

# Controlling musical emotionality: an affective computational architecture for influencing musical emotions

---

**Steven R. Livingstone<sup>1,3</sup>, Ralf Mühlberger<sup>1</sup>, Andrew R. Brown<sup>2,3</sup> and Andrew Loch<sup>1,3</sup>**

<sup>1</sup> The University of Queensland, Australia

<sup>2</sup> Queensland University of Technology, Australia

<sup>3</sup> The Australasian CRC for Interaction Design, Australia

srl@itee.uq.edu.au  
a.brown@qut.edu.au  
ralf@itee.uq.edu.au  
loch@itee.uq.edu.au

## Abstract

Emotions are a key part of creative endeavours, and a core problem for computational models of creativity. In this paper we discuss an affective computing architecture for the dynamic modification of music with a view to predictably affecting induced musical emotions. Extending previous work on the modification of perceived emotions in music, our system architecture aims to provide reliable control of both perceived and induced musical emotions: its emotionality. A rule-based system is used to modify a subset of musical features at two processing levels, namely score and performance. The interactive model leverages sensed listener affect by adapting the emotionality of the music modifications in real-time to assist the listener in reaching a desired emotional state.

Keywords: affective computing, emotion induction, emotionality, music emotion, rule-based system

## 1 Introduction

Arguably music's most enduring feature in the Western tradition is the communication of emotion (Juslin 2001). Music has the capacity to engender affect in listeners in powerful and often deeply personal ways. While the form this affect takes is still a matter of debate (Gabrielsson 2001; Kivy 2002; Scherer 2004), the body of scientific evidence in support of this hypothesis is considerable. Music's operation as an affective medium has in turn seen it employed in numerous domains, notably in cinema (Cohen 2001). Since the turn of the 20<sup>th</sup> century, empirical music psychologists have sought to determine what it is within music that can elicit or communicate emotion. With a growing body literature we are now at the point where these relationships between musical features and emotional connotations can be used to modify a musical work's emotionality with consistent, repeatable and quantifiable outcomes.

A computational methodology permits an examination of the musical emotionality precept in greater detail. Modifications to the music experienced by the listener can be expressed at two levels, a change in the work's score and the performative interpretation. Such changes to a musical work's emotionality can be gathered from the audience in

real-time with physiological response capture techniques and other monitoring methods. While post-testing or self-reporting response techniques are also an excellent measure for empirical analysis, the temporal dynamic afforded by certain physiological metrics can be leveraged in the form of an adaptive feedback mechanism.

In this paper we outline an affective computing architecture designed to influence induced musical emotions through modification of a work's score and performance. These modifications are brought about through the application of a set of rules, triggered by monitoring and analysis of the audience's emotional states. A rule engine applies modifications to the structure and performance of musical features, such as tempo, mode, and dynamics that previous research has shown are associated with particular emotions (Livingstone and Brown 2005).

## 2 Musical emotions

### 2.1 Induction, expression and perception of musical emotions

While conceptually musical emotion is often treated as a rather amorphous quantity, the reality is that it can in fact be discussed in a simple and precise terminology when viewed as a transmissible entity in a multi-stage process; or life cycle. The distinction between expression and induction of musical emotions has its roots in the early philosophical differentiation of cognitivism and emotivism

(Kivy 1989). Emotivists contend that music can induce emotion in listeners in response to musical stimuli, while cognitivists adjudge that music can only express or represent emotion. It is our belief that musical emotion exists in both roles—as an expressive representation and as an affective response—along with a third important form, perceived expressive intent, as described in the Music-emotion life cycle of Figure 1.

The significance of such distinctions in Figure 1 has been underscored by music educators, and others, when describing musical epistemology (Swanwick 1988; Brown 2000). This minimalist representation does not detail the processes behind each step nor their potential for overlap or conflation. For example, whether the performer was previously also a listener, how improvisation conjoins many steps, how external factors influence emotion induction (including attitudes, physiological state, see Livingstone et al. 2005a; Livingstone et al. 2005b), or how step 1 emotion came about and was then translated into step 2 emotion. For the discussions in this paper it is important to emphasise that step 4 is self-induced emotion prior to performance, and differs to step 7 which is emotion induced due to the listening process.

