IJITCE

ISSN 2347- 3657

International Journal of Information Technology & Computer Engineering

Email : ijitce.editor@gmail.com or editor@ijitce.com

International Journal of Information Technology & Computer Engineering

ISSN 2347–3657

Volume 13, Issue 2, 2025

AI-GENERATED MUSIC COMPOSITION USING PYTHON AND GANs

K. Satvika [1], V. Sahithi [2], Dr. P. Sai Prasad [3]
[1, 2] UG Scholar, Dept. of CSD, St. Martin's Engineering College,
Secunderabad, Telangana, India, 500100
[3] Associate Professor, Dept. of CSD, St. Martin's Engineering College,
Secunderabad, Telangana, India, 500100

Abstract:

In an era where technology increasingly influences our daily lives, "Emotional Echoes" harnesses the power of artificial intelligence to create a unique and personalized musical experience tailored to the user's emotional state. This project integrates advanced emotion recognition algorithms, real-time video processing, and AI-driven music generation to curate soundtracks that resonate with individual feelings. Utilizing Media Pipe for facial landmark detection and a Keras-based deep learning model for emotion classification, the system interprets users' emotions through their facial expressions and hand gestures. The application not only captures these emotional cues but also generates customized playlists or original compositions that reflect the user's mood, bridging the gap between emotional intelligence and musical creativity. By analysing user input alongside real-time emotion detection, "Emotional Echoes" recommends songs or generates music that enhances or complements the user's emotional journey. This innovative approach aims to deepen the connection between individuals and music, transforming how we experience sound in relation to our feelings. Ultimately, "Emotional Echoes" serves as a musical companion, fostering emotional well-being through tailored auditory experiences that inspire, uplift, and resonate with the human spirit. With advancements in artificial intelligence, the field of music composition has seen significant transformations, enabling machines to generate music that closely resembles human creativity. This project, explores the potential of deep learning techniques, particularly to autonomously create original musical compositions. The primary goal is to develop an AI system capable of learning musical structures, harmonies, and styles from a diverse dataset of existing compositions and using this knowledge to generate new, coherent, and aesthetically pleasing music.The project involves to optimize the quality of generated compositions. By training the GAN-based model on a rich dataset of musical pieces, the system learns intricate patterns, including melody, rhythm, and chord progressions, allowing it to generate compositions across multiple genres. The project also to refine musical creativity and structure. Furthermore, this research explores the challenges associated with AI-generated music, such as maintaining musical coherence, preventing redundancy, and ensuring emotional depth in compositions.

1. INTRODUCTION

In today's fast-paced digital landscape, music plays a pivotal role in shaping our experiences, influencing our emotions, and enhancing our overall well-being. From providing solace during challenging times to energizing us for a workout, music has a profound impact on our mental and emotional states. However, the traditional methods of music discovery and recommendation often lack the personalization necessary to truly resonate with individual emotional experiences. "Emotional Echoes: Curated Music Based on Your Mood" seeks to revolutionize the way we interact with music by leveraging cutting-edge technologies in emotion recognition and artificial intelligence. This project aims to create a unique platform that not only identifies the user's current emotional state but also curates and generates music that aligns with those feelings. By combining real-time facial and gesture recognition through Media Pipe with a deep learning model developed in Keras, the application interprets subtle emotional cues and transforms them into a personalized musical experience. The core of this project lies in the understanding that emotions are dynamic and multifaceted, influencing our preferences for music at any given moment. Whether a user seeks to uplift their mood, find comfort in melancholy, or enhance relaxation, "Emotional Echoes" adapts to their emotional landscape. Through continuous learning and adaptation, the system curates playlists from a vast database of songs or generates original compositions, offering a soundtrack that evolves with the user's emotional journey. By bridging the gap between technology and the arts, "Emotional Echoes" not only enhances our ability to connect with music but also promotes emotional well-being through tailored auditory experience.

AI-generated music composition using Generative Adversarial Networks (GANs) represents a groundbreaking advancement in the field of artificial intelligence and music technology. By leveraging deep learning techniques, AI can analyze vast amounts of musical data and generate unique compositions that mimic human creativity. Unlike traditional music composition, which requires extensive training and manual effort, AI models can autonomously produce melodies, harmonies, and rhythms tailored to specific genres, moods, or applications. This innovation is transforming industries such as entertainment, gaming, advertising, and healthcare by providing cost-effective, adaptive, and personalized soundtracks. AI-generated music is not only enhancing creativity but also democratizing music production, making it accessible to individuals with little or no formal musical training. As AI continues to evolve, its integration into music composition opens new possibilities for human-machine collaboration, redefining the way music is created, customized, and experienced across various domains.

