

AffectMachine-Pop: A controllable expert system for real-time pop music generation

Kat R. Agres^{1*}, Adyasha Dash^{2*}, Phoebe Chua³, and Stefan K. Ehrlich⁴

¹Centre for Music and Health, Yong Siew Toh Conservatory of Music, National University of Singapore, Singapore, Singapore

²Yong Siew Toh Conservatory of Music, National University of Singapore, Singapore, Singapore

³Department of Information Systems and Analytics, National University of Singapore, Singapore, Singapore

⁴SETLabs Research GmbH, Munich, Germany

Abstract

Music is a powerful medium for influencing listeners' emotional states, and this capacity has driven a surge of research interest in AI-based affective music generation in recent years. Many existing systems, however, are a black box which are not directly controllable, thus making these systems less flexible and adaptive to users. We present *AffectMachine-Pop*, an expert system capable of generating retro-pop music according to arousal and valence values, which can either be pre-determined or based on a listener's real-time emotion states. To validate the efficacy of the system, we conducted a listening study demonstrating that *AffectMachine-Pop* is capable of generating affective music at target levels of arousal and valence. The system is tailored for use either as a tool for generating interactive affective music based on user input, or for incorporation into biofeedback or neurofeedback systems to assist users with emotion self-regulation.

Introduction

Artificial intelligence (AI)-based music generation has rapidly gained traction in recent years, driven by advances in machine learning, computational creativity, and the growing availability of music datasets. Such systems have broad applications, from enhancing artistic creativity and aiding music composition through co-creative systems, to providing interactive entertainment in gaming and VR scenarios, and even person-centered, adaptive therapeutic applications (Ji, Luo, and Yang 2020; Agres et al. 2021; Dash and Agres 2024; Civit et al. 2022; Wang et al. 2024).

Advancements in generative AI, specifically large language models (LLM) and natural language processing (NLP) technology, have led to groundbreaking advancements in the field. Generative models such as AudioLM (Borsos et al. 2023), MusicLM (Agostinelli et al. 2023), and MusicGen (Copet et al. 2024) can be prompted with either text descriptions or melodies to produce high-quality audio samples that mimic diverse genres, styles and sentiments. Aside from text-based prompting, previous approaches have also achieved music generation by sampling latent spaces encoding distinct musical attributes, styles, and sentiments (Tan and Herremans 2020; Gupta et al. 2023).

Within the field of music generation, there has been significant interest in AI-based *affective* music generation (AI-AMG) systems, which aim to create music with specific emotional qualities (Grekow and Dimitrova-Grekow 2021; Sulun, Davies, and Viana 2022; Cui, Sarmento, and Barthet 2024; Dash and Agres 2024). Through their ability to induce, mediate and enhance listener's emotional states, AI-AMG systems widen the scope of algorithmic music generation by enabling tailored emotional engagement in applications such as mental health interventions (Elliott, Polman, and McGregor 2011; Stewart et al. 2019; Agres et al. 2021), assistive technology (Hagerer et al. 2015), and physical rehabilitation (Ashoori, Eagleman, and Jankovic 2015).

While generative AI has brought about important advancements in the development of AI-AMGs, specifically regarding the high quality and naturalness of the generated music, most approaches relying on generative AI (especially neural network-based approaches) suffer from several drawbacks, namely: 1) the generation of musical samples from generative AI models is typically performed at the song level and usually requires several seconds up to minutes; 2) near real-time inference for instantaneous music generation is rare; and 3) the music generation process is non-deterministic, rendering fine-grained control of the affective qualities of the music challenging, and the expected emotional impact on the listener non-trivial. Furthermore, the majority of generative AI models rely on pre-existing music for training, which has sparked controversies surrounding copyrighted training data, making approaches *not* relying on pre-existing music increasingly attractive.