The Music-emotion life cycle highlights a distinction between three forms of emotion: induced, expressed and perceived. Delineating musical emotion into these three subsets allows for a more flexible grammar, both for the purposes of discussion and computation.




Figure 1. The Music-emotion life cycle illustrates how the roles of composer, performer and listener equate with different aspects of emotionality in a musical work. Within the life cycle, steps 1, 4, & 7 may or may not operate, and are dependant upon the music and actor instances.

Induced emotion is that felt by, or elicited in, an observer in response to a given stimulus or event (for example, memory). For illustrative purposes consider the following example: a smiling clown may induce happiness in a child. Expressing emotion is the embodiment of emotion through audio, visual, or other forms of stimulus; the clown expresses happiness by smiling. Lastly, perceiving emotion is the act of sensing/detecting expressed emotion in the stimulus (whether or not the emotion was intended or genuine); the child perceives the clown to be expressing or feeling happiness. Importantly, each form of emotion is distinct as the link between each is not always transitive. That is, a happy person may not always express happiness in response to an event, due to for example social norms; conversely, a seemingly happy event may not result in happiness being felt, due to for example a strong prior state of depression. Similarly, an observer may not always correctly perceive the intended emotion being expressed. For example “is the person crying out of joy or sadness?”

With a conceptual model of emotion transmission in place, we can move now to a discussion of a representational technique for computational implementation.

## 2.2 Measurement and representation

For a computational architecture to leverage musical emotions a representation is required that is both consistent and quantitative. One such form is the ‘Two dimensional emotion space’ (Schubert 1999), and is used throughout the system architecture. For the remainder of this discussion we will refer to this representation as 2DES. A critique and justification of 2DES has been discussed previously by the authors, and we refer the reader to Livingstone and Brown (2005) and Livingstone et al. (2005a). In 2DES emotions are mapped onto a two-dimensional plane, with the vertical axis representing arousal (energy level), and the horizontal axis as valence (pleasantness); see Figure 2.

The 2DES is an effective, reductionist representation for describing perceived emotions. As discussed previously, judgement accuracy of perceived emotion has been found to be more robust and accurate than that of induced emotion (Livingstone et al. 2005a). Coupled with this, current physiological response techniques are restricted largely to the capture of arousal, while ignoring valence. For these reasons and others the music-emotion rule engine used in the system is built upon rules and test results designed to influence perceived musical emotions. As discussed though, there exists a clear divide between the perception of emotion in music, ‘that piece sounds happy’, and the inducement of emotion ‘that piece makes me feel happy’. There is growing support however for the hypothesis that the perception of a particular emotion can often translate into a similar induced emotion. In a survey of 141 participants, 74% stated that they always or often felt the same emotion as that which was being expressed by music (Juslin and Laukka 2004). In a similar study of 45 listeners, 70% of participants stated that they felt the same or similar emotion as to that which was being expressed in the music (Evans and Schubert, in press). These results suggest that it would be advantageous to investigate the effectiveness of the rule engine to influence induced emotion.

## 3 A computational music-emotion rule engine

### 3.1 Collation and analysis

For over a century music psychologists have asked the question “what it is within music that allows it to convey emotion?” Many of these investigations involve the creation of rules, with musical experts first identifying a set of musical features and then correlating these with the musical work’s emotionality as described in listener self-reporting surveys. Music-emotion rules can be separated generally into two groups: structural and

performative. Structural rules are those used by the composer and relate to the music score itself, see step 2 in Figure 1. Performative rules are those used by the performer in their interpretive expression of the score, see step 5 of Figure 1, and share many similarities. For a more detailed discussion on this separation, see Gabrielsson (2001) and Livingstone and Brown (2005).

Previous work by the authors collated a series of music-emotion structural rules, aided by recent summative works (Schubert 1999; Gabrielsson and Lindstrom 2001; Gabrielsson and Juslin 2003; Juslin and Laukka 2003). The index consists of music rules grouped into their respective emotional octants, where each octant (quadrant subdivision) is associated with a position on the 2DES plane. For a complete listing of the music-emotion rules see Livingstone and Brown (2005).

In Figure 2 a subset of the structural rules for two of the 2DES octants are highlighted. The numerical value in parenthesis beside each rule type is the number of independent studies that reported that result. For example, 15 studies reported that works described as sad, or similar, had a slow tempo; 7 studies reported that such works were typically written in a minor mode, and so on.