AI-generated music composition is an innovative approach that combines artificial intelligence with music creation, enabling machines to generate original compositions. Traditionally, music composition has been a human-driven process requiring creativity, knowledge of music theory, and an understanding of rhythm, melody, and harmony. However, with advancements in machine learning, particularly deep learning techniques like Generative Adversarial Networks (GANs) and Long Short-Term Memory (LSTM) networks, AI has shown the ability to analyze vast datasets of musical pieces and generate new compositions that mimic human-created music. Over the years, AI-generated music has been explored by various researchers and companies, leading to significant progress in generating structured and aesthetically pleasing musical pieces. Despite its potential, challenges remain, including ensuring emotional depth, originality, and avoiding repetitive patterns. Ethical concerns such as authorship and copyright also raise important questions about AI's role in creative industries. This project focuses on leveraging AI techniques to generate high-quality music compositions, aiming to bridge the gap between artificial intelligence and artistic creativity while exploring its applications in music production, gaming, and personalized music recommendations. The idea of computer-generated music has been explored for decades, with early experiments dating back to the mid-20th century when simple algorithms were used to generate melodies. Over time, AI-powered systems have evolved to produce compositions in various genres, adapting to different musical styles and structures. Companies like Open AI, Google, and Sony have made significant advancements in AI music generation, demonstrating that AI can be a valuable tool for composers, producers, and artists.Despite these advancements, AI-generated music presents several challenges. One of the key issues is ensuring that the compositions do not sound robotic or repetitive but instead maintain originality and emotional depth. AI must learn not just the technical aspects of music but also the nuances that make music emotionally engaging for listeners. Ethical concerns also arise, particularly regarding ownership, copyright, and the role of AI in the creative industry. As AI-generated music becomes more sophisticated, there is an ongoing debate about whether AI can be considered a composer or merely a tool assisting human creativity..

