Curious Whispers: An Embodied Artificial Creative System

Rob Saunders[1], Petra Gemeinboeck[2], Adrian Lombard[1], Dan Bourke[1], and
Baki Kocabali[1]

[1] Faculty of Architecture, Design and Planning, University of Sydney, Australia
[2] College of Fine Arts, University of New South Wales, Australia

Abstract. Creativity, whether