Real-Time Emotional Adaptation in Automated Composition

Maia Hoeberechts and Jeffrey Shantz
Department of Computer Science
University of Western Ontario
London, Ontario, Canada
hoebere@csd.uwo.ca, jshantz4@csd.uwo.ca

Abstract. A feature shared by films, restaurants, video games, stores and dinner parties is that music is used to create ambiance and set a mood. Music conveys emotions and evokes responses from the listener. The challenge we address in this research is to create an automated composition system which can adapt generated music in real-time to desired emotional characteristics. This paper describes the underlying theoretical foundations for the design of adaptive components and their implementation in AMEE (Algorithmic Music Evolution Engine). We survey related work, and establish a foundational standpoint for our research.

1 Introduction

The motivation for our research is to create an automated composition system which can adapt composed music to desired emotional characteristics in real-time. We established a theoretical foundation for our work in the field of emotion perception in music. Background research was conducted in music psychology to determine relationships between alterations to low-level musical elements and changes in perceived emotion. We have a fully functional composition engine—AMEE (Algorithmic Music Evolution Engine)—which creates novel emotionally adaptive compositions in real-time. Our primary target application is video game music, but we have considered other applications such as ringtone generation, rapid musical prototyping and standalone composition tools.

In our previous work, we described an initial prototype and proof of concept for AMEE [13]. In the present research, we have greatly expanded the theoretical background underlying the emotional adaptation, extended the supported moods, and implemented real-time support. Recently, two aspects of AMEE's design have been patented [12]: the pipelined architecture (see Section 4.1) and emotion mapper (see Section 4.2).

The component of the system which is the focus of this paper is the emotion mapper. Based on the desired mood characteristics of the piece, the emotion mapper makes adjustments to low-level components to effect emotional adaptation of the music. The system alters six musical dimensions—mode, tempo, pitch, consonance, volume and articulation—which are used to produce ten supported emotions—happy, sad, triumphant, defeated, excited, serene, scared, ominous, angry and crazy. An important feature of the emotion mapper is that adjustments can be made in real-time during the composition process; for example, a sudden change from serene to ominous music can occur when a character in a game enters a cave fraught with danger. Furthermore, the same piece can be reinterpreted under different moods, enabling altered compositions containing recognizable musical themes to be created.

In this paper, we survey related work (Section 2), establish the theoretical foundations of our project and summarize the background research for the emotion mapper (Section 3), and describe implementation of real-time emotional adjustments in AMEE (Sections 4 and 5). We conclude by suggesting directions for further research (Section 6).

2 Related Work

A number of approaches for affective composition have emerged in the literature that can be roughly divided into the categories of stochastic processes, which incorporate random choice or probability in composition; rule- or knowledge-based methods, where input parameters are mapped to musical features based on well-established rules; evolutionary/genetic techniques, which create new patterns through mutation; recombinatorial approaches that combine short, prearranged music segments to produce longer pieces; state machine or grammar driven approaches in which composition is controlled though state or symbol transitions respectively; and hybrid approaches.

MySong [24] uses a stochastic, machine learning approach, and serves as the underlying technology for Microsoft Research Songsmith [19]. Songsmith allows a user to sing a melody track, and will automatically generate a full band accompaniment according to one of a number of user-selectable styles, such as *50's Rock* or *Big Band*. Users can tune the emotional quality of generated music by adjusting a *Happy* slider, which alters the mode of the piece, and a *Jazzy* slider, which affects how closely selected chords fit the notes of the melody. The system employs a Hidden Markov Model (HMM) that is trained using a database of melodies, along with their associated chord progressions.

Livingstone et al. make use of the rule-based Two-Dimensional Emotion Space (2DES) [23], which maps the input parameters of valence (positive/negative charge) and arousal (energy level) to a set of both structural and performative musical features, such as tempo, mode, accents, and slurs [17]. Emotions can be expressed by selecting appropriate values for these parameters. For instance, selecting high values for both valence and arousal might produce a happy composition with increased tempo and a major mode employed. A comprehrensive treatment of 2DES and its capacity for musical feature transformations is given by Schubert [23]. A similar approach is used by the emotion mapper in AMEE, but the two systems differ at the level of parametric abstraction employed. In AMEE, emotional content – such as the level of happiness, anger, or serenity – can be directly specified, rather than being the result of a combination of more abstract parameters.

