COMPUTER-GENERATING EMOTIONAL MUSIC: THE DESIGN OF AN AFFECTIVE MUSIC ALGORITHM

Isaac Wallis, Todd Ingalls, Ellen Campana

Arts, Media, and Engineering Program,
Arizona State University,
Tempe, Arizona 85281, USA
iwallis@asu.edu

ABSTRACT

This paper explores one way to use music in the context of affective design. We've made a real-time music generator that is designed around the concepts of valence and arousal, which are two components of certain models of emotion. When set to a desired valence and arousal, the algorithm plays music corresponding to the intersection of these two parameters. We designed our algorithm using psychological theory of emotion and parametrized features of music which have been tested for affect. The results are a modular algorithm design, in which our parameters can be implemented in other affective music algorithms. We describe our implementation of these parameters, and our strategy for manipulating the parameters to generate musical emotion. Finally we discuss possible applications for these techniques in the fields of the arts, medical systems, and research applications. We believe that further work will result in a music generator which can produce music in any of a wide variety of commonly-perceived emotional connotations on command.

1. INTRODUCTION

In recent years a paradigm has emerged, affective design[1], where things are designed in part to elicit emotional responses in users. This paper explores the possibility of using computer music for this purpose. The elicitation of emotion is an intuitive application of music: one can make the argument that the elicitation of emotion is exactly what most composers and musicians spend their lives trying to accomplish. Music psychologists have studied music in terms of perception of emotion, and computer musicians have been creating algorithmic music for well over half a century—but there have only been a few cases where the two worlds collide[2,3].

We focus on algorithmic techniques for the purpose of evoking specific emotions. For example, in our engine a user can select a desired value for both valence and arousal; these two scales can be used to describe many human emotions. The algorithm will then generate piano music which feels similar in quality to the desired affect settings. This work crosses multiple disciplines including music, psychology, and digital media—and members of each of these disciplines might find this work useful.

Musicians can benefit from this work because creation of this algorithm demands study of music structural features, including abstract features such as rhythmic roughness, to determine their effect on emotion. Knowledge of a music feature's affect can help composers and performers make creative decisions with emotion in mind, just as they already do with well understood music structures like tempo and mode.

Psychologists benefit because these algorithmic techniques could be valuable in future study of emotion. Consider a hypothetical situation where a researcher needs to determine the effects of mood upon a specific task: one way to go about this research might be to use an affective music algorithm to influence the mood of the test subjects. Such an algorithm could also be useful in the study of affective disorders or autism.

The field of digital media has arguably the most to gain from these techniques because of the multi-modal nature of today's media. An algorithm like ours, built using affective principles, could serve as an automatic soundtrack generator for movies or video games, or be used in interactive systems designed for more serious purposes like education, rehabilitation, or art.

2. EMOTION THEORY

Music and its relationship to emotion has been studied by scientists since the late 19th century, and by musicians for much longer. Starting in the 1930's, music psychologists such as Hevner[4] and Gundlach[5] have been able to identify and quantify some of the emotional effects of specific features of music such as tempo, rhythm, and mode.

Since then, the study of emotion and music has evolved rapidly. In the experiments of Hevner and Gundlach, emotions were considered to be discrete elements: feelings such as anger and sadness were not really related. Since then, various models of emotion have been developed, many of them mapping emotion to continuous multidimensional space. One of the first of these models is the circumplex model of emotion [6], which has some similarities with the emotion groupings used by Hevner. The circumplex model considers any emotion to be part of a two-dimensional space, in which the x-axis is related to the emotional valence and the y-axis is related to emotional arousal.

Valence[1] is a term with meaning shared across many domains: in electronics, for example, valence describes whether particles are positively or negatively charged. In emotion theory the meaning is similar. Positive emotions such as happiness, peacefulness, and love are all considered to have a positive valence. Similarly, less pleasant emotions such as anger, sadness, and fear all have negative valence. Arousal[2] is a description of the amount of energy in an emotion. Emotions like fury, panic, and excitement all have high arousal, while emotions like depression, contentment, or solemnity have low arousal. In the circumplex model, each emotion is described as the intersection of a certain valence with a certain arousal—a point in the valence/arousal space.