From this investigation a series of recurrent rules were identified which featured in each emotional octant with some prominence (3 or more independent studies reporting the same result). When graphed visually on the 2DES, their importance in the control of perceived musical emotions was immediately evident, see Figure 3. They were subsequently labelled the Primary music-emotion structural rules after René Descartes, who in 1649 first enumerated the set of primary emotions in his work *Les passions de l'ame* (Soloman 2002).

The diagram in Figure 3 shows each of the primary rules as a concentric ring. Rules near the periphery are based on the greatest number of studies, while those near the centre on the fewest number of studies. For example,




Figure 2. The 'Two dimensional emotion space' (2DES), and music emotion structural rules grouped into their corresponding positions on the plane. Reprinted with permission from (Livingstone et al. 2005a).

'happy' music, which exists in quadrant 1, typically has a fast tempo (5<sup>th</sup> ring from centre, Q1, 20 studies), a major mode (4<sup>th</sup> ring from centre, Q1, 19 studies), above average loudness (3<sup>rd</sup> ring from centre, Q1, 10 studies), and so on for each ring that exists in the quadrant for that emotion.

What is most striking about the figure is the reflective symmetry of the rules around the axes. For example, works described as having a positive emotion (happy, excited,




Figure 3. The Primary music-emotion structural rules graphed on the 'Two dimensional emotion space'. Reprinted with permission from Livingstone and Thompson (2006), adapted from Livingstone and Brown (2005).

tender) typically had a major mode and simple harmonies (4<sup>th</sup> ring from centre, Q1 and Q4); conversely, works described as negatively emotive were generally written in a minor mode with complex harmonies (4<sup>th</sup> ring from centre, Q2 and Q3). In this example, we see the use of opposing rules reflected around the vertical axis, which separates positive from negative emotions. This form of symmetry strongly suggests their use as a communication mechanism. We will return to this point shortly.

### 3.2 System prototype: testing the rules

A system prototype operating on a symbolic music representation was constructed to test the effectiveness of the primary music-emotion structural rules on influencing the perceived emotions of music. MIDI encodings of two musical works were selected, one from the Western classical music genre, the other from a popular computer game. For each work participants were first played the original unaltered piece, followed by a modified version in which a set of rules were applied. For example, those which shift the emotion towards quadrant 1 (happier): increasing the tempo, increasing the loudness, transposing to the major mode, raising the pitch, increasing the level of staccato articulation. The participants then marked on a 2DES sheet how they perceived the overall emotion of the work had shifted between the two versions. Results were encouraging for influencing music towards quadrants 1, 3 and 4; loosely described as ‘happier’, ‘sadder’ and more ‘relaxed’. It was later found that methodological issues were largely responsible for lower than expected quadrant 2 results.

Post-testing feedback also revealed that, at times, participants had trouble reporting the emotional changes, stating that the performances of the works were “too mechanical”. Indeed this is a common complaint against MIDI performances generated from symbolic scores. In our testing only the structural

information of the work was maintained, while all the performative information was stripped. Indeed, the separation of musical rules into structural and performative is largely an ontological one, and cedes nothing to the final audio experienced by the listener where the two rule sets are combined. Thus it has been identified that the presence and control of both structural and performative features is required for effective influencing of musical emotions.

The significance of performative aspects has been illustrated recently in a survey investigating the similarities between the expression of emotion in speech and music performance. An analysis of 104 studies on vocal emotion, and 41 on music performance found a strong similarity both in the communication accuracy and the patterns of emotion rules used to express particular emotions (Juslin and Laukka 2003). These patterns, or groupings, are similar to those expressed in Figure 2. In an extension of these results, Livingstone aligned the previously collated structural music-emotion rules with these results, finding a 96% degree of 3-way rule agreement between expression of emotion in speech, music performance and music structure. These findings offer some of the strongest evidence yet that music is, at its heart, a form of emotional communication used in affective engagement (Livingstone and Thompson 2006). Based on these findings, our music modification system has been extended to incorporate a series of performative rules. These additional rules serve two purposes, firstly as a means of ‘humanising’ the work such that it sounds less ‘mechanical’, and secondly to provide additional methods of influencing musical emotions. Research into ‘humanising’ the computer performance of MIDI works has been the aim of the KTH rule system for the last twenty years (Friberg 1995; Friberg 2006). The KTH rule collation and implementation has assisted in the development of performative and humanising rules in this work. A complete list of rules implemented by the system can be found in Table 1.

