



The 2nd DeepDive Challenge

Music Genre Classification

Ordibehesht 1402



Classification is a fundamental problem in machine learning, where the goal is to assign a label or category to a given input based on its features. Music genre classification is a specific instance of this problem, where the goal is to assign a genre label to a given song based on its audio features. This task is challenging due to the subjectivity and ambiguity of music genres, as well as the complexity and variability of musical styles.

The "Prediction of Music Genre" dataset includes 50,000 songs from 10 different genres, along with various audio features extracted from each song. The goal of this challenge is to develop a machine learning model that can accurately classify songs into different genres based on their audio features.

However, music genre classification is a challenging problem due to several factors. First, music genres are inherently subjective and can vary depending on cultural and personal preferences. Second, the boundaries between different genres can be blurry, and some songs may contain elements of multiple genres. Finally, genre labels assigned by artists or distributors may not always be accurate or consistent.

Despite these challenges, music genre classification has several potential benefits. It can help music streaming services, radio stations, and music marketers to provide better recommendations and targeted advertising to users. It can also help music researchers to analyze trends and patterns in music over time and understand the evolution of musical styles.

The objectives of this task are to perform exploratory data analysis (EDA) on the "Prediction of Music Genre" dataset, and to develop and evaluate a machine learning model for music genre classification. The EDA will involve analyzing the distribution and correlations between the various audio features in the dataset, as well as identifying any outliers or missing values. The modeling and evaluation phase will involve developing a machine learning model that can accurately classify songs into different genres based on their audio features, and evaluating the model's performance using standard metrics such as accuracy, precision, recall, and F1 score.

The "Prediction of Music Genre" dataset includes the following features:

- **instance_id**: A unique identifier for each song instance in the dataset.
- **artist_name**: The name of the artist who performed the song.
- **track_name**: The name of the song track.
- **popularity**: A measure of the song's popularity, based on the number of times it has been played and downloaded.
- **acousticness**: A measure of the song's acoustic characteristics, with higher values indicating a more acoustic sound.
- **danceability**: A measure of the song's danceability, based on factors such as tempo, rhythm stability, and beat strength.
- **duration_ms**: The duration of the song in milliseconds.
- **energy**: A measure of the song's energy, with higher values indicating a more energetic sound.

- **instrumentalness**: A measure of the song's instrumental content, with higher values indicating a more instrumental sound.
- **key**: The key of the song, represented as a character between A and G, with additional values for sharps and flats.
- **liveness**: A measure of the song's liveliness, based on the presence of audience sounds or other indications of a live performance.
- **loudness**: A measure of the song's loudness, in decibels (dB).
- **mode**: The mode of the song, represented as either minor or major.
- **speechiness**: A measure of the song's spoken word content, with higher values indicating a more spoken word-heavy sound.
- **tempo**: The tempo of the song in beats per minute (BPM).
- **obtained_date**: The date on which the song data was obtained.
- **valence**: A measure of the song's emotional valence, with higher values indicating a more positive or happy sound.
- **music_genre**: The genre label assigned to the song, which is the target variable for the machine learning model.