These factors render state-of-the-art neural networks-based generative AI music systems sub-optimal for real-time music generation, which is essential for interactive entertainment in gaming and VR scenarios, as well as human-in-the-loop adaptive therapeutic systems. The challenge of *narrative adaptability* (Dash and Agres 2024) – the ability of AI-AMG systems to instantaneously generate music seamlessly adapting to a given emotional narrative or sequence, while still maintaining musical coherence – has not been explored in great depth. However, real-time adaptation of musical output, which can in turn influence the listener's emotions, is a key area given music's significant impact on human emotion (Juslin, Sloboda et al. 2001). Few systems have addressed this challenge. For example, En-

*These authors contributed equally.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

gels, Tong, and Chan (2015) employed hierarchical Markov models to generate original music for video games, capable of mimicing the variations in musical structure often found in longer human-composed pieces. Similarly, Ehrlich et al. (2019) developed a rule-based probabilistic music generation algorithm as part of a closed-loop brain-computer interface system. The algorithm generates a stream of MIDI events, with the type and occurrence of events modulated by emotional arousal and valence parameters. Recent work by Agres, Dash, and Chua (2023) introduced an AI-AMG system called *AffectMachine-Classical*, capable of creating non-repetitive, continuous affective music in a classical style. This system was validated in a study with human participants, confirming the emotional qualities of the generated musical and supporting its utility for the aforementioned use cases.

Building upon *AffectMachine Classical*, we present here *AffectMachine-Pop*, a novel AI-AMG system designed to generate continuous affective music in a pop style. Compared with classical music, pop music is characterized by simpler, more repetitive forms that resonate more broadly with the general public (Chan, Qu, and Mak 2009). However, generating multi-track pop music presents the technical challenge of maintaining melodic, harmonic and rhythmic coherence across several musical layers (Ren et al. 2020; Zhu et al. 2018). Therefore, unlike most modern pop music AI-AMGs, which predominantly rely on data-driven methods and high-performance AI architectures, e.g. (Ren et al. 2020; Huang and Yang 2020), our approach takes a different path. Existing pop systems often depend heavily on pre-existing music datasets and suffer from limited transparency, functioning as black-box models. In contrast, our method features a compositional framework guided by established musical rules, deliberate instrumentation, and precise timbral arrangements. These elements are manually refined by expert musicians and composers to achieve a balance of technical accuracy and artistic expressiveness. In addition, our approach features near real-time adaptability, enabling the system to seamlessly adjust its affective content based on the user's inputs or physiological states, by traversing the valence-arousal space (Russell 1980) while preserving musical and stylistic coherence. The system generates musical samples with controllable emotional qualities, i.e., music at specific levels of arousal and valence, as demonstrated by empirical evidence from the listening study presented here.

Below we describe *AffectMachine-Pop*, our expert system for affective automatic music generation in a popular music style. We also highlight the ways in which the system compares and differs from *AffectMachine-Classical*. Finally, we present the methodology and results of a listening study conducted to validate the system's ability to produce music at different target levels of arousal and valence.

AffectMachine-Pop system description

Here we describe the parameters and design of our system, *AffectMachine-Pop*, for computationally-generating affective pop music. Unlike the majority of AI-AMG systems, which have focused on generating music in a classical genre (e.g., Agres, Dash, and Chua (2023); Ehrlich et al. (2019);

Wallis et al. (2011)), the present system composes music in a retro-pop genre. This makes it particularly well-suited for middle-aged and older adults, and anyone who enjoys popular music from the 60s and 70s.

Figure 1 depicts the basic building blocks of our system. There are four major components, namely, *Harmonic parameters*, *Timbral & loudness parameters*, *Pitch characteristics of voices*, and *Time & rhythm parameters*. Each of these blocks contains different musical rules for tailoring features of the music (i.e., tempo, chord, etc) to convey specific levels of emotion. The emotion parameters (valence and arousal values) are given as input to the system, and then the system generates emotional music (output) through a complex interplay between these blocks. In the following section, we describe each of these blocks in greater detail.

We focus on highlighting the *differences* between *AffectMachine-Pop* and *AffectMachine-Classical*, which has already been described in-depth by Agres, Dash, and Chua (2023). Where the designs overlap, we offer only a brief summary.