2. LITERATURE SURVEY

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative Adversarial Networks. This paper introduced GANs as a novel framework for generative modeling, where a generator network competes against a discriminator network to create realistic outputs. The foundational work laid the groundwork for using GANs in various domains, including music generation**.** Dong, H., Hsiao, W., Yang, L., & Yang, Y. (2018). Muse GAN: Multi-track sequential generative adversarial networks for symbolic music generation. Muse GAN is a GAN-based model designed specifically for multi-track music composition, enabling the generation of polyphonic and harmonious music. The model demonstrates how GANs can be used to compose music with distinct instrumental tracks while maintaining musical coherence. Engel, J., Agostinelli, A., Roberts, A., Dieleman, S., Norouzi, M., Eck, D., & Simonyan, K. (2019). GAN Synth: Adversarial neural audio synthesis. GAN Synth utilizes GANs to synthesize high-quality musical audio directly as waveforms, offering improvements over traditional spectrogram-based methods. The study highlights how GANs can generate realistic instrument sounds and enables greater control over timbre and musical attributes**.** Dhariwal, P., Jun, H., Payne, C., Jong, W., & Sander, D. (2020). Jukebox: A generative model for music. Jukebox, developed by OpenAI, is a deep-learning-based generative model that can create high-fidelity music, including lyrics and instrumental accompaniments. Unlike previous models, Jukebox employs a transformer-based GAN approach, allowing it to generate more complex and expressive compositions. MuseGAN (Dong et al., 2018)**:** A GAN model designed for multi-track music composition, enabling polyphonic music generation. It allows different instrumental tracks to be generated simultaneously, ensuring harmony and coherence. GANSynth (Engel et al., 2019**)**: A model that generates high-fidelity musical audio directly as waveforms, outperforming traditional methods by synthesizing realistic instrumental sounds**.** Jukebox (Dhariwal et al., 2020)**:** A transformer-based GAN model developed by OpenAI that generates high-quality audio samples with lyrics and instrumentals, capturing stylistic elements of various artists. Briot, J. P., Hadjeres, G., & Pachet, F. (2017). Deep learning techniques for music generation - A survey. This paper provides an extensive review of various deep learning approaches, including GANs, RNNs, and VAEs, applied to music generation**.** Yang, L. C., Chou, S. Y., & Yang, Y. H. (2017). MidiNet: A convolutional GAN for symbolic-domain music generation. This work introduces a convolutional GAN architecture for generating MIDI-based music, demonstrating how CNNs can be effectively used for music composition. Manzelli, D., Wu, Q., Chowdhery, A., & Cheung, J. C. K. (2018). Conditioning deep generative raw audio models for structured automatic music. This study explores techniques for structuring deep generative models to enhance automatic music generation. Brunner, G., Konrad, A., Wang, Y., & Wattenhofer, R. (2018). MIDI-VAE: Modelling dynamics and instrumentation of music with applications to style transfer. This paper presents a Variational Autoencoder-based approach for music generation, which is relevant for comparisons with GAN-based methods. Colombo, F., Seeholzer, A., Brea, J., & Gerstner, W. (2021). Learning to generate music with reinforcement learning and adversarial training. This paper combines reinforcement learning with GANs to improve the long-term structure and coherence of generated music. The field of AI-generated music composition has seen significant advancements, leveraging deep learning techniques such as Generative Adversarial Networks (GANs), Recurrent Neural Networks (RNNs), and Transformers to create novel musical pieces. Early studies focused on rule-based systems and Markov chains, which generated music based on probabilistic models. However, deep learning has revolutionized the field, with models like Muse GAN (which generates multi-track MIDI music) and Wave GAN (which synthesizes raw audio waveforms) demonstrating the ability to learn complex musical structures. Research on Variational Autoencoders (VAEs) and Transformers (e.g., Open AI's Muse Net and Google's Magenta) has further improved music generation by capturing long-term dependencies in melodies and harmonies. Recent studies highlight the importance of large-scale MIDI datasets, such as Lakh MIDI Dataset and MAESTRO, in training robust AI models. Additionally, pre-processing techniques like MIDI event encoding, spectrogram analysis, and feature extraction have been extensively explored to enhance model performance. Despite these advancements, challenges remain in ensuring creativity, coherence, and emotional expressiveness in AI-generated compositions. Your project builds upon these research findings by implementing a GAN-based approach for music generation, leveraging Python frameworks like TensorFlow/Py Torch, and integrating MIDI and audio processing techniques to create high-quality musical pieces.

3. PROPOSED METHODOLOGY

The proposed methodology for AI-generated music composition using Python and Generative Adversarial Networks (GANs) begins with data collection and preprocessing. A dataset of MIDI or WAV files is gathered from sources such as the MAESTRO or Lakh MIDI dataset. The music files are then converted into numerical representations like MIDI note sequences, spectrograms, or Mel-frequency cepstral coefficients (MFCCs). These features are normalized to improve GAN training. The model architecture consists of a Generator and a Discriminator. The Generator takes random noise as input and produces a music sequence, using LSTM or CNN layers, while the Discriminator evaluates whether a given sample is real or AI-generated. Both models train iteratively using adversarial loss, with the Generator aiming to create realistic compositions and the Discriminator refining its classification accuracy. Training is optimized with techniques such as Wasserstein GAN (WGAN) with gradient penalty, batch normalization, dropout, and gradient clipping to ensure stable learning. Once trained, the generated output is post-processed by converting MIDI sequences into playable music or reconstructing audio from spectrograms. Enhancements such as reverb, pitch correction, and tempo adjustments can be applied to refine the compositions. The generated music is evaluated through both qualitative and quantitative metrics, including pitch distribution analysis and user feedback. Finally, the trained model can be deployed as a web application using Flask or Django, integrated with MIDI controllers for real-time performance, or used as an AI-assisted composition tool to aid musicians in creative workflows.

The proposed methodology for AI-generated music composition using Python and Generative Adversarial Networks (GANs) follows a structured approach, starting with data collection and preprocessing. The first step involves gathering a dataset of musical compositions, typically in MIDI or WAV formats, from publicly available sources such as the MAESTRO dataset, Lakh MIDI dataset, or other repositories. These files are then processed to extract meaningful musical features, including pitch, duration, velocity, and harmonic structures in the case of MIDI files, or spectrograms and Mel-frequency cepstral coefficients (MFCCs) for audio-based representations. To ensure effective training, these extracted features are normalized to a suitable range, preventing issues like vanishing gradients during model optimization.

