Experiments on collective decision-making and musical metacreative performances with audience participation

Marcelo Gimenes
Interdisciplinary Centre for Computer Music Research - University of Plymouth
marcelo.gimenes@plymouth.ac.uk

Abstract

This paper introduces the performance-based research led by the author in the area of collective decision-making and musical metacreative performances with audience participation. A series of six experiments (or interactive music performances) has been created in order to test alternatives to the implementation of features in metacreative systems aiming to foster engagement and improve interaction between performers and audience. This paper describes the evolution of the solutions sequentially adopted in these experiments with regard to control data, sound generation and user experience.

1. Introduction

This paper introduces design approaches adopted by the author in the areas of collective decision-making and metacreative performances. The aim is to describe how these approaches were progressively tested in a series of six interactive performances. Each one of these performances gradually informed the introduction of improvements with respect to the implementation of a distributed system aiming to support collective decision-making and control of an intelligent generative music machine. This practice-based research draws on the experience reported by many other artists and researchers in the context of metacreative systems. At the same time, it proposes combinations of technologies that have not yet been tested (to the best of the author's knowledge). One of these is the possibility of using human intelligence (via audience participation) alongside artificial intelligence techniques embedded in computer systems.

One of the background questions of the research reported in this paper are the apparent divisions between the roles of composers, performers and audience. Distinctions between these roles seem to be relatively well defined but, in fact, it is often difficult to consider them individually without acknowledging a good number of exceptions. The European common practice period, from where much of today's western musical practices inherit, seem to be at one end of the spectrum. It is fairly easy to see, in this case, where composers, performers and audiences sit. The same distinction is not as clear, however, in practices and traditions such as African dance and music, which are largely functional and participatory in nature. Improvisation, broadly speaking, is another example where these roles are not necessarily distinct as composition (in the sense of a successive generation of musical ideas) and performance occur (almost) simultaneously.

Given the above difficulties, it is relevant to consider how metacreative systems are situated within this spectrum of possibilities. Aiming at a particular spot on this spectrum is often a question of design and, generally speaking, each individual artist and researcher will have his/her own view and approaches about where to place generative and performative components.

The third element of the equation (the audience), however, seems oblivious to this discussion and is often overlooked. The work reported in this paper tries to redefine the role of the audience as an element of participatory design. In the past as much as today, audience participation has been mediated by all sorts of technologies, from simple push buttons (and sliders, etc.) to more advanced motion tracking devices (such as infrared cameras, and the Microsoft Kinect). Often interactive systems are evaluated with respect to creativity or the simulation of specific genres, among other criteria (Agres, Forth et al. 2016). In this case, the importance of the audience's feedback is essential.

The literature reports a wealth of research in this area, which showcases many creative, technical and conceptual approaches. Paul (2003) and Ouzounian (2008) provide a comprehensive survey of systems and approaches to audience participation especially in sound installations. Trying to survey the area is an arduous task that goes beyond the scope of this paper. However, a few examples must be mentioned in order to better situate the research introduced here. Freeman and Godfrey (2010), for instance, introduced Flock, a complex work for saxophone quartet, dancers, audience participation, video and electronic sound. Location of participants, in this case is tracked with a computer vision system. Game elements (such as competitiveness) are also present as well as music notation generated in real time. NOMADS, introduced by Burtner, Kemper et al (2012) addresses audience participation through a mobile interactive computer ensemble distributed across a network in multiple performance venues. Here, a "single emergent sound and visual structure" is generated based on the input from "thousands of simultaneous users" (*ibid.*). Quintet.net (Hajdu 2005) also supports real-time interactive performance distributed on local networks as well as the Internet.

This paper addresses a particular set of technologies applied to audience participation, in particular the ones resulting from the overwhelming advances in communication occurred during the last decade or so. These include smartphones and other wearables, such as smart watches and earphones. Belonging to an increasingly connected world has become essential in people's lives today. In fact, even considering differences of ownership in terms of geographic and economic contexts, smartphones are widely used almost everywhere.