These two scales, by themselves, do not do a good job of distinguishing all types of emotion. Emotions like fear and anger are perceptually quite different, for example, but are likely to in-

1 Sometimes called *pleasure*.
2 Sometimes called *activity* or *intensity*.

habit the same regions of the valence/arousal space. This fact has led to the development of various conflicting higher-dimensional models of emotion[7,8], the most commonly used of which is the PAD model. The PAD model includes the valence and arousal axes but adds a third dimension called dominance, which is a measure of approach and avoidance response. The PAD model would do a better job of distinguishing the two emotions used in the above example, because anger is a high-dominance emotion, while fear is a low-dominance emotion.

It is clear that music is capable of generating more finely-grained emotions than the valence and arousal scales can distinguish[9], and eventually we hope to study music generation in terms of three or more axes. The valence and arousal model is perfectly valid, however, and is advantageous in terms of interface design. For example, a GUI has been designed for our engine in which a single mouse click anywhere on a two-dimensional grid is used to set the desired valence and arousal of the music (Figure 1). Similar grid interfaces have already been tested and validated in terms of evaluating emotion[10].

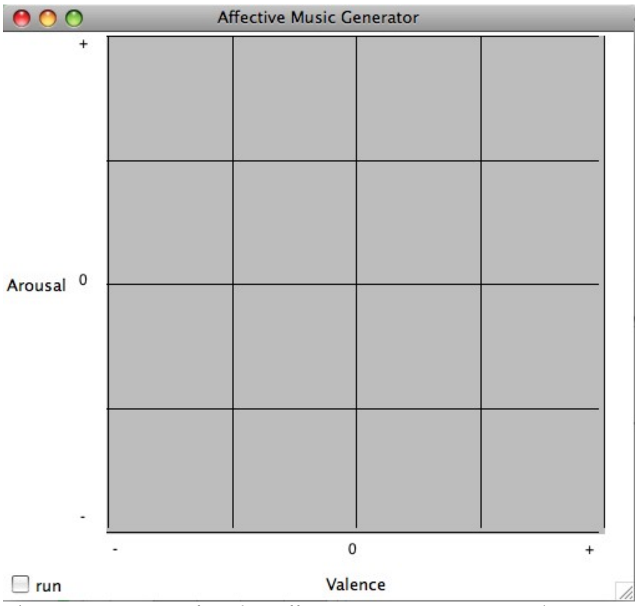

Figure 1*: A GUI for the affective music engine. The "run" checkbox turns the music on and off, and the emotion quality of the generated music is set by clicking anywhere in the "Arousal/ Valence" grid.*

3.ALGORITHM STRUCTURE

We feel that the most important part of our work is in the definition of music parameters which can be implemented in other algorithms. Music parameters, in the context of this paper, indicate music features which can be quantified in some way, and compared using less-than and more-than relationships. Any aspect of music can be a parameter, and new parameters can be contrived, as in the example of the rhythmic roughness feature to be discussed later.

Nothing about this particular music generator has anything to do with emotion other than the fact it manipulates music parameters which are based on well-studied musical features; for this reason, description of the algorithm structure seems almost irrelevant. Nevertheless, it would be impossible to describe some of the generalizable music parameters without their implementation within this engine's structure. Figure 2 is a flow chart of the current algorithm. It consists of three interdependent modules: one that controls timing, one that controls the harmonic framework, and one that selects notes and transmits them via MIDI to an external sound engine.

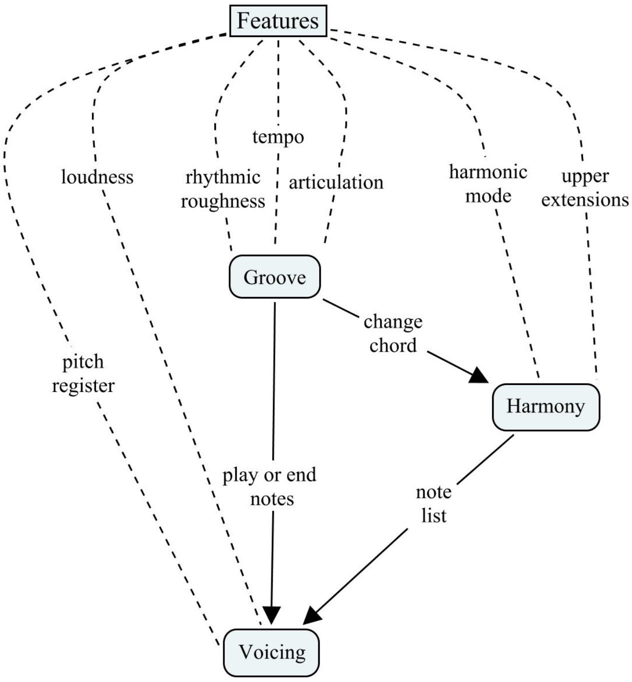

Figure 2: *Algorithm flow chart. Solid edges indicate control or information flow, dashed edges indicate an "implemented by" relationship.*