| Rule                    | Modification               | Type | Description   | Emotional Effect   | Progress |
|-------------------------|----------------------------|------|---|--|----------|
| Tempo                   | Increase, decrease         | S    | A change in the overall played speed of the work  | Arousal shifting   | C        |
| Mode                    | Major, minor               | S    | A major harmonic mode, or minor harmonic mode   | Valence shifting   | C        |
| Loudness                | Increase, decrease         | S    | Overall change in volume of the work  | Arousal and valence shifting   | C        |
| Articulation            | More staccato, more legato | S    | Notes are sharpened through decreasing the ratio of $\left(\frac{\text{duration note 1 onset to note 1 offset}}{\text{duration of note 1 onset to note 2 onset}}\right)$ referred to as $\left(\frac{d_{io}}{d_{ij}}\right)$  | Arousal and valence shifting   | C        |
| Pitch                   | Raise, lower               | S    | Overall raising of the pitch for each note in the work  | Arousal and valence shifting   | C        |
| Harmony                 | Simplify (partial)         | S    | Simplifying the harmony by removing notes in the harmony section using a stochastic algorithm. Rule was removed as it was found to produce an unmusical result.   | Valence shifting   | R        |
| Expressive contour      | Increase, Decrease         | P    | Expressive Contours, the pattern of $\left(\frac{d_{io}}{d_{ij}}\right)$ used in a series of notes, see (Juslin 2001). These patterns are naturally grouped into emotional types ("happy", "sad" etcetera).   | Arousal and valence shifting   | D        |
| Tempo Variation         | Increase, Decrease         | P    | Tempo variations between sections varies in intensity, sharpness of application and frequency.  | Arousal and valence shifting.  | D        |
| Tone Attacks            | Increase, Decrease         | P    | Gradient of Attack in ADSR envelope (Attack Duration Sustain Release)   | Arousal and valence shifting   | D        |
| Stable Note Accent      | Increase, Decrease         | P    | Accent notes considered to be stable in their position in the harmonic scale. These include typically, tonic, dominant and subdominant. Accent achieved using the Note Accent form described below.   | Increases intensity of arousal and valence.  | D        |
| Phrase Arch             | -                          | H    | A highly complex rule that uses quadratic equations applied at various phrase levels to mimic performer acceleration and ritard over music phrases. Included is a rise in dynamic at peak curve. Implementation is a simplification of that discussed in (Friberg 1995). These can be changed real-time to indicate changing emotional intensity. | More 'humanistic'. Mimics the phrase rubato and timings used by performers in phrasing. Can shift arousal and valence intensity, or reflect different musical genre. | C        |
| Pedal Accent            | Increase, Decrease         | P    | Pedal off before higher-level phrases start and finish  | Increases intensity of arousal and valence.  | D        |
| Originality             | -                          | H    | Great performers are known to exaggerate rule applications. Stochastic application of various rules around phrase boundaries.   | More humanistic, can also increase intensity of arousal and valence.   | D        |
| Stochastic Fluctuations | -                          | H    | Random fluctuations on note start time, length, loudness etc help to mimic natural performance variability.   | More humanistic.   | C        |
| Chord Asynchrony        | Increase, Decrease         | P    | Accents notes, often done through delaying onset time of melody notes   | More humanistic, can also increase intensity of arousal and valence.   | D        |
| Melody Accent           | Increase, Decrease         | P    | Accent the melody notes to increase perceptual acuity. Achieved using note accent rule types.   | More humanistic, can also increase intensity of arousal and valence.   | C        |
| Beat Accent             | Increase, Decrease         | P    | Change loudness of notes that occur on major beats, generally in the form of increasing harmony note loudness   | More humanistic, can also increase intensity of arousal and valence.   | D        |
| Note Accent             | Increase, Decrease         | P    | A generalised rule set that can be used for accenting. Play preceding note detached, lengthen played note and closely tied to succeeding note, pedal off prior to playing and pedal down at end, increase dynamic of note, soften dynamic of other chord notes.   | More humanistic, can also increase intensity of arousal and valence.   | D        |
| Slurs                   | Increase, Decrease         | P    | Mimicry of 'down-up' motion of a slur. The intensity of down up decreases as the number of notes within the slur increases.   | More humanistic, can also increase intensity of arousal and valence.   | D        |