What makes these technologies particularly appealing as facilitators of audience participation in musical metacreative systems is the intrinsic value derived from their many sensors, which provide a "huge potential to gather precise, objective, sustained, and ecologically valid data on the real-world behaviours and experiences of millions of people" (Miller 2012, p. 221). In addition, smartphones are used in such a physical and personal way that they are often regarded as extensions of peoples' own minds and bodies.

A formal analysis of how smartphones facilitate interactive performances has been proposed by Essl and Rohs (Essl and Rohs 2007, Essl and Rohs 2009). The Open Symphony system (Hayes, Barthet et al. 2016) uses a "smartphone-friendly web-based application for audience voting" along with a "quasi real-time visualization of a symbolic score for performers generated based on the audience's votes". Internet messaging services (Wang, Oh et al. 2011), among other approaches and techniques, have also been used for musical applications.

A performance-based research

It is obvious, given all the above-mentioned examples, that using smartphone technologies in the design of metacreative systems can provide a plethora of new creative possibilities. In 2016, the general theme of the Peninsula Arts Contemporary Music Festival (in Plymouth, UK), "Frontiers: expanding musical imagination" set the initial motivation for a series of six performances:

- Embodied iSound (Peninsula Arts Contemporary Music Festival, Plymouth)
- Performance without Borders (Peninsula Arts Contemporary Music Festival, Plymouth)
- Smart iSound (klingt gut!, Hamburg)
- Sound Games (International Computer Music Conference, Utrecht)
- Sound Games 2 (Off the Lip, Plymouth)
- Vox Populi (International Conference on Digital Research in the Humanities and Arts, Plymouth)

Figure 1 contains a very succinct overview of these performances:

Date	Title	App	Movement	Game
27/02/16	Embodied iSound	Levinsky	y	n
28/02/16	Performance without Borders	Sherwell	n	n
27/05/16	Smart iSound	Levinsky	y	n
12/09/16	Sound Games 1	Levinsky	y	y
22/10/16	Sound Games 2	Levinsky	y	y
08/11/17	Vox Populi	Levinsky	n	n

Figure 1: Summary of performances

The mind-set initially established was to progressively experiment with different approaches to audience participation in order to foster engagement and improve interaction between performers and audience. Members of the audience would become active participants (as opposed to passive listeners), decision-makers, performers and game players, sharing the responsibility of defining the end sound result.

2. First experiments and main systems

At the heart of the above-mentioned experiments is a set of systems, in particular a smartphone app, which collects the bulk of raw control data generated by the audience (the 'app') and a desktop application (the 'server'), which retrieves data from the apps, processes and maps (the resulting) calculations to sound generation.

Two main sets of systems were implemented: Sherwell (Gimenes 2016c) and Levinsky (Gimenes 2016b). Even though the app remained essentially the same throughout all the performances, on each occasion the server was incrementally adapted in order to accommodate specific needs. Progressively, new features were added whilst previous solutions abandoned, from approaches to control data and sound generation techniques, to spacialization and embodiment.

The first two experiments, described below, occurred during the Peninsula Arts Contemporary Music Festival at Plymouth University.

Performance 1: Performance without Borders

Performance without Borders (Gimenes, Largeron et al. 2016) is an improvisational piece for acoustic instruments (piano and violin) performed by musicians (on stage) and electronics controlled by audience on a traditional space (concert hall) as well as remotely via the Internet. The initial theme and general musical reference, performed on piano and violin is the composition “Dreams”, by John Cage, around which improvisation follows. Other sounds and musical elements are progressively added, based on decisions taken by the audience using the Sherwell app.

Sherwell is essentially a voting platform whereby (interface shown in Figure 2) members of the audience can choose from 5 different parameters and cast votes (for each one of these parameters) by tapping on a central square:

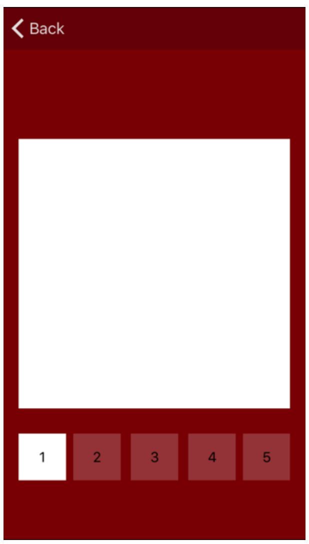

Figure 2: The Sherwell app main screen

The actual meaning of parameters and position of votes inside the square (‘x’ and ‘y’ coordinates) is part of the composition design and is explained to the audience before the performance starts. Data generated by the app (running on each one of the participant’s smartphones) is immediately sent to a webserver in the Cloud and, at regular intervals (every ‘n’ milliseconds), retrieved by the server running on a desktop computer. This feature allows the participation of audience that is not only physically present in the concert hall, but also in other parts of the world, through the Internet.

