Analysing Group Creativity: A Distributed Cognitive Study of Joint Music Composition

Shahin Nabavian (shahin@dcs.qmul.ac.uk)
Nick Bryan-Kinns (nickbk@dcs.qmul.ac.uk)
Department of Computer Science: Queen Mary University of London
London E1 4NS, UK

Abstract

We investigated the processes involved in an instance of group creativity by conducting a pilot study that looked closely at a group of expert musicians creating musical compositions over a period of time. We used Distributed Cognition principles to investigate how information propagates and transforms and how it impacts the compositional process. We identified three key processes (attainment, experimentation and structuring) that help the group achieve a successful operation in creating compositions together over a period of time.

This approach has implications for how we investigate group creativity in general. By focussing on the propagation and transformation of information, we can gain a more systematic understanding of how groups come to create products together. Analysis from this pilot study demonstrates that an individual's creative musical idea possibly bears less influence on group creativity than the group's ability to transform that idea in a desirable manner. Therefore, information processing within Joint Music Composition has a major impact in how groups create compositions. This view sheds light on how we view the notion of group creativity. It can inform how we design tools to support it and how we design experiment to investigate it.

Keywords: Group Creativity; Distributed Cognition; Music Composition

Introduction

Many researchers have discuss the contested nature of what constitutes creativity and how it can be viewed in different domains (Boden 1992, 1994; Csikszentmihalyi, 1996; Sternberg, 1999; Schneiderman, 2000; Sawyer, 2003; Paulus and Nijsatd, 2003). However, what is yet to materialise is a practical application of theories that can be used to investigate the phenomenon of individual or group creativity.

Group Creativity Research

The term group creativity can be attributed to a number of diverse situations in which a group of people communicate and/or work together. Our research defines group creativity as situations that are not scripted beforehand and therefore have an improvisational element where members of the group collaborate together to shape the emerging flow of interaction (Sawyer, 2003) and where the creation of a product requires a distribution of cognitive and physical labour.

Studying musical groups provides an opportunity to investigate group creativity. Sawyer studied improvisational music and theatre as it "exaggerated the key characteristics of all group creativity: process, unpredictability, intersubjectivity, complex communication and emergence". The principle idea of improvisation is that the process is the product. Sawyer's theory of analysing improvisation is based on *semiotic mediation,* which relates to linguistic ideas of *deictics* and *indexical entailments.* This theory helps illustrate how a person makes some presumption about a future action and is therefore a suitable way of understanding the improvisational process, be it in conversation or music. However, there are numerous collaborative situations such as Joint Music Composition (JMC), in Western contemporary music, where synchronous interaction such as improvisation plays a major part, yet it is not the product. Indeed the product in JMC, where 2 or more musicians collaborate to compose, is a song (composition) that retains a structure once it is deemed complete. JMC is product creativity that develops over time and in a group context. At this stage, our research is not focussed on understanding how individuals come to make predictions on future actions. We are interested in the role of information processing and how it impacts the creative process.

Distributed Cognition

The principles of Distributed Cognition (DC) have been around for over a decade and have primarily been utilised in workplace settings (Nardi, 1996; Hutchins, 1995a, 1995b; Heath and Luff 1997). More recently DC has been employed in new directions. For example, Kirsh's (2004) design principles for a visual e-learning environment are based on the notion that metacognition can sometimes be associated with external processes; hence metacognitive decisions are sometimes distributed between the internal and external and can be affected by a visual tool. We illustrate the process of JMC using Marr's (1982) computation, algorithmic and implementation levels of description in similar manner to how Hutchins (1995a) and Flor and Maglio (1997) utilised it. Our findings are based on a pilot study that was conducted with a group of expert musicians who were given the task of writing songs over a period of three weeks. At the computational level we describe the constraints that need to be satisfied to achieve a successful operation in JMC. At the algorithmic level we specify three levels that help encode the propagation and transformation of information for key processes within the system. At the implementation level we shall illustrate details of how the representations are actually realised in the system.

We conclude the paper by discussing the implications of our approach and findings for research into JMC and group creativity.

Computation of JMC

Our research is based on studies of Western contemporary music groups (i.e., Rock, Pop, Folk etc.) and therefore our definition of JMC is based on work situations governing these genres of music.