Evolutionary approaches have also demonstrated promising results in adaptive music composition. Legaspi et al. asked human subjects to rate a set of musical scores according to several bipolar emotional adjectives using a web-based tool [16]. These ratings are used by their system to generate rules to identify musical structures that express various emotions. Composition occurs by means of a genetic algorithm which employs a fitness function accounting for both adherence to musical structure and desired emotional expression when evaluating mutations.

Eladhari et al. use a neural network approach to generate music tailored to suit individual characters in video games [7]. The system employs a *Mind Module* that consists of a spreading activation network, comprised of interconnected nodes modelling traits, emotions, moods, and sentiments. The mood of a character is measured in two dimensional space, with one axis representing an inner mood scale (depressed to happy) and the other representing an outer mood scale (angry to exultant). Inner mood influences the harmonic qualities of the composition, such as the scale employed, while outer mood affects the time signature of the piece. Real-time affective changes occur when a sentiment felt by a character exceeds a given threshold, triggering an emotional response in the activation network, and influencing the values of its inner and outer mood scales.

Comair et al. combine a finite state machine with a recombinatorial approach to produce adaptive audio for video games [3]. Pre-composed blocks of music are associated with the states of the machine. For instance, one state might be assigned a relaxed music block, while another state might be assigned a more excited block. An *interactivity parameter* modelling some aspect of gameplay determines when to transition between states through comparison to numeric entry and exit values assigned to each state. Although composition in the system does not occur in real-time, each player experiences a different composition since gameplay changes the order in which blocks of music are played.

Affective audio has also penetrated the commercial video games market, with an increasing number of games adopting at least some form of emotional adaptation to enhance gameplay. A *mod* developed for the popular *Unreal Tournament 2004* changes music based on events occurring within the game, as well as a player's status [5].

Automated composition in AMEE is accomplished using a hybrid approach comprising recombination, stocastic and rule-based processes. The adaptive component (the emotion mapper), which maps ten input emotion parameters to musical features based on established theoretical knowledge, is rule-based.

3 Music and emotions

The perspective on emotional expression in music which is the foundation for our research comprises several key points:

1. There is a set of moods which can be convincingly expressed in music.
2. The moods conveyed by a given piece of music will be identified relatively consistently by listeners with a similar cultural background.
3. Perceived moods in a composition can be linked to specific characteristics of the piece such as mode, articulation, tempo, volume, etc. Changes to these characteristics alter perceived mood.

These statements are not without controversy. In the following sections, we summarize some key theoretical perspectives on music and emotion, and clarify our standpoint in the context of the surveyed positions.

3.1 Overall perspective

That music is capable of evoking and/or expressing emotions is not a universally accepted idea. Stephen Davies summarizes some of the problems as follows: "Given that music is nonsentient, how could emotions be expressed in it? Are musical phrases understood as utterances about emotions? [. . .] Or are they interpreted as gestures betraying the composer's or performer's emotions? Or are they expressive insofar as they move the listener?" ([6], p. x). We also face the question of whether the emotions are merely *perceived* in the composition by the listener, or whether they are *felt* — that is, experienced more fundamentally than simple cognitive identification [15]. For a survey of key research perspectives and issues, see [14].

The major theoretical positions regarding emotion and music are distinguished as follows [1, 22]:

Absolutist vs. Referentialist
An absolutist holds that the meaning of a musical work is contained within the work itself, and can be conveyed through the perception of the work in isolation. A referentialist maintains that music's meaning involves more than the notes itself; for example, a story or spiritual state is necessary as background context for understanding the music.

Formalist vs. Expressionist
The absolutist viewpoint can be further classified as either formalist or expressionist. A formalist holds that the meaning of a piece of music is primarily intellectual and is derived from perception of the structural relationships between elements in a composition (for example, the musical form or harmonic progressions). An expressionist believes that the principal sources of meaning of a composition are the emotions evoked in a listener as a result of listening to (or perceiving the structural relationships within) the piece.

For our purposes, we adopt an expressionist viewpoint since we are designing an interface which allows the user to choose characteristics that the generated music should be capable of expressing. Both the absolutist and referentialist perspectives are useful to us. We can aim to generate music which "stands on its own" in the sense that the generated music contains elements designed to evoke specific emotions in the listener. However, in the context of a video game where there is a story line, we might also use a referentialist perspective in which the game players' associations between the storyline and the background music serve to evoke emotions.