The server component of the Sherwell system contains the algorithms that continuously runs statistical analyses of the data received from the Cloud (including timing and coordinates of the votes related to each of the five parameters) and summarizes them into commands that are then sent to other components (of the distributed system) responsible for generating sound. In the case of the Performance without Borders, this was accomplished via sound synthesis and sound clips triggered in Max/MSP. Figure 3 shows a photograph taken during the performance with a projection in the background of a feedback screen which helped audience and performers on stage to communicate.

Figure 3: Author performing at the Performance Without Borders

As a first experiment, the evaluation of the performance based on reports from the performers and the audience confirmed positive and negative aspects of the voting approach to audience participation of the Sherwell system. On the positive side, participants knew exactly what they were voting for, as instructions were simple and easy to follow. On the negative side, averaging the coordinates for each one of the parameters meant that participants had difficulty to follow their individual decisions. This particular approach was abandoned in the following experiments, which sought a more direct correspondence between control data and sound output.

Performance 2: Embodied iSound

The second performance, named Embodied iSound, was also premiered at the Peninsula Arts Contemporary Music Festival in Plymouth and used Levinsky. This distributed system implements audience participation from an entirely different perspective, namely, direct control. In addition to Levinsky (app and server), other components of the composition included an electroacoustic composition in Logic and algorithmically generated sound synthesis in Pure Data.

The controls implemented in Levinsky are primarily derived from participants’ movements. For that purpose, the app (main screen shown in Figure 4) grabs data from the smartphone’s gyroscope, accelerometer, digital compass and position (location) in the performance space where 6 iBeacons (Bluetooth low energy emitters) are installed. Audience is invited to walk around the performance space and also gesticulate with their smartphones.

Some performance parameters are controlled by the collective behaviour of the audience and others by individual participants. For instance, the sound output used a quadraphonic surround sound system and, in order to dynamically control the volume of each one of the loudspeakers, Levinsky calculates the number of participants next to each one of them. The more participants are closer to one particular (iBeacon and) loudspeaker, the softer the sound of this loudspeaker is (relative to the other ones).

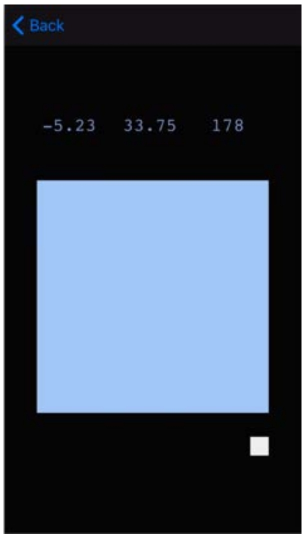

Figure 4: Levinsky app main screen

In addition to observing collective behaviour during the performance, Levinsky also hands direct and individual control of some performance parameters over to a limited number of participants. For instance, given that, say, three participants are next to loudspeaker 'A', the system randomly selects one of them to directly control the positioning (surround angle) of a particular track (or group of tracks) in Logic. Data from the smartphones that control this parameter are generated by the gyroscope's 'yaw'. This mapping is actually very intuitive given that, as the participant moves the device towards a particular spot, the position of the sound (in the surround space) follows the movement to that spot.

In addition to sound positioning, Levinsky also allows direct control of an algorithmic sound generation machine that runs in Pure Data. This is performed through a button on the app interface: one tap on the button starts the sound machine and another one stops it. Following a similar strategy, pitch and roll angles of the smartphones' gyroscope control sound synthesis parameters that are set in Pure Data.