When writing songs, a group often starts with fragments of musical ideas that are manifested through the instruments individual members play, along with verbal, written and gestural communication. Over time, through various forms of interaction between group members and artefacts in their environment, a composition emerges. The primary task of the group is to create and co-ordinate what each musician plays and when. JMC is a classic DC work situation where the cognitive and physical labour of the task is distributed (Hutchins, 1995a) but with the added dimension of time. By this we mean compositions can take several work sessions, at different points in time (i.e., several days or weeks), before being completed.

JMC is an informal work setting in the sense that there are no manuals written on how to conduct work. We would classify JMC as an "ill structured" system where the role of the participants, the processes and artefacts are unspecified or under defined (Perry, 1999). However, there are some structures that help musicians co-ordinate action in JMC. Each musician plays a musical instrument and each instrument has a function within a song. Musical and behavioural conventions, relating to genres of various Western contemporary music, can also help musicians co-ordinate actions and possibly expectations. In addition, musicians commonly use a variety of techniques in overcoming the cognitive burden associated with performing songs (Flor and Maglio, 1997). Therefore, though we state that JMC is ill-structured, it is clear that some structures exist.

Musical Conventions

In most situations, the musical role of each musician is greatly influenced by conventions associated to the genre of music within which the group is composing. In simple terms, musical choice can be constrained to certain actions depending on the genre. For example, within standard Jazz the set of chords/musical notes and the structure that underlie the composition, is known as the *form*. Awareness of the form and knowledge of the conventions of a genre, such as standard Jazz, often dictates the choices made by musicians in selecting notes and scales they play over the form.

Form

The concept of the form is important to understand as we believe it plays a major role in how a group develop a composition. In JMC, the form is often developed within a group context especially as it may not be fully created in the beginning. This means that the group have to extend the form and agree whether the new form is acceptable to the group. In most cases of JMC, the initial musical idea that is proposed for composition can be thought of as musical information for certain parts of the form and possibly for certain members of the group. To complete a composition, all members should play sequences that adhere to a common structure and that are linked in some way (i.e., key, scales, harmonies, rhythms etc.). At the computational level, the most basic constraint that needs to be satisfied for a successful operation in JMC is the emergence of a composition that redeems a form after a period of development. In Western contemporary music such as pop or rock, other compositional features (for example, melodies and solos that are played over the form) are often expected to be redeemed for future performances. We use the terms redeem or retain to mean the process by which individuals within the system store and reproduce existing musical sequences. This process involves both internal and external representations, which are often interlinked. For example, a guitarist can remember the notation of a composition by recording the labels of chords and notes in their own memory (internal representation) and/or recording a representation externally, for example creating written notes and audio recordings. A guitarist can also use their guitar a as memory aid by recalling the finger positions that they use when playing the composition.

It is important to note that though musical conventions play a major role in the decision making process of the individual, a song can still be written without the individuals sharing knowledge of the same conventions. Therefore, by our definition, the constraints for a successful operation in JMC to be satisfied are not always down to conventions. JMC constraints are satisfied if the group manage to create a form that they can reproduce and that they themselves recognise as a composition.

In order to investigate group creativity, we needed to create a description of the algorithmic and implementation levels in JMC. In particular we were interested in how information propagates and transforms and how this related to the compositional process.

Pilot Study

A pilot study was conducted from 1st - 22nd November 2004, at Queen Mary University of London (QMUL), which entailed an analysis of a group engaged in JMC.

Aim

Since this study was the first of its kind, our aim was to identify observable information that appeared relevant to the work of the group; what representations were used and why; what was the affect of physical environment; what was the role of the artefacts in the process; what was the affect of time in the process of work?

Observation Set Up

The Group The group consisted of musicians who had not worked together before. They were all expert musicians in the sense that they could play their instruments proficiently; two were classically trained, one had some form of formal training and the fourth had been involved in writing music in Western contemporary music bands for over 10 years. The group consisted of two males and two females all in the 20-30 age range. Three members were research students in different departments at QMUL, and the one was an artist and animator.

One of the participants was the researcher of the pilot study. The main reason for being directly involved was to follow outside interaction between group members as it may impact the compositional process. Another reason for this participation was for the researcher to have some form of understanding of the process from within the system. The other members were aware of the dual role of the researcher. Based on the observations that we have made of other musical groups in their work environment, we believe that this participation did not appear to impact the process being analysed (i.e., we did not identify any side effects).