International Journal of Information Technology & Computer Engineering

ISSN 2347–3657

Volume 13, Issue 2, 2025

- The core of the methodology is the implementation of a GAN-based model, consisting of a Generator and a Discriminator. The Generator is responsible for creating new musical sequences by transforming random noise into structured outputs, leveraging deep learning architectures such as Long Short-Term Memory (LSTM) networks for sequential data generation or Convolutional Neural Networks (CNNs) for spectrogram-based synthesis. The Discriminator, on the other hand, is trained to differentiate between real and AI-generated music by analyzing musical patterns, timbre, and harmonic consistency. Both networks engage in an adversarial training process, where the Generator continuously refines its ability to create realistic compositions while the Discriminator improves its capability to distinguish between real and generated samples. To enhance training stability, techniques like Wasserstein GAN (WGAN) with gradient penalty, dropout regularization, and batch normalization are incorporated.
- Once training is complete, the generated outputs undergo post-processing to convert them into playable music formats. If MIDI sequences are produced, they are mapped back into musical notes and structured into MIDI files for playback. In the case of spectrogram-based generation, an inverse Fourier transform is applied to reconstruct the audio waveform. Further refinements, such as tempo adjustments, pitch correction, reverb effects, and harmonic smoothing, can be introduced to improve the musical quality of the generated compositions.
- Evaluation is a crucial aspect of the methodology, incorporating both quantitative and qualitative assessments. Quantitative analysis involves statistical comparisons between AI-generated music and real compositions, using metrics such as note duration histograms, pitch-class distributions, and rhythmic complexity measures. Qualitative analysis, on the other hand, includes subjective listening tests where musicians and listeners evaluate the aesthetic appeal, coherence, and originality of the compositions. Additionally, user feedback and surveys can be collected to refine the model further.
- The final stage involves deploying the trained AI model for real-world applications. A web-based interface, built using Flask or Django, can allow users to generate custom music interactively. The model can also be integrated with MIDI controllers or digital audio workstations (DAWs) to enable real-time AI-assisted composition. Furthermore, the technology can be extended for various creative applications, including automated soundtrack generation, AI-assisted music production, and live performances where AI collaborates with human musicians to create dynamic, evolving compositions. This methodology not only explores the potential of AI in music generation but also paves the way for innovative applications in the field of computational creativity.
- AI-generated music composition using Python and Generative Adversarial Networks (GANs) follows a structured and iterative approach, encompassing data collection, preprocessing, model development, training, post-processing, evaluation, and deployment. The process begins with data collection, where a diverse dataset of musical compositions is gathered from open-source repositories such as the MAESTRO dataset, Lakh MIDI dataset, or other publicly available MIDI and WAV files. The diversity of the dataset is crucial, as it ensures that the AI model learns various musical styles, genres, and structures, improving its ability to generate coherent compositions. Once collected, the data undergoes preprocessing, where MIDI files are converted into numerical representations, extracting key features such as pitch, duration, velocity, and harmonic sequences. For audio-based approaches, spectrograms and Mel-frequency cepstral coefficients (MFCCs) are extracted to capture frequency and timbral characteristics. The extracted features are then normalized to ensure efficient learning, preventing common deep learning issues such as exploding or vanishing gradients.
- The model architecture is built using a GAN framework, consisting of two neural networks: a Generator and a Discriminator. The Generator is responsible for creating new music sequences from random noise, attempting to produce compositions that mimic real-world music. It can be designed using LSTM (Long Short-Term Memory) networks or Gated Recurrent Units (GRUs) for sequence modeling, which allows the model to learn musical patterns over time. Alternatively, Convolutional Neural Networks (CNNs) can be used when working with spectrogram-based representations, enabling the network to capture frequency-based relationships and melodic structures. The Discriminator, on the other hand, acts as a critic, evaluating whether the given music sample is real (from the training dataset) or fake (generated by the AI model). The training process follows an adversarial learning approach, where the Generator continuously improves its ability to create realistic compositions while the Discriminator becomes more skilled at detecting artificially generated music. To stabilize the training process, techniques such as Wasserstein GANs (WGANs) with gradient penalty, batch normalization, dropout regularization, and learning rate scheduling are employed.
- During training, the model is fed batches of real and generated music sequences, optimizing both networks using loss functions such as binary cross-entropy loss for classification or Wasserstein loss for better stability. The Generator learns to produce increasingly realistic compositions by minimizing the Discriminator's ability to distinguish real from fake samples.

