising, or enabling deep focus.

- Before classification, I used the Librosa Python library to extract RMS Energy, MFCC means, Spectral Flux, Spectral Centroid, and Zero Crossing Rate (ZCR), as these features showed meaningful correlations with EEG activity in my analysis. I used a Google Gemini API because it is free to use, and provides satisfactory responses.

- I tried multiple prompts until finalising one that gave accurate classifications and casual responses that were not overly technical. For instance, “Based on its features, this tune is perfect music to focus on your work or studies. The song has a quick, repetitive rhythm that can help you stay on task without being distracting. Its overall sound is smooth and mellow, avoiding any loud or sharp noises that might break your concentration.”

The patterns observed in the EEG analysis directly shaped the conceptual design of the NeuroHarmony application. From the EEG audio correlations, I identified which musical properties most strongly influenced neural activity, allowing me to define the three core music categories, calm, focus, and energy boost, in a way that was grounded in neuroscience rather than subjective mood labels. For instance, higher tempo, greater spectral flux, and increased

RMS energy were linked with elevated Beta and Gamma band power, indicating heightened focus and energy. Conversely, slower tempos, smoother spectral transitions, and lower RMS energy were associated with higher Alpha power, reflecting a calm but alert state.

These findings informed two key components of the application:

1. **Music Classification Model** - Although implemented using a CNN trained on Mel spectrograms, the choice of categories and the interpretation of their underlying patterns were guided by EEG findings.

2. **Likeness Prediction Model** - This combined explicit EEG features from temporal channels with corresponding audio features, directly applying the EEG analysis methods to predict whether a listener would like or dislike a song.

Web Application Images:

Home Page

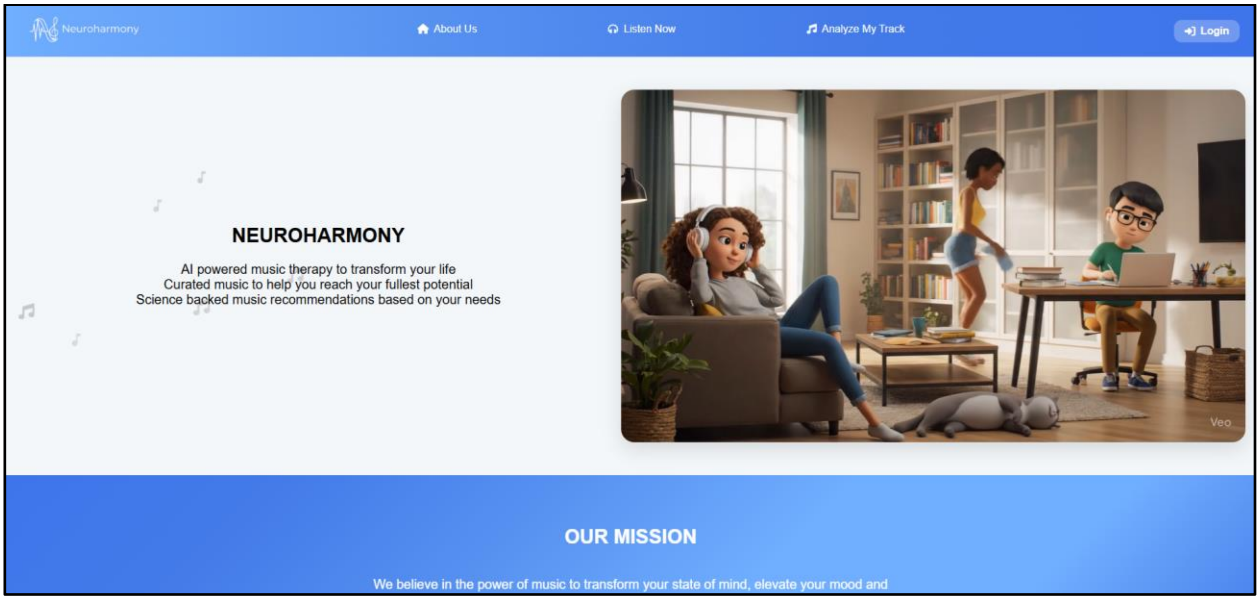

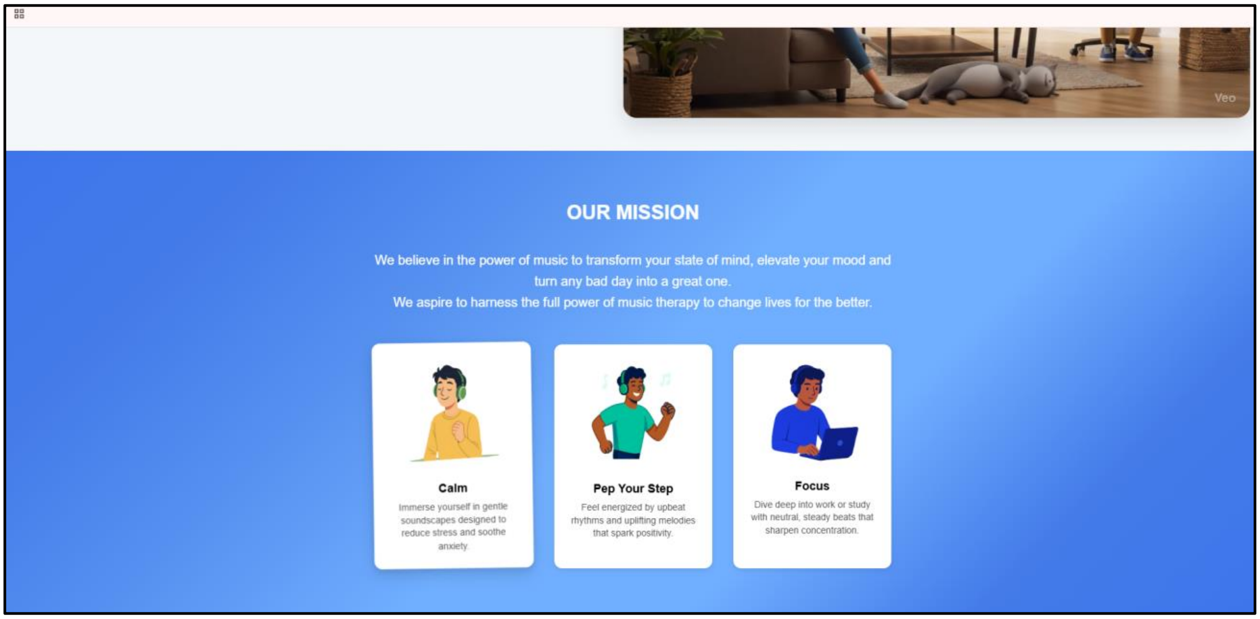

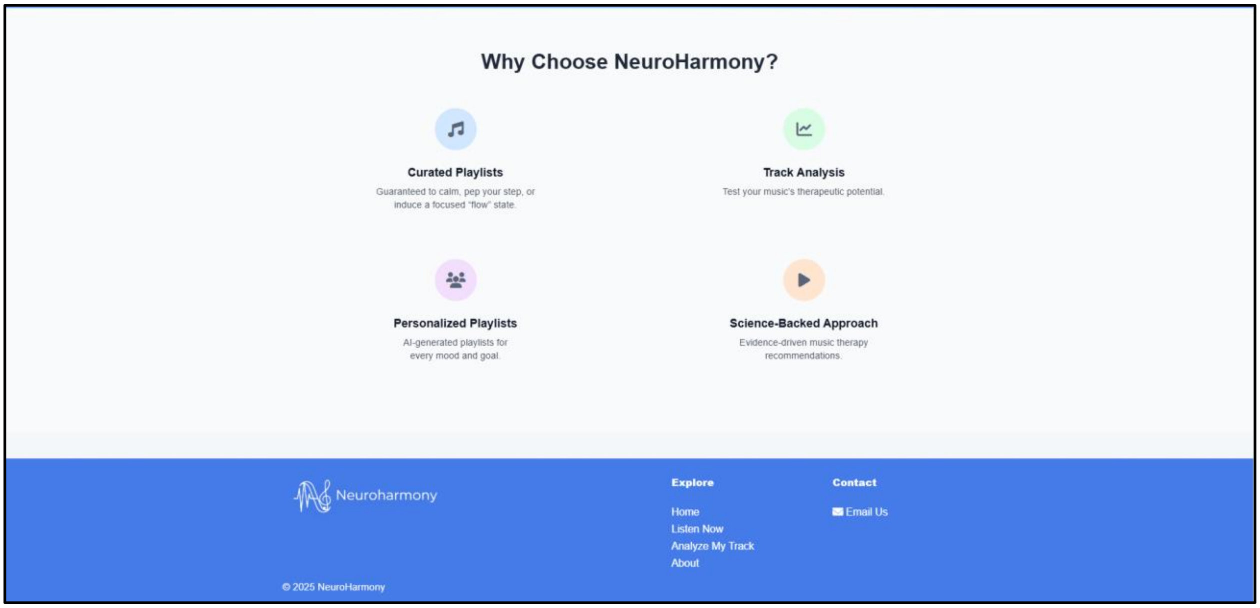
Why Choose NeuroHarmony?
Curated Playlists
Guaranteed to calm, pep your step, or induce a focused "flow" state.
Track Analysis
Test your music's therapeutic potential.
Personalized Playlists
AI-generated playlists for every mood and goal.
Science-Backed Approach
Evidence-driven music therapy recommendations.
Neuroharmony
Explore