3.2 Types of emotions expressed in music

In the context of psychology of music, emotional response to music is divided into several different categories. Here are the most important ones (definitions are based on those in [1] and [22]):

- *Affect/emotion/mood*: This is the "usual" interpretation of emotional response. While listening to a given piece, a person feels an emotion such as joy or anxiety.
- *Aesthetic*: An aesthetic response has both an affective and intellectual component. If we think a piece of music is "beautiful" or "well-structured", these would be aesthetic evaluations of it.
- *Preference*: This term encompasses like or dislike of a particular piece of music, or preference of one over another. To distinguish preference from the previous two categories, imagine a piece which makes you sad, that you think is beautiful, but that you do not like.
- *Taste*: Musical taste reflects long term attitudes toward music, developed from experience. The content of your CD collection would reflect your musical taste.

For the purposes of the engine design, we are most interested in the first category: How does a given piece of music make a listener feel? Preference is also relevant since we probably want to produce music that our listener likes, however this is a secondary concern. Likewise, we are not so interested in aesthetic response and taste. Therefore, we are justified in concentrating our research on papers which deal with emotional/mood/affective response to music.

Returning now to our foundational perspective (Section 3), we are justified in adopting point 1. from a expressionist standpoint. Regarding point 2., studies confirmed that within a fixed cultural group, emotions are identified relatively consistently by listeners of different ages [25], and varied musical expertise [25, 9, 4].

3.3 Emotion perception studies

The methodology used in studies on music and emotion varies widely. This is due in part to the lack of a standard theoretical framework [18]. Juslin and Västfjäll's recent work [15] shows promise in addressing this problem: the authors argue that musical emotions must be investigated with regard to their underlying mechanisms such as brain response and musical expectancy. They claim that some disparities in experimental results can be explained by the researchers' failure to take these mechanisms into account.

In the studies used as a basis for our research, there is variation in the music or musical stimuli which are used. We find studies analyzing whole pieces of classical music or popular music songs [10, 8], segments of pieces [25], altered pieces (order of sections, ending, tempo, harmony etc. changed) [10, 26], pieces composed specifically for the study to represent a desired musical quality [9, 20, 8], and even single line melodies or single intervals [4].

The studies which consider a musical element in isolation are the easiest for us to interpret, since they state a result about one clearly defined musical element. However, it is unclear whether this methodology produces musically meaningful results since some researchers deny that there is musical significance to an element devoid of musical context [10]. In creating our software, we incorporated results from any of these various studies, focusing most strongly on those results which have been corroborated in studies with different methodologies.

In addition, there are different perspectives on the appropriate metric for measuring emotional response:

Physiological response
Some researchers maintain that the mechanism through which music is capable of evoking emotions is by causing a physiological reaction, such as changes to heart rate or muscle tension, which is in turn perceived by the subject as an emotion. Bartlett's work [2] contains a comprehensive survey of more than 130 papers in this area. The advantage of studying emotional response in this way is that the response is quantifiable and easily measured; however, in psychology research, the link between physiological reaction and subjective emotion is not clear-cut. The step from the physiological reaction to emotional experience is not the concern of our research.

Neurological response
PET scans, fMRI, EEG and other brain imaging techniques have been used to visualize neurological response to musical stimuli. Studies comparing responses from non-musicians and musicians, music listening and playing, perception of harmony, rhythm, melody etc. have been done [21]. However, as with physiological response, analyzing the link between neurological response and emotion is beyond the scope of our research.

Adjective descriptors
Adjective descriptors are the traditional means for evaluating emotional response to music. Normally, subjects are asked to check off adjectives from a list which they feel best describe their feelings while listening to music or samples. Hevner's adjective circle [9, 10], is an early example of this research method, and has served as the basis of much subsequent research. There have been many other attempts to compile complete and appropriate adjective lists; for example, in [25], an experiment was conducted in which a list of 100 mood indicators were rated on their usefulness in capturing musical meaning. This study contributed to the choice of supported emotions in AMEE. However, a problem with using adjective descriptors is that they do not allow for any degree of response.