Considering that the sound output of the Embodied iSound performance is done through four loudspeakers, there is the possibility of having one participant directly controlling each one of them (four simultaneously). On the other hand, given that during the performance there can be more than four participants in total, Levinsky has to determine which (four) participants are in control at any given time. Therefore, at regular intervals (every 'n' seconds), the system calculates which participants are closer to each one of the loudspeakers (using the iBeacon signal) and then randomly selects the ones who will be in control in the next round. Levinsky then sends a message to the participants that have been chosen to let them know they are in control: the background colour of the screen in their apps changes to a different colour and the device vibrates. At the next round, if the same participant is not chosen, the background of the app goes back to its original colour.

Given the particular mapping between data control and synthesis, it was clear during the performance that specific gestures made by members of the audience with the movement of their arms and hands had the same or similar effect on the sound produced. Facial expressions of enjoyment and surprise were observed throughout the performance, as the audience explored the space and sound possibilities derived from their gestures. In addition, during the performance, participants often shared their enthusiasm amongst them. On the negative side, participants reported that the interval initially defined for individual control (5 seconds) was too small in order to achieve further engagement. Finally, conversations among the participants after the initial 5 minutes of performance indicated that from this moment onward the level of engagement decreased considerably. This particular aspect led the author to experiment with a game approach described in Performance 4, below.

Performance 3: Smart iSound

After these two initial experiments and the feedback from participants, a number of adjustments were introduced in the Levinsky system. Control time, initially set to 5 seconds, was increased to 15 seconds. A new sound generative machine was created in and entirely handled by Pure Data. The objective this time was to reduce the number of parameters and facilitate the perception of control. Essentially, participants do not have time to rehearse before the actual performance begins.

Other components of the performance remained fundamentally the same. Participants were invited to move and gesticulate around the performance space. The experimental and improvisational general approach was maintained but this time without the interference of pre-composed soundtracks.

As in the previous experience, control was handed to four participants simultaneously. This time, however, each one of them controlled a specific instrument in Pure Data. The timbre of these instruments was very distinct from each other as they used different synthesis techniques. The tap button on the Levinsky app was used to start and stop the synthesis of the controlled instruments and the gyroscope (pitch, roll and yaw) to modulate synthesis parameters. As expected, movement of arms and hands directly influenced the quality of the sound synthesis. For instance, it became possible to define and predict what type of sound the acceleration/deceleration of the smartphones would produce in large/short, slow/quick movements.

A new element was introduced, however, as the author decided to use a music keyboard to choose the fundamental frequencies of each one of the synthesizers in Pure Date during the performance. The objective was to have a sort of 'harmonic anchor' and use it to give shape (an overall structure) to the music experience.

The aim and initial assumption of the design described above was that, given a reduced number of variables, the participants would be able to have a better understanding of the mappings between control and sound generation and, consequently, achieve a better interaction during the performance. Better interaction in this case would be understood as the attainment of cohesiveness, and the spontaneous practice of interactive protocols (imitation, continuation, etc.).

These assumptions were not entirely fulfilled, however. Participants reported that, effectively, correspondences between control and sound production were easily understood. The general impression, though, was that they were more engaged with the technology itself (trying, for instance, to explore different movements during the performance) than with interacting with each other. In addition, apart from the fact that staying next to a beacon determined the choice of who would be in control in the next round, keeping in constant movement during the performance did not something they were particularly keen to do.

3. Further experiments: gameplay

Once again, lessons learned in the previous performances helped to shape the following experiments, named "Sound Games", presented at the International Computer Music Conference – ICMC (Utrecht, 2016) and the Off the Lip Conference (Plymouth, 2016). This time, a new form of interaction was implemented in order to provide a more engaging experience for the participants. The idea of a gameplay occurred naturally as most of the elements of what would look like a 'sound game' were already present. These included the concept of 'participants' (that could be easily translated into 'gamers') and the gameplay itself: a narrative derived from the execution of specific tasks. These tasks were part of the instructions given to the gamers and included the action of grabbing an instrument and moving around the performance space. Game rules, however, had to be more explicitly defined as well as a score system.