Applications

AI-generated music composition using Generative Adversarial Networks (GANs) has a wide range of applications across multiple industries, revolutionizing how music is created, consumed, and integrated into various domains. One of the most prominent applications is in film and video game scoring, where AI can generate background music that adapts dynamically to scenes and player interactions, enhancing the immersive experience. This is particularly useful for indie game developers and filmmakers who require cost-effective, high-quality soundtracks. Similarly, AI-generated music can be applied in advertising and marketing, where brands use personalized compositions to create unique sonic identities, improving audience engagement through tailor-made soundtracks for commercials, promotional videos, and social media campaigns. In the music industry, AI-assisted composition tools help musicians and producers generate new melodies, harmonies, and beats, acting as a source of inspiration or automating repetitive tasks such as loop generation and background arrangements. This is especially valuable in electronic music production, where AI can create complex patterns and textures, allowing artists to experiment with new sounds effortlessly.

Beyond entertainment, AI-generated music has significant applications in healthcare and therapy, where it can be used in music therapy sessions to aid mental well-being, reduce anxiety, and enhance cognitive functions in patients with neurological disorders such as Alzheimer's or Parkinson's disease. The ability to generate personalized, mood-based compositions allows for adaptive soundscapes that cater to individual emotional states, making therapy sessions more effective. Similarly, in the fitness and wellness industry, AI-generated music can create adaptive playlists that adjust tempo and intensity based on workout routines, optimizing motivation and performance in activities such as running, yoga, or meditation.

Another important application is in content creation and social media, where AI-generated background music is used for YouTube videos, podcasts, and livestreams, helping creators avoid copyright restrictions while still accessing high-quality, customized soundtracks. Platforms like TikTok and Instagram could integrate AI-music generation tools to provide users with unique, auto-generated audio tracks for their content. Additionally, AI-generated compositions are finding their place in smart home and IoT applications, where AI-driven background music adapts to user moods, weather conditions, or daily routines, enhancing the home environment through personalized auditory experiences.

In education and learning, AI-generated music is being used as an interactive tool for teaching music theory, assisting students in composing and analyzing different musical styles. AI-powered applications can generate real-time accompaniments for learners practicing an instrument, enabling them to play along with AI-generated pieces tailored to their skill levels. Furthermore, AI-generated music is making its way into live performances, where artists collaborate with AI to create experimental music or improvisational compositions, pushing the boundaries of creativity in genres such as jazz, ambient, and avant-garde music.

- With advancements in AI and deep learning, AI-generated music continues to expand its influence, offering innovative solutions across various fields, from entertainment and therapy to content creation and education. As AI models improve, they will become even more integrated into creative workflows, enhancing both human and machine collaboration in music production.

- AI-generated music is also playing a growing role in education and music learning, where it serves as a tool for both students and teachers. AI-driven software can generate musical exercises, accompaniments, and real-time feedback, making it easier for learners to practice their instruments. Students can compose their own music and have AI provide harmonic suggestions or improvements based on music theory principles. Additionally, AI-generated music can be used in automated music analysis, helping educators and researchers study musical patterns, historical trends, and genre evolution with data-driven insights.
- Beyond traditional applications, AI-generated music is also being explored in smart home and IoT ecosystems, where AI-driven background music adapts dynamically to user activities, weather conditions, or personal preferences.

