Electric Sheep: Designing Improvised Musical Play

Alistair McLean
Faculty of Fine Arts & Music
University of Melbourne
alistair.mclean@unimelb.edu.au

Robert Jarvis
Exertion Games Lab
RMIT University, Australia
bob@exertiongameslab.org

CHI PLAY '18 Extended Abstracts, October 28–31, 2018, Melbourne, VIC, Australia

ACM 978-1-4503-5968-9/18/10☐$15.00
https://doi.org/10.1145/3270316.3272057

Abstract

This paper outlines the development, creation and initial presentation of a computer program designed to direct human performers in improvised musical performance. The work, titled "Electric Sheep", is grounded in models of play-based improvisation of the late 20th century, focused around American composer John Zorn's "game-pieces" of the 1980s. It seeks to overcome some of the technical limitations of previous game-pieces whilst also providing a functioning example of player-computer interaction (PCI) in improvised music practice. Utilising an iterative rehearsal and development process we were able to isolate and highlight the importance of non-verbal and non-musical communication between improvising musicians and offer suggestions for incorporating this kind of feedback into future systems. Through this work, we will highlight the value of exploring the intersection of PCI and musical play as a valuable method of forming insight into rich PCI interactions.

CSS concepts

• **Human-centered computing → Ubiquitous and mobile computing design and evaluation methods** • *Human-centered computing → Interaction design*

Author Keywords

Improvisation; music; musical play; performance

Figure 1. John Zorn Cobra, North Sea Jazz, Rotterdam / NL, 2009' by William Veerbeek available at https://www.flickr.com/photos/william_veerbeek/3713911785 under a Creative Commons Attribution-NonCommercial 2.0. Full terms at https://creativecommons.org/licenses/by-nc/2.0/.

Introduction

The use of computing in the production of music is widespread, and has been formative in many of the stylistic, technical and artistic developments of the later 20th and early 21st centuries [7]. The dominant paradigm throughout this has been the use of computing to assist human action: to automate processes, enable faster and more intensive editing, and in the creation of sound via synthesis [7].

We find that comparatively little focus has been given to the agency of computing, and exploring aspects of creation, creativity and aesthetics generated by computers within musical practice. Recent developments in machine learning and deep learning have demonstrated potential for seemingly artistic practice generated by computers, including in musical applications [5, 15, 18]. However, these have largely focused on autonomous sound synthesis or on non-real time composition (often based on established notation or musical data languages) [5, 12, 14, 15, 19, 20, 22, 24].

Inspiration: Zorn's *Cobra*

One of the most significant experimental music improvisation works of the 20th century is American composer and saxophonist John Zorn's 1984 game-piece *Cobra* [3, 25] (fig. 1). Game-pieces are playful yet formal structures for improvisation often reminiscent of board games, involving rules designed to challenge conventional musical practice. *Cobra* is a complex game-piece, based loosely on the rules of a tactical military simulation game of the same name released in 1977 by games company TSR, itself based on a historical battle in WWII, featuring eight full pages of complicated rules and allowable moves.

This framework was intended to generate novel and strange interactions between performers as a way to explore not only experimental music practice, but also human interaction and the psychology of groups and group performance.

In *Cobra,* an ensemble of unspecified size (although often between 10-20 musicians) works in collaboration with a prompter who delivers commands via a set of cue-cards (fig. 1). Whilst the prompter delivers these commands, the commands are chosen by individual members of the ensemble, who suggest these to the prompter via a series of physical gestures, often utilising a specified number of fingers held in front of the performer's face.

These instructions suggest actions or modifiers for the performers to interpret, and are mostly focused on structural aspects of music making: who is or is not playing, when to start, stop or change an action, and when the piece begins or ends. They do not specify any melodic, harmonic or thematic information, and allow for great freedom in an individual musician's decision making, despite overlaying a strict form and direction to the piece.