Semantic differentials
Semantic differentials are pairs of bi-polar adjectives for which a rating can be chosen by the subject (e.g. [20]). An example would be a 10 point scale labelled "sorrow" at one end and "joy" at the other. This technique overcomes the main limitation of adjective descriptors in that one can specify the degree to which a given emotion is experienced. However, it has a disadvantage—there are implied semantics in the choice of opposite terms for each differential. To illustrate, consider the pair excited-calm. Suppose you find a musical sample not to be exciting—with this differential, you would be forced to rate it as calm, despite the fact that you might feel that some other adjective such as "stately" better describes it.

In AMEE, we combine advantages of adjective descriptors and semantic differentials with the following solution: adjective descriptors are used, but a degree of intensity of each desired emotion between 0 (emotion not present) and 100 (extreme emotion) can be specified. There are two significant challenges in determining the relative contribution of musical characteristics to perceived mood. First, characteristics affect more than one mood. For example, increased tempo is associated with increased perception of happiness, excitement, and anger. Second, combinations of desired moods pose problems since the expected changes to characteristics sometimes conflict. For instance, tempo should be increased to evoke happiness, but decreased for serenity. Hence, we encounter a conflict when attempting to produce a mood which is both happy and serene. The resolution of these problems in AMEE (dimension adjustments) is presented in Section 4.2.1.

The emotion perception studies discussed in this section and elsewhere in the paper studies present evidence for points 1. and 3. in our foundational perspective (Section 3).

4 Emotional Adaptation in AMEE

4.1 System overview

The high-level structure in AMEE is a pipelined architecture, which builds a piece in stages by calling on *producers* to create high-level musical elements such as sections, blocks and musical lines. The producers rely on *generators* to create the low-level musical elements such as harmonic patterns and motif patterns. In order to instantiate composition, a *performer* object must be created which contains the pipeline and all the musicians participating. Music generation commences by calling the *play()* method on the performer. Figure 1 summarizes the pipeline process, and Figure 2 provides an overview of the system architecture.

I. The section producer supplies a section.

II. The block producer supplies a block and uses a harmonic generator to create a harmonic pattern for the block.

III. The line producer uses motif generators to fill the block with a musical line for each musician.

IV. The completed block is sent to the output producer for playback.

Figure 1: Overview of pipeline process

The generators can operate in a library based mode (in which musical elements are reused), or create new elements. The musical lines can either be composed in one global style (Ensemble mode), or each virtual musician's line can reflect a personal style (Jam Session mode). The system includes both a programming interface (allowing it to be embedded in another software system), and a graphical user interface (allowing it to be used as a standalone tool). For more details on AMEE's design, see [13].

At every step in the generation process, the musical components are adjusted by an *emotion mapper* to reflect the currently desired emotional characteristics of the composition. The pipeline process can be interrupted in real-time when a change to the composition is requested. The EventDispatcher shown in Figure 2 initiates real-time changes to the composition pipeline.

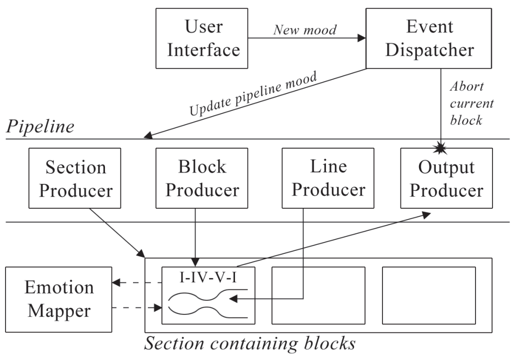

Figure 2: Overview of AMEE system architecture

Some important design considerations behind the system were to permit flexibility in the music composition process, to provide an easily extensible architecture with a simple API (application programming interface), to allow musical motifs and harmonic sequences to be stored and reused, and to permit the composition process to be either fully automatic, or human and/or software controlled.

4.2 Emotion mapper implementation

The emotion mapper acts on low-level musical elements (volume, tempo, mode, harmonic patterns, and motif patterns). Motif patterns are sequences of harmony and mode independent notes: for instance, a descending sequence of 3 whole tones. Within a composition, the motif will be resolved to actual mode tones such as B, A, G or $F\#, E, D$, depending on the current mode and the harmonic function of the motif (for more details on motif resolution see [13]). The emotion mapper adjusts motif patterns through alterations to pitch, consonance and articulation.