Table 1. List of music emotion rules and their implementation progress in the rule engine: complete (C), in development (D), removed (R). Rule types include structural (S), performative (P) and humanising (H). Humanising rules are a superset of performative rules, however only performative rule values can be modified in realtime when a change in emotion is desired; humanising rules are applied pre run-time based on the work's structure. This rule set is a significant enhancement over the previous system iteration.

The history of empirical music analyses has yielded surprising insight into the forms and techniques used by composers and performers in the expression of emotion. The rules detailed in this paper offer a simple and relatively powerful method of influencing musical emotions. The use of these rules is however but one small step in the understanding and influencing of musical emotion. Like Gabrielsson (1999), we believe that the application of these rules is intimately connected with the composer's and performer's intentions for emotional expression of a work. In the present system iteration the majority of rules are applied without any local or global musical context; indeed without such analyses higher-order 'intentions' can never be ascertained. Future examinations could begin with the concept of musical expectations (Meyer 1956), a cornerstone of modern music emotion psychology (Lerdahl 2001; Narmour 1990), and discussed previously by the authors (Livingstone et al. 2005b).

## 4 System architecture

**4.1 Describing an affective architecture**  
Music is an immensely affective, abstract medium. It has the power to express and induce a broad range of emotions in both the listener and performer. With the low cost and simplicity of reproduction technologies, music is now ubiquitous in everyday life. From public transport to concert halls, operating theatres to elevators, it is now unusual for two hours to go by where the average person isn't exposed to music in some way (Sloboda and O'Neill 2001). With such a broad range of applications, the capability to influence how music affects a listener is a tantalising prospect. Through the modification of a musical work's structure and method of performance, such a tool is possible. The previous system iteration had focused on the influencing of perceived musical emotions through a modification of structural features. As we have just outlined

there are a number of new factors that can justifiably be included in the architecture of our computational music system. These include the growing support for a positive relationship between perceived and induced emotion, an enlarged rule set which includes performative aspects, and encouraging results for the partial sensing of induced emotional state.

There are two categories of implementation for our music emotionality system, the architectures of which are displayed in Figure 4. In the autonomous model, the system can operate without affective feedback, where changes to the music are triggered via external system events. The first example of such an application would be as an adaptive music system for a computer game. A second type of situation, known as the interactive model, is where an external agent monitors the actions of a human user, and the analysis of these actions is used as a basis for music modification as in, for example, the use of music for controlling stress or arousal in critical work places or therapeutic sessions. We have discussed this model in previous work (Livingstone et al. 2005a). The interactive model is where a human user's response is monitored, and score and performance changes are made in real-time to compensate for or escalate the emotionality.

**4.2 Automated capture techniques**  
User feedback to the system can be provided in a number of ways, including the monitoring of interaction patterns such as keyboard and game controller usage, or physiological measurement such as skin conductance; a demonstrated measure of music affect response (Rickard 2004). These can provide the system with an indication of the listener's current arousal level, one of the two emotional dimensions in the 2DES representation. This emotional status information can be utilised by the system, allowing it to adapt in real-time to modify the emotionality of the music, to help influence the audience to reach the desired emotional state.

Research is ongoing into ways of eliciting a person's affective state from a diverse range of human characteristics and behaviours. Additional techniques that can be utilised in the music listening environment, and have not been discounted include: facial recognition/pupil dilation, body posture/gesture/movement, and multimodal (combination of these). Similarly with empirical psychology, these methods are restricted largely to the capture of arousal. To overcome this, Wang et al. (2004) combined elements of self-report with automated methods (such as facial recognition), in an attempt to measure valence.

Studies have shown that humans derive high levels of affective information from facial expressions; the accuracy of which can be increased significantly when used in combination with body gestures (Gunes et al. 2003). While it has been shown that the facial expressions of performers influences the perceived affective judgements of the music by listeners (Thompson et al. 2005), there is little work investigating the capture of such expressions made by listeners (Witvliet and Vrana 1996).

