

Visualising Mutual Engagement

Nick Bryan-Kinns
Queen Mary University of London
Interactional Sound and Music, Centre for Digital Music
School of Electronic Engineering and Computer Science
London, E1 4NS. UK.
nickbk@eecs.qmul.ac.uk

Jennifer G. Sheridan
Togeva Ltd
London
E2 7BB
Jennifer@togeva.com

Understanding what goes on in collective creativity is difficult. This paper presents work on visualising mutual engagement in musical micro-creativity and shows how musical memes can be visualised and tracked through extensive logs of creative interaction.

Mutual Engagement, Musical Memes, Micro-Creativity, Collective Creativity, Exploratory Analysis.

1. INTRODUCTION

Being creative is a key part of a healthy lifestyle, contributes to our wellbeing, and is an inherently social activity (Csikszentmihalyi, 1991). There are many definitions and understandings of creativity (see Sternberg, 1991, and Boden, 1994 for an overview) ranging from cognitive and computational models (e.g. Boden, 2003; Wiggins, 2006) which help us to understand the mechanics of human creativity through to practical approaches to increasing creativity (cf. De Bono, 1992). Similarly, there are many studies of creative practice (e.g. Candy and Edmonds, 2002), creative processes (e.g. Shneiderman, 2000), and collaborative creativity (e.g. Paulus and Nijstad, 2003) which help us to understand the social and inter-personal nature of creativity.

In this paper we are particularly interested in *collective creativity* - creative acts by non-experts which rely on social interaction as a key form of *everyday creativity* (cf. Bryan-Kinns, 2009). Furthermore, our view of everyday creativity is that it involves amateurs in the domain, as opposed to workplace creativity involving experts such as designers, artists, musicians, logicians, and so on. We draw on the body of creativity research to characterize collective creativity as involving acts which:

- Are autotelic (an end in itself; Csikszentmihalyi, 1991)
- Involve innovation of ideas with social value (cf. Boden, 2003),
- Involve social dissemination (cf. Shneiderman, 2000).

Examples of collective creativity include: brainstorming, playground games (cf. Sheridan, 2011), creating online photo collections with friends, family bake-offs, knitting circles, flash-mobs, and so on. Identifying and creating new ways to support everyday creativity is key to the development of a strong and dynamic innovation culture and economy.

The emergence of new forms of digital social infrastructure including social networks such as Facebook (cf. Lampe et al. 2008) blogging tools (cf. Nardi et al., 2004) and micro-blogging such as Twitter (cf. Zhao and Rosson, 2009) clearly illustrate the populist potential of digital technologies to provide increased opportunities for collective creativity where the emphasis is on the enjoyment of being creative together rather than focusing on completing work-oriented tasks. We refer to this form of creativity as *micro-creativity* which have the following features:

- Digitally-mediated creative activities
- Carried out as a number of fleeting collaborative interactions
- Happen over an extended period of time
- Take place in a wide range of interaction contexts from desktop computers to mobile phones.

For example, people use micro-blogging tools to play word games over periods of weeks, or to engage in collective (micro) drawing by uploading and iteratively editing small shared sketches, or even to programmatically create music (Graham-Rowe, 2009). However, current approaches to micro-creativity have predominantly focused on textual interaction between people and have not been driven any principled understanding of

creativity or human centred design, typically emerging ad-hoc from specialist interest groups. Moreover, we have little understanding of how micro-creativity emerges and is sustained over time. This lack of principled design and evaluation hinders the growth and development of more radical creativity support, relying instead on whimsical and incremental development of the genre

Music is collective creativity without words; it requires skill to produce beautiful music, yet anyone can be emotionally touched by music. It is both inclusive and exclusive, and acts as a cultural memory. Digital technologies now provide everyday folk with the ability to make and share music wherever they may be, and provide unique opportunities to explore new forms of music making as illustrated by developments in the field of New Interfaces for Musical Expression (NIME; Poupyrev et al., 2001). For example, Ocarina (Wang, 2009), and Daisyphone (Bryan-Kinns, 2004) are mobile phone Apps for social music making. Ocarina supports micro-creativity by allowing people to create short pieces of music using a simple four key interface. These musical contributions are then shared with a global community of users who can tag and rate them. However, picking up and modifying others' contributions is more problematic and requires skill in listening and playing. In contrast, Daisyphone allows direct co-editing of short loops of music, but little support is provided for social interaction beyond the music.