Performances 4 and 5: The Sound Games

Sound Games, therefore, became the evolution (and to some extent, a summary) of all the previous systems and performances. This time, in order to foster competitiveness, an additional element was introduced: the distinction between (and use of) personal (or standard) and leading (or special) virtual instruments.

The gameplay involves the following steps: initially (when connecting to the system), all gamers are given a personal instrument. There is a fixed number of personal instruments available which are allocated sequentially every time a new player connects to the game. Therefore, if there are more players than the number of available personal instruments, some participants will end up playing the same instrument. There is also a limited number of leading instruments (four), each one corresponding to one iBeacon placed in the game space (previously named 'performance space'). In other words, one specific iBeacon corresponds to one specific leading instrument.

Both leading and personal instruments are virtual synths pre-programmed in Logic. The sounds of the leading instruments are 'brighter' and 'punchier' as opposed to the personal instruments, which tend to be 'softer' and 'duller'. Instruments are played by tapping on the app's central button. The pitch played by the instruments (both leading and personal) is defined by the 'current note set' (a set of integers corresponding to MIDI note numbers), that continuously change as the game progresses. This idea derives from the 'harmonic anchor' defined in the above-mentioned previous experiment.

With regard to note sets, two possibilities are implemented. The first one is an 'automatic mode' whereby the note set changes every 'n' seconds according to a pre-programmed sequence (of note sets). The second is a 'manual mode' whereby who manages the game (typically the author) plays a MIDI controller, the input of which defines the note set. Choosing between these two options is done through the server's user interface

In terms of mapping, every time players tap on the app's central button, the next note in the current note set is played. If players move the smartphone on the roll axis of the gyroscope, the note changes octave. Moving the smartphone to the right increases the octave and, to the left, decreases the octave. The pitch axis of the gyroscope controls the volume of the instrument. An upward movement increases the volume and a downward movement decreases it.

Only one player plays a leading instrument at a time. Every 'n' seconds the dice is rolled (through a random number generator) and the leading instrument is given to participants who are next to each one of the beacons. When the player is not performing a 'leading instrument', he/she will get back his/her 'personal instrument'. When the player of a leading instrument is chosen, the background of the app changes colour and the smartphone vibrates. When the player goes back to the personal instrument, the background also goes back to its previous colour.

At the start of the game, the instruction is to go grab a leading instrument. This instruction holds a strong (unintentional) parallel with the location-based game Pokémon Go (Niantic 2016) which was launched around the same time the Sound Games was being implemented.

In terms of scoring, playing leading instruments gives more points than playing personal instruments. Playing different leading instruments in a row gives more points than playing the same leading instrument all the time. Playing 'in sync' (with other gamers) also gives more points (than playing 'out of sync') as the system tries to detect similar gestures amongst gamers within a defined time window.

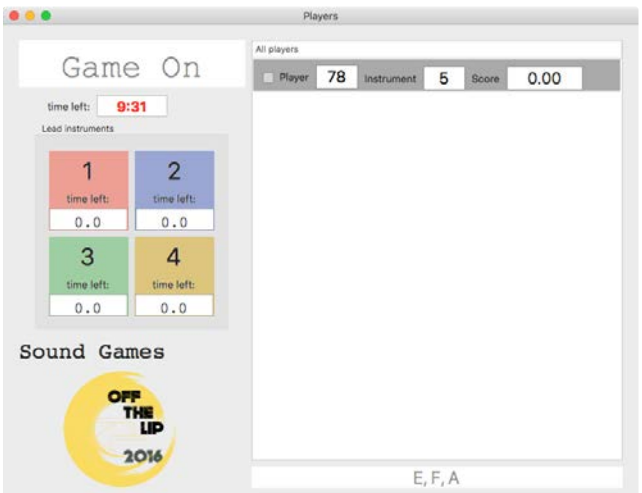

Figure 5: Scoreboard of the Sound Games

Figure 5 shows the state of development of the Sound Games scoreboard at the time of the performance at the Off the Lip Conference. On the right-hand side, the figure shows a list of players in descending order of points. For the purposes of this paper, the figure shows only one. The row informs, in addition to the score, the ID of the player and the number of the instrument the player is currently performing. At the bottom, the letters "E, F, A" indicate the 'current note set'. On the left-hand side the interface shows the amount of time left to end the game and the amount of time left for each one of the leading instruments.