Making *Sheep* with a computer

Whilst *Cobra* succeeds in creating a framework that encourages musicians to enter new situations and perform novel actions, the requirement for a human prompter can lead to some fundamental technical limitations, chiefly that it is difficult for multiple commands to occur at the same time. Similarly, the need for commands to be displayed on static, pre-written cards limits the number of possible actions, as well as requiring extensive memorisation of rules and

Figure 2. The first live performance of *Electric Sheep*

Figure 3. The first live performance of *Electric Sheep.*

symbols prior to the performance. Furthermore, as a performative action, the detailed and complicated rules often led to audiences being unable to follow the direction of gameplay, which can serve to mask the rationale for the musical and non-musical decisions being made by the performers.

These limitations led us to investigate how software could be utilised in the creation of "game-pieces" for improvising ensembles, and how these works could both overcome limitations of previous game-pieces and demonstrate a model for human-computer co-creation in experimental music performance.

We anticipated that by removing the limitations that human prompting places on a performance it would be possible to realise a game-piece that was both easier to learn due to a smaller number of rules and more intuitive commands, while providing a similar or increased variety of experimentation and diverse musical actions to those of other game-pieces. Historically, experimentation and diversity have been highly sought after in this field of musical practice [17], so increasing their occurrence would be an indicator of an artistically successful work.

The initial concept was to develop a computer program that could deliver individual commands to an ensemble of eight improvising musicians utilising visual cues. Considerable research and development has been made into systems that improvise audio alongside performing musicians with varying degrees of success [5, 6, 12, 18, 22, 24], but there has been little investigation into computer led music performance involving human performers. Interestingly, a core concept in much previous research is the idea of player-computer interaction (PCI) occurring with both parties conceived of as equals, often discussed as collaboration [4, 5, 6, 18]. Similarly, there is a large focus on the technical aspects of creating successful human-computer interactions, including pitch, rhythmic and gestural analysis, modification and response [4, 5, 6, 8, 11, 18, 22], but little research on how group behaviour or human responses are affected in music performance involving PCI.

In this instance we made a conscious effort to explore the direction of musicians by a computer as opposed to a 'computer-as-musician' model, to more closely follow the tradition of Zorn's *Cobra* and related 'game-pieces' which rely on a director or conductor, and in the process investigate a model of PCI in music practice where humans assist a computer program to realise its "artistic wishes". This creates a model where the interest and drama, to both audience and performers, comes from the idiosyncratically "unhuman" commands that are given and how the group must interpret them to realise them in a successful and aesthetically pleasing or interesting manner. This leads to situations where performers must not only interact with each other, but also with a dynamic non-human director, realising Zorn's desire for game pieces to "*deal with form, not with content, with relationships, not with sound*." [25].

Development

The instigator for the development of *Electric Sheep* was a performance entitled *no new noise* as part of the 2017 Melbourne Festival (fig. 2, fig. 3). Presented by the Australian Creative Music Ensemble (ACME), *no new noise* explored the use of artificial intelligence and creative technology within contemporary music

Figure 4. The player commands for *Electric Sheep.*

practise. It asked a number of composers to respond to the idea of artificial intelligence and generative technology replacing the roles of composers and musicians.

We sought to design a model where a computer program could work with an ensemble of human performers to collaborate on musical performance. To do so, we established a number of criteria:

- There must be interaction or dialogue between the human and computer participant/participants
- The model must allow human performers a degree of creative freedom, to ensure there is true collaboration and not merely human performance of computer creativity
- The work must be artistically engaging for the performers, and not an example of "tech for tech's sake"

These criteria led us to a number of early decisions as to the shape of the work.

Firstly, we decided that the computer's output should mostly be decided in real-time, and not pre-composed. This allows for the potential of input data to affect its decisions, as well as ensuring it is operating on the same "time-scale" as the human performers. Secondly, due to the exploratory nature of the project and the desire for real time creativity and problem solving, we established that we would work with musicians with an extended practice in improvisation and experimental music. Lastly, and most significantly in the design of this project, we identified that the majority of current musical HCI and PCI projects utilised sound as a shared language, however historically many musical pieces have been communicated or developed via other means, including verbally, graphically and via instruction [1, 13, 21]. We identified this as an interesting and previously unexplored application of PCI in a musical context.