To the authors' knowledge, skin conductance is the only successful measure of a listener's physiological response to music. As such it is the primary capture technique employed in the system. In future work we hope to investigate a multimodal technique, which incorporates facial responses with skin conductance.

#### 4.3 Discussion

The system architectures displayed in Figure 4 are quite similar. We will first discuss the interactive model. Prior to each performance a suitable baseline data sample is collected from the selected user monitoring system (Rickard 2004), notated as 'Baseline (once)' in the audience affect component. Once the music performance begins, listener response metrics in the form of arousal events are passed to the emotion agent. For each musical work there

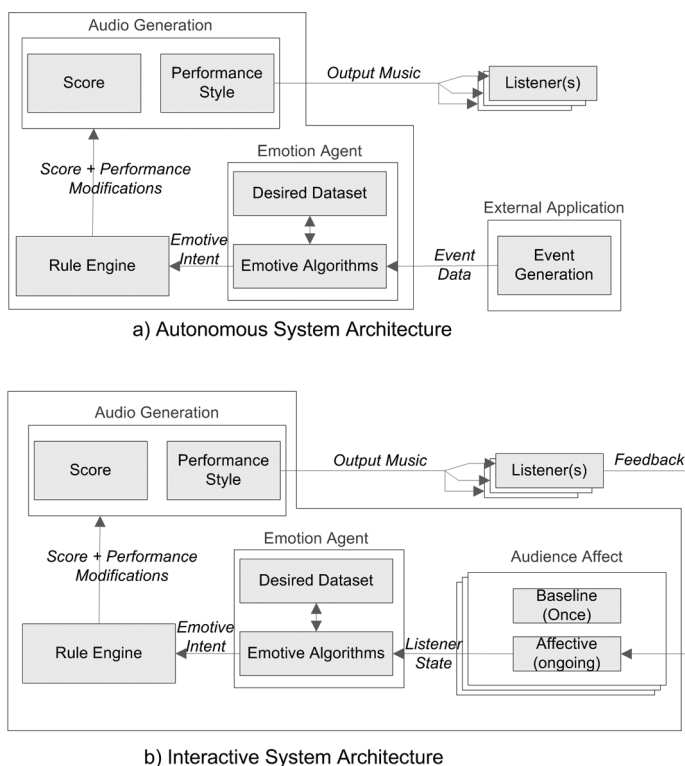


Figure 4. The architecture for an autonomous and interactive music emotionality system.

is a pre-constructed dataset of arousal-valence values that describes the desired audience affect responses in time to the work's performance. This desired dataset is stored in the emotion agent component, and is used in concert with the emotive algorithms to determine the emotional changes to be made to the music.

While the listener's state feedback is described by the single dimension, arousal, the emotion agent passes intent in the form of a two-dimensional data point, arousal-valence, as required by the rule engine. The emotive algorithms extrapolate the necessary change in valence based on that which occurred in arousal. From the 'Primary music-emotion structural rules' graph shown in Figure 3, it can be theorised that the two dimensions are interconnected for music emotion rules. For example, an increase in tempo would



presumably have an effect both on arousal and valence. This degree of interplay on induced emotion is one such research outcome to be investigated by later iterations of this project.

Upon receiving the desired and detected arousal-valence data pairs, the rule engine then calculates the necessary changes to be made to the music's structure and performance to bring about the desired emotional state. For example, the measured arousal-valence values may be in the range corresponding to 'moderately happy', but the desired state may be in the mid-upper range, which is closer to 'very happy'. In this case an increase in the structural rules: tempo, loudness, pitch, and articulation, and the performative rule: expressive contour, could be employed. It is the task of the audio generation unit to apply these modifications, with the changes made directly to the MIDI source to be played by the computer.

The dynamic MIDI scheduling, emotion agent, rule engine and audio generation unit have all been written in *Impromptu*, a dynamic programming music environment for Mac OS X that employs the *Scheme* programming language (Sorensen 2005). In the first case model, autonomous, input events can be in the form of either arousal-valence, or simply arousal. As with the interactive architecture, the task of the emotive algorithms is the same. The desired dataset component is maintained in the case where a pre-sought path of listener affect is required, such as in the previously described therapeutic or work-place control situations. This model allows for an external sensing agent to be employed, which could leverage detection techniques beyond those described here.