The participants were named 'H' (violinist), 'C' (guitarist), 'S' (bassist) and 'A' (keyboardist).

Task The musicians were asked to write at least one song over the course of three weekly sessions. They were told that these sessions would be filmed. No description was given of what constitutes a song. They were asked to be prepared to write a song with people whom they would meet in the session. They were invited to bring compositional ideas if they wished to do so. Apart from the researcher, the participants were paid £5 per session.

Physical Setting The sessions took place in a section of the Electronic Engineering (EE) Lab in the Engineering building at QMUL. We positioned the chairs for people to have access to the equipment that they were to use. 'A' was positioned behind the keyboards; 'C' next to the guitar amp; 'S' next 'C' and within distance of the bass amp; 'H' was positioned in a way that was in the line of vision to all members and at least one of the two cameras recording the pilot study. In essence we attempted to recreate as natural a rehearsal setting as possible with the assumption that visual and aural channels were important to the process of work.

Data Capture The study consisted of filming 1.5 hours of the musicians working together in the EE lab. Each session lasted between 1.5 and 2 hours. Copies of the written notes that musicians created were collected after the final session. E-mail exchanges that the researcher had access to, were also treated as data

Method of Analysis

To start with we had to define what constituted *information* in this environment. We classified information in terms of what we observed as possible inputs and outputs to and from a musician. Primarily, we focussed on verbal communication, some gestural communication (gaze, nodding, smiling etc.), sounds produced by musical instruments and written notes created by the musicians during the process of work.

We then proceeded to investigate where information comes from, how it propagates and how is it transformed. We then correlated this with key events in order to map the relationships between interactions and consequences for the compositional process. Examples of key events are: instances where musical ideas were proposed to the group; the beginning of a session when the group had to recommence working on compositions from a previous session; times when perceptions of representation relating to the composition (i.e., labels for musical structures and musical parts) appeared to be different for members of the group especially when divergence in perception caused a breakdown in communication or music playing.

We transcribed the communications to highlight areas of misunderstanding or areas where we wanted to demonstrate a certain process of work. We analysed the information created within the written notes and tracked the timings of when the written notes were created and when they were referenced.

Summary of Findings

For the purposes of this paper we will describe three algorithmic levels of description that we feel reflect the key group processes within JMC. It must be stated that these are high-level descriptions and constitute some but not all algorithmic levels of description available in JMC; this is beyond the scope of this paper. Primarily, we use these levels to classify individual actions within a specific context of group work to illustrate the fundamental processes that help the group to achieve a successful operation in JMC.

Key Algorithmic Levels in JMC

(Attainment) This is characterised by gathering information of core musical units of principle idea, through various mediums. Members interact with each other and the artefacts to determine whether the states of representation have propagated appropriately.

(Experimentation) Typically at this level, the core musical units have propagated to musicians and they then contribute their own knowledge to the attained information in order to extend the form. The information can transform into playing an instrument or it may prompt a verbal contribution. This is primarily the stage during which the group experiments and verifies what they want to retain.

(Structuring) At this level, a structure is created based on the retained ideas, as a way to co-ordinate cues for musical changes. Individuals can employ different algorithms in structuring musical information but would need to achieve the same output in order to perform in unison.

It is our theory that a group cannot develop a composition from a stage where the form is incomplete (i.e., musical information is not fully available or created for the form for all members in the group) without employing these algorithms. The composition cannot develop if the states of representation relating to the core musical units that constitutes the form (i.e., notes, the structure, the tempo, the rhythm etc.) does not propagate across the members of the group. *Attainment* refers to the process by which members obtain core musical units relating to the form. The form is rarely formulated from the outset. Some information may exist for certain members but in the end all members would have to conduct work on the basis of the core musical units. At the same time to create a composition, the form needs to be extended. This is the *experimentation* process. Once it is extended, the group need a system of co-ordination to reproduce what they have extended. We refer to this process as *structuring.* Of course, there are other fundamental processes such as *verification* and *retention*. We shall touch on these in the course of the paper.

It would be beyond the scope of this paper to create details beyond general descriptions of the three levels. However, we shall attempt to illustrate how the **implementational level** relates to the algorithmic levels in relation to compositional process by using some excerpts from the pilot study.

