

Evaluating Human-Robot Interaction with Embodied Creative Systems

Rob Saunders

Design Lab
University of Sydney
NSW 2006 Australia
rob.saunders@sydney.edu.au

Emma Chee

Small Multiples
Surry Hills, Sydney
NSW 2010 Australia
emma@small.mu

Petra Gemeinboeck

College of Fine Art
University of NSW
NSW 2021 Australia
petra@unsw.edu.au

Abstract

As we develop interactive systems involving computational models of creativity, issues around our interaction with these systems will become increasingly important. In particular, the interaction between human and computational creators presents an unusual and ambiguous power relation for those familiar with typical human-computer interaction. These issues may be particularly pronounced with embodied artificial creative systems, e.g., involving groups of mobile robots, where humans and computational creators share the same physical environment and enter into social and cultural exchanges. This paper presents a first attempt to examine these issues of human-robot interaction through a series of controlled experiments with a small group of mobile robots capable of composing, performing and listening to simple songs produced either by other robots or by humans.

Introduction

Creativity is often defined as the generation of novel and valuable ideas, whether expressed as concepts, theories, literature, music, dance, sculpture, painting or any other medium of expression (Boden 2010). But creativity, whether or not it is computational, doesn't occur in a vacuum, it is a situated activity that is connected with cultural, social, personal and physical contexts that determine the nature of novelty and value against which creativity is assessed. The world offers opportunities, as well as presenting constraints: human creativity has evolved to exploit the former and overcome the latter, and in doing both, the structure of creative processes emerge (Pickering 2005).

There are three major motivations underlying the research of developing computational creativity: (1) to construct artificial entities capable of human-level creativity; (2) to better understand and formulate an understanding of creativity; and, (3) to develop tools to support human creative acts (Pease and Colton 2011). The development of artificial creative systems is driven by a desire to understand creativity as interacting systems of individuals, social groups and cultures (Saunders and Gero 2002).

The implementation of artificial creative systems using autonomous robots imposes constraints upon the hardware and software used. These constraints focus the development process on the most important aspects of the computational

model to support an embodied and situated form of creativity. At the same time, embodiment provides opportunities for agents to experience the emergence of effects beyond the computational limits that they must work within. Following an embodied cognition stance, the environment may be used to offload internal representation (Clark 1996) and allow agents to take advantage of properties of the physical environment that would be difficult or impossible to simulate computationally, thereby expanding the behavioural range of the agents (Brooks 1990).

Interactions between human and artificial creators within a shared context places constraints on the design of the human-robot interaction but provides opportunities for the transfer of cultural knowledge through the sharing of artefacts. Embodiment allows computational agents to be creative in environments that humans can intuitively understand. As Penny (1997) describes, embodied cultural agents, whose function is self reflexive, engage the public in a consideration of the nature of agency itself. In the context of the study of computational creativity, this provides an opportunity for engaging a broad audience in the questions raised by models of artificial creative systems.

The 'Curious Whispers' project (Saunders et al. 2010), investigates the interaction between human and artificial agents within creative systems. This paper focuses on the challenge of designing one-to-one and one-to-many interactions within a creative system consisting of humans and robots and provides a suitable method for examining these interactions. In particular, the research presented in this paper explores how humans interacting with an artificial creative system construe the agency of the robots and how the embodiment of simple creative agents may prolong the production of potentially interesting artefacts through the interaction of human and artificial agents. The research adopts methods from interaction design to study the interactions between participants and the robots in open-ended sessions.

Background

Gordon Pask's early experiments with electromechanical cybernetic systems provide an interesting historical precedent for the development of computational creativity (Haque 2007). Through the development of "conversational machines" Pask explored the emergence of unique interaction protocols between the machine and musicians. MusiColour,

seen in Figure 1, was constructed by Gordon Pask and Robin McKinnon-Wood in 1953. It was a performance system comprising of coloured lights that illuminated in conjunction with audio input from a human performer.