4. Merging realities

The last experiment reported in this paper was run in the context of the Vox Populi performance (Gimenes 2017b). The general design of direct control supported by the Levinsky system was preserved from previous experiments. This time, however, participants were allowed to control some parameters of an artificial intelligence-based interactive computer music system previously developed by the Author, the Interactive Music Environments - iMe (Gimenes 2017a).

The full description of the iMe system, already reported in the literature, is outside the scope of this paper. In a nutshell, iMe is a comprehensive platform (or central hub) designed to afford a number of experiments with computational creativity. In its origins, the system focused primarily on the study of music evolution through computational modelling (Gimenes 2008a). It also aimed at exploring "interactive musical systems as a method to model and achieve the transmission of musical influence in artificial worlds and between humans and machines" ... and ... "to experiment with artificial and alternative developmental musical routes in order to observe the evolution of musical styles" (*ibid*.). This approach was called the 'ontomemetical model' due to the emphasis on a developmental path ('ontogenesis') and a 'memetic' basis of music evolution (Dawkins 1976).

The iMe system is highly inspired by the real world and incorporates a number of components that tries to mirror human perceptive and cognitive abilities (learning and creative models), as well as social interaction, life cycle (reproduction) and decision making, among others.

For the purposes of this paper, in order to give an appropriate context to the experiment introduced in this section, it suffices to mention that the iMe system models interactivity from the perspective of artificial agents that interact with each other within the artificial environment where they live as well as the with outside (real) world via offline music data and real-time (MIDI) input. Broadly speaking, interaction is achieved via the execution of a series of musical tasks (listening to, composing, improvising music, etc.).

Soon after the implementation of iMe's initial prototype, it became clear that this system would also be efficient for real time interactive improvisational sessions (Gimenes 2008b). Learning and generative algorithms initially implemented to deal with MIDI files were re-implemented in order to process real time input from a connected MIDI controller. Learning and generative algorithms of these two 'modes', however, are essentially the same.

In addition, the longer simulation sessions required for the exploration of music evolution scenarios, were reduced, in the case of real-time performances, to one single iteration (time cycle). The rationale of the system is that at every time cycle, agents evaluate (decide) what they do (which music task to perform) next. In the case of a real-time interactive performance, this feature is not necessary as just one single cycle is used to execute one task. As a result, operating under this 'mode', iMe behaves in a similar way as other real time interactive music systems such as the Continuator (Pachet 2002), OMax (Assayag 2018), GenJam (Biles 2013), and CIM (Brown, Gifford et al. 2016), among others.

Performance 6: Vox Populi

In summary, after its first public real-time appearance at the Peninsula Arts Contemporary Music Festival (Gimenes 2008b) the iMe system continued to be used in other public artistic and academic events (Gimenes 2015, Gimenes 2016a, Gimenes 2018). More recently, iMe and Vox Populi became the catalysts that allowed the integration of the two above-mentioned lines of research into an artificial intelligence-based performance with audience participation.

Vox Populi is essentially a free improvisation between a human and an artificial agent, both playing an electronic keyboard. At the beginning of the performance (or in the system's jargon, the 'simulation') the memory of the agent is empty. After a few seconds listening to the human player, the agent starts the musical dialogue, replicating some of the human's musical structures, modifying others, etc. The learning and generative algorithms of this system are introduced in (Gimenes and Miranda 2011).

Audience participation in this context happens by means of giving the participants the ability to control some of the agent's performance parameters. This is a radical change in relation to the original iMe system, which is essentially autonomous. Here, instead of trying to learn every single musical feature played by the human musician, the agent obeys to some decisions taken in real time by the audience.