Home
Listen Now
Analyze My Track
About
Contact
Email Us
© 2025 NeuroHarmony

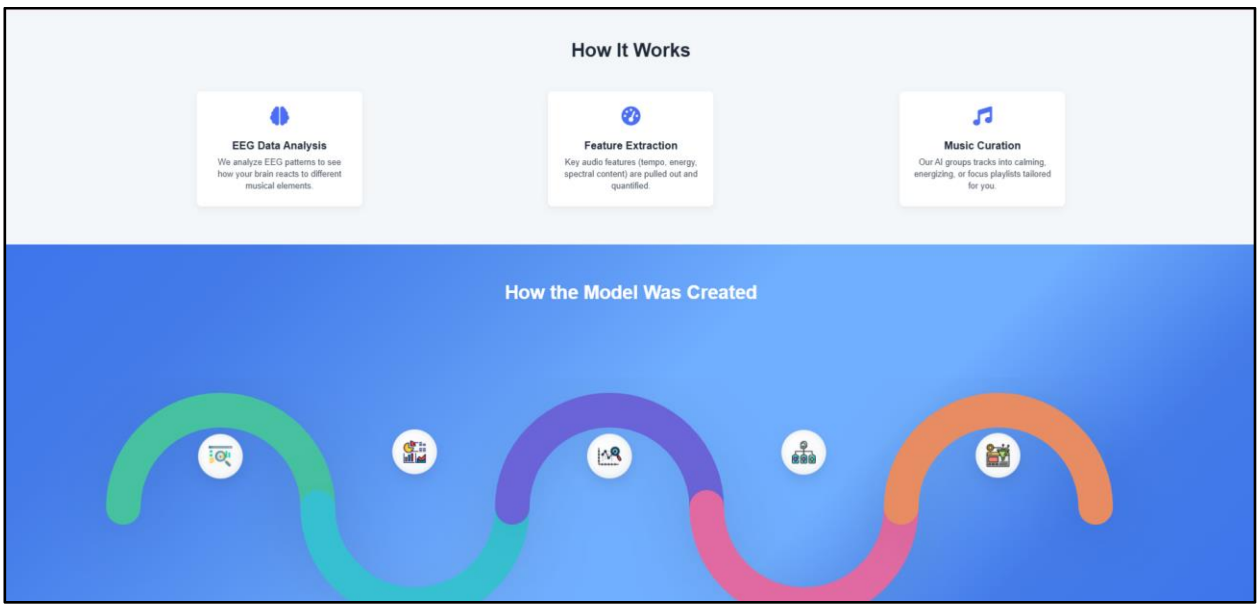
How It Works
EEG Data Analysis
We analyze EEG patterns to see how your brain reacts to different musical elements.
Feature Extraction
Key audio features (tempo, energy, spectral content) are pulled out and quantified.
Music Curation
Our AI groups tracks into calming, energizing, or focus playlists tailored for you.
How the Model Was Created