Initially we focused on the design of a set of commands and iconography to control the piece, focusing on a number of design goals.

- Rules must be simple to learn and follow in a performance situation.
- Rules must be open ended enough to offer a wide range of performance outcomes, but not so open-ended as to offer no useful instruction.
- Iconography should be designed so a non-musical audience can follow.

Following these goals, we designed five player commands with accompanying icons (fig. 4). In addition to these commands, a player's command icon could animate by spinning around its axis directing the player to immediately alter what they are playing. Whilst there are a number of similarities between the commands in *Electric Sheep* and those of prior pieces of game-music, care was taken to ensure that commands were included based on their usefulness in gameplay and not based on precedent.

Continuing with our design goal of providing a visual interface that was understandable for both performers and audiences, we took care to design a simple but visually exciting visual display. Animation was utilised

Figure 5. The group commands for *Electric Sheep.*

to draw viewer's attention to cues as they were triggered, and visual interest was generated through smooth animated transitions between commands, which also gave the musicians a period of time to prepare for the upcoming command.

Initial concepts involved providing individualised commands to each performer via their own screen, but through discussion it was decided that a single visual output that could be shown to all performers and the audience would provide the more engaging interface, as well as ensuring a system that was more easily constructed and ported between different performance venues (fig. 6).

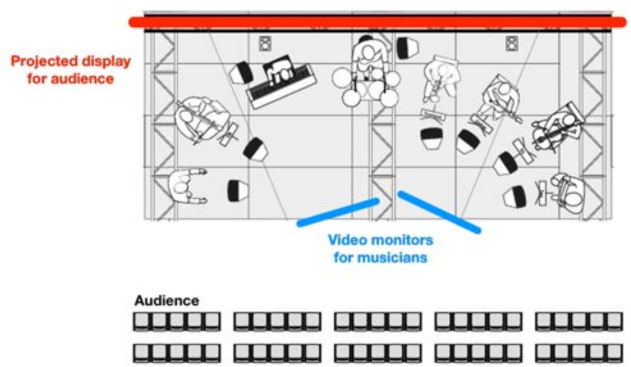

Figure 6: Stage layout for Electric Sheep showing positions of screens for musicians and audience.

Working with a standard 16:9 ratio display, the screen was divided into a three-by-three grid, with each performer allocated a single square. Each musician's cell displayed an iconic representation of their instrument, and commands were overlaid on top (fig. 7).

Player 1	Player 2	Player 3
Player 4	**Group Commands**	Player 5
Player 6	Player 7	Player 8

Figure 7: Screen layout for Electric Sheep, showing individual performer squares as well as area for group commands

The centre square was reserved for "group commands", these are commands that applied to all performers, and acted as further modifiers for their actions. These commands were intended to provide moments of cohesion, where all performers would align to play in the same tonality, mood or tempo, and also to inform performers of the end of a piece (fig. 5).

Technology

We identified a weighted probabilistic model as the most suitable option for generating commands as it would allow us to easily modify to the likelihood of each action occurring, and therefore provide an avenue to shape the decisions made by the program towards a

Figure 8. An example of the visual output of the final program. It shows four individual commands, one with an ostinato modifier, plus a group mood command.

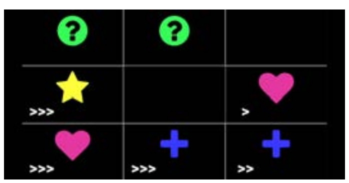

Figure 9. A screenshot of the early prototype used to begin the iterative development process.

desired aesthetic. To provide an overall shape to each performance, we implemented an extra parameter referred to as the "density curve". This was a user-drawn curve created immediately before a performance that controlled the overall duration of the piece, as well as providing approximate control over how many players should have an active command at any point in time.