2. MUTUAL ENGAGEMENT

Mutual engagement occurs when people creatively spark together (Bryan-Kinns & Hamilton, 2009). It is imprinted in the moment-by-moment contributions of people when they are engaged in micro-creativity. Music is an artistic activity which relies heavily on mutual engagement - the riffing and jamming on musical ideas generates new forms of creative expression which convey emotion without words or pictures. However, we do not have effective techniques for identifying the birth, development, sustenance, and propagation of musical ideas (or memes, cf. Dawkins, 1976) in mutually engaging micro-creativity, making it difficult to effectively design and evaluate engaging user interfaces.

The nature of human auditory perception allows us to easily identify repeating patterns in music (Bregman, 1994). However, the temporal nature of music also means that manually identifying, tagging and tracking musical memes is extremely time consuming in large data sets. Furthermore, simply listening to hundreds of hours of musical interaction to spot repeated patterns of interaction is simply

impractical. This work shows how extensive logs of musical micro-creativity can be visualised to explore micro-creativity in the time domain, visually foregrounding musical memes across users, time, and spatial proximity. This provides a representation which allows researchers, musicians, and other interested parties to identify and explore how musical memes emerge and propagate in mutually engaging experiences. In this paper eight years of logs (approximately 10m points of interaction) of micro-creativity through the online music making tools *Daisyphone* and *Daisyfield* (Bryan-Kinns, 2011) are analysed for repeated and manipulated patterns of musical memes - points at which mutual engagement occurs between people.

3. DAISYPHONE AND DAISYFIELD

Daisyphone and Daisyfield (referred to collectively as *Daisy**) allow co-editing of short loops of music by co-located and online groups of users through web and iOS interfaces in focused sessions or over extended periods of time. Daisyphone was launched in October 2003 (Bryan-Kinns, 2004), and at launch it received between 4 and 18 players per day from all over the world. Logs of interaction have been collected since its launch, and there are now 160Mb of log files, or approximately ten million individual interactions.

Daisyphone provides a shared loop of music (5 seconds; 48 beats) which can be edited by up to 10 people at once via a network connection. Each participant can create notes using four different voices, can edit any notes, and a shared graphical annotation is provided around and within the musical loop, or Daisy to allow for some social communication beyond the music. Figure 1 illustrates Daisyphone in use on an Apple iPhone - the score is represented by the circle of dots, and



