Music Organisation Using Colour Synaesthesia

Michael Voong
School of Computer Science
University of Birmingham
Birmingham, UK B15 2TT

m.voong@cs.bham.ac.uk

Russell Beale
School of Computer Science
University of Birmingham
Birmingham, UK B15 2TT

r.beale@cs.bham.ac.uk

CHI 2007, April 28 – May 3, 2007, San Jose, USA
ACM 1-xxxxxxxxxxxxxxxxxx.

Abstract
The movement of music from physical discs to digital resources managed on a computer has had an effect on the listening habits of users. We explore using the potential of the innate synaesthesia that some people report feeling between colour and mood in a novel interface that enables a user to explore their music collection and create musical playlists in a more relevant way.

We show that there is a reasonable degree of consistency between users' associations of colour and music, and show that an indirect descriptor can aid in the recall of music via mood, making playlist generation a simpler and more useful process.

Keywords
Interaction, music, playlists, SOM, visualisation

ACM Classification Keywords
H.5.2 User Interfaces, H.1.2 User/Machine Systems, H.5.5 Sound and Music Computing, H.3.3 Information Search and Retrieval, I.2.1 Applications and Expert Systems

Introduction

The development of P2P (peer-to-peer) and the rising popularity of online music stores such as iTunes have led to an increase in the amount of music that people have on their computers and portable music players. People are now exploring more and more genres of music [6], made easier with the advent of new music recommendation systems such as Last.fm [1] and Pandora [2]. These systems are often described as being contributors to the new social music revolution ([8], [15]); where people contribute their listening preferences. The systems gives music recommendations using this data, with the assumption that people would like to listen to music similar to the ones they listen to already.

Our studies [5] have suggested that changes in modern listening habits have not been reflected on in modern music player UIs, and that there is a need to adapt the behaviour of playlist generation to a user's mood. The studies showed that many choose to create playlists grouped by mood, for example, by naming a playlist "relaxing". We found that users tend to switch between these playlists as their mood changes. The rationale for our system design was that this categorisation task is cumbersome. Many audio processing systems have been developed to attempt to detect mood in an audio track [12] but fail due to the subjective nature of music. We decided to use colour to provide the expressivity and simplicity required. Colour is often associated with emotion and moods, supported by existing psychological theories ([7], [3], [11]) and it provides a natural and simple way of expressing something about a song [9]. Colour psychology as a field is still relatively immature, and it is not surprising to find that different cultures can have radically different interpretations of colours [3].

Initial User Study

An initial user study was performed by a survey aiming to find out user habits in terms of how they used their music players and whether using colour as a music descriptor could work. We hoped to find evidence to support whether users' music associations correlated with each other.

An online questionnaire was created, and 91 people responded. We kept the system anonymous, so have no demographic or social information about the respondents, but owing to the nature of the distribution of the URL, we expect most to be in the range 18-35, with their own Internet access. Administering a questionnaire in this form allowed us to stream music clips to the respondent. In one question, the respondent would listen to five clips and indicate which colour, out of a selection of ten, was triggered in their mind. One visualisation was made for each of the music clips, showing colour responses of the respondents. An example is shown in figure 1.

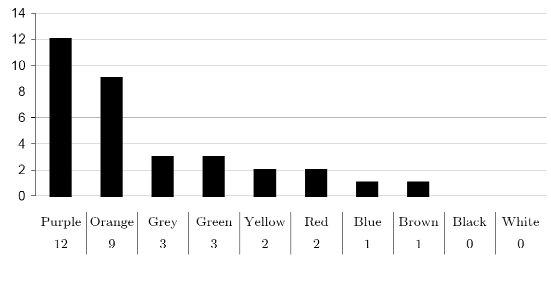

figure 1: Strong responses to two colours are a typical response when hearing a music track.

The Colour Player UI

We show a fully working prototype of a music player (figure 2), based on the iTunes interface, which includes two additional interactive modes to the traditional ones ("artist/album/search"). Nielsen argued that user interface consistency is important when you want to increase the usability of a system [13]. Basing the interface on an existing popular application used by millions ensured that it was easier to learn, increasing the chances of a wider participation in the study. As there was a completely new way to interact with a user's music collection, making the choice to augment existing, well-known interaction methods was an obvious one.

The following is an overview of the features supporting the new system of tagging tracks:

Assigning Colours to Tracks

Colour associations are assigned by dragging the mouse over the colour picker, situated on the left of the UI.