Emotion inputs

Russell's (1980) two-dimensional model of affective space, which represents emotions along the dimensions of arousal (intensity or energy level) and valence (degree of pleasantness) (Russell 1980), has inspired many AI-AMG systems (Dash and Agres 2024; Huang, Chen, and Yang 2024; Agres, Dash, and Chua 2023; Ehrlich et al. 2019), including *AffectMachine-Pop*. This model has been widely adopted (and adapted) in emotion research, as it allows for the representation of a broad range of emotions, from high-arousal, positive states (e.g., joy) to low-arousal, negative states (e.g., sadness). By mapping musical features to this affective space, *AffectMachine-Pop* can generate music that meets target arousal and valence levels, which are either predetermined, or based on real-time inputs, such as physiological data reflecting a listener's emotion states (such as HR or EEG data).

Harmonic parameters

Mode *AffectMachine-Pop* uses a bespoke probabilistic chord progression matrix that takes inspiration from retro-pop music, e.g., English and Mandarin pop songs from the 1960s-1970s, such as The Beatles' *Hey Jude* and Theresa Teng's *Thousands of Words*. Specifically, the chord progression matrix can be modeled as a directed probabilistic graph, $G(V, E)$ with vertices V and edges E . Each vertex represents a chord, and edges represent transition probabilities between chords. The generated chord sequence is guided by a pre-defined sequence or array of valence values; each valence value corresponds to one bar in the progression, such that the array's length matches the number of bars in the generated music. The chords and probabilities in our expert system were tuned by hand based on previous empirical findings, as well as the musical insights from our interdisciplinary team of musicians and composers. Figure 2 shows a representation of our chord progression algorithm as a directed probability graph for a 4-bar musical excerpt with a valence parameter of [0.1, 0.8, 0.7, 0.8].




Figure 1: Architecture of AffectMachine-Pop.




Figure 2: Sample Chord Progression as a probabilistic graph.

The musical structure governing the length of the generated chord progressions differs from that of AffectMachine-Classical. In AffectMachine-Pop, the music follows a 32-bar musical structure comprised of two 8-bar sections (labelled section “A” and “B”, as per convention) that repeat in an AABB pattern. (This AABB pattern is distinct from the AffectMachine-Classical system.) Each bar in the structure is associated with a fixed chord function, while the specific chords for the bar and their associated transition probabilities are determined by the valence level.

Timbral and loudness parameters

Instrumentation AffectMachine-Pop uses five virtual instruments, played over different audio tracks, to generate affective polyphonic music that resembles “retro” pop music from the 1960s-1970s. The five instruments are a percussion kit, bass guitar, electric guitar, violin section and French horn. Each instrument is played using a single audio track in the DAW (Ableton Live) except for the electric guitar, which is played back over two audio tracks. The violin section and French horn receive identical MIDI information and play the main melodic line. The doubling of the violin section by French horn helps to smooth and balance the timbre and enhance the prominence of the melodic line. The electric guitar provides harmonic accompaniment by strumming chords, which are played back on one audio track. At lower levels of arousal ($aro < 0.7$) where the music is slower and less dense with minimal percussion, a plucked electric guitar is also used to add texture, to convey a desired level of valence when note activations were very sparse. The strummed and plucked electric guitars are recorded on separate audio tracks to facilitate audio mixing.

Velocity Volume (or dynamics) is frequently used to convey emotional expressiveness in music (Gabrielsson and Juslin 2003). Increased sound levels are associated with higher-arousal emotions such as “happy” or “angry”, while lower sound levels tend to be associated with lower-arousal emotions such as “sad” (Bresin and Friberg 2011). To enable real-time adaptation of volume, we mapped arousal values to attack velocity in MIDI following Eq. 1. Attack velocity indicates the force with which a note is struck; aside from controlling volume, higher velocity values also produce brighter and harder timbres which contribute to the perception of increased arousal (Agres, Dash, and Chua 2023).

$$velocity = 60 + (aro * 15) \quad (1)$$

Pitch characteristics of voices

Voice leading Voice leading – the linear progression of melodic voices and their interaction to create harmonies – was added to the system for musical sophistication and coherence. Separate voice leading logic was developed for each instrument and track in AffectMachine-Pop. The bass guitar plays the roots and fifths of the current chord, a common note pattern used by bass players in pop music. The root is selected with higher probability ($p = 0.9$) on the first beat of each bar to help establish the chord. Next, the plucked electric guitar is used primarily to add texture at lower levels of arousal; hence, it plays a randomly selected sequence of chord tones from the current chord (with all chord tones being equiprobable). This allows the current chord quality to be expressed clearly even at lower levels of arousal with little additional dissonance.