Log In Page:

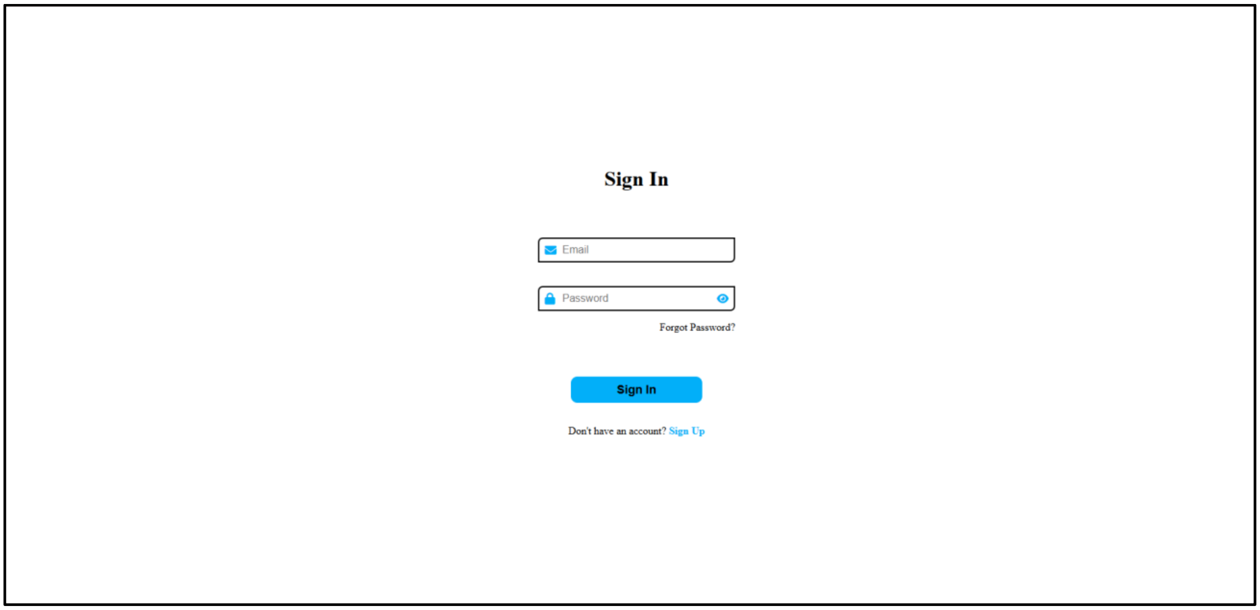

About Us Page:

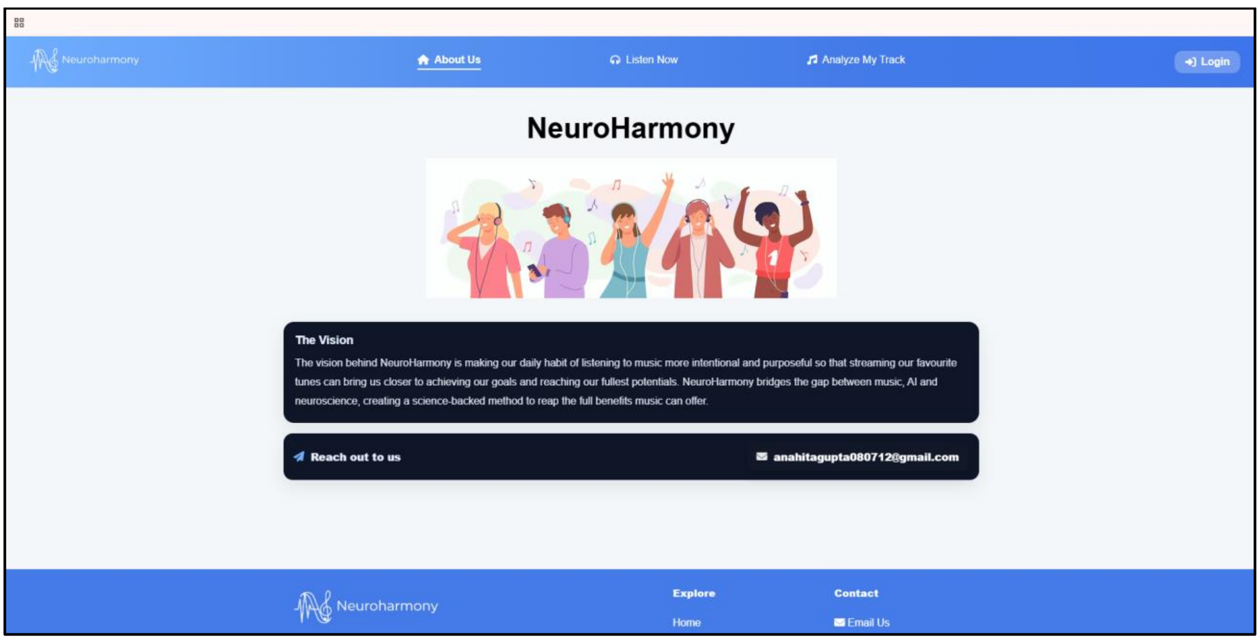

Listen Now Page:

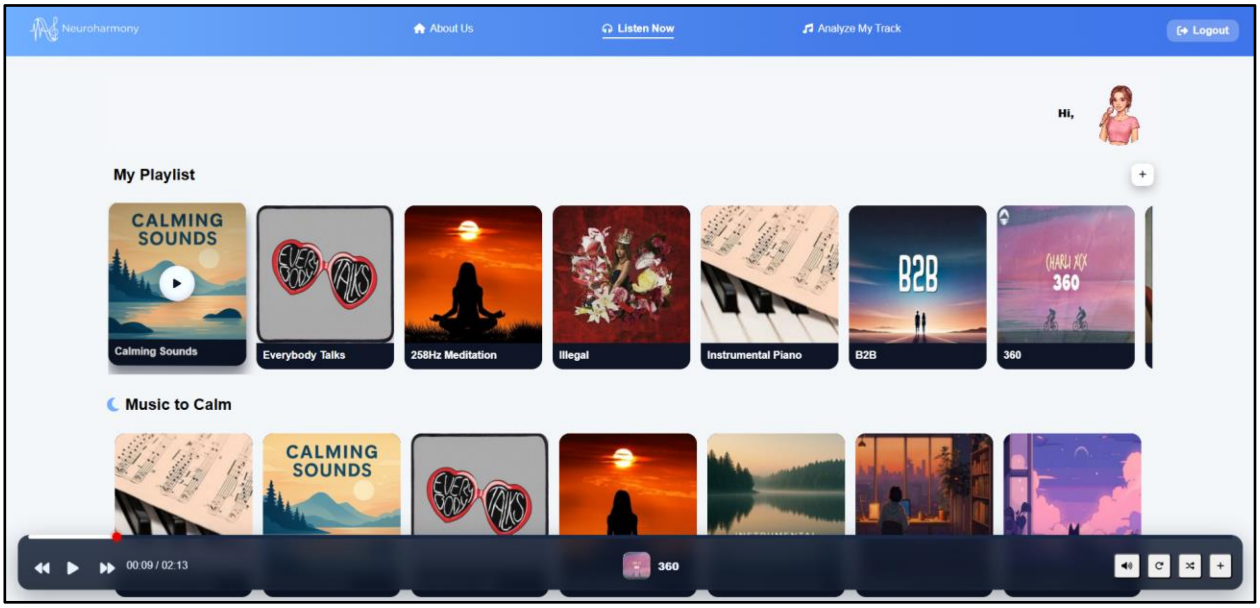

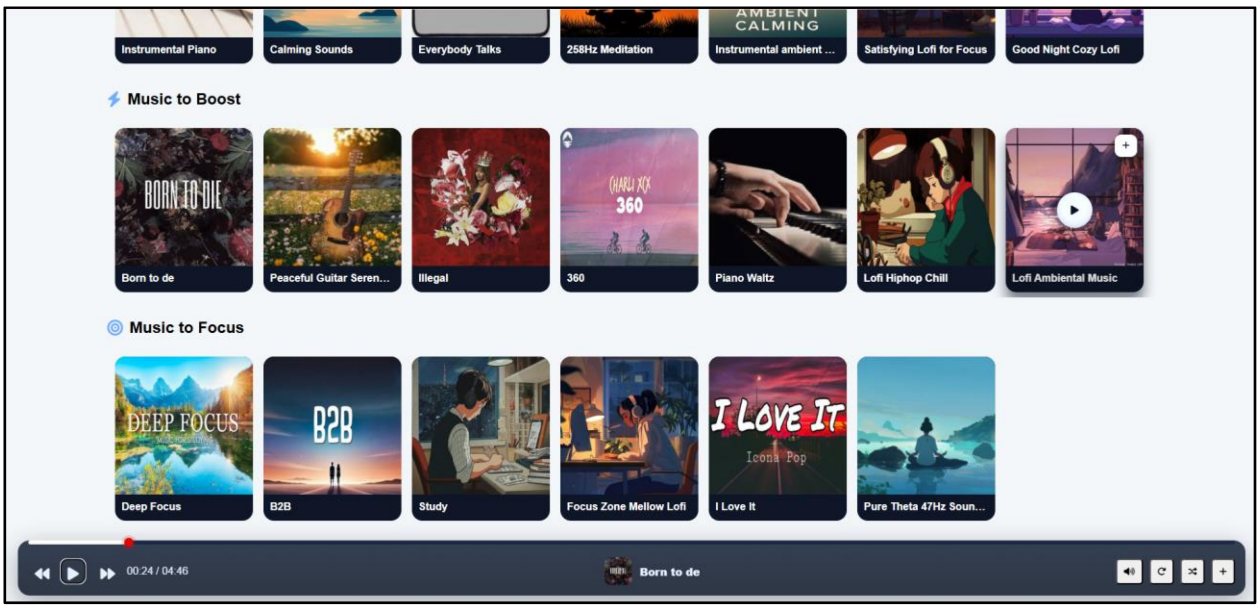

Analyze My Track Feature:

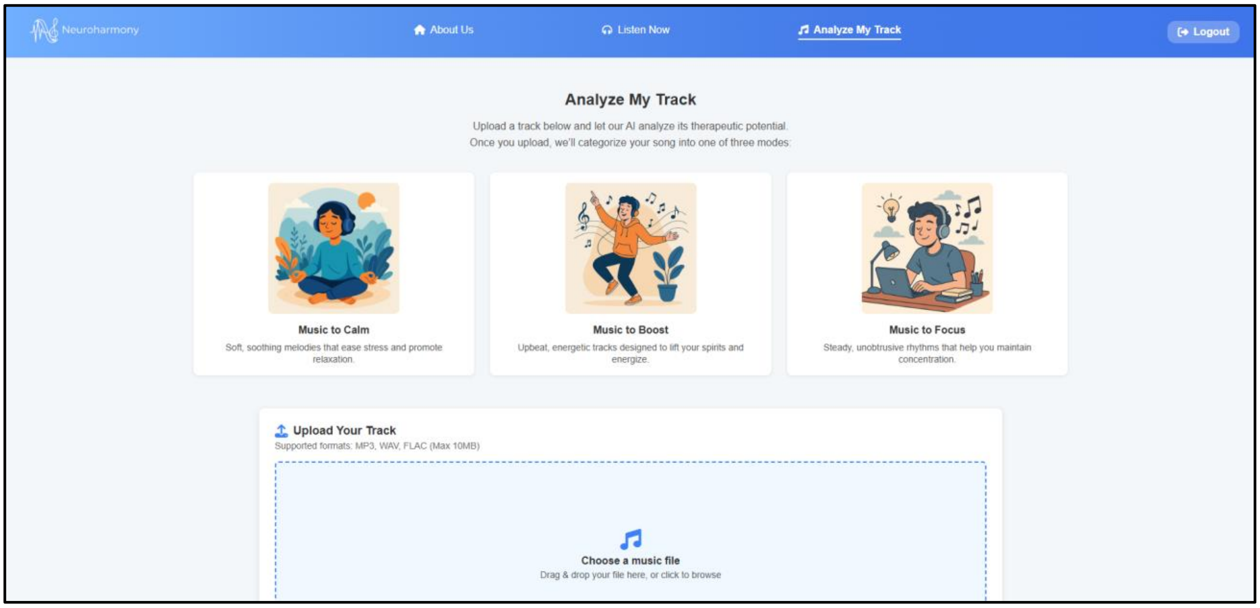

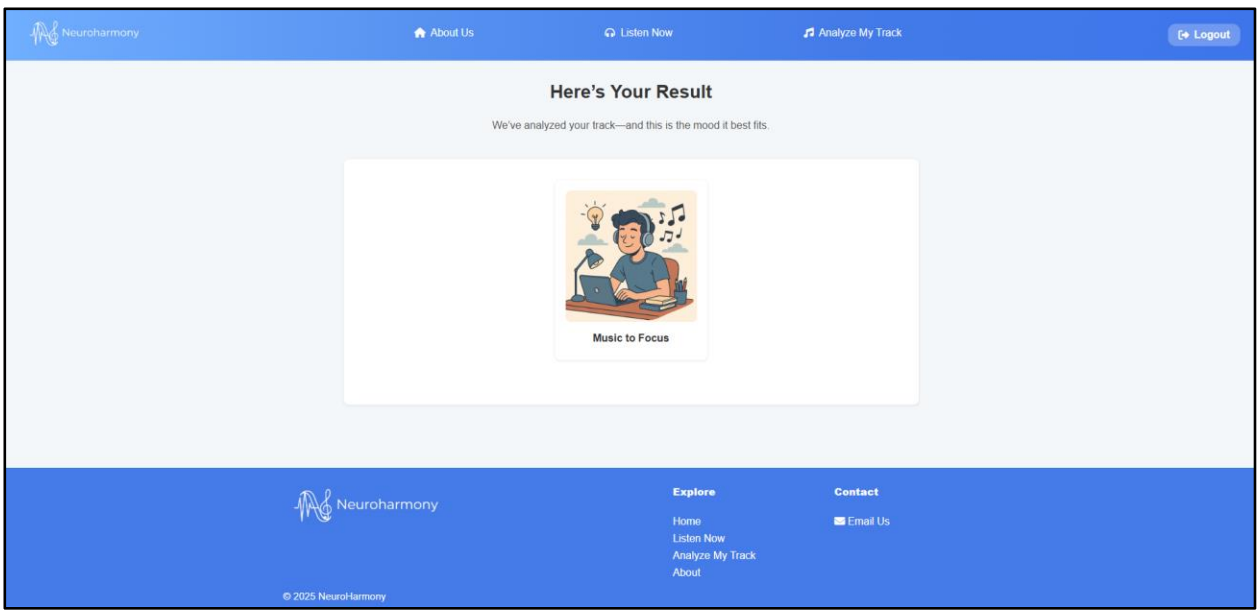

Link to Application:

https://neuroharmony-988798124110.asia-south1.run.app

Resources:

For my project I used the NMED-T dataset collected by Stanford University. This dataset contained EEG recordings from 20 participants as they listened to 10 different songs, using 125 electrode channels. Alongside the EEG data, it included participant ratings that reflected their subjective impressions of each song across various parameters.

I was guided throughout the project by my mentors at The Innovation Story. To deepen my understanding of EEG data collection and interpretation, I met with neurologist Dr. Swati Jeh. Her guidance helped me resolve several uncertainties. For example, she explained that I could focus specifically on the electrode channels connected to the temporal lobe, which was valuable because the dataset's 125 channels per participant made the data dense and challenging to process.

Dr. Jeh also clarified a common misconception I had about EEG graphs. I initially thought that the shape and amplitude of the waves were the most important features, but she explained that wave frequency is a much more reliable indicator of brain activity levels. Based on her advice, I focused on alpha waves (slower waves) and beta waves (faster, high-frequency waves), included theta waves to a lesser extent, and excluded delta waves, as they are primarily associated with deep sleep and unconsciousness.