## 5 Conclusion

In this paper we have presented a computational architecture that enables the emotionality of music to respond in real-time to human affect feedback and external programmatic

events. This system enables the real-time modification of music according to rules for use in adaptive circumstances including games and entertainment, musical therapies, and computer music performance. For this end, we have developed a list of salient performance rules to accompany our previously published list of structural rules; the importance and application of which has been discussed in Table 1.

The basis for a theory of music emotionality that is inclusive of expressed, induced and perceived emotions has also been presented. The relationships between these different aspects of emotionality are expressed in the Music-emotion life cycle diagram, which plays a key role in the modelling of creative behaviour. This, combined with our own previous work on structural music emotion rules for controlling adaptive music systems, has made it evident that both performance and structural rules are required for successfully influencing the emotionality of music in a computational system.

## Acknowledgements

The work reported in this paper has been funded in part by the Cooperative Research Centre Programme through the Australian Government's Department of Education, Science and Training.

## References

- Brown, A. R. (2000) 'Modes of compositional engagement', *The Australasian Computer Music Conference*, Brisbane, QLD, pp. 8–17.
- Cohen, A. J. (2001) 'Music as a source of emotion in film' in *Music and emotion: theory and research*, eds P. Juslin and J. Sloboda, Oxford University Press, Oxford, pp. 249–274.

- Evans, P. and Schubert, E. (in press) 'Quantification of Gabriellsson's relationships between felt and expressed emotions in music', *9th International Conference on Music Perception and Cognition*, Bologna, Italy.
- Friberg, A. (1995) 'Matching the rule parameters of phrase arch to performances of Träumerei: a preliminary study', *KTH Symposium on Grammars for Music Performance*, Stockholm, Sweden, pp. 37–44.
- Friberg, A. (2006) 'pDM: an expressive sequencer with real-time control of the KTH music-performance rules', *Computer Music Journal*, vol. 30, no. 1, pp. 37–48.
- Gabriellsson, A. (1999) 'The performance of music' in *The psychology of music*, 2nd edn, ed. D. Deutsch, Academic Press, San Diego, pp. 501–602.
- Gabriellsson, A. (2001) 'Emotions in strong experiences with music' in *Music and emotion: theory and research*, eds P. Juslin and J. Sloboda, Oxford University Press, Oxford, pp. 431–449.
- Gabriellsson, A. and Juslin, P. N. (2003) *Emotional expression in music*, Series in Affective Sciences, Oxford University Press, New York, pp. 503–534.
- Gabriellsson, A. and Lindstrom, E. (2001) 'The influence of musical structure on emotional expression' in *Music and emotion: theory and research*, eds P. Juslin and J. Sloboda, Oxford University Press, Oxford, pp. 223–248.
- Gunes, H., Piccardi, M. and Jan, T. (2003) 'Face and body gesture recognition for a vision-based multimodal analyzer', *Proceedings of Workshop on Visual Information Processing*, Sydney, NSW, pp. 19–28.
- Juslin, P. N. (2001) 'Communicating emotion in music performance: a review and a theoretical framework' in *Music and emotion: theory and research*, eds P. Juslin and J. Sloboda, Oxford University Press, Oxford, pp. 309–340.
- Juslin, P. N. and Laukka, P. (2003) 'Communication of emotions in vocal expression and music performance: different channels, same code?', *Psychological Bulletin*, vol. 129, no. 5, pp. 770–814.
- Juslin, P. N. and Laukka, P. (2004) 'Expression, perception, and induction of musical emotions: a review and a questionnaire study of everyday listening', *Journal of New Music Research*, vol. 33, no. 3, pp. 217–238.
- Kivy, P. (1989) *Sound sentiment: an essay on the musical emotions*, Temple University Press, Philadelphia, PA.
- Kivy, P. (2002) *Introduction to a philosophy of music*, Oxford University Press, USA.
- Lerdahl, F. (2001). *Tonal pitch space*, Oxford University Press, New York.
- Livingstone, S. R. and Brown, A. R. (2005) 'Dynamic response: real-time adaptation for music emotion', *Australasian Conference on Interactive Entertainment*, Sydney, NSW.
- Livingstone, S. R., Brown, A. R. and Mühlberger, R. (2005a) 'Influencing the perceived emotions of music with intent', *Third International Conference on Generative Systems*, Melbourne, VIC.
- Livingstone, S. R., Mühlberger, R. and Brown, A. R. (2005b) 'Playing with affect: music performance with awareness of score and audience', *Australasian Computer Music Conference*, Brisbane, QLD.
- Livingstone, S. R. and Thompson, W. F. (2006) 'Multi-modal affective interaction: a comment on musical origins', *Music Perception*, vol. 21, no. 1, pp. 89–94.
- Meyer, L. B. (1956) *Emotion and meaning in music*, The University of Chicago Press, Chicago.
- Narmour, E. (1990) *The analysis and cognition of basic melodic structures*, The University of Chicago Press, Chicago.
- Rickard, N. S. (2004) 'Intense emotional responses to music: a test of the physiological arousal hypothesis', *Psychology of Music*, vol. 32, no. 4, pp. 371–388.
- Scherer, K. R. (2004) 'Which emotions can be induced by music? what are the underlying mechanisms? and how can we measure them?', *Journal of New Music Research*, vol. 33, no. 3, pp. 239–251.
- Schubert, E. (1999) 'Measurement and time series analysis of emotion in music', PhD, University of New South Wales.
- Sloboda, J. A. and O'Neill, S. A. (2001) 'Emotions in everyday listening to music' in *Music and emotion: theory and research*, eds P. Juslin and J. Sloboda, Oxford University Press, Oxford, pp. 415–429.