Figure 1: Daisyphone

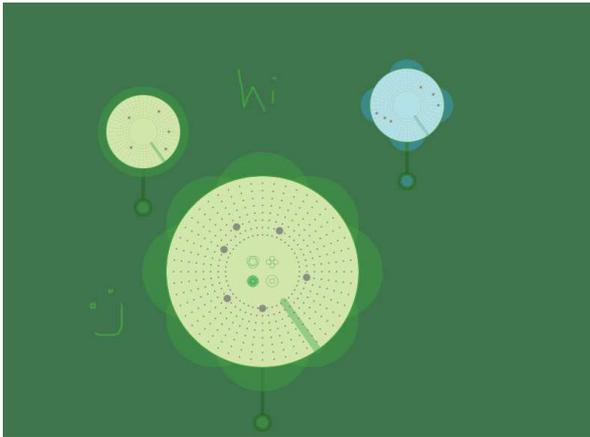


Figure 2: Daisyfield

the currently played set of notes is indicated by the grey line radiating from the centre which rotates clockwise over the period of 5 seconds. Shapes represent different kinds of sound (ambient electronic sound palette in C major scale consisting of bass, lead, wash, and percussion), and the shapes in the centre allow participants to select which sound type and volume they create notes with. Colours are uniquely assigned to users to provide a simple sense of identity. In 2009 Daisyfield was developed which allows multiple loops (i.e. Daisys) to be played concurrently, and for participants to arrange Daisys in a two dimensional space which can be graphically annotated. Again, a single shared score of 48 beats is created from the individual Daisys placed in the space, and each participant hears the same audio output. Figure 2 illustrates the Daisyfield interface with three Daisys shown, the larger one is opened for editing. The aim of this interface is to provide a richer musical and communicative user interface for exploring mutual engagement in micro-creativity.

4. VISUALISING MUTUAL ENGAGEMENT

In order to track the birth and growth of musical memes in micro-creativity in Daisy* we need to be able to algorithmically identify musical phrases. Dannenberg and Hu (2002) provide an overview of approaches to patterns matching for discovering musical structure and repeating patterns. Similarly, pattern matching algorithms such as edit distance matching (cf. Gusfield, 1997) and melody matching algorithms such as Uitenbogerd and Zobel (1999) can be used to identify musical memes as they evolve over time. Following these approaches we developed a tool to analyse logs from Daisy* and identify musical memes in Daisy* as having:

- At least three notes (not pauses), and
- No more than two pauses between notes.

From this definition of musical memes we can collapse the data across the following dimensions to provide useful visualisations of the mutually engaging interaction as discussed in the rest of this paper:

- **Time:** When the meme was made and also when it was modified. Varying the granularity of the time dimension allow memes to be tracked as they develop.
- **User:** Who started the meme, and who contributed to its on going creation. More co-editing indicates greater mutual engagement between participants.
- **Spatial location:** Where the meme is located in the user interface - closer memes indicate more mutual engagement (Bryan-Kinns, 2012).
- **Musical range** of the meme: whether the meme is monophonic or polyphonic, and whether there is a wide dynamic of notes in the meme.
- **Density of musical** meme: what percentage of musical pauses are used in the meme. This allows us to track musical style in the music.

4.1 Development of Memes

Removing the spatial dimension, and laying musical memes out sequentially allows us to visualise the development of musical memes over time. Figure 3a illustrates a musical meme identified by the our visualisation tool consisting of four notes in a descending sequence, whereas figure 3b shows four notes of the same pitch in a sequence. The colours used in the figures represent different users, and time for the music goes from left to right.

Figure 4 illustrates the development of a musical meme by one user as identified by our tool. The meme starts from a sequence of 2 pairs of notes in the leftmost box, after 4 iterations, the meme has become an interesting musical meme composed of



Figure 3a: Descending notes

3b: Same notes



Figure 4: Development of a meme

a 3 descending notes followed by 3 rising notes. Figure 5 illustrates the memes identified by our visualisation tool in 15 minutes of interaction between 4 users of Daisyfield. In the figure several memes are seen emerging (time is laid out from left to right, top to bottom), for example the memes illustrated in Figures 3a, 3b, and 4 are clearly seen in Figure 5 as they develop and are repeated. Figure 5 also shows the persistence, or popularity (as participants can delete notes they are unhappy with), of certain memes such as the descending meme illustrated in Figure 3a which is repeated throughout most of the 15 minute interaction.

Figure 5 also illustrates points at which participants

compliment each others' contributions and co-create musical memes together. For example, on the fifth row down the green and red participant co-create a musically harmonious meme together as summarised in Figure 6. This shows high levels of mutual engagement as the participants feed off each others' contributions. Other examples are shown in the second and third rows to the bottom of figure 5.