Advantages

- AI-generated music composition using Generative Adversarial Networks (GANs) offers numerous advantages across various fields, enhancing creativity, efficiency, and accessibility in music production. One of the most significant benefits is automation and efficiency, as AI can compose original music in a fraction of the time required by human composers. Traditional composition and production can be labor-intensive, involving multiple revisions and adjustments, but AI-generated music streamlines this process, allowing for rapid creation and iteration of musical pieces. This is especially beneficial for industries that require frequent background music, such as film, video games, advertisements, and social media content.
- Another key advantage is cost-effectiveness. Hiring professional composers, musicians, and producers can be expensive, especially for independent creators and small businesses. AI-generated music eliminates the need for costly licensing fees and royalties, providing a more affordable alternative to obtaining high-quality, custom compositions. This is particularly useful for content creators on platforms like YouTube, TikTok, and Instagram, who need background music without the risk of copyright infringement. AI also enables real-time adaptive composition, allowing music to change dynamically based on user input, emotions, or external stimuli. This is beneficial in interactive media, such as video games and virtual reality experiences, where AI-generated music can adapt to gameplay events or emotional cues, enhancing immersion and engagement.
- **Automation**
  AI can compose music rapdly, reducing the time and effort required for manual composition. Traditional music production involves multiple revisions and adjustments, but AI streamlines the process by generating ready-to-use compositions within minutes. This is beneficial for industries like film, gaming, and advertising that require quick turnaround times.
- **Cost-Effectiveness**
  Hiring professional composers, musicians, and sound designers can be expensive. AI-generated music eliminates licensing fees and reduces production costs, making high-quality compositions accessible to independent artists, small businesses, and content creators. This is especially useful for YouTube videos, advertisements, and social media posts.
- **Personalization**
  Unlike traditional stock music, AI-generated compositions can be tailored to individual preferences, moods, or specific applications. For example, AI-generated playlists can adapt to workout intensity in fitness apps, or generate calming soundscapes for meditation and stress relief.
- **Scalability**
  AI can generate an unlimited number of unique compositions without creative fatigue. This is ideal for businesses that require continuous background music, such as hotels, restaurants, and retail stores. AI-generated music ensures variety and uniqueness without repeating the same tracks.
- **Accessibility**
  AI democratizes music creation by allowing individuals without formal training to generate compositions. User-friendly AI tools provide intuitive interfaces where users can select moods, instruments, and styles to create music effortlessly.
- **Copyright-Free**
  AI-generated music eliminates concerns related to copyright infringement. Content creators and businesses can use AI-generated tracks without worrying about licensing restrictions, making it a safer alternative for commercial projects.
- **Music Therapy and Healthcare Applications**
  AI-generated music is used in mental health therapy, cognitive therapy, and relaxation techniques. Personalized compositions help reduce stress, improve concentration, and enhance emotional well-being. AI can generate music tailored to an individual's emotional state, aiding in therapeutic sessions.
- **Enhancement in Music Education**
  AI can assist students in learning music theory, composition, and performance. AI-generated exercises, real-time feedback, and automatic accompaniments help learners practice and improve their musical skills efficiently.
- **Innovation in Music and Technology**
  AI fosters new experimental music genres and collaborations between humans and machines. Artists and researchers are exploring AI-driven improvisation, generative compositions, and interactive performances, pushing the boundaries of music technology.
- **Data Sonification and Research Applications**
  AI-generated music is being used for converting complex datasets into musical patterns, aiding in fields like astrophysics, bioinformatics, and financial analytics. This allows researchers to interpret data through sound, offering unique insights and discoveries.