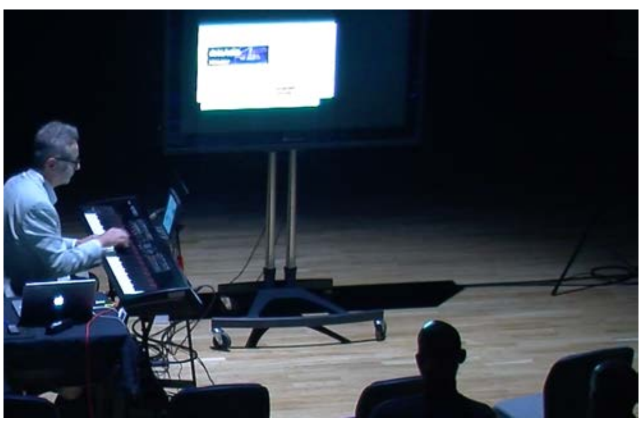

Figure 6: Author pays the Vox Populi with artificial agent

As in previous experiments, the Levinsky system provides the physical control to the audience. This time no one walks in the performance space (see Figure 6) as the audience holds their smartphones in their seats. There is only one leading participant at any given time who is randomly chosen from the members of audience connected to the system. Participants take turns every 'n' seconds in this role.

The octave the agent plays is determined by the 'leading participant' in the audience. A movement to the right on the roll axis of the smartphone's gyroscope increases the octave; to the left, decreases it. A movement upwards on the pitch axis increases the volume; downwards decreases it. Finally, a tap on the central button makes the agent play; another tap silences the agent.

Figure 7 shows a window that was implemented to provide a simple visual feedback for the audience to follow during the performance. Essentially, the window allows other participants (who are not controlling the performance) to know about the decisions of the leading participant. It shows the ID of the leading participant, whether participation is enabled, the octave and the volume the agent is playing at, and whether the agent is playing or not.

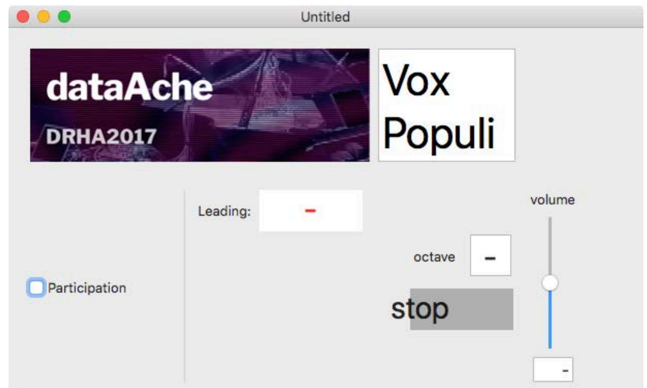

Figure 7: Vox Populi feedback window

5. Conclusion

This paper presents an overview of the performance-based research led by the author in the area of collective decision-making and musical metacreative performances with audience participation. One of the general aims of this investigation, exploratory in its nature, is to achieve a better understanding and integration with audiences in contemporary music.

In each one of these experiments, specific design approaches to a number of features were tested, from compositional ideas to technical implementations of control data, sound production and the connections between them. Voting systems and direct control were tested in a number of contexts, including algorithmic music generation, direct mapping and artificially-based interactive systems.

Under a critical perspective, considering the author's personal judgement in addition to discussions with and feedback from participants, these approaches were progressively evaluated and, eventually, some demonstrated to be more applicable to specific circumstances than others. In summary, there is often a sensitive trade-off between complexity of design and engagement. Straightforward solutions such as voting systems are relatively easy to grasp but (at least in the context of the above-mentioned experiments) fall short in terms of engagement. Direct control seem to provide better engagement but, on its own, not sufficiently to keep it for longer periods of time. A game approach to interaction can provide an additional element to a more prolonged engagement but, again, on its own, not enough to ensure stylistic consistency. These initial conclusions represent an important step towards an upcoming series of experiments and performances aiming at systematically testing metacreative performances with audience participation.

Given the above, the benefits of the reported approaches to the community of musical metacreative performers are numerous. Firstly, it is noticeable the enthusiasm with which audience adheres to the idea of using their personal mobile devices to participate in music performances. Secondly, it represents an alternative to the (very often) arduous implementation of specific artificial intelligence-based features in metacreative systems. It is true that sharing control of specific performance parameters with the audience removes precision or autonomy from the perspective of the system designer. On the other, however, it can potentially increase complexity and, hence, interest. Instead of embedding intelligence in the system, one should consider 'borrowing intelligence' from the audience.