The strummed electric guitar provides harmonic progression, and its voice leading logic follows the heuristic outlined in Wallis et al. (2011), which is that musicians tend to voice new chords in a manner that is as similar as possible to the previous chord, in terms of interval and placement on the instrument. We implement this by selecting the voicing for the subsequent chord that is least dissimilar to the current chord (details of the dissimilarity calculation can be found in Agres, Dash, and Chua (2023)).

The main melodic line is played by the violin section and French horn. For these instruments, we employ a mix of composed melodic motifs and probabilistic voice leading logic. Specifically, the instruments play a melodic pattern

probabilistically determined by the voice leading logic for the first four bars of each 8-bar section. Like AffectMachine-Classical, this voice leading logic is encoded in the form of transition matrices. However, in AffectMachine-Pop, we divide the valence range into low ($V \leq 0.5$) and high valence ($V > 0.5$) regions and compose one matrix for each region to generate appropriate melodies for each level of valence. In the second four bars, the instruments play a probabilistically-selected composed rhythmic motif. The set of composed rhythmic motifs were developed with reference to representative pieces from the Western and mandarin retro-pop music canon during the 1960s-1980s, and include typical motifs such as pentatonic patterns and arpeggiation.

Pitch register Generally, higher pitches tend to be associated with positive emotions such as excitement and serenity (Collier and Hubbard 1998), while lower pitches correlate with negatively-valenced emotions such as sadness. We use chord inversions to gradually shift the pitch register of the strummed and plucked electric guitar. Regardless of the current level of valence, there is always a constant probability ($p = 0.6$) that the chord voicing does not change, to prevent the pitch register of the music from increasing or decreasing too quickly. As valence increases, the probability that the inversion of the current chord voicing increases by one (e.g., from root to first inversion) also increases according to the formula $p(inv + |inv \neq 0) = val$, while the probability that the inversion of the current chord voicing decreases by one (e.g., from first to root inversion) decreases according to the formula $p(inv - |inv \neq 0) = 1 - val$.

Register is not relevant for the percussion instrument, and for the remaining instruments (French horn and violin section), the register is fairly consistent as they return to the composed melodic motive in the second four bars of each 8-bar section, and there are typically no obvious register changes due to voice leading logic in the remaining four bars.

Time and rhythm parameters

Rhythm We composed a set of three equiprobable 8-bar rhythmic patterns for the bass guitar, and randomly select a new rhythmic pattern at the beginning of each 8-bar section. For the strummed electric guitar, we divided the arousal range into three regions: low ($aro < 0.40$), moderate ($0.4 \leq aro < 0.70$) and high ($aro \geq 0.70$), and composed several rhythmic patterns (and their associated probabilities) for each region. A new rhythmic pattern is selected for strummed guitar at the beginning of each bar. For the percussion kit, we composed several 8-bar rhythmic patterns that gradually reduce in density as arousal decreases, with greater use of quieter instruments and techniques such as rim clicks. A new rhythmic pattern for percussion is selected at the beginning of each 8-bar section. A single pattern was composed for low ($aro \leq 0.30$) and moderate ($0.3 < aro \leq 0.70$) levels of arousal, while a set of three rhythmic patterns were composed for high ($aro > 0.70$) levels of arousal.

For the main melodic instruments (string section and French horn), we employ a mix of composed rhythmic mo-

tives and rhythmic roughness. Specifically, during the first four bars of each 8-bar section, the instruments play a rhythmic pattern probabilistically determined by the roughness parameter, which is a measure of how irregular the rhythm of a piece of music is. Music with smooth, regular rhythms are typically perceived as higher in valence and lower arousal. In AffectMachine-Pop, we use note density as a proxy for rhythmic roughness (Wallis et al. 2011) in the following manner: as arousal increases, roughness decreases in a linear fashion and note density increases. For example, if roughness is set to 0, each bar would be populated with eight notes of equal length. Because this tends to sound overly dense (since the tempo is also faster at higher levels of arousal), we set a lower bound for the roughness parameter.