Figure 6: Complimenting contributions

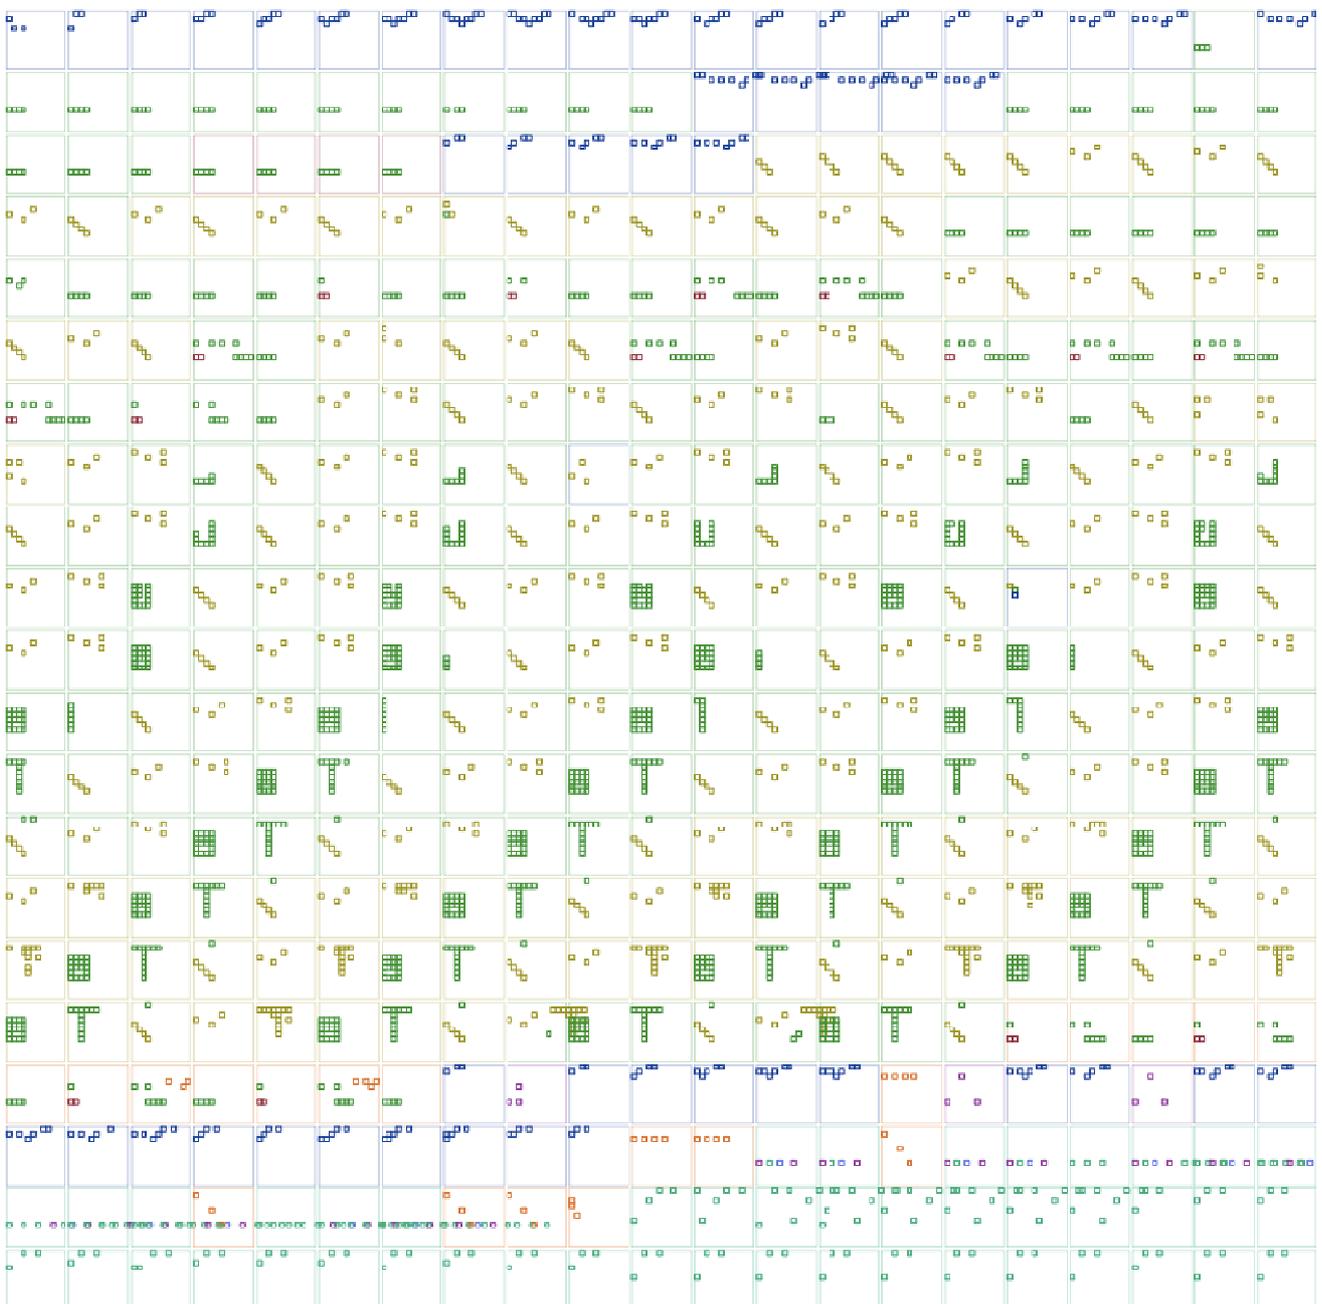


Figure 5: 15 minutes of interaction analysed as memes

4.2 Collapsing Time and Space

The linear layout of memes illustrated in Figure 5 allows us to inspect the detail of how memes develop and propagate over short amounts of time. The visualisation tool also provides a visualisation which foregrounds the spatial location of contributions, allows variable temporal depth, and lays out memes in terms of their musical range and density as illustrated in this section.

Figure 7 shows 45 minutes of memes laid out spatially in terms of where they were created on the user interface (x and y location map to x and y axes on Figure 7), and coloured to represent which user

created them. From this visualisation it is clear that the blue participant mostly interacted by themselves in the centre of the screen whereas the green and orange participants predominantly interacted together in the bottom left quarter of the screen. The layering of the memes gives a sense of the areas of most activity on the screen.

Figure 8 illustrates how development over time can be shown by - foregrounding the manipulation of the memes whilst hiding the details of the location and actual editing. In Figure 8 time is represented on the x axis with coloured stripes indicating activity. The interactive nature of the visualisation tool allows the analyst to explore visualisations in