3.1.Groove: A Timing Module

The groove module controls all temporal aspects of this algorithm. It tells the harmony module when to change chords, and it tells the voicing module when to play and end notes. It implements the music parameters of tempo, rhythmic roughness, and articulation, all to be discussed in detail later. These parameters have all been linked to the perception of musical arousal.

3.2.Harmony: A Harmonic Framework Module

Discussions of harmonic matters may require some background knowledge of music theory for understanding. Persichetti[11] is a useful resource on these topics.

The harmony module ensures that the generated music fits into a western modal music context. It implements the musical parameters of upper extensions and harmonic mode, each of which is associated with valence perception. A chord progression is contained within this module which is written, not with explicit chords as is the norm, but in terms of chord functions:

In modal and tonal music, each chord serves a specific function within the harmonic framework. Although there have been numerous functions identified through analysis, they can be boiled down to three categories: dominant functions, subdominant functions, and tonic functions. Chords of similar function within the mode can usually be interchanged without affecting the quality of the chord progression.

Because of this fact, one can write a chord progression using harmonic functions rather than specific chords. Our harmony module contains such a chord progression: tonic, subdominant, dominant, tonic, repeat. Whenever a chord change is requested by the groove module, the harmony module first proceeds to the next chord in the cycle, then translates that chord into a list of notes using methods to be described later, in the discussions of the upper extensions and harmonic mode parameters.

Once a note list has been obtained for the chord, that list is sent to the voicing module. The voicing module can only select and play notes that are in the list. Note that this means our music engine cannot play non-chord tones.

3.3. Voicing: A Note Selector and MIDI Interface

The voicing module is the last module in the generation algorithm, the one which transmits MIDI information on to the sound engine [12]. It also implements the music parameters of loudness and pitch register through a probabilistic note selection scheme.

Whenever the groove module tells the voicing module to play, a probability table is constructed on the notes from the harmony module. The weights for the probability table are based on three rules:

Rule 1: Generated notes tend toward a central range, keeping the music from going too high or low. The range can be transposed up or down in pitch.

Rule 2: New notes will not play on top of notes which are already playing. In addition, each currently playing note is surrounded by a range that new notes tend to avoid. This feature allows a simple adjustment of the openness or closedness of chord voicings. Voicings are more open more in the lowest pitch ranges for psychoacoustic reasons.

Rule 3: After notes are released, new notes tend to play near the same place. This rule enforces the concept of voice-leading, with which pianists and composers are familiar. Voice-leading is a strategy by which the movement of individual notes is minimized over chord changes.

The resulting probability table is used to probabilistically select each note. Next, the volume for each note is determined by a gaussian random number generator. Finally, the notes are sent on to the sample player via MIDI.

4. MUSICAL PARAMETERS

Determining the parameters was done through condensing, extrapolating from, or uniting the work of various affective musical psychologists. Of special interest were psychologists who performed experiments controlling for specific musical features in regards to emotion. Gabrielsson and Lindstrom [13] have a fairly comprehensive review of relevant work.

Many more features have been studied than can be coded into any algorithm, so we tried to focus on features that had the greatest correlation with perceived valence or arousal. We ended up with tempo, rhythmic roughness, harmonic mode, upper extensions, loudness, articulation, and pitch register.

Although each of these features has been studied by several music psychologists, often our parameters are not constructed the same way as used in the experiments. For example, in many experiments, when the scientists needed to quantify the music features they were controlling for, their preferred method was to bring in musical experts to grade each piece of music in terms of the desired features. While this is a perfectly valid quantification method, it offers no information to aid in algorithm design, so in those cases the created parameters were designed from scratch.

4.1. Rhythmic Roughness

In separate experiments, Hevner and Gundlach each studied the affective value of rhythm through the use of an obvious contrivance: reducing the entirety of rhythm to a single parameter having to do with how perceptually smooth or rough a rhythm is. They also had different definitions of smoothness and roughness. Nevertheless, their results showed that their made-up parameters could be mapped to perceived musical affect. Generally, rougher rhythms have higher arousal values, although there is entanglement with tempo which, in our case, serves to reduce this effect.