There are three factors which contribute to the adjustment made to a specific musical characteristic: the desired mood, the initial value of the characteristic, and the degree of effect of the mood on that characteristic. For the ten supported emotions in AMEE (happy, sad, triumphant, defeated, excited, serene, scared, ominous, angry and crazy), the *emotional intensity*, a percentage ranging from 0 (emotion not present) to 100 (extreme emotion), can be specified by the user. Any combination of the supported emotions is permitted. The *mood* is the collection of emotion intensities. Mood is considered in two contexts in the system: global mood, which represents the overall feeling of the composition, and individual moods, which characterize the mood of each musician in the performance. Depending on the desired outcome, a mood from either of these categories could be sent to the emotion mapper.

Figure 3 describes the steps executed to emotionally adjust a musical element. This process is initiated by one of the producers, and therefore could occur at each step in Figure 1.

i. If necessary, the producer uses a generator to create a low-level element.

ii. The producer calls the emotion mapper to adjust the element according to the desired mood.

iii. The emotion mapper determines adjustments to each dimension (mode, tempo, pitch, consonance, volume, and articulation) based on mood.

iv. The relevant features of the low-level element are adjusted.

Figure 3: Emotion mapper process

4.2.1 Dimension adjustments

At Step iii. in Figure 3, the relative contributions of each emotion in the desired mood are combined to determine the requisite effect on the musical characteristics. As an example of relative contributions, consider the influence of happy, sad, and scared on consonance as summarized in Table 1. The references used to establish each result are also shown.

Emotion	Effect	[refs]	Contribution
happy	**o** +	[4] [10]	+0.75
sad	**-**	[4, 10]	-0.75
scared	**-**	[4, 10, 20]	-1.0

o: no effect, **-:** more dissonant, **+:** more consonant

Table 1: Relative contribution example

The "contribution" of each emotion is a real number between -1.0 and 1.0 which reflects the polarity and strength of the influence of that emotion on the respective musical characteristic. In Table 1, we see that happy has a strong positive impact on consonance, whereas scared has a very strong negative impact on consonance (therefore producing dissonance).

The dimensional contributions are combined using a formula to produce a final adjustment number for each musical characteristic. For dimension i, calculate its adjustment $adjust_i$ from emotions $1..\#emotions$ by combining the intensity of emotion j ($intensity_j$) with its relative contribution to dimension i ($contrib_{(j,i)}$):

$$adjust_i = \frac{1}{100}\left(\sum_{j=1}^{\#emotions} contrib_{(j,i)} \times intensity_j\right)$$

where $\#emotions$ is the number of emotions (presently, 10). In the case that there are several emotions with extreme intensities, $adjust_i$ might fall outside the range $[-1.0, 1.0]$. In this case, it is set to -1.0 or 1.0 as required (maximal influence).

To illustrate emotional adjustment, suppose a musical change is requested for a character setting out on a quest far away from home (a mixture of happy and sad), who suddenly is startled by a creature in a forest (scared). The mood in AMEE could be set as follows: happy=30, sad=20, triumphant=0, defeated=0, excited=0, serene=0, scared=50, ominous=0, angry=0 and crazy=0. Using the above formula and contributions in Table 1, the calculated dissonance contribution is:

$$\frac{1}{100}(0.75 * 30 + -0.75 * 20 + -1.0 * 50) = -0.425$$

As the result is negative and relatively large, a strong influence toward more dissonance would be heard in the music. The adjustments for other dimensions are calculated analogously.

4.2.2 Musical element adjustments

As the final step, the emotion mapper uses the calculated dimension adjustments to make changes to the low-level musical elements. The three types of changes AMEE supports are listed below, together with the characteristics to which they apply.

1. **Discrete alteration:** for characteristics with discrete values (e.g. mode and consonance). Depending on the final dimension adjustment and thresholds for change, the characteristic may be altered to a different value. For example, the number of dissonant notes in a motif might be increased or decreased.
2. **Continuous alteration:** for characteristics with a range of possible values (e.g. tempo, pitch, articulation). The existing characteristic is increased or decreased by a percentage depending on the dimension adjustment. For example, the tempo is reduced to 75% of its present value.
3. **Effect alteration:** to portray the desired emotion, the characteristic should vary in time (e.g. volume). For example, for defeated, we gradually decrease the volume, and for crazy, there are sudden changes in volume in the composition.