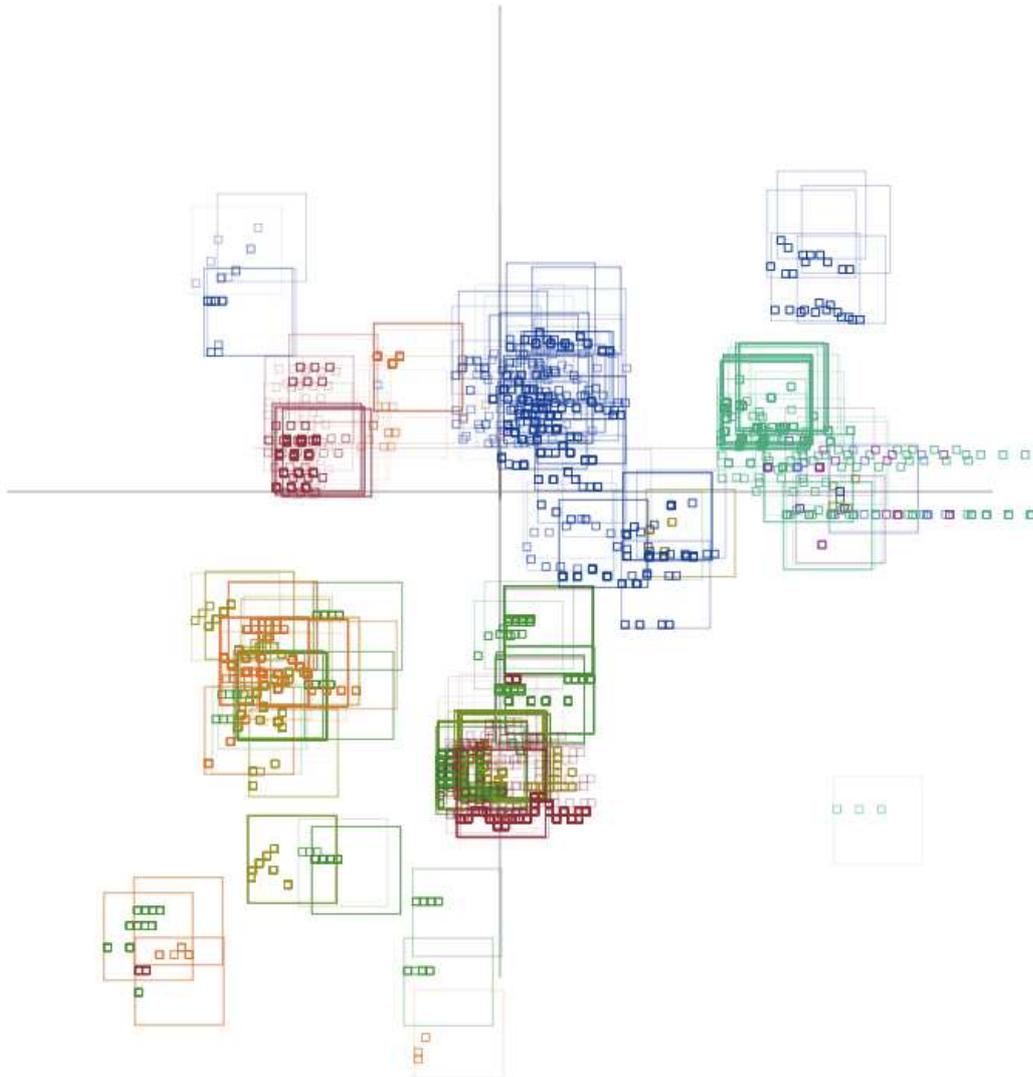


Figure 7: 45 Minutes of memes laid out spatially

real-time, affording exploratory data analysis.

4.3 Musical Features

The visualisation tool also allows memes to be laid out in terms of musical range and density. Figure 9 shows the same set of memes illustrated in Figure 7, but with x mapped to musical *range*, and y mapped to musical *density*. This provides a visualisation which allows analysts to spot styles of musical contribution, and to identify whether these are tied to specific participants. In Figure 9 it is clear that the orange participant has the most unique musical style - quite musically complex with only a few pauses in their memes. The dark green participant predominantly made memes with no pauses (top left quadrant), whereas the light green participant predominantly made memes with half the beats as pauses (bottom left quadrant). Overall, it can be seen that most participants created quite monophonic memes (containing notes predominantly of the same pitch). This is may in part be use to the ambient nature of the music which may favour simple memes - it would be interesting to explore whether the same pattern is found with other musical styles, or other levels of