They each thought of rhythmic roughness as a discrete parameter. For example, Gundlach, whose feature definition was better for our uses, separated all rhythms into three roughness categories: 1) smooth, where all notes are of equal length, 2) uneven, where the pulse is maintained but some notes are subdivided by half, and 3) rough, where there are multiple note lengths.

Because we are interested in creating a music-emotion space which seems continuous, we could not use a rhythmic roughness parameter with only three states. Therefore, our rhythmic roughness is on a smooth to rough continuum. At maximum smoothness, each measure is divided into sixteen events of equal length. As the roughness increases, randomly selected event pairs are joined into single events. At the maximum roughness, nine junctions will have occurred (Figure 3). The value of nine was chosen because with nine randomly-selected junctions it is not possible for a rhythmic pattern consisting of all equal lengths to result.

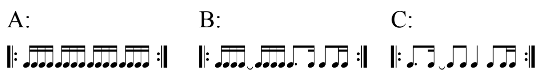

Figure 3: *(A) Smoothest rhythm, consisting of sixteen equal-length rhythmic events. (B) Rougher rhythm, in which four of the sixteen events are joined at random. (C) Roughest rhythm, in which five additional events (nine total) are joined.*

4.2. Tempo

Tempo is a feature, strongly correlated with arousal, that seems ubiquitous: everyone knows it as a synonym for "musical speed." Unfortunately, in the context of the mind, it is just not that simple. There are actually multiple different kinds of tempo. In the context of composition there is score tempo, which is easily quantified in beats-per-minute. This is what musicians think of when they think of tempo. It is also what is used in this music engine: it works by expanding or contracting the note durations in the rhythm pattern generated by the rhythmic roughness parameter.

Unfortunately, score tempo as written is not always as perceived, which is why two other definitions exist for tempo, called preferred tempo and perceived tempo. Preferred tempo is the tempo that listeners tap their foot to; and often this is half or double the score tempo. It tends to land in a moderate range of

60 to 120 beats-per-minute[14], no matter what the score tempo is.

Perceived tempo is not easily quantified by a number. It is less related to the speed of the musical pulse, and much more related to the sparsity or density of musical events. For instance, a composer might write a passage at 40 bpm, and most musicians would agree that this is a slow score tempo. However, if this passage is densely composed using all thirty-second notes, listeners will hear fast music. This loose description of musical speed, "fast", is its perceived tempo.

Perceived tempo, not score tempo or preferred tempo, is used by the mind to assign an emotional quality to the music. This is slightly problematic from an algorithmic design standpoint, because it means that tempo is affectively entangled with rhythmic roughness. However, one advantage to our rhythmic roughness parameter design is that the entanglement is easily understood: the smoothest rhythm is also the densest, so it sounds the fastest. The roughest is also sparsest, so it sounds the slowest. In some applications a fixed score tempo is needed because of aesthetic or synchronization reasons. In those cases we have found that rhythmic roughness alone is sufficient to control perceived tempo.

4.3. Articulation

Musical passages are sometimes notated as either staccato or legato. These are descriptions of musical articulation, which is equivalent to saying they are descriptions of note length. Staccato music has short, almost percussive, notes; while legato music is played in such a way that every note blends together as smoothly as possible without discordant overlapping. The former tends to be high in arousal, the latter is low.

Rather than having only two states, our parameter has a continuous parameter from short to long articulations. Between the beginnings of any two subsequent events there is a length of time called the inter-onset interval. The shortest articulations produce note lengths which are only one-fifth of the inter-onset intervals. The longest produce note lengths which are double the inter-onset intervals. In the long articulations, notes overlap, which could allow the possibility of inadvertent discordant intervals; but since our algorithm disallows non-chord tones, this likelihood is very small.

4.4. Harmonic Mode

Mode is a harmonic structure that has a strong relationship with musical valence. Often when young music students are being taught to distinguish between the commonly used minor and major modes[3], they are told that one sounds happy while the other sounds sad. These two modes seem to be the only two which have ever been studied scientifically in terms of affect.

However, other modes do exist. Much western music is based on a single eight note pattern of pitch intervals called the diatonic scale. The pattern of intervals is always found in the same order, but it can be rotated by moving the first interval to the end or the last interval to the front. Because of rotation, there exist seven different configurations of the diatonic scale. These configurations are called modes. The seven modes, according to music theorists such as Persichetti, can be ordered from brightest to darkest: Lydian, Ionian, Mixolydian, Dorian, Aeolian, Phrygian, and Locrian. The difference between any two consecutive modes on this scale is a single rotation of the diatonic scale.