We broke the system into separate front and back-ends to decouple distinct tasks. The backend was written in Max 7, a visual programming environment for art and multimedia, with most of the logic written in JavaScript. This backend was responsible for generating all commands and had its own graphical interface for adjusting probability weights, drawing the density curve and controlling the start and stop of the piece (fig. 10). This enabled us to adjust the decision's output by the system as the piece was running, allowing for very rapid development during rehearsals as well as providing an option for human intervention during a performance if required.

The frontend was written in TouchDesigner, a visual programming language focussed on interactive 3D graphics (fig.8). Communication between the frontend and backend was accomplished using Open Sound Control (OSC) allowing either end to be developed independently and making the system easily modifiable.

Refinement

Once we had developed an initial working prototype (fig. 9), a period of iterative development with human performers commenced. This period encompassed five sessions over a ten-week period, involving eight improvising musicians. In these sessions, feedback was sought from the performers as to their experience and interaction with the program, but also as to whether the commands were suitable for generating performances that were perceived as being "musical".

Initial responses suggested that the commands were generated too often, and had too high a level of instruction, which both hindered performers' choices and often led to them feeling overwhelmed.

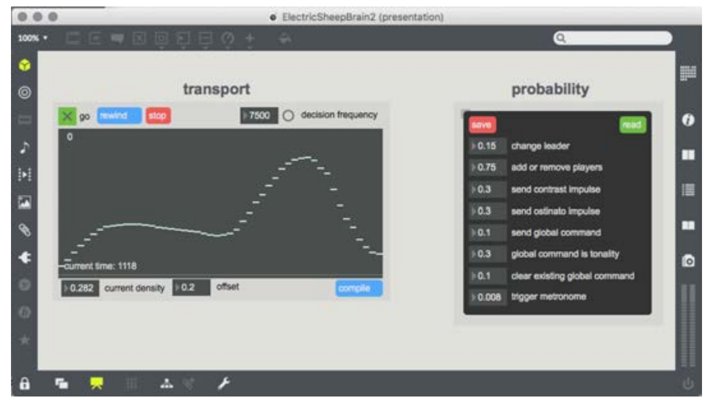

Figure 10: The Max 7 interface providing parameterized user control to the decision-making model and the density curve.

During this development phase we discovered our initial assumptions about appropriate inputs for musical decision making were incorrect. We had assumed that a more sophisticated version of this application could incorporate audio input from each musician, analysing pitch and amplitude over time and using this information to inform upcoming commands (for example see [12]). Interestingly, whilst our desire in future iterations of the piece was to involve audio analysis, the development process suggested that monitoring and responding to interpersonal feedback was more important than musical analysis when determining whether a command was perceived of as appropriate or "musical" by the performers.

Similarly, when engaged in performance, the group's collective sense of humour would often steer the piece one way or another. This emerged as musical jokes such as small phrases being repeated or reinterpreted, external musical references such as familiar choruses or themes, and musical and gestural references to interactions that occurred outside of the current performance (such as at a previous rehearsal). As a human participant in musical play, this factor is intuitive, and recognition is often based on non-verbal cues or shared cultural knowledge, factors that rely on very different inputs than those of audio signals and MIDI values that are most readily associated with an interactive musical system.

Aside from humour, other non-verbal communication also emerged as a repeated oversight in the software's decision making. Improvising musicians rely heavily on eye contact, conventional and intuitive gestural cues [17]. It is notable that all performers involved in this project were trained musicians and all were very aware of the importance of communication during improvised musical play, yet this particular oversight was not initially apparent.

Finally, equity of performance opportunity arose numerous times, where the weighted probabilities did not adequately stop certain performers from dominating iterations of the piece, whilst other performers did not get a chance to play. Whilst this could be seen as a valid outcome of a probabilistic model, and indeed having certain performers featured more at times than others is often musically interesting, it all also led to situations where performers felt the need to apologise after performances for dominating the game, and it was felt that the program was not delivering "fairness" in its commands.

Viewed together, these findings may stand as the most significant take away for the project – whilst audio and musical inputs are important considerations in improvised music, when designing a system for computer mediated musical play, subtle, interpersonal communication plays an equally fundamental role and needs to be accounted for.

Future Work

Electric Sheep as it stands lacks real-time feedback and is unaware of the actions of the human performers. Although we think an ideal implementation of the system would incorporate sophisticated feedback and contextual awareness, we found the simplicity of the system led us more directly to understanding the subtleties of the game-piece. As such, refraining from closing this loop in at least initial future iterations could offer yet more insight into PCI and musical play.

We are investigating redeveloping the system as a mobile or web application for wider distribution. By incorporating some form of continuous or prompted user response we may be able to gather enough human assessment of model-generated musical decisions to feed into a reinforcement learning process. Assessing the musicality of such abstract commands presents challenges but the success of our own reflective development process leads us to think this approach may be fruitful.

We also intend to incorporate basic audio analysis. Simply implementing a measure of amplitude could enable additional game mechanisms and goals such as directing the ensemble to play very loud for a set amount of time or rapidly change dynamics.

Enabling *Electric Sheep* to have a reliable sense of social, cultural and musical context, and then responding appropriately and with some amount of originality could also be an avenue for future work.

Although the output of this and similar systems will likely not reach the standards of performers or expectations of their audiences in the foreseeable future, it can provide contrasting examples that expose the gaps in our understanding of PCI and forms of creative play.

Conclusion

This project has demonstrated a model for the creation and development of a computer program to direct human performers in improvised musical performance. Whilst grounded in a tradition of improvised musical 'game-pieces', through utilising current technology it improves upon prior models of human-led performance, offering a richer and more diverse range of possibilities, and presents a working example of PCI in improvised music performance.

Through description of the iterative development process, which involved refinement of the program in collaboration with human performers, the common use of shared cultural knowledge, humour and non-verbal communication amongst human performers of improvised music is highlighted, and demonstrated to be an important future area of research in the development of PCI in interactive musical situations.

Acknowledgements

We would like to thank the musicians of ACME, Matthew Hoy and the Melbourne Festival, The Substation and the Alan C. Rose Memorial Scholarship for their assistance in the development and initial public performance of Electric Sheep. We would also like to thank Rohit Khot, Florian 'Floyd' Mueller and The Exertion Games Lab for their feedback and support, and Dr. Robert Vincs and the Jazz & Improvisation department at the University of Melbourne.

References

1. Brett Boutwell. 2012. Morton Feldman's Graphic Notation: Projections and Trajectories. *Journal of the Society for American Music* 6, 4(2012), 457ăĂŞ482. DOI: http://dx.doi.org/10.1017/S1752196312000363

2. John Bowers and Sten Olof Hellström. 2000. Simple interfaces to complex sound in improvised music. In *CHI'00 Extended abstracts on Human factors in Computing Systems*. ACM, 125–126.

3. John Brackett. 2010. Some Notes on John Zorn's Cobra. *American Music* 28, 1 (2010), 44–75.

4. Federica Bressan, Tim Vets, and M Leman. 2017. A multimodal interactive installation for collaborative music making: from preservation to enhanced user design. In *Proceedings of the European Society for Cognitive Sciences of*

Music (ESCOM) Conference. 23–26.

5. Benjamin Leigh Carey. 2016. _*derivations and the performer-developer: co-evolving digital artifacts and human-machine performance practices*. Ph.D. Dissertation.

6. Marcelo Cicconet, Mason Bretan, and Gil Weinberg. 2012. Visual Cues-based Anticipation for Percussionist-robot Interaction. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI '12)*. ACM, New York, NY, USA, 117–118. DOI: http://dx.doi.org/10.1145/2157689.2157713

7. Brian Eno. 2013. The Studio as a Compositional Tool. In *Audio culture: Readings in modern music*, Christoph Cox and Daniel Warner (Eds.). Bloomsbury, New York, 115-126.

8. Alexandre François, Elaine Chew, and Dennis Thurmond. 2007. Visual feedback in performer-machine interaction for musical improvisation. In *Proceedings of the 7th international conference on New interfaces for musical expression*. ACM, 277–280.

9. Alexandre François, Isaac Schankler, and Elaine Chew. 2013. Mimi4x: An interactive audio-visual installation for high-level structural improvisation. 6 (01 2013), 138–151.

10. Garth Griffin and Robert Jacob. 2013. Priming Creativity Through Improvisation on an Adaptive Musical Instrument. In *Proceedings of the 9th ACM Conference on Creativity & Cognition (C&C '13)*. ACM, New York, NY, USA, 146–155. DOI: http://dx.doi.org/10.1145/2466627.2466630

11. Tomoko Hashida, Takeshi Naemura, and Takao Sato. 2007. A System for Improvisational Musical Expression Based on Player's Sense of Tempo. In *Proceedings of the 7th International Conference on New Interfaces for Musical Expression (NIME '07)*. ACM, New York, NY, USA, 407–408. DOI: http://dx.doi.org/10.1145/1279740.1279839

12. William Hsu. 2007. Design issues in interaction modeling for free improvisation. In *Proceedings of the 7th international conference on New interfaces for musical expression*. ACM, 367–370.

13. Evarts John. 1968. The New Musical Notation: Graphic Art?. *Leonardo*. 1, 4: 405-412

14. Laewoo Leo Kang, Steven J Jackson, and Phoebe Sengers. 2018. Intermodulation: Improvisation and Collaborative Art Practice for HCI. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 160.

15. Tang-Chun Li. 1999. Who or what is making the music: music creation in a machine age. In *Proceedings of the 3rd conference on Creativity & cognition*. ACM, 57–62.

16. Duri Long, Hannah Guthrie, and Brian Magerko. 2018. Don't Steal My Balloons: Designing for Musical Adult-child Ludic Engagement. In *Proceedings of the 17th ACM Conference on Interaction Design and Children (IDC '18)*. ACM, New York, NY, USA, 657–662. DOI: http://dx.doi.org/10.1145/3202185.3210795

17. Ingrid T. Monson. 1996. *Saying something : jazz improvisation and interaction.* Chicago : University of Chicago Press, 84, 115, 124-125,

18. Fabio Morreale and Antonella De Angeli. 2016. Collaborating with an Autonomous Agent to Generate Affective Music. *Comput. Entertain.* 14, 3, Article 5 (Dec. 2016), 21 pages. DOI: http://dx.doi.org/10.1145/2967508

19. Tomás Pérez-García, José M. Iñesta, Pedro J. Ponce de León, and Antonio Pertusa. 2011. A Multimodal Music Transcription Prototype: First Steps in an Interactive Prototype Development. In *Proceedings of the 13th International Conference on Multimodal Interfaces (ICMI '11)*. ACM, New York, NY, USA, 315–318. DOI: http://dx.doi.org/10.1145/2070481.2070540

20. Isaac Schankler, Alexandre RJ François, and Elaine Chew. User-directed improvisation and play with meta-compositions for Mimi4x.

21. Dieter Schnebel. 1973. Composition 1960: La Monte Young. *Musique En Jeu*. 11, 6-16.

22. Belinda Thom. 2000. BoB: An Interactive Improvisational Music Companion. In *Proceedings of the Fourth International Conference on Autonomous Agents (AGENTS '00)*. ACM, New York, NY, USA, 309–316. DOI: http://dx.doi.org/10.1145/336595.337510

23. Konstantinos Vasilakos and others. 2016. *An evaluation of digital interfaces for music composition and improvisation*. Ph.D. Dissertation. Keele University.

24. Gil Weinberg and Scott Driscoll. 2007. The interactive robotic percussionist: new developments in form, mechanics, perception and interaction design. In *Proceedings of the ACM/IEEE international conference on Human-robot interaction*. ACM, 97–104.

25. John Zorn. 2013. The game pieces. In *Audio culture: Readings in modern music*, Christoph Cox and Daniel Warner (Eds.). Bloomsbury, New York, 196-200.