But MusiColour did more than transcode sound into light, it manipulated its coloured light outputs such that it became a co-performer with the musician, creating a unique (though non-random) output with every iteration (Glanville 1996). The sequence of the outputs not only depended on the frequencies and rhythms but also repetition: if a rhythm became too predictable then MusiColour would enter a state of ‘boredom’ and seek more stimulating rhythms by producing and stimulating improvisation. As such, it has been argued that MusiColour acted more like a jazz co-performer might when ‘jamming’ with other band members (Haque 2007).

The area of musical improvisation has since provided a number of examples of creative systems that model social interactions within creative activities, e.g., GenJam (Biles 1994), MahaDeviBot (Kapur et al. 2009). The recent development of Shimon (Hoffman and Weinberg 2010) provides a nice example of the importance of modelling social interactions alongside the musical performance.




Figure 1: MusiColour: light display (left) and processing unit (right) (Glanville 1996).

‘Performative Ecologies: Dancers’ by Ruairi Glynn is a conversational environment, involving human and robotic agents in a dialogue using simple gestural forms (Glynn 2008). The Dancers in the installation are robots suspended in space by threads and capable of performing ‘gestures’ through twisting movements. The fitness of gestures is evaluated as a function of audience attention, independently determined by each robot through face tracking. Audience members can directly participate in the evolution by manipulating the robots, twisting them to record a new gesture. Successful gestures, i.e., those observed to attract an audience, are shared between the robots over a wireless network.

The robotic installation ‘Zwischenräume’ employs embodied curious agents that transform their environment through playful exploration and intervention (Gemeinboeck and Saunders 2011). A small group of robots is embedded in the walls of a gallery space, they investigate their wall habitat and, motivated to learn, use their motorised hammer to in-



Figure 2: Performative Ecologies: Dancers (Glynn 2008)

roduce changes to the wall and thus novel elements to study. As the wall is increasingly fragmented and broken down, the embodied agents discover, study and respond to human audiences in the gallery space. Unlike the social models embodied in MusiColour and Performative Ecologies, the social interactions in Zwischenräume focus on those between the robots. Audience members still play a significant role in their exploration of the world but in Zwischenräume visitors are considered complex elements of the environment.

In ‘The New Artist’, Straschnoy (2008) explored issues of what robots making art for robots could be like. In a series of interviews, the engineers involved in the development of The New Artist expressed different interpretations of the meaning and purpose of such a system. Some questioned the validity of the enterprise, arguing that there is no reason to construct robots to make art for other robots. While others considered it to be part of a natural progression in creative development “We started out with human art for humans, then we can think about machine art for humans, or human art for machines. But will we reach a point where there’s machine art for machines, and humans don’t even understand what they are doing or why they even like it.” — Interview with Jeff Schneider, Associate Research Professor, Robotics Institute, Carnegie Mellon (Straschnoy 2008)

The following section describes the current implementation of the ‘Curious Whispers’, an embodied artificial creative system. The implemented system is much simpler than those described above, i.e., the robots employ a very simple generative system to produce short note sequences, but it provides a useful platform for the exploration of interaction design issues that arise with the development of autonomous creative systems involving multiple artificial agents.

Implementation

The current implementation of Curious Whispers (version 2.0) uses a small group of mobile robots equipped with speakers, microphones and a movable plastic hood, see Figure 3. Each robot is capable of generating simple songs, evaluating the novelty and value of a song, and performing those songs that they determine to be ‘interesting’ to

other members of the society – including human participants. Each robot listens to the performances of others and if it values a song attempts to compose a variation. Closing their plastic hood, allows a robot to rehearse songs using the same hardware and software that they use to analyse the songs of other robots, removing the need for simulation.




Figure 3: The implemented mobile robots and 3-button synthesiser.