In terms of the last four bars of each 8-bar section, the instruments play a selected composed rhythmic motive. We divided the arousal range into three regions: low ($aro < 0.30$), moderate ($0.3 \leq aro < 0.60$) and high ($aro \geq 0.60$), and composed a set of three equiprobable rhythmic patterns for each region. Finally, rhythmic patterns for the plucked electric guitar are determined probabilistically by the roughness parameter.

Tempo Tempo refers to the rate or speed of the music, often measured in terms of beats-per-minute (bpm). In our system, tempo is controlled by the arousal value. Specifically, tempo is governed by a logarithmic relationship with arousal, and has a range of $tempo \in [36, 130]$. We found the logarithmic relationship and limited range to be useful, as changes in tempo were more perceptually apparent when the tempo was lower, and using overly fast tempos sounded unpleasant and unnatural.

AffectMachine-Pop Listening study

To evaluate the effectiveness of AffectMachine-Pop in generating music at specific levels of arousal and valence, we conducted a music listening study. This study assesses the system's ability to convey the intended emotional content of the generated music to listeners (i.e., perceived emotion). Future work will examine AffectMachine's capacity to reliably induce emotions in listeners.

Musical Stimuli Generation AffectMachine-Pop is designed to compose music that conveys emotions corresponding to any point on the valence-arousal (V-A) plane. To capture a range of different emotions, we generated musical excerpts from 13 different points on the V-A plane. These points were selected to span a variety of different emotions, including the corners, the middle of each quadrant, and the neutral point at the centre of the space. For more details on the selection of points in A-V space, please refer to Agres, Dash, and Chua (2023). To minimize bias towards any single musical excerpt and improve generalizability, three musical stimuli were generated for each of these 13 points, resulting in a total of 39 musical excerpts. The average duration of the excerpts was 32.6 seconds. The musical stimuli were composed based on a 4-bar, 8-bar, or 16-bar progression. For stimuli corresponding to low arousal levels, the duration was limited to 4 bars (ranging from 27-34 sec per stimulus).

Stimuli with higher levels of arousal contained 8 or 16 bars, as 4 bars were too brief for these stimuli with a faster tempo.

Experimental Procedure We conducted the listening study with 24 participants (average age = 23.0 yrs, SD = 2.6 yrs; male = 9, female = 15). All participants were given verbal and written instructions about the study prior to providing their written consent. The study was approved by the Institutional Review Board (IRB) of the National University of Singapore (NUS).

The experiment took place in a quiet room free from audio-visual distractions. The experimenter explained the procedure to each participant, who then provided written informed consent to participate in the study. Participants completed the study individually. Before starting the listening task, participants provided demographic information such as their age, gender, and ethnicity.

At the beginning of the study, participants listened to two practice trials to become familiar with the procedure. Subsequently, the 39 musical excerpts were presented in randomized order. After listening to each stimulus, participants were asked to indicate their perceived emotion, i.e., the emotion they perceived in the musical excerpt. Responses were collected using the Self-Assessment Manikin (SAM), a visual 9-point scale ranging from “Very unpleasant” (1) to “Extremely pleasant” (9) for valence, and from “Calm” (1) to “Excited” (9) for arousal. Participants were allowed to listen to each stimulus only once, but could take as much time as needed to provide their ratings. The experiment lasted approximately 40 minutes, and participants were compensated with \$6 SGD for their time.

The valence and arousal ratings collected across participants were then analyzed to evaluate the efficacy of our system in conveying specific emotions through its generated music. The results of the analyses are presented below.

Results and Discussion

To evaluate the efficacy of AffectMachine-Pop in generating music that expresses a desired emotion, we analyzed the user ratings collected during the listening study. The primary goal was to determine whether the music generated by AffectMachine-pop successfully conveys the intended levels of valence and arousal to listeners. Specifically, we performed a comparative analysis between the average user ratings for the musical stimuli and the valence or arousal parameter settings used during the music generation process. First, we normalized the perceptual ratings for each listener using Eq. 2. Here, the $MaxValence$ refers to the maximum possible valence rating (i.e., 9) and $MinValence$ refers to the minimum possible valence rating (i.e., 1). A similar normalization procedure was applied for arousal. The resulting normalized valence and arousal ratings, which range from 0 to 1, were used for further analysis. In the remainder of the article, the normalized valence and normalized arousal ratings will be referred to as valence and arousal ratings, respectively.