This model could support other types of alterations, such as pitch bend effects, distortion, rhythm alterations, etc.

4.3 Real-time interaction

AMEE allows the mood of composed music to be altered in real-time by changing the values of its ten emotional intensities. Depending on the application in which AMEE is embedded, changes can be made directly by a user, or can be the result of an event occurring, such as the start of a battle in a video game.

An application developed to demonstrate the engine allows the user to alter the desired mood by dragging sliders to change emotion parameters. AMEE generates and outputs blocks of music one at a time, and thus once a mood change has been made, an event is dispatched to the pipeline, playback of the current block stops, and it is discarded. Generation and playback of the next block then immediately occurs with the new mood applied to it (see Figure 2).

AMEE saves the state of the pipeline before generating each block. When a mood change event occurs and a block is discarded, the state of the pipeline is restored to the checkpoint taken before the block was generated, and the new block is generated using the restored pipeline. The system also tracks the precise time at which the old block was aborted, and restarts playback of the new block at that same position to ensure continuity.

5 Discussion of AMEE's approach

Portability
We designed AMEE with the intent of embedding it within an application such as a video game or assisted composition system. Care was taken to minimize programming language dependence and to decouple the engine from the underlying music output technology.

Intuitive API
Changes to the mood of a given AMEE composition are made by selecting intensities for specific emotions, such as *triumphant* or *scared*. Working directly with emotional adjectives is intuitive to end users and programmers who may not be sufficiently familiar with the theory of emotion to use more abstract parameters such as valence and arousal.

Extensibility
Additional musical styles can be readily added to AMEE's repertoire by creating pattern databases based on existing pieces. This is in contrast to machine learning approaches, for instance, in which extensive training is necessary to adopt a new style with palatable results, and composition quality is highly dependent on the selection of sufficiently representative training data.

Ease of collaboration
One of the goals governing the design of AMEE has been to enable collaboration between composers who might wish to use the system to generate ideas and share interesting pieces with others. AMEE facilitates this process by allowing users to create and export musical styles, musicians, and generated compositions, which can then be shared with other composers.

6 Future development

The engine is fully functional, and has been tested extensively on a Windows system. The next step in our research will be to embed the engine in a casual game in order to showcase its features and to conduct further testing in a realistic application setting.

We list here a few other refinements which are currently planned or in progress.

6.1 Improvements to musicality

- **Exaggerate emotion effects**. Informal feedback has indicated that some emotional changes are not as easily perceived in AMEE as others (for example, altered articulation is often not noticed).
- **Changes to instruments**. Presently, we do not change instrumentation based on emotion. Some instruments are conventionally associated with moods; for instance, a melancholy saxophone. An argument for not changing instrumentation is that AMEE attempts to preserve the original composition and convey changed mood through changes to low-level elements only. An alteration to the sound of the piece might be too drastic to fulfill this goal. Further thought and research is needed on this topic.
- **Transitions between moods**. Changes are presently made abruptly to the music when a mood alteration is requested. It would be desirable to allow the user to request a gradual change.
- **Refinements to emotions**. Further literature research, refinements to existing emotions and possible extensions to other supported emotions and musical alterations are in progress.

6.2 Technical improvements

- **Database creation tools**. A tool for creating pattern databases based on existing pieces has already been created, and we are currently adding support for creating motif pattern databases.
- **Instrument quality**. While artificial instrument quality is a common complaint amongst MIDI users, we are researching possibilities for creating custom MIDI soundbanks to improve the realism of instruments used in compositions.
- **Mobile edition**. In response to the increasing popularity of mobile games, we are considering porting AMEE to a mobile platform. To our knowledge, no affective composition engine currently exists in the mobile domain.

6.3 Empirical study

To further validate the implementation of real-time emotional adaptation in AMEE, we are considering conducting a formal survey of perceived emotions in the engine's compositions. This would aid in determining directions for future development and suggest improvements to the present framework.

Acknowledgements

This research was supported by the NSERC Idea to Innovation (I2I) Program (Principal Investigator, Dr. Micheal Katchabaw, University of Western Ontario). The authors thank Michael Katchabaw and the Audio Mostly reviewers for their suggestions.

References

[1] Harold F. Abeles and Jin Won Chung. Responses to music. In Hodges [11], pages 285–342.