There seems to be a direct relationship between Persichetti's mode brightness and the perceived valence of the music. This has been scientifically proven only in the cases of the two most commonly heard modes of Aeolian and Ionian; the other modes have historically been used only in the context of specific musical genres, so affective scientists have overlooked them. Nevertheless, we want our mode parameter to seem continuous, so it must have more than two states. Therefore our algorithm uses all the modes excepting Locrian, which for harmonic reasons is rarely used in musical practice—with exceptions in some modern music styles and jazz, where Locrian is used fleetingly.

Each mode has a different set of chords serving its tonic, dominant, and subdominant functions, but each mode has several of each. Whenever the harmony module cycles to the next chord function in its progression, the first step in translating that function into a chord list is to select, at random, one of the chords in the current mode which fulfills the desired function. This dynamic function-based chord generation means our music engine's harmonies will adapt to do the same thing harmonically across multiple different modes.

4.5. Upper Extensions

One well-studied music feature we were interested in implementing was harmonic complexity. It had a strong inverse relationship with valence, but was very loosely defined, making it hard to parametrize. We spent a significant amount of time implementing a harmonic complexity parameter into our algorithm, using a complex system of theoretically-based chord substitutions. Then we realized we had a problem: the harmonic complexity parameter was too entangled with the mode parameter.

By our definition, harmonically complex music is music that either: 1) contains a high level of chromaticism, 2) is difficult to ascribe a single mode to, or 3) both. Our parametrization succeeded in creating harmonically complex music, but caused the mode parameter to be less reliable. Therefore, for the time being, we have stopped using the harmonic complexity parameter. However, the work we did was not a total loss: part of our implementation could be retained. This part is the upper extensions parameter.

The simplest type of chord is the triad, consisting of three notes spaced in intervals of thirds. More complex chords can be constructed from triads, however, through the addition of notes called sevenths and upper extensions. Jazz and twentieth-century music, in particular, makes common use of these extra notes.

At the minimum upper extensions value, all chords will be triads. As the upper extensions value begins to increase, sevenths will begin to be included. At the maximum upper extensions value, all non-discordant sevenths and upper extensions will be included in the chord.

It seems reasonable to assume that this parameter will, like harmonic complexity, have an inverse relationship with musical valence, so this is how we use the parameter. Once we have the capability, however, it is important that we study this parameter in isolation from the other parameters in order to determine if our assumptions are correct.

4.6. Loudness

Loudness has been, quite intuitively, associated with musical arousal through scientific experimentation. There are differences between studies, however, leading us to implement loudness in

3 Henceforth, we will refer to the major and minor modes with their alternate names, the Ionian and Aeolian modes.

more than one way. In Hevner's experiments, a piano player performed all the music excerpts for her test subjects, so in her case the louder music was caused by louder playing: each individual note was played louder. Other scientists, however, used recorded orchestral music. In orchestral music, the loudness changes with the number of players at any one time. If orchestra members are resting, the music is quieter than when all are playing at once.

We implemented both types of loudness. Individual note loudness is manipulated by adjusting the mean of the gaussian random number generator used in the voicing module for determining note volume. At the same time, we have another, more indirect loudness manipulation. During quiet music, only two notes will play at any one time. As the loudness increases, the sizes of the chords increase as well, until at maximum eight notes are played per event—nearing the limits of practicality for a real piano performance. This second loudness manipulation may someday, after some experimentation, be split off into its own parameter. But for now, since we do not know how it affects emotion other than that it corresponds to loudness, we will leave it entangled with the more standard definition of loudness.

4.7.Pitch Register

This parameter is used to describe, in a holistic way, the highness or lowness of a musical passage. As studied by the various psychologists, it has a weak relationship with valence, quite dependent on context. In the future, when we begin to use these algorithmic techniques to study emotion models with more than two dimensions, we may find that pitch has a greater correspondence with a dominance axis: this would be an intuitive mapping due to the tendency of low pitches in nature to originate from large sources.

The pitch register mapping is implemented in a simple way. As you will recall, the voicing module's first note selection rule involves tending to generate notes in a central range. We manipulate the pitch register by transposing that range up or down in pitch.