Acknowledgements

The author thanks Mr. Frank Loesche for implementing the Android version of the Levinsky app.

References

Agres, K., J. Forth and G. A. Wiggins. 2016. "Evaluation of musical creativity and musical metacreation systems." Computers in Entertainment (CIE) 14(3): 3.
Assayag, G. 2018. "The OMax Project Page." Retrieved 11/03/2018, from https://goo.gl/KvE6mS
Biles, J. A. 2013. Performing with Technology: Lessons Learned from the GenJam Project. MUME 2013 Workshop.
Brown, A. R., T. Gifford and B. Voltz. 2016. "Stimulating Creative Partnerships in Human-Agent Musical Interaction." Computers in Entertainment (CIE) 14(2): 5.
Burtner, M., S. Kemper and D. Topper. 2012. "Network socio-synthesis and emergence in nomads." Organised Sound 17(1): 45-55.
Dawkins, R. 1976. The Selfish Gene. Oxford, Oxford University Press.
Essl, G. and M. Rohs. 2007. ShaMus: A Sensor-Based Integrated Mobile Phone Instrument. International Computer Music Conference.
Essl, G. and M. Rohs. 2009. "Interactivity for Mobile Music-Making." Organised Sound 14(02): 197-207.
Freeman, J. and M. Godfrey. 2010. "Creative collaboration between audiences and musicians in Flock." Digital Creativity 21(2): 85-99.
Gimenes, M. 2008a. An Approach to Machine Development of Musical Ontogeny. Doctor of Philosophy, University of Plymouth.
Gimenes, M. 2008b. "Stella by Starlight." Plymouth. University of Plymouth. Interactive music performance. Available on https://youtu.be/40CC89YLbsE
Gimenes, M. 2015. Artificial Worlds and the Simulation of Music Evolution. Study Day on Computer Simulation of Musical Creativity. University of Huddersfield.
Gimenes, M. 2016a. Hands On Computer Simulation of Music Evolution. 1st Conference on Computer Simulation of Musical Creativity. University of Huddersfield.
Gimenes, M. 2016b. "Levinsky Music (iOS app)." Apple App Store. Apple Inc. Available on https://goo.gl/kggCBd
Gimenes, M. 2016c. "Sherwell (iOS app)." Apple App Store. Apple. Available on https://goo.gl/Tve4Hw
Gimenes, M. 2017a. "Interactive Music Environments - iMe (for macOS) v2.0." Computer Music Lab.
Gimenes, M. 2017b. "Vox Populi (computer based interactive music performance)." Retrieved 11/03/2018, from https://marcelogimenes.blog/2017/09/06/vox-populi/
Gimenes, M. 2018. "Extending self." Plymouth. University of Plymouth. Interactive Performance.
Gimenes, M., P.-E. Largeron and E. R. Miranda. 2016. Frontiers: Expanding Musical Imagination with Audience Participation. New Interfaces for Musical Expression, Brisbane, Australia.
Gimenes, M. and E. Miranda. 2011. Emergent Worldviews: An Ontomemetic Approach to Musical Intelligence. A-Life for Music: On Music and Computer Models of Living Systems. E. Miranda. Middleton, Wisconsin, A-R Editions.
Hajdu, G. 2005. "Quintet. net: An environment for composing and performing music on the Internet." Leonardo 38(1): 23-30.
Hayes, K., M. Barthet, Y. Wu, L. Zhang and N. Bryan-Kinns. 2016. A participatory live music performance with the Open Symphony system. Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems, ACM.
Miller, G. 2012. "The smartphone psychology manifesto." Perspectives on Psychological Science 7(3): 221-237.
Niantic. 2016. "Pokémon Go." Available on https://www.pokemongo.com/
Ouzounian, G. 2008. Sound art and spatial practices: situating sound installation art since 1958, University of California, San Diego.
Pachet, F. 2002. The continuator: Musical interaction with style. International Computer Music Conference, Gothenburg, Sweden
Paul, C. 2003. Digital art, Thames & Hudson London.
Wang, G., J. Oh, S. Salazar and R. Hamilton. 2011. World Stage: A crowdsourcing paradigm for social/mobile music. International Computer Music Conference.