Excerpt 1

0:43:02 *A:<plays keyboards> "E flat?" (Looks to 'C')*
0:43:09 *C:"So that erm B flat major over a B flat chord I am sure there is better name for it"*
H:"it's diminished"
C:"yeah half diminished or something?"
A:<plays keyboard> "isn't it just major 7^{th}*s?"*
C:"yeah erm" (looks at his guitar carefully as he plays notes)
0:43:43 *C:"eh it has a ahh what it is it's a" <plays the same chord twice> "it's kind of ambiguous because it doesn't have a* 3^{rd} *in it does it?"*
A:(puts her hands on the keyboard) "oh does it not?" <plays keyboard> (not quite in unison with C)
0:44:02 *C:"its got a oh its got a 9 in it".*

This excerpt is typical of attainment level interaction when musicians aligned their understandings of what each other played. The excerpt demonstrates several important representational and interactional features of all three algorithmic levels: 1) musical and verbal modes of communication were the most prevalent and it created a fluid form of expression where verbal communication was intertwined with musical sounds, 2) the majority of representation used in the interaction was transient in nature, 3) musicians perceived the same musical information slightly differently, 4) musicians used their instruments to map internal and external representations, 5) the fluid nature of the work situation played an integral part in the cognitive processes of the musicians (i.e., instant feedback and access to other sources of information influenced how they worked), 6) the open channel of communication invited people outside of direct interaction to participate, in this example H became involved.

In the pilot study, individuals transformed propagated information into some form of action, like playing their own instruments. At certain points, the group *jammed* (performed the composition in its present state). The performance can be seen as an attempt by the group to co-ordinate the individuals' transformation of propagated information, as manifested by the instruments they play. This type of activity helps the group to verify whether the states of representation relating to the compositional information have propagated appropriately.

Excerpt 2

'S' plays a set of notes; 'C' replicates it on his guitar. Once one set is replicated, 'S' plays the next set of notes. After replicating (or attempting to replicate) what 'S' plays, 'C' compiles a set of chords. 'C' asks 'S' "so the chords you are outlining is" <plays a chord> "that" <plays a chord>.

Once the information propagated and the group achieved a common ground of understanding (Clarke and Brennan, 1991), they extended the form together. Often the group verified what each member added to the composition. This meant that members of the group assessed whether musical sequences met the criteria of what they were attempting to compose. Notice that in certain genres, like standard Jazz, verification could be based solely on *objective external verifiers* such as musical conventions. Therefore, conventions can be used to create a criteria for verification. The pilot study group appeared to set their own boundaries for verification, which is something that we have seen in previous observations of JMC (Nabavian, 2002). The verification may have been influenced by conventions each musician had knowledge of, however verbal judgement was rarely based on external verifiers (i.e., conventions of what scales should be played); s*ubjective verification* played a bigger role. Subjective verification can be seen as musicians basing judgement on what they feel sounds good rather than whether the sequence belonged to a certain convention.

Excerpts 1 and 2 show that the attainment of the core musical units rely heavily on the interaction between the musicians and between the musicians and the musical instruments. We therefore believe the transformation of compositional information, in context of the work in JMC, is a group process. In particular, the feedback between musicians plays a major role in defining whether there is a common ground of understanding between the group. This can exist through out the three algorithmic levels set out because of two main reasons: 1) the responsibility or decision-making process is often at system level and it requires constant feedback from members to determine whether propagation of information and the transformation appears agreeable and 2) the representations used within the group do not always yield the desired transformation.

We state that the decision making process is often at system level as musicians can influence what each other plays. Even though the information propagates and transforms in a manner deemed suitable by the individual, it does not mean that it is deemed suitable by others. For example, 'H' asks 'A' to make her notes “sparser than that...because otherwise I can’t put anything on top”. 'C' asks 'A' to play her notes higher up the keyboard “because there isn't much happening at the top”. 'C' asks 'S' to change his bass line to play more in the same rhythm as his guitar part. This is more the type of processes that we associate with the experimentation levels of description.

There is sufficient evidence in the pilot study and previous studies (Nabavian, 2002) to suggest that the representations used within the JMC do not always yield the desired transformation.