A simple 3-button synthesiser allows participants to play songs that the robots can recognise and if a robot considers a participant’s songs to be interesting it will adopt them. Using this simple interface, humans are free to introduce domain knowledge, e.g., fragments of well-known songs, into the collective memory of the robot society. For more information on the technical details of the implementation see Chee (2011).

Methodology

To investigate the interactions between robots and human participants we adopted a methodology from interaction design and employed a ‘technology probe’. Technology probes combine methods for collecting qualitative information about user interaction, the field-testing of technology, and exploring design requirements. A well-designed technology probe should balance these different disciplinary influences (Hutchinson et al. 2003). A probe should be technically simple and flexible with respect to possible use: it is not a prototype but a tool to explore design possibilities and, as such, should be open-ended and explicitly co-adaptive (Mackay 1990). The probe used in this research involved three observational studies exploring different aspects of the human-robot interaction with the embodied creative system.

The observational studies were conducted with different arrangements of robots and human participants, allowing us to observe how interaction patterns and user assessments of the system changed in each configuration. Each session was video recorded and at the end of each session the participants were interviewed using a series of open-ended questions. The interview was based on a similar one developed by Bernsen and Dybkjær (2005) in their study of conversational agents. Employing a ‘post-think-aloud’ method at the end of each session the participants were first asked to describe their experiences interacting with the robot. A similar

method was used in the evaluation of the Sonic City project (Gaye, Mazé, and Holmquist 2003). The video recordings were transcribed and interaction events noted on a timeline. The ‘post-think-aloud’ reports were correlated with events in the video recordings where possible.

Six participants were observed in the studies. The participants came from a variety of backgrounds and included 2 interaction designers, 2 engineers, 1 linguist, and 1 animator. All participants were involved in the 1:1 (1 human, 1 robot) observation study. Two participants (Participant 5 and 8) went on to be part of the 1:3 (1 human, 3 robots) observation study, the other four (Participant 6, 7, 9 and 10) were involved in the 2:3 (2 humans, 3 robots) observation study.

1:1 Interaction Observation Study The purpose of the first study was to observe the participants behaviour whilst interacting with a single robot. Each participant was given a 3-button synthesiser to communicate with the robot and allowed to interact for as long as they wished, i.e., no time limit was given.

1:3 Interaction Observation Study The second observational study involved each participant interacting with the group of 3 robots to examine how participants interacted with multiple creative agents at the same time and how the participants were influenced by the interactions between robots. This study involved 2 participants, both participants had previously completed the first observation study.

2:3 Interaction Observation Study The third observational study involved pairs of participants interacting with the system of 3 working robots. This study allowed for the participants to not only interact and observe the working system but to also interact with each other to share their experiences. This study involved 4 participants working in two groups of two. The 4 participants were chosen from those who completed the 1:1 study but were not involved in the 1:3 observation study.

Results

This section presents a brief summary of the observational studies, a more detail account can be found in Chee (2011).

1:1 Interaction The 1:1 interaction task allowed the participants to form individual theories on how single robots reacted to them, most learned that the robots did not respond to individual notes but sequences of them. Participants spent between 2 and 4 minutes interacting with the robot, much of that time was spent experimenting to determine how the robot reacted to different inputs: “[I] first tried to see how it would react, pressed a single button and then tried a sequence of notes” (Participant 6). Several of the participants learned to adopt a turn-taking behaviour with the robots, e.g., “when it started to play I stopped to watch, I only tried to play when it stopped” (Participant 5). Some of the participants interpreted the opening and closing of the hood as a cue for when they could play a song for the robot to learn, as Participant 9 commented: “I played a noise and it took that song and closed up and was like ‘alright I’m gonna think of

something better'. It sounded like it was repeating what I did but like a bit different. Like it was working out what I'd done." Most of the participants assumed the role of teacher and attempted to get the robot to repeat a simple song. But in the case of Participant 8 the roles were reversed as the participant began copying the songs played by the robot.