$$NormalizedValence = \frac{RatedValence - MinValence}{(MaxValence - MinValence)} \quad (2)$$

Comparative Analysis of User Ratings To access whether our system accurately expresses the desired emotions through its generated music, we performed a comparative analysis between the mean emotion ratings and the valence and arousal settings used to generate the music. For example, the mean arousal ratings for stimuli generated with settings valence, arousal = [{0,1}, {0.5,1}, {1,1}] were used to evaluate the system’s performance when arousal was set to its maximum value. Figure 3 presents the average ratings along with their respective standard errors.

As shown in the figure, the valence ratings exhibit a positive linear trend but tend to plateau when $V >= 0.75$. According to literature, the middle range of psychometric rating scales often receives a higher density of responses, whereas the extremes of the scale typically receive fewer responses (Leung 2011). This phenomenon could explain why raters found extreme values of valence, such as valence values > 0.75 and < 0.25 , to be less distinguishable.

In contrast, arousal values showed better correspondence with the average user ratings across most of the scale, except around an arousal value of 0.5. This may be due to the fact that the excerpts were generated using three random instantiations (e.g., these particular stimuli may have deviated from the system’s typical output at $A = 0.5$). It is also possible that the tempo around $A = 0.5$ is too slow. Future work will explore a faster tempo for moderate levels of arousal, and will test a larger number of excerpts.

To quantify the positive correlation between average user ratings and parameter settings, we performed linear regression-based curve fitting. The coefficient of determination was $R^2 = 0.93$ ($p < 0.01$) for valence and $R^2 = 0.86$ ($p < 0.05$) for arousal, indicating good correspondence between user ratings and target values and thereby confirming the effectiveness of our system to convey the target emotion.




Figure 3: Mean arousal (a) and valence (b) ratings (error bars represent standard error).

We also investigated whether the perception of arousal is influenced by variations in valence and vice versa. To this end, we examined the dependence of perceived emotion user

ratings on both the valence and arousal parameter settings simultaneously. Figure 4 illustrates this relationship, showing the interpolated average arousal ratings (Fig 4a) and valence ratings (Fig 4b). The stars correspond to the 13 points in the valence-arousal plane used to generate the musical stimuli.

Generally, there is a strong correspondence between parameter settings and emotion ratings. As shown in the figure (Fig 4b), perceived valence is lower than the intended valence parameter setting (for values of $V > 0.8$) when arousal is set below < 0.6 . This implies that excerpts intended to convey high valence are perceived as having only moderate valence when arousal value is low. Conversely, stimuli generated at low valence settings (values of $V < 0.2$) are perceived to express higher valence when arousal is set high (values of $A > 0.7$). This suggests that music intended to express low valence is perceived as having higher valence when the arousal value is higher. These observations can be partially attributed to the effect of tempo in generated music, where slower tempos are associated with lower arousal values and vice versa.

In contrast, we observed a more consistent correspondence between the arousal parameter settings and arousal ratings (except for $A = 0.5$), and found that the arousal ratings show no dependence on valence settings. This observation aligns with findings from the literature (Wallis et al. 2011), which reported an asymmetrical “crossover” effect between arousal and valence. Specifically, Wallis et al. (2011) observed that while perceived valence correlates with intended arousal, perceived arousal does not significantly correlate with intended valence. The observed crossover effect is most likely due to the non-orthogonal relationship between the valence and arousal dimensions in relation to one or more musical features. Alternatively, it could also arise from uncontrollable factors, such as cultural influences. Nevertheless, this crossover effect does not greatly impact the intended functioning of our systems, as the results within each dimension align well with expectations.

These findings confirm the potential of AffectMachine-Pop to reliably generate emotion-infused music that expresses varying levels of arousal and valence in a controlled manner. However, we note that the results may be influenced by the limited sample size, which could affect generalizability. Further in-depth investigations will be useful to expand on the above observations.

Conclusion

We have presented *AffectMachine-Pop*, an expert system designed for generating affective music in a pop style. This system is capable of real-time music generation, offering a novel and innovative tool for creating music tailored to convey specific emotion states.