5.MANIPULATION OF FEATURE VECTOR

So far, we have described how the algorithm is structured, and we have described each affective parameter we use in it. However, every one of those parameters is contextual. It is often the case in real music that some feature which would normally serve to move the affect one direction is overpowered by other features. For example, there are numerous examples of music which, despite being in a minor harmonic mode, can only be described as upbeat and cheerful.

Therefore, a strategy must exist for controlling the parameters as a group. For now, we use the simplest possible strategy: the parameters move in parallel with one another. If, on our interface, we choose to hear music at the maximum arousal, we will hear the maximum tempo, maximum roughness, maximum loudness, and shortest articulations. If we desire minimum arousal, the opposite will result. Similarly, if we ask for the highest valence, we will hear music in the brightest mode, at the highest pitch range, with the least upper extensions, and vice versa if we choose to hear music with low valence.

The parallel strategy is wonderful for tuning the parameters: one can listen to music that is meant to sound angry and say to oneself, “This does not sound fast or discordant enough to be angry”; then go back and re-scale the parameters. Unfortunately, it means that if the engine is left playing the same valence/arousal setting for a long time, there might not be as much musical variety as desired. One way to alleviate this could be to change to a probabilistic strategy, where the desired valence and arousal are approximate values that change from time to time within a small, centrally weighted range. This might help—somewhat—with the musical self-similarity.

We are currently in the process of developing a far better strategy, however. Right now the algorithm is calibrated to our perception of musical emotion. We are currently performing user studies in order to calibrate it to the average person's perception. Once this has been accomplished, we can begin a process of incremental studies isolating and controlling for each individual musical parameter, as well as any new parameters we might want to implement.

The result of each incremental study will be a vector or curve in the emotion space for the studied feature. Given that data, we will no longer need to move the parameters in blocks: we can use the change in one parameter to offset the change in another parameter. We will be capable of much more musical variety while leaving our music generator on only one setting.

6. PRELIMINARY APPLICATIONS

Because of the newness of the algorithm and the concepts, we are only just beginning to use this work in digital media applications. However, we have begun to explore it for artistic use with dancers, by using a motion capture system to drive the valence and arousal controls. This possibility could lead to further affective computing research, where we try to extract the emotional quality from dance using motion capture, so the interactive music can automatically follow the emotion of the dancer.

Also using motion capture, we have begun to explore this music engine in the context of a Parkinson's Disease rehabilitation system. Physical therapy sessions for Parkinson's sufferers use some repetitive exercises in order to help the patient retain balance and mobility. We can provide affective music feedback to these patients during their sessions. This research path leads to some interesting questions, like: Which emotions are more conducive to rehabilitation? Is it better for the patient to be excited or relaxed? Does each patient have different affective needs, to which the system must adapt?

Because of the way music, emotion, and movement are all linked together neurologically[15], affective music feedback has the potential to greatly improve Parkinon's patients' benefit and enjoyment of rehabilitation. Also, Parkinson's sufferers have been shown to respond positively to active music therapy[16], so there is some precedent for our optimism in regards to this application.

7.CONCLUSION

This paper has focused on the design details of our music generation engine, with emphasis on music feature parametrization. We designed our parameters as individual elements, around music features which have been tested by affective psychologists. This results in a music generation algorithm which can be set to generate music corresponding to a specific arousal and valence.

Our work in defining affective musical parameters in algorithmic ways, as opposed to the often subjective way in which the features were originally studied, was specifically done so that the parameters could be generalizable and repeatable in other, stylistically and musically different affective music engines. For that matter, they might even be useful as guidelines for musicians who compose in the traditional way without computers.

The possibilities for applying these algorithmic techniques in digital media are quite vast. We are of the opinion that the video game industry, for example, could make great use of this type of music generation. Currently, most video game music is composed out of audio loops, a technique which can grow monotonous and is not interactive. An affective algorithm need never repeat itself, can easily be designed to interact and synchronize with the game, and serves the purpose game designers need from music by producing the desired emotion for the scene.

Other more serious digital media systems can make use of these techniques as well, as evinced by the Parkinson's Disease rehabilitation system that was discussed. Also, scientists could use an affective music generator in experiments when trying to learn more about disorders like depression or autism. We are currently in the process of collecting data to ensure that our algorithm corresponds well to the greater population's perception of music and emotion. Once that is completed, this algorithm can be a baseline for experimentation on a variety of topics, including any music parameters we want to define in terms of affect—especially the ones which are already implemented.