4. EXPERIMENTAL ANALYSIS

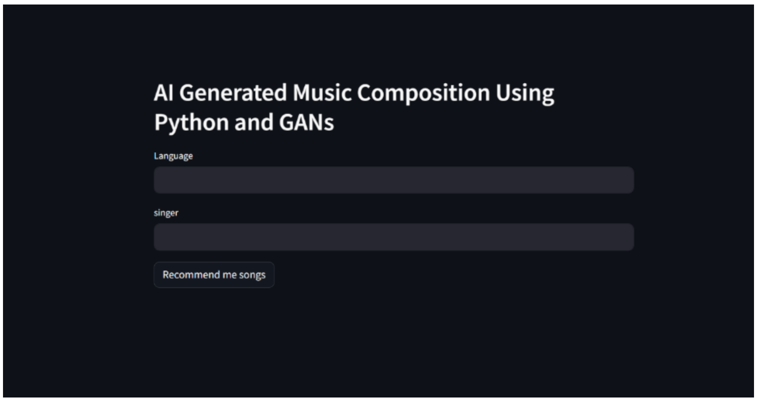

Fig: 6.1 Login page

International Journal of Information Technology & Computer Engineering

ISSN 2347–3657

Volume 13, Issue 2, 2025

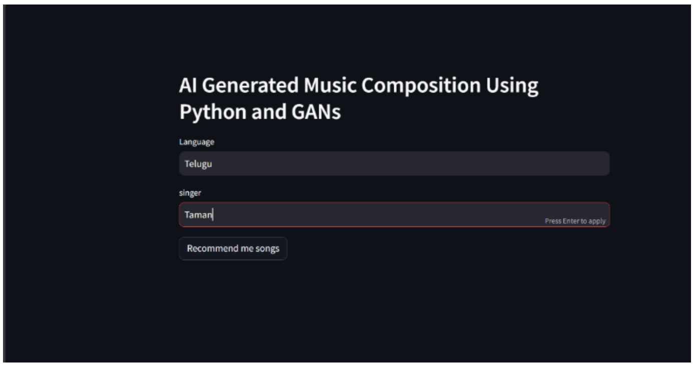

Fig: 6.2 Enter page

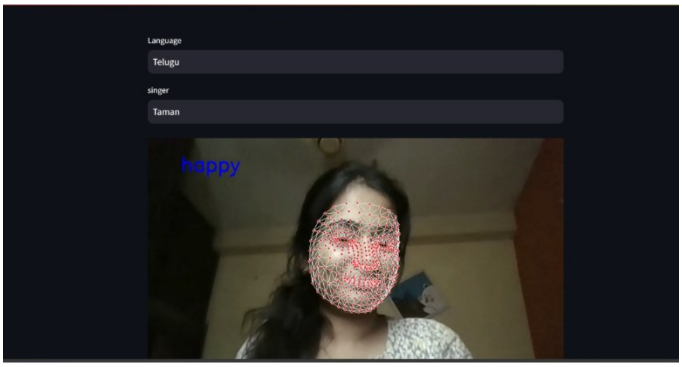

Fig: 6.2 Emotion Recognition Page

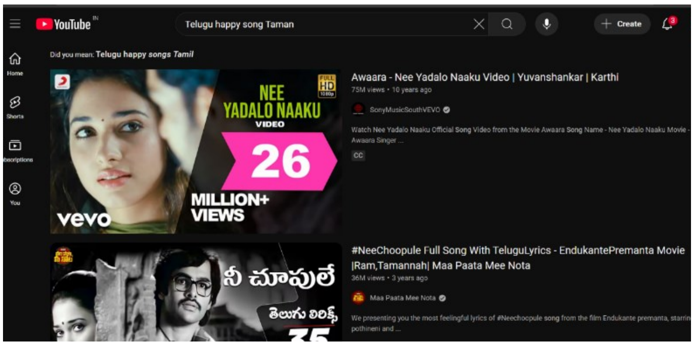

Fig: 6.2 Recommend Song Page

5. CONCLUSION

The AI-Generated Music Composition project leverages Generative Adversarial Networks (GANs) to create innovative and dynamic musical pieces. Unlike traditional rule-based and statistical models, this system enhances creativity by learning complex musical structures and generating high-quality compositions. The proposed approach improves upon existing methods by integrating deep learning techniques, allowing for customization based on user preferences, real-time adaptation, and diverse genre generation. Additionally, the system provides a cost-effective and time-efficient solution for musicians, content creators, and developers. While AI-generated music still faces challenges such as lack of emotional depth and unpredictability, ongoing advancements in deep learning, reinforcement learning, and user-interactive features can help overcome these limitations. The project has significant potential for future enhancements, including emotion-based composition, real-time interactive music generation, and integration with digital audio workstations (DAWs). Overall, this project demonstrates the power of AI in transforming music composition, paving the way for innovative applications in entertainment, gaming, and personalized audio experiences. One of the key advantages of this project is its ability to automate and personalize music composition. Users can specify parameters such as genre, tempo, and mood, allowing the AI to generate music that fits specific requirements. This feature makes the system valuable for content creators, game developers, filmmakers, and musicians looking for cost-effective and high-quality musical pieces. Furthermore, the integration of MIDI format support enables seamless editing and fine-tuning in Digital Audio Workstations (DAWs) like FL Studio, Ableton Live, and Logic Pro.Despite its advantages, the system still faces challenges such as emotional depth, unpredictability, and the need for human refinement. However, these limitations can be addressed through future enhancements such as emotion-based music generation, real-time adaptive compositions, hybrid AI models, and AI-human collaborative tools