In a music listening study, we investigated whether the system is able to generate music that effectively expresses desired emotions (specified via specific levels of arousal and valence). The results showed a strong correspondence between the given valence and arousal parameter settings and listeners’ averaged valence and arousal values, with $R^2 = 0.93$ for valence and $R^2 = 0.86$ for arousal. These results confirm the system’s effectiveness.




Figure 4: Mean (interpolated) arousal (a) and valence (b) ratings as a function of the valence and arousal parameters. Vertical color bars represent the colors corresponding to different values of normalized average ratings over the range 0.1 to 0.9.

Future work will evaluate the system’s ability to *induce* emotions in listeners, extending beyond the current focus on perceived emotion. Additionally, we aim to address current limitations, including the need for a larger and more diverse group of participants. The long-term goals of this line of research are to establish the system’s potential to reliably generate affective music, and explore its applications in emotion self-regulation. One promising future direction is the integration of AffectMachine into a Brain-Computer-Interface (BCI), which aims to assist users in achieving desired emotional or mood states through adaptive, real-time music generation guided by their neural activity.

Acknowledgments

This work was supported by the RIE2020 Advanced Manufacturing and Engineering (AME) Programmatic Fund (No. A20G8b0102), Singapore. We wish to thank the composer Cliff Tan for his musical contributions to the development of AffectMachine-Pop, and our former research assistant, Ting Yuan Tay, for her help with data collection for the listener study.

References

- Agostinelli, A.; Denk, T. I.; Borsos, Z.; Engel, J.; Verzetti, M.; Caillon, A.; Huang, Q.; Jansen, A.; Roberts, A.; Tagliasacchi, M.; et al. 2023. Musiclm: Generating music from text. *arXiv preprint arXiv:2301.11325*.
- Agres, K. R.; Dash, A.; and Chua, P. 2023. AffectMachine-Classical: a novel system for generating affective classical music. *Frontiers in Psychology*, 14: 1158172.
- Agres, K. R.; Schaefer, R. S.; Volk, A.; van Hooren, S.; Holzapfel, A.; Dalla Bella, S.; Müller, M.; De Witte, M.; Herremans, D.; Ramirez Melendez, R.; et al. 2021. Music, computing, and health: a roadmap for the current and future roles of music technology for health care and well-being. *Music & Science*, 4: 2059204321997709.
- Ashoori, A.; Eagleman, D. M.; and Jankovic, J. 2015. Effects of auditory rhythm and music on gait disturbances in Parkinson’s disease. *Frontiers in neurology*, 6: 234.

Borsos, Z.; Marinier, R.; Vincent, D.; Kharitonov, E.; Pietquin, O.; Sharifi, M.; Roblek, D.; Teboul, O.; Grangier, D.; Tagliasacchi, M.; et al. 2023. Audiolfm: a language modeling approach to audio generation. *IEEE/ACM transactions on audio, speech, and language processing*, 31: 2523–2533.

Bresin, R.; and Friberg, A. 2011. Emotion rendering in music: range and characteristic values of seven musical variables. *cortex*, 47(9): 1068–1081.

Chan, W.-Y.; Qu, H.; and Mak, W.-H. 2009. Visualizing the semantic structure in classical music works. *IEEE transactions on visualization and computer graphics*, 16(1): 161–173.

Civit, M.; Civit-Masot, J.; Cuadrado, F.; and Escalona, M. J. 2022. A systematic review of artificial intelligence-based music generation: Scope, applications, and future trends. *Expert Systems with Applications*, 209: 118190.

Collier, W. G.; and Hubbard, T. L. 1998. Judgments of happiness, brightness, speed and tempo change of auditory stimuli varying in pitch and tempo. *Psychomusicology: A Journal of Research in Music Cognition*, 17(1-2): 36.

Copet, J.; Kreuk, F.; Gat, I.; Remez, T.; Kant, D.; Synnaeve, G.; Adi, Y.; and Défossez, A. 2024. Simple and controllable music generation. *Advances in Neural Information Processing Systems*, 36.