- Soloman, R. C. (2002) 'Back to basics: on the very idea of "basic emotions"', *Journal for the Theory of Social Behaviour*, vol. 32, no. 2, pp. 115–144.
- Sorensen, A. (2005) 'Impromptu: an interactive programming environment for composition and performance', *Australasian Computer Music Conference*, Brisbane, QLD.
- Swanwick, K. (1988) *Music, mind and education*, Routledge, London.
- Thompson, W. F., Graham, P. and Russo, F. A. (2005) 'Seeing music performance: visual influences on perception and experience', *Semiotica*, vol. 156, no. 1, pp. 177–201.
- Wang, H., Prendinger, H. and Igarashi, T. (2004) 'Communicating emotions in online chat using physiological sensors and animated text', *Proceedings of CHI 2004*, ACM Press, pp. 1171–1174.
- Witvliet, C.V. and Vrana, S.R. (1996) 'The emotional impact of instrumental music on affect ratings, facial EMG, autonomic response, and the startle reflex: effects of valence and arousal', *Psychophysiology*, vol. 33 (Suppl. 1), no. 91.

**Steven R. Livingstone** is a doctoral student at The University of Queensland, and Australasian CRC for Interaction Design. He received his B.Sc. in physics and B.Inf.Tech. (hon.) from The University of Queensland. His thesis explores the use of computers for influencing the perceived and induced emotions of music. He has published in affective computing, music psychology, evolutionary origins of music, and emotion psychology. As a Research Assistant he examined new media collaboration, content management systems and DRM. Steven is also an accomplished classical pianist and composer.

**Andrew R. Brown** is an Associate Professor at the Queensland University of Technology where he teaches music and sound and is the Digital Media Program Manager for the Australasian CRC for Interaction Design (ACID). Dr. Brown's expertise is in technologies that support creativity, computational music and art, and the philosophy of technology. His current research focuses on the aesthetics of process and adaptive music for interactive entertainment. He is an active computer musician and a builder of software tools for dynamic content creation.

**Ralf Mühlberger** teaches computer science at The University of Queensland. Dr. Mühlberger's background is in distributed information management using integrating technologies such as workflow management systems and peer-to-peer computing, particularly from a modelling perspective. A member of the Interaction Design Research Division, he is now looking at engagement models, affective computing, and various study and design methods from the human-computer interaction (HCI) and computer support for collaborative work (CSCW) fields.

**Andrew Loch** is a doctoral student at The University of Queensland, and Australasian CRC for Interaction Design (ACID). His doctoral thesis, entitled 'Responsive media spaces for affective computing environments' is in the area of affective computing and interaction design.