Evaluation:

The NeuroHarmony project successfully combined neuroscience, audio signal processing, and AI to explore how music affects brain activity and listener engagement.

- **Music Classification Model:**
  The CNN-based classifier achieved 92% accuracy when classifying songs into three categories. The use of Mel spectrograms helped capture both spectral and temporal characteristics of music, leading to robust performance across genres.

- **Listener Preference Model:**
  A Random Forest Classifier trained on both EEG features and audio features achieved an accuracy of 71% in predicting whether a listener liked or disliked a song. This demonstrated the potential of combining brain and music data for understanding personal preferences.

- **"Analyze My Track" Feature:**
  Implemented using the Google Gemini API, this feature provided clear, concise summaries of the mood or cognitive purpose a song serves. The average latency was

around 30 seconds per track, depending on file size and processing load. Prompts were optimized to ensure accurate and user-friendly responses.

Reflection and Learning

When I started working on NeuroHarmony, I thought it would simply be about analysing EEG signals to understand how people respond to music. Very quickly I realised that brain data alone could not tell the full story. Music is more than just sound; it is rhythm, texture and emotion. I needed a way to capture those elements as well. That led me to include audio features such as tempo, timbre and loudness, and combine them with the EEG data.

This meant stepping into several fields I had never fully explored before. From my mentors I learned the more complex aspects of EEG processing and machine learning. On my own I studied audio analysis, experimenting with different methods to extract meaningful features from songs and match them with the brain data.

Every choice I made influenced the results. Selecting the right filters, identifying the most useful audio features and deciding on the models to train all played a part. One of the most important moments was when I combined song categories with both EEG and audio features. The predictions immediately became more accurate and felt more connected to how people actually experience music.

The most rewarding part of the project was how it brought together neuroscience, music and coding into one approach.

Ethical and Safety Concerns:

The project uses publicly available, anonymized EEG data, so there are no privacy concerns. NeuroHarmony does not collect or store any personal or medical information. It is not a medical tool and is not intended to diagnose or treat any condition.

Future Enhancements:

<u>Larger song catalogue</u>

- To improve the web application NeuroHarmony, I would apply the Music Classification model to a wider range of songs of varied genres. Having a larger catalogue of music would enhance the experience of using NeuroHarmony, cater to different individuals' personal tastes, and improve the users' experience.

Understanding more varied benefits of listening to music

- While I do believe the three fundamental purposes of music I selected are representative of most benefits of listening to music, I also recognise that each individual's experience of listening to music is unique. Different individuals may listen to music for different reasons, and the Analyse My Track model could thus be improved to categorise songs into a larger range of purposes.

Conclusion:

The NeuroHarmony project shows that combining EEG research with audio feature analysis can deepen our understanding of how music affects mood and engagement. While EEG was only used in the research stage to identify the most impactful musical features, those findings now power an application that works solely from audio input, making it widely accessible without specialised hardware.

The implications extend far beyond personal playlists. In music therapy, the system could help practitioners select songs with proven mood-enhancing properties. In education and workplace settings, it could support the creation of focus-oriented soundscapes. Public spaces could use it to design sound environments that promote calmness or alertness.

By grounding music recommendations in both neuroscience findings and audio signal processing, NeuroHarmony bridges the gap between scientific insight and everyday listening, opening the door to more beneficial uses of music in daily life.

Bibliography:

https://www.sciencedirect.com/science/article/pii/S1110016824008652 -

https://www.researchgate.net/publication/379983748_MUSIC_AND_THE_EFFECTS_ON_THE_BRAIN_AND_EMOTIONS

https://ieeexplore.ieee.org/abstract/document/8614971

https://www.obrela-journal.gr/index.php/obrela/article/view/306

https://ieeexplore.ieee.org/abstract/document/8416667

https://www.scirp.org/html/5-3900795_98060.htm

https://neuraldatascience.io/7-eeg/mne_python.html

https://pypi.org/project/librosa/

https://pypi.org/

https://flask.palletsprojects.com/en/stable/

https://firebase.google.com/