AI-generated music composition using Generative Adversarial Networks (GANs) is revolutionizing the way music is created, customized, and consumed across various industries. By leveraging deep learning techniques, AI can generate high-quality, adaptive, and personalized compositions, making music creation more efficient, cost-effective, and accessible to a wider audience. From film scoring, gaming, and advertising to music therapy, education, and content creation, AI-generated music is enhancing creativity, personalization, and innovation. Its ability to generate unlimited compositions without copyright concerns provides immense value to businesses, artists, and independent creators. Furthermore, AI's role in adaptive music, interactive experiences, and experimental sound design pushes the boundaries of traditional composition, fostering new possibilities in both artistic and technological domains. While AI does not replace human creativity, it serves as a powerful tool that complements and enhances the creative process, enabling artists and industries to explore new frontiers in music. As AI continues to evolve, its integration into music composition will expand, further bridging the gap between technology and artistic expression, shaping the future of sound in unprecedented ways.

REFERENCES

1. **Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014).** Generative Adversarial Networks. This paper introduced GANs as a novel framework for generative modeling, where a generator network competes against a discriminator network to create realistic outputs. The foundational work laid the groundwork for using GANs in various domains, including music generation.
2. **Dong, H., Hsiao, W., Yang, L., & Yang, Y. (2018).** MuseGAN: Multi-track sequential generative adversarial networks for symbolic music generation. MuseGAN is a GAN-based model designed specifically for multi-track music composition, enabling the generation of polyphonic and harmonious music. The model demonstrates how GANs can be used to compose music with distinct instrumental tracks while maintaining musical coherence.
3. **Engel, J., Agostinelli, A., Roberts, A., Dieleman, S., Norouzi, M., Eck, D., & Simonyan, K. (2019).** GANSynth: Adversarial neural audio synthesis. GANSynth utilizes GANs to synthesize high-quality musical audio directly as waveforms, offering improvements over traditional spectrogram-based methods. The study highlights how GANs can generate realistic instrument sounds and enables greater control over timbre and musical attributes.
4. **Dhariwal, P., Jun, H., Payne, C., Jong, W., & Sander, D. (2020).** Jukebox: A generative model for music. Jukebox, developed by OpenAI, is a deep-learning-based generative model that can create high-fidelity music, including lyrics and instrumental accompaniments. Unlike previous models,

International Journal of Information Technology & Computer Engineering

ISSN 2347–3657

Volume 13, Issue 2, 2025

Jukebox employs a transformer-based GAN approach, allowing it to generate more complex and expressive compositions.

5. **MuseGAN (Dong et al., 2018)**: A GAN model designed for multi-track music composition, enabling polyphonic music generation. It allows different instrumental tracks to be generated simultaneously, ensuring harmony and coherence.
6. **GANSynth (Engel et al., 2019)**: A model that generates high-fidelity musical audio directly as waveforms, outperforming traditional methods by synthesizing realistic instrumental sounds.
7. **Jukebox (Dhariwal et al., 2020)**: A transformer-based GAN model developed by OpenAI that generates high-quality audio samples with lyrics and instrumentals, capturing stylistic elements of various artists.
8. **Briot, J. P., Hadjeres, G., & Pachet, F. (2017).** Deep learning techniques for music generation - A survey. This paper provides an extensive review of various deep learning approaches, including GANs, RNNs, and VAEs, applied to music generation.
9. **Yang, L. C., Chou, S. Y., & Yang, Y. H. (2017).** MidiNet: A convolutional GAN for symbolic-domain music generation. This work introduces a convolutional GAN architecture for generating MIDI-based music, demonstrating how CNNs can be effectively used for music composition.
10. **Manzelli, D., Wu, Q., Chowdhery, A., & Cheung, J. C. K. (2018).** Conditioning deep generative raw audio models for structured automatic music. This study explores techniques for structuring deep generative models to enhance automatic music generation.
11. **Brunner, G., Konrad, A., Wang, Y., & Wattenhofer, R. (2018).** MIDI-VAE: Modeling dynamics and instrumentation of music with applications to style transfer. This paper presents a Variational Autoencoder-based approach for music generation, which is relevant for comparisons with GAN-based methods.
12. **Colombo, F., Seeholzer, A., Brea, J., & Gerstner, W. (2021).** Learning to generate music with reinforcement learning and adversarial training. This paper combines reinforcement learning with GANs to improve the long-term structure and coherence of generated music.