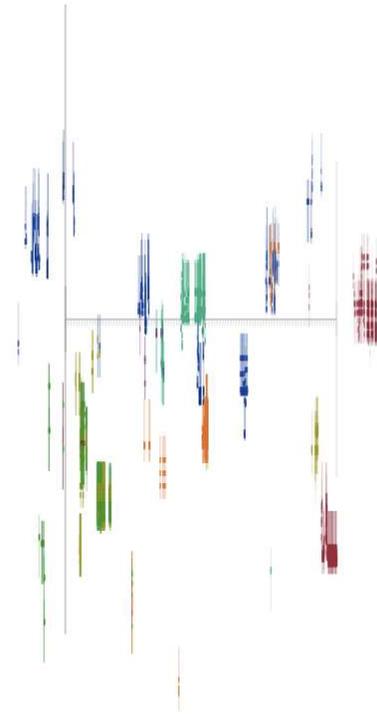


Figure 8: Time of memes

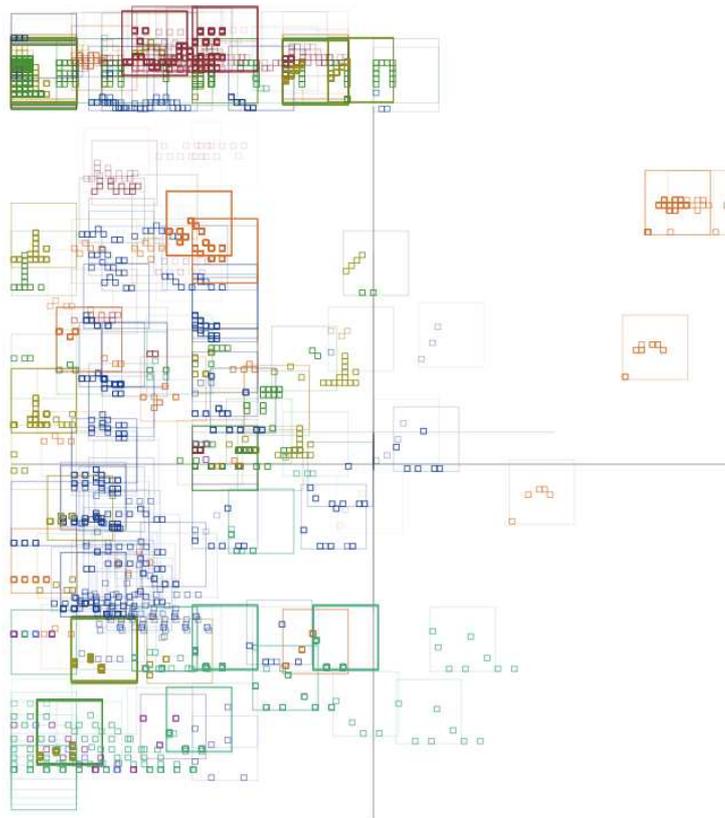


Figure 9: Musical Range and Density of memes

musical experience of participants. The interesting aspect of such pieces of interaction is that we can see that participants are making different *kinds* of memes, and if we switch to spatial layout in the visualisation tool (as illustrated in Figure 7), we see that they are making these contributions in different spatial locations, but at the same time. So, space and musical features are similarly differentiated between participants. The interactive nature of the visualisation tool allows the analyst to easily explore such connections.

The visualisation tool also allows musical features to be viewed over time to explore how features of the memes change with participant contributions. This form of visualisation provides the opportunity to explore concurrent participant activity.

5. SUMMARY

This paper illustrated the use of an interactive visualisation tool to explore logs of collective creativity in music making. The tool allows analysts to explore large data sets across the dimensions of time, participant, spatial location, and musical features of memes. Such a tool provides windows into the process of micro-creativity - how participants feed off each others' contributions, and how they use both time, space, and musical style to differentiate and harmonise with each other. The interactive nature of the tool provides opportunities for exploratory data analysis of the memes as they emerge, are modified, and are co-edited by participants. Developing more sophisticated pattern matching will allow larger data sets to be more easily explored and will make the connections between participants, their memes, and their actions clearer.

6. ACKNOWLEDGEMENTS

Work reported here has been supported by EPSRC EP/H042865/1, EP/E045235/1, GR/S81414/01.

7. REFERENCES

Boden, M. (ed). (1994). *Dimensions of Creativity*. Cambridge, MA: MIT Press, 119-142.

Boden, M. (2003). *The creative mind: Myths and mechanisms*, 2nd edition. London: Routledge.

Bregman, A. S. (1994). *Auditory Scene Analysis: The Perceptual Organization of Sound*, MIT Press.

Bryan-Kinns, N. (2004). Daisyphone: The Design and Impact of a Novel Environment for Remote Group Music Improvisation. In *Proceedings of DIS 2004*, Boston, USA, 135-144.

Bryan-Kinns, N., (2009). Everyday Creativity. In *Proceedings of ACM Creativity and Cognition 2009*, ACM Press, New York, iii-iv.

Bryan-Kinns, N., and Hamilton, F. (2009). Identifying Mutual Engagement. *Behaviour & Information Technology*.

Bryan-Kinns, N. (2011). Annotating Distributed Scores for Mutual Engagement in Daisyphone and Beyond. *Leonardo Music Journal*, Vol. 21.

Bryan-Kinns, N. (2012). Mutual Engagement and Collocation with Shared Representations. *International Journal of Human Computer Studies*. DOI: <http://dx.doi.org/10.1016/j.ijhcs.2012.02.004>

Candy, L., and Edmonds, E. (2002). *Explorations in Art and Technology*. London: Springer-Verlag.

Csikszentmihalyi, M. (1991). *Flow: The Psychology of Optimal Experience*. New York, NY: Harper Collins.

Dannenbergh, R., and Hu, N. (2002). Pattern Discovery Techniques for Music Audio. In *Proceedings of ISMIR 2002*, Paris, IRCAM, 63-70.

Dawkins, R. (1976). *The Selfish Gene*. New York, New York: Oxford University Press.

De Bono, E. (1992). *Serious Creativity*. New York, N.Y.: Harper Collins.

Graham-Rowe, D. (2009). Best of Twitter tunes album released. *New Scientist: Tech*, 19 November 2009.

Gusfield, D. (1997). *Algorithms on Strings, Trees, and Sequences*. Cambridge, USA.

Lampe, C., Ellison, N. B., and Steinfield, C. (2008). Changes in use and perception of facebook. In *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work*, 721-730

Nardi, B. A., Schiano, D. J., and Gumbrecht, M. (2004). Blogging as social activity, or, would you let 900 million people read your diary?. In *Proceedings of the 2004 ACM conference on Computer supported cooperative work*, November 06-10, 2004, Chicago, Illinois, USA.

Paulus, P. B., & Nijstad, B. A. (2003) (Eds.). *Group creativity: Innovation through collaboration*. New York: Oxford University Press.

Poupyrev, I., Lyons, M. J., Fels, S., and Blaine, T. (2001). New interfaces for musical expression. In *CHI '01 extended abstracts on Human factors in computing systems*, Seattle, USA, 491-492.

Sheridan, J. G. (2011). Developing an Open Source Exertion Interface for Two-Handed 3D and 6DOF Motion Tracking and Visualisation. In *Proceedings of BCS-HCI*, Newcastle, UK.

Shneiderman, B. (2000). Creating Creativity: User Interfaces for Supporting Innovation. *ACM*

Transactions on Computer-Human Interaction, 7(1), March 2000, 114–138.

Sternberg, R. (ed), (1999). *Handbook of Creativity*. Cambridge Univ. Press, Cambridge, UK.

Uitenbogerd, A., and Zobel, J. (1999). Melodic Matching Techniques for Large Music Databases. In *Proceedings of ACM Multimedia '99*, Orlando, USA, 57-66.

Wiggins, G. A. (2006). Searching for computational creativity. *New Generation Computing*, 24(3), 209–222.

Wang, G. (2009). Designing Smule's iPhone Ocarina. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. Pittsburgh.

Zhao, D. and Rosson, M. (2009). How and why people Twitter: the role that micro-blogging plays in informal communication at work. In *Proceedings of the ACM 2009 international Conference on Supporting Group Work*, 243-252.