Cui, W.; Sarmento, P.; and Barthet, M. 2024. Mood-LoopGP: Generating Emotion-Conditioned Loop Tablature Music with Multi-granular Features. In *International Conference on Computational Intelligence in Music, Sound, Art and Design (Part of EvoStar)*, 97–113. Springer.

Dash, A.; and Agres, K. 2024. AI-Based Affective Music Generation Systems: A Review of Methods and Challenges. *ACM Computing Surveys*, 56(11): 1–34.

Ehrlich, S. K.; Agres, K. R.; Guan, C.; and Cheng, G. 2019. A closed-loop, music-based brain-computer interface for emotion mediation. *PloS one*, 14(3): e0213516.

Elliott, D.; Polman, R.; and McGregor, R. 2011. Relaxing music for anxiety control. *Journal of music therapy*, 48(3): 264–288.

Engels, S.; Tong, T.; and Chan, F. 2015. Automatic real-time music generation for games. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 11, 220–222.

Gabrielsson, A.; and Juslin, P. N. 2003. *Emotional expression in music*. Oxford University Press.

Grekow, J.; and Dimitrova-Grekow, T. 2021. Monophonic music generation with a given emotion using conditional variational autoencoder. *IEEE Access*, 9: 129088–129101.

Gupta, C.; Kamath, P.; Wei, Y.; Li, Z.; Nanayakkara, S.; and Wyse, L. 2023. Towards controllable audio texture morphing. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.

Hagerer, G. J.; Lux, M.; Ehrlich, S.; and Cheng, G. 2015. Augmenting affect from speech with generative music. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, 977–982.

Huang, J.; Chen, K.; and Yang, Y.-H. 2024. Emotion-driven Piano Music Generation via Two-stage Disentanglement and Functional Representation. *arXiv preprint arXiv:2407.20955*.

Huang, Y.-S.; and Yang, Y.-H. 2020. Pop music transformer: Beat-based modeling and generation of expressive pop piano compositions. In *Proceedings of the 28th ACM international conference on multimedia*, 1180–1188.

Ji, S.; Luo, J.; and Yang, X. 2020. A comprehensive survey on deep music generation: Multi-level representations, algorithms, evaluations, and future directions. *arXiv preprint arXiv:2011.06801*.

Juslin, P. N.; Sloboda, J. A.; et al. 2001. Music and emotion. *D. DEUTSCH (Org.)*.

Leung, S.-O. 2011. A comparison of psychometric properties and normality in 4-, 5-, 6-, and 11-point Likert scales. *Journal of social service research*, 37(4): 412–421.

Ren, Y.; He, J.; Tan, X.; Qin, T.; Zhao, Z.; and Liu, T.-Y. 2020. Popmag: Pop music accompaniment generation. In *Proceedings of the 28th ACM international conference on multimedia*, 1198–1206.

Russell, J. A. 1980. A circumplex model of affect. *Journal of personality and social psychology*, 39(6): 1161.

Stewart, J.; Garrido, S.; Hense, C.; and McFerran, K. 2019. Music use for mood regulation: Self-awareness and conscious listening choices in young people with tendencies to depression. *Frontiers in psychology*, 10: 1199.

Sulun, S.; Davies, M. E.; and Viana, P. 2022. Symbolic music generation conditioned on continuous-valued emotions. *IEEE Access*, 10: 44617–44626.

Tan, H. H.; and Herremans, D. 2020. Music fadernets: Controllable music generation based on high-level features via low-level feature modelling. *arXiv preprint arXiv:2007.15474*.

Wallis, I.; Ingalls, T.; Campana, E.; and Goodman, J. 2011. A rule-based generative music system controlled by desired valence and arousal. In *Proceedings of 8th international sound and music computing conference (SMC)*, 156–157.

Wang, L.; Zhao, Z.; Liu, H.; Pang, J.; Qin, Y.; and Wu, Q. 2024. A review of intelligent music generation systems. *Neural Computing and Applications*, 36(12): 6381–6401.

Zhu, H.; Liu, Q.; Yuan, N. J.; Qin, C.; Li, J.; Zhang, K.; Zhou, G.; Wei, F.; Xu, Y.; and Chen, E. 2018. Xiaoice band: A melody and arrangement generation framework for pop music. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 2837–2846.