The Colour Map

After associating colours to tracks, the colour visualisation updates dynamically to reflect the changes. The map represents the user's entire collection of colour-tagged music.

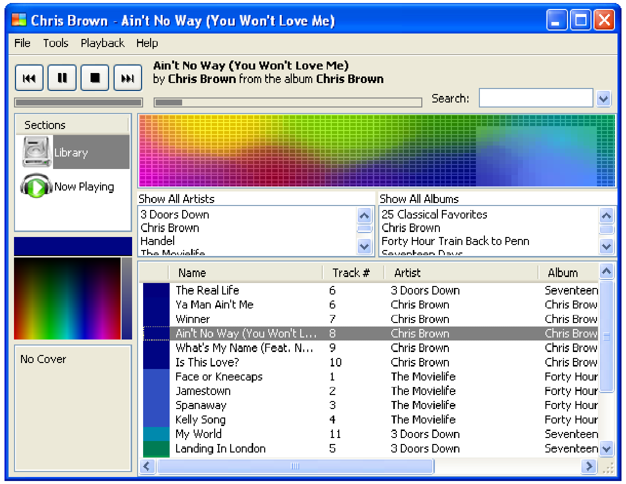

figure 2: The user interface

Interaction with the Colour Map

There are two modes of interaction on the colour map:

1. Marquee mode: The user drags a box around the areas of the colour they are targeting. The track listing updates automatically to show only tracks that match the colours selected.
2. Arrow mode: This was developed so that playlists can be generated to indicate a dynamic change of mood. Movement through the colour map is indicated with a line, and reflected on dynamically by the change in songs put into the ongoing playlist as the tracks are played.

Updating the Colour Map

The program utilises a self-organising map (SOM) [10], also known as a Kohonen map, to cluster and organise the music in the 2D representation.

The SOM is a clustering algorithm that takes in vectors of numbers as input. As the algorithm runs, the input vectors that are similar to each other (using the Euclidean measure) are represented as being distributed near to each other on the output map, which is arranged as a grid of a set size. It is called an unsupervised learning method as it organises the output solely on the input. So, in our implementation, the input vectors are a set of RGB (red, green, blue) values between 0 and 255. The output is a grid of vectors that represent the input space – the set of colour associations tagged to tracks by the user.

The SOM has a number of useful properties [4]. It is able to provide a representation of an arbitrary number of songs into a relatively compact 2D space, and it clusters related colours, and hence, related songs near to each other. Due to the nature of the algorithm, colours that are very different tend to move away from each other – and hence songs that are very different in mood get positioned far away from each other.

The SOM is also sensitive to the frequency of classification, in that a relatively larger area is used to represent an area of the input space that has more data points in it. In practical terms for us, this means that a larger 2D area in the representation is used to represent a smaller range of colours, so that if a large number of songs are clustered around a common colour they are spread out across the map. This provides a greater degree of disambiguation, using more space to represent a high density of data points and compacting areas in which very few songs are represented.

To ensure that the user experience is as smooth as possible, the SOM is retrained every time a new colour association has been made, dynamically updating the Colour Map. Due to the simplicity of the SOM algorithm, these changes are processed without any pause, allowing the system to remain responsive.

Gathering Data for Evaluation

The software was developed using an iterative, user-centered design process [14]; there was constant interaction with users through questionnaires and comment forms on the website, and they were encouraged to send ideas, criticisms and feedback. The program was constantly updated to reflect new points made in feedback.

The software included a feature that showed a dialog box to the user when they first started it, streaming the same music clips to them that were used in the original online questionnaire. By doing this, not only could users quickly get an idea of how the colour association system worked, we could also build up a database of associations made for a given artist and title pair, and investigate the constancy of the synaesthetic response. We are still collecting data on this and it is too premature to report here.

User Evaluation

The first user evaluation involved users recording their evaluation of the system on a Likert scale (1–5 equivalent to *Very Difficult - Very Easy or Strongly Disagree - Strongly Agree* as appropriate). Most users found it easy to create playlists in their current application, but found that the Colour Map speeded this up. Seven of these nine would use this approach as an alternative to their existing one. They found it neither hard not easy (mean = 3.11, standard deviation = 1.05) to associate colours with music, supporting the earlier findings that users did not expect it to be easy though these experiments suggest that in practice they seem to do better than they expect and find it useful.