Excerpt 3

21:18	*S:“So the 1st verse sounded alright we could do it 8-2-4-2”*
	C <looks at S and nods>
21:19	*H:“8”*
	S:“2-4-2”
21:30	*H:“you are talking in terms of bars?”*
	A to S: “what do you mean?”
21:31	*S to H: “bars”*
	A:“Oh because you are doing it longer anyway”
21:34	*S: “So one repetition – I call that a bar”*
21:37	*A:“because the 2nd time <plays something> <laughs>*
21:42	*H:“that's the chorus”*
21:46	*H:“The chorus is”*
21:48	*S:“The 2nd part” <plays notes>*
	H:“Twice as long”
	S: “The chorus”
	A:“Twice as long”
21:54	*C <leans over to look at A's notes>*
21:58	*A:“because the first one” <plays notes> “and the chorus goes” <plays notes>*
	S <plays with 'A'>
22:08	*S:“yeah I see what you mean”*

This excerpt illustrate how representations can be perceived differently even after multiple sessions of work with the same composition. In this instance, the system that musicians employed in counting cycles of music appeared to be different. They managed to conduct work because the output was the same (i.e., they all managed to change at the right places). S and C appear content with using the term bar number to signify a certain number of patterns to be played. The representation “8-2-4-2” signified a structure for them. 'A' did not appear to use this method. In fact, there are other places where she plays her sequence in the wrong places. The concept of what constituted a bar was not something that 'H' and 'A' shared with 'S' immediately. This excerpt was taken from the final session of the pilot study. The group had performed the composition with various structures without realising their notions of the structure of the form may be different. If musicians can perform whilst having different states of representation when expressing the structures, why do they create them?

In theory structures are created in order to retain co-ordination of the playing of existing composition. They can serve as a basis to co-ordinate cues for the changes of musical parts during the performance of composition. Structures also help musicians to breakdown the task of recalling what they have to retain. Flor and Maglio (1997) refer to this type of activity as ‘chunking’. Essentially, structures segment the serial composition into sections with labels where each musician plays a certain part for an agreed length. An example of a Western contemporary pop song structure can contain the following labelled sections: *Introduction, Verse, Chorus, Bridge, Middle 8*. When musicians reach the end of the length of the section, they change to the next segment without actually needing to know what someone else is playing; they just need the cue to change. Musicians can utilise a number of methods to track changes when performing to a structure. For example, if a musician loses count of the number of bars that s/he has played, s/he can look for gestural cues (like nods from other musicians) or musical events within the composition (like a particular sequence in the song) in order to establish when to change. In this way, they co-ordinate the playing of retained ideas over a period of time without attempting to remember everything that occurs in the system.

Discussion

We have illustrated that a musical group can create compositions without sharing knowledge of the same musical conventions. We have illustrated that, at the implementational level, there can be more than one type of representation utilised in achieving the same output (for example, cues to change by counting bar number or looking for particular musical sequences or events). We have also shown that, in JMC, members of the group rely on constant feedback from the system to overcome any differences in interpretation that might be caused by using different representations in the process work.

In terms of the development of the composition, the group needed a common understanding of the fundamental underlying information that constitutes the composition (i.e., the form). There are occasions (i.e., excerpt 1), when the group literally spell out each musical note and there are periods when they do not discuss or query what each person plays. Therefore, we state that the states of representation of certain musical units (for example, information about the form) must remain the same through out the system for the composition to develop. We have also illustrated that much of the transformation, in relation to what is produced and retained for the composition, is dependant on the group verification.

We have specified three algorithmic levels of description (attainment, experimentation and structuring) that help the group in satisfying the constraints of JMC at the computational level (i.e., creating a composition that redeems a form after a period of development). In other words, without attaining information about the form, extending the form and structuring what has been created in order to reproduce it, the constraints for a successful operation in JMC would not be satisfied.

Structuring the analysis of natural observations of JMC in this way allows us to assess the impacts of the propagation and transformation of information at various points during the development of the composition. This is important because it illustrates how information processing within JMC impacts group creativity. Information may propagate in a number of ways and create a number of transformations. Each of these transformations will play a role in the group's ability to achieve a desired outcome for the composition. In simple terms, group creativity is influenced by how the states of representation of musical ideas (as manifested by the musical information such verbal labels of musical notes, hand movements on instruments, written musical notation etc.) propagate and transform within the group. This means that the individual's creative musical idea possibly bears less influence on group creativity than the group's ability to transform the information contained in that idea in a desirable manner (i.e., achieving a transformation that is deemed suitable in the context of composition). Therefore, it is not just the individual's idea but what the group makes of the idea that impacts group creativity.