[2] Dale Bartlett. Physiological responses to music. In Hodges [11], pages 343–385.

[3] Claude Comair, Rory Johnston, Lawrence Schwedler, and James Phillipsen. Method and apparatus for interactive real time music composition. United States Patent 6,822,153, November 2004.

[4] Marco Costa, Pio Enrico Ricci Bitti, and Luisa Bonfiglioli. Psychological connotations of harmonic musical intervals. *Psychology of Music*, 28:4–22, 2000.

[5] Creative Technology Ltd. Unreal Tournament 2004 Spotlight. <http://www.soundblaster.com/technology/welcome.asp?1=sbg&j2=ut2004>, accessed May 2009.

[6] Stephen Davies. *Musical Meaning and Expression.* Cornell University Press, 1994.

[7] Mirjam Eladhari, Rik Nieuwdorp, and Mikael Fridenfalk. The soundtrack of your mind: mind music - adaptive audio for game characters. In *ACE '06: Proceedings of the 2006 ACM SIGCHI international conference on Advances in computer entertainment technology*, page 54, New York, NY, USA, 2006. ACM.

[8] Alf Gabrielsson and Patrick N. Juslin. Emotional expression in music performance: Between the performer's intention and the listener's experience. *Psychology of Music*, 24:68–91, 1996.

[9] Kate Hevner. The affective character of the major and minor modes in music. *The American Journal of Psychology*, 47(1):103–118, January 1935.

[10] Kate Hevner. Experimental studies of the elements of expression in music. *The American Journal of Psychology*, 48(2):246–268, April 1936.

[11] Donald A. Hodges, editor. *Handbook of Music Psychology.* IMR Press, second edition, 1996.

[12] Maia Hoeberechts, Ryan Demopoulos, and Michael Katchabaw. A flexible music composition engine. PCT Patent WO 2009/036564.

[13] Maia Hoeberechts, Ryan J. Demopoulos, and Michael Katchabaw. A flexible music composition engine. In *Audio Mostly 2007 Conference Proceedings*, Ilmenau, Germany, September 27–28 2007.

[14] Patrik N. Juslin and John A. Sloboda. *Music and Emotion: Theory and Research.* Oxford University Press, 2001.

[15] Patrik N. Juslin and Daniel Västfjäll. Emotional responses to music: The need to consider underlying mechanisms. *Behavioral and Brain Sciences*, 31(5):559–621, 2008.

[16] Roberto Legaspi, Yuya Hashimoto, Koichi Moriyama, Satoshi Kurihara, and Masayuki Numao. Music compositional intelligence with an affective flavor. In *IUI '07: Proceedings of the 12th international conference on Intelligent user interfaces*, pages 216–224, New York, NY, USA, 2007. ACM.

[17] Steven R. Livingstone, Ralf Mühlberger, Andrew R. Brown, and Andrew Loch. Controlling musical emotionality: an affective computational architecture for influencing musical emotions. *Digital Creativity*, 18(1):43–53, 2007.

[18] Patrick T. McMullen. The musical experience and affective/aesthetic responses: A theoretical framework for empirical research. In Hodges [11], pages 387–400.

[19] Microsoft Corporation. Microsoft Research Songsmith, 2009.

[20] Sören Nielzén and Zvonimir Cesarec. On the perception of emotional meaning in music. *Psychology of Music*, 9(2):17–31, 1981.

[21] Isabelle Peretz and Robert Zatorre, editors. *The Cognitive Neuroscience of Music.* Oxford University Press, 2003.

[22] Rudolf E. Radocy and J. David Boyle. *Psychological Foundations of Musical Behavior.* Charles C. Thomas Publisher Ltd., fourth edition, 2003.

[23] Emery Schubert. *Measurement and time series analysis of emotion in music.* PhD thesis, University of New South Wales, Sydney, 1999.

[24] Ian Simon, Dan Morris, and Sumit Basu. Mysong: automatic accompaniment generation for vocal melodies. In *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 725–734, New York, NY, USA, 2008. ACM.

[25] Mark Meerum Terwogt and Flora van Grinsven. Musical expression of moodstates. *Psychology of Music*, 19:99–109, 1991.

[26] Barbara Tillman and Emmanuel Bigand. Does formal musical structure affect perception of musical expressiveness? *Psychology of Music*, 24:3–17, 1996.