A Second User Evaluation

Six subjects were studied in order to find quantitatively how much faster users find creating playlists using colour compared to traditional methods in iTunes. The main parts of the test were as follows:

1. Ask the test subject to think of three different moods or states of mind they might be in during the day that might affect their choice of music.
2. Time the process of creating three playlists in iTunes—one for each mood, containing three tracks each.
3. Ask the test subject to assign a different colour in their mind to each of the three moods they selected
4. Record the amount of time it takes the test subject to assign each colour to three tracks each using Colour Player.
5. Record how long it takes for the test subject to create a playlist mixing two moods together: a) in iTunes, and b) in Colour Player.
6. Request that the test subject tries to recall what colour the nine tracks were associated with. The tracks are played back if requested.

The most similar method in both programs for achieving the goal of organising a list of tracks by some mood descriptor was used as a measurement of performance. The colour recall test was designed to find out whether people could easily recall the colour associations they made with music.

The results are summarised in Table 1. The first conclusion we could draw from this test is that Colour Player is an easy to use music player. It only took the users 16% longer to add in the colour information, despite little exposure to the player, and we would expect that time to reduce as they became more familiar with it. The second conclusion that can be drawn is that users can create playlists of tracks combining moods much quicker in Colour Player than in iTunes. Users quickly realised the best way of doing this.

			Merging Playlists (s)	
Subject	Build iTunes Playlist (s)	Assigning Tracks (s)	iTunes	Colour Player
1	121	91	16	1
2	218	268	36	12
3	124	184	86	4
4	120	238	34	5
5	171	141	19	4
6	94	80	18	1
Mean	141	167	35	5
Std.	45.09	76.88	26	4

Table 1: Results of users performing timed tasks, comparing iTunes and Colour Player.

The colour recall test (test 7) showed significant results: 8 of the 9 test subjects recalled all of their colours correctly. Several of the subjects said that hearing the music aided in their colour recall, but said it was not essential. On asking their opinion on whether they would remember the association if they waited a day, every subject replied "yes".

Conclusions

We have described the design process of creating a music player with a new interaction element and the subsequent evaluation of the system with users. The primary focus of these experiments was to find out the nature of users' correlations of music-mood synaesthesia, and whether the new interface made it easier to create playlists. The results were positive - users liked the new interface, and found it a useful way to organise their collections of music, tagging them by mood by associating mood with colour. We found that users reason about their colour associations with a high degree of accuracy, showing that the tagging system aids recall of music. Using the Kohonen map provides an effective way of visualising large quantities of data in a compact space in a comprehensible manner.

References

[1] *Last.fm*. [cited 2006 February]; Available from: http://www.last.fm/.

[2] *Pandora Internet Radio*. [cited; Available from: http://www.pandora.com/.

[3] *Colour Wheel Pro*. 2005 [cited 2005; Available from: http://www.color-wheel-pro.com/color-meaning.html.

[4] Beale, R. and T. Jackson, *Neural computing: an introduction*. 1990: IOP.

[5] Beale, R. and M. Voong. *Managing online music: attitudes, playlists, mood and colour*. in *HCI 2006*. 2006.

[6] Brown, B.A.T., E. Geelhoed, and A.J. Sellen. *The Use of Conventional and New Music Media: Implications for Future Technologies*. in *Interact '2001*. 2001. Tokyo, Japan.

[7] Campbell, I.G., *Basal Emotional Patterns Expressible in Music.* The American Journal of Psychology, 1942.

[8] Filter^, T. *The Social Music Revolution?* 2006 [cited 2006 September]; Available from: http://thefilter.blogs.com/thefilter/2006/07/the_social_musi.html.

[9] Kawanobe, M., M. Kameda, and M. Miyahara, *Corresponding affect between music and color*, in *Paper presented at the Systems, Man and Cybernetics, IEEE International Conference*. 2003.

[10] Kohonen, T. *The self-organizing map*. in *Proceedings of the IEEE*. 1990.

[11] Levkowitz, H., *Color Theory and Modeling for Computer Graphics, Visualization, and Multimedia Applications*. 1997: Kluwer.

[12] Li, T. and M. Ogihara. *Detecting emotion in music*. in *The International Symposium on Music Information Retrieval*. 2003.

[13] Nielsen, J., *Coordinating user interfaces for consistency.* SIGCHI Bull., 1989. **20**(3): p. 63-65.

[14] Norman, D. and S. Draper, *User Centered System Design; New Perspectives on Human-Computer Interaction*. 1986: Lawrence.

[15] Thermo[SAT]. *The Social Music Revolution...* 2006 [cited 2006 September]; Available from: http://www.thermosat.qc.ca/index.php/2006/03/13/the-social-music-revolution/.