1:3 Interaction For the 1:3 interaction studies the group of robots were placed on a table in a quiet location, as shown in Figure 4. The participants interacted with the group of robots for approximately 5 minutes. Both participants already knew the robots were responsive to them from the 1:1 study, but they found it difficult to determine which robot they were interacting with: "you knew you could interact but you were not really aware of the reaction as a group" (Participant 5). The participants noticed that the robots were different: "the green robots song was slightly different to blue and purple" (Participant 5); and, that they exhibited social behaviour amongst themselves: "Noticed they didn't rely just on the [synthesiser], the 3 of them were communicating. I thought they sang in a certain order as one started and the others would reply" (Participant 8). Both participants came to realise that system would continue to evolve new songs without their input and spent time towards the end of their sessions observing the group behaviour.




Figure 4: An example of the interaction in the 1:3 study.

2:3 Interaction Working together the participants in the third study quickly arrived at the conclusion that they needed to take turns in order to interact with the robots. Participant 6 saw that the robots moved towards Participant 7 and asked to be given one of the robots. Participant 7 replied "No, they have to go to you on their own", suggesting that Participant 7 recognised that the robots could not be commanded. Later, the participants became competitive in their attempts to attract the robots away from each other. As the participants shared observations about the system, they explored the transference of songs. By observing the interactions between Participant 7 and the robots, Participant 6 was able to determine that the robots responded to songs of exactly 8 notes and that the robots would repeat the song 3 times while it learned. At one point Participant 9 commented: "...when I pressed it like this 'beep beep beep beep' it went 'beep beep boop beep' so it was like changing what I played". These observations suggest that over time the participants were able

to build relatively accurate 'mental' models of the processes of the robotic agents.



Figure 5: An example of the interaction in the 2:3 study.

Discussion

Unlike traditional interactive systems that react to human participants (Dezeuze 2010), the individual agents within artificial creative systems are continuously engaged in social interactions: the robots in our study would continue to interact and share songs without the intervention of the participants. While initially confusing, participants discovered through extended observation and interaction that they could inject songs into the society by teaching them to a single robot. Participants sometimes also assumed the role of learner and copied the songs of the robots and consequently adopted an interaction strategy more like that of a peer.

The autonomous nature of the embodied creative system runs counter to typical expectations of human-robot interactions; making interacting with a group of robots is significantly more difficult than interacting with one. The preliminary results presented here suggest that simple social policies in artificial creative systems, e.g., the turn-taking behaviour, coupled with cues that indicate state, e.g., closing the hood while practicing and composing songs, allow for conversational interactions to emerge over time.

Conclusion

The development of embodied creative system offers significant opportunities and challenges for researchers in computational creativity. This paper has presented a possible approach for the study of interaction design issues surrounding the development of artificial creative systems.

The Curious Whispers project explores the possibility of developing artificial creative systems that are open to these types of peer-to-peer interactions through the construction of a 'common ground' based on the expression and perception of artefacts. The research presented has shown that even a simple robotic platform can be designed to exploit its physical embodiment as well as its social situation, using easily obtained components.

The implemented system, while simple in terms of the computational ability of the agents, has provided a useful

platform for studying interactions between humans and artificial creative systems. The technical limitations of the robotic platform place an emphasis on the important role that communication plays in the evolution of creative systems, even with the restricted notion of what constitutes a 'song' in this initial exploration. Above all, the technology probe methodology used in our observational studies have illustrated the usefulness of implementing simple policies in artificial creative systems to allow human participants to adapt to the unusual interaction model.

Acknowledgements

The research reported in this paper was supported as part of the Bachelor of Design Computing Honours programme in the Faculty of Architecture, Design and Planning at the University of Sydney.