In short, we have created a novel musical interface, and hope that it leads to open-ended research on emotion, music, and digital media. We also hope that the concepts used in designing the engine can proliferate, leading to a new way of thinking about algorithmic music design; where concepts like tempo and loudness aren't just implementation problems, but are considered in terms of emotion as well.

8. AUDIO EXAMPLES

Here are some examples of the algorithm's sound output:

Angry: High arousal, low valence
http://ame2.asu.edu/students/riwallis/Angry.mp3

Sad: Low arousal, low valence
http://ame2.asu.edu/students/riwallis/Sad.mp3

Joyful: High arousal, high valence
http://ame2.asu.edu/students/riwallis/Joyful.mp3

Glad: Low arousal, high valence
http://ame2.asu.edu/students/riwallis/Glad.mp3

9. REFERENCES

1] C. Reynolds and R.W. Picard, “Designing for Affective Interactions,” in *Proc. of 9th Intl. Conf. on Human-Computer Interaction,* New Orleans, LA, United States, Aug. 5-10, 2001.

2] R. Legaspi, Y. Hashimoto, K. Moriyami, S. Kurihara and M. Numao, “Music Compositional Intelligence with an Affective Flavor,” in *Proc. Intl. Conf. on Intelligent User Interfaces,* Honolulu, HI, United States, Jan. 28-31, 2007, pp. 216-224.

3] S. Dubnov, S. Mcadams and R. Reynolds, “Structural and Affective Aspects of Music from Statistical Audio Signal Analysis,” *Journal of the American Society for Information Science and Technology,* vol. 57, no. 11, pp. 1526-1536, 2006.

4] K. Hevner, “Experimental Studies of the Elements of Expression in Music,” *American Journal of Psychology,* vol. 48, pp. 246-268, 1936.

5] R.H. Gundlach, “Factors Determining the Characterization of Musical Phrases,” *American Journal of Psychology,* vol. 47, No. 4, pp. 624-643, 1935.

6] J.A. Russell, “A Circumplex Model of Affect,” *J. Pers. Soc. Psychol.,* vol. 39, no. 6, pp. 1161-1178, Dec. 1980.

7] J.A. Russell and A. Mehrabian, “Evidence for a Three-Factor Theory of Emotions,” *Journal of Research in Psychology*, vol. 11, no.3, pp. 273-294, 1977.

8] R.J. Fontaine, K.R. Scherer, E.B. Roesch and P.C. Ellsworth, “The World of Emotions is not Two-Dimensional,” *Psychological Science,* vol. 18, no. 12, pp. 1050-1057, Dec. 2007.

9] G.L. Collier, “Beyond Valence and Activity in the Emotional Connotations of Music,” *Psychology of Music,* vol. 35, no. 1, pp. 110-131, 2007.

10] J.A. Russell, A. Weiss and G.A. Mendelsohn, “Affect Grid: A Single-Item Scale of Pleasure and Arousal,” *J. Pers. Soc. Psychol.,* vol. 57, no. 3, pp. 493-502, Dec. 1989.

11] V. Persichetti, *Twentieth-Century Harmony: Creative Aspects and Practice*, chapter W.W. Norton & Co. Inc., NY, USA, 1961.

12] M. Haylock, “VSamp,” Available at http://www.vsamp.com, Accessed March 23, 2008

13] A. Gabrielsson and E. Lindstrom, Music and Emotion: Theory and Research, chapter The Influence of Musical Structure on Emotional Expression, pp. 224-248, P.N. Juslin and J.A. Sloboda Eds. Oxford UP, 2001.

14] R.A. Duke, “Musicians' Perception of Beat in Monotonic Stimuli,” *Journal of Research in Music Education*, vol. 37, no. 1, pp. 61-71, 1989.

15] I. Molnar-Szakacs and K. Overy, “Music and Mirror Neurons: From Motion to 'E' Motion,” *Social Cognitive and Affective Neuroscience, Special Issue: Genetic, Comparative and Cognitive Studies of Social Behavior,* vol. 1, no. 3, pp. 235-241, Dec. 2006.

16] C.Pacchetti, F. Mancini, R. Agliere, C. Fundaro, E. Martignoni and G. Nappi, “Active Music Therapy in Parkinson's Disease: An Integrative Method for Emotional Rehabilitation,” *Psychosomatic Medicine*, vol. 62, no. 3 pp. 386-393, 2000.