We are currently investigating whether we can help musicians enhance the group's ability to transform information in a more uniform manner by creating constraints for the propagation of the states of representation relating to the three algorithmic levels that we have identified. We would like to test the possibilities of attempting to superimpose a structure, through a computer-mediated tool, that would constrain the possibility of interpretation of representations during the process of work. Musical scores can be thought of as system that imposes certain constraints. However, formal scores do not appear to be used in JMC, as seen in the pilot study. There are numerous reasons for the absence of scores in JMC: 1) musicians all need the knowledge in creating and reading scores, 2) it is a time consuming task, 3) composition is dynamic, changes can be frequent and therefore to make written notes of everything will disrupt flow of work and 4) scores do not always help to produce the exact transformations (in terms of how musical notes are phrased). Computer mediated tools may be one way that enable groups to create representations that help attainment, retention and structuring of musical ideas in a group context fluidly. We seek to investigate the features that are required to help maintain the states of representation and examine whether they help achieve a transformation of information better than or compliment current representations.

Acknowledgements

This research is funded by the Engineering and Physical Sciences Council (EPSRC). Grant reference: GR/P00130/01.

References

Boden, M. A. (1992) *The Creative Mind.* London: Abacus.

Boden, M. A. (1994). *Dimensions of Creativity.* London: MIT Press.

Clark, H.H. and Brennan, S.E. (1991). Grounding in Communication. pp.127-149 In Resnick, L.B., Levine, J and Behrend, S.D. (eds.) *Perspectives on Socially Shared Cognition*. Washington DC.: American Psychological Association.

Csikszentmihalyi, M. (1996). *Creativity: Flow and the Psychology of Discovery and Invention*, HarperCollins

Flor, N. V. and Maglio, P. P. (1997) Emergent Global Cueing of Local Activity: Covering in Music, *CSLC Proceedings,* 45-52.

Heath, C. and Luff, P. (1997) Reconfiguring Media Space: Supporting Collaborative Work. in Finn, K.E., Sellen, A.J. and Wilbur, S., (eds.) *Video-Mediated Communication*. Lawrence Earlbaum Associates. Mahwah, New Jersey. pp.323-347.

Hutchins, E. (1995a). *Cognition In the Wild.* Cambridge, MA: MIT Press.

Hutchins, E. (1995b). How A Cockpit Remembers Its Speed, *Cognitive Science,* 19, 265-288.

Kirsh, D. (2004). Metacognition, Distributed Cognition and Visual Design. In Gardinfas, P. and Johansson, P. (eds.) *Cognition, Education and Communication Technology*, Lawrence Erlbaum**.**

Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information* . San Francisco, CA: Freeman

Nabavian, S., (2002) *The Emergence of composition from improvisation*, AMSc. Thesis, Dept. of Computer Science, QMUL.

Nattiez, J. J. (1990). *Music and Discourse: Toward a semilogy of music* (Carolyb Abbate, Trans.). Princeton, NJ: Princeton University Press.

Nardi, B. A. (1996). 'Studying Context: A Comparison of Activity Theory, Situated Action Models and Distributed Cognition.' In Nardi, B. (ed.) *Context and Consciousness: Activity Theory and Human-Computer Interaction.* MIT Press.

Paulus, P. B., and Nijstad, B. A. (2003). *Group Creativity*, NewYork: Oxford University Press

Perry, M., (1999). The application of individually and socially distributed cognition in workplace studies: two peas in a pod?, *European Conference in Cognitive Sciences, Siena, Italy* – p. 87-92.

Ruwet, N. (1972). Linguistics and musicology. *International Social Science Journal, 19,* 79-87.

Schneiderman, B. (2000). Creating Creativity: User Interfaces for Supporting Innovation, *ACM Transactions on Computer-Human Interaction, Vol. 7, No. 1.*

Sawyer, K. R. (2003). *Group Creativity: Music, Theater, Collaboration.* New Jersey: LEA.

Sternberg. R. J. (1999). *Handbook of Creativity.* Cambridge University Press.