References

- Bernsen, N., and Dybkjær, L. 2005. User interview-based progress evaluation of two successive conversational agent prototypes. In Maybury, M.; Stock, O.; and Wahlster, W., eds., *Intelligent Technologies for Interactive Entertainment*, volume 3814. Springer Berlin / Heidelberg. 220–224.
- Biles, J. A. 1994. Genjam: A genetic algorithm for generating jazz solos. In *Proceedings of the International Computer Music Conference*, 131–137.
- Boden, M. A. 2010. *Creativity and Art: Three Roads to Surprise*. Oxford: Oxford University Press.
- Brooks, R. 1990. Elephants don't play chess. *Robotics and Autonomous Systems* 6:3–15.
- Chee, E. 2011. Curious Whispers 2.0: Human-robot interaction with an embodied creative system. Honours Thesis, University of Sydney, Australia. Available online at http://emmachee.com/Thesis/Emma_Chee_Thesis_2011.pdf.
- Clark, A. 1996. *Being There: Putting Brain, Body, and World Together Again*. Cambridge, MA, USA: MIT Press.
- Dezeuze, A. 2010. *The 'do-it-yourself' artwork: Participation from the Fluxus of new media*. Manchester: Manchester University Press.
- Gaye, L.; Mazé, R.; and Holmquist, L. E. 2003. Sonic city: the urban environment as a musical interface. In *Proceedings of the 2003 Conference on New interfaces For Musical Expression*, 109–115.
- Gemeinboeck, P., and Saunders, R. 2011. Zwischenräume: The machine as voyeur. In *Proceedings of the First International Conference on Transdisciplinary Imaging at the Intersections between Art, Science and Culture*, 62–70.
- Glanville, R. 1996. Robin mckinnon-wood and gordon pask: A lifelong conversation. *Journal of Cybernetics and Human Learning* 3(4).
- Glynn, R. 2008. Performative Ecologies: Dancers, <http://www.ruairglynn.co.uk/portfolio/performative-ecologies/>.
- Haque, U. 2007. The architectural relevance of Gordon Pask. In *4d Social: Interactive Design Environments*. Wiley & Sons.
- Hoffman, G., and Weinberg, G. 2010. Shimon: an interactive improvisational robotic marimba player. In *CHI '10 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '10, 3097–3102. New York, NY, USA: ACM.
- Hutchinson, H.; Mackay, W.; Westerlund, B.; Bederson, B. B.; Druin, A.; Plaisant, C.; Beaudouin-Lafon, M.; Conversy, S.; Evans, H.; Hansen, H.; Roussel, N.; and Eiderbäck, B. 2003. Technology probes: inspiring design for and with families. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '03, 17–24. New York, NY, USA: ACM.
- Kapur, A.; Eigenfeldt, A.; Bahn, C.; and Schloss, W. A. 2009. Collaborative composition for musical robots. *Journal of Science and Technology of the Arts* 1(1):48–52.
- Mackay, W. 1990. *Users and Customizable Software: A Co-Adaptive Phenomenon*. Ph.D. Dissertation, Massachusetts Institute of Technology.
- Pease, A., and Colton, S. 2011. On impact and evaluation in computational creativity: A discussion of the turing test and an alternative proposal. In *Proceedings of the AISB symposium on AI and Philosophy 2011*.
- Penny, S. 1997. Embodied cultural agents: At the intersection of art, robotics, and cognitive science. In *Socially Intelligent Agents: Papers from the AAI Fall Symposium*, 103–105. AAAI Press.
- Pickering, J. 2005. Embodiment, constraint and the creative use of technology. In *Freedom and Constraint in the Creative Process in Digital Fine Art*.
- Saunders, R., and Gero, J. S. 2002. How to study artificial creativity. In *Proceedings of Creativity and Cognition 4*.
- Saunders, R.; Gemeinboeck, P.; Lombard, A.; Bourke, D.; and Kocabali, B. 2010. Curious whispers: An embodied artificial creative system. In *International Conference on Computational Creativity 2010, 7–9 January 2010*.
- Straschnoy, A. 2008. The New Artist, <http://www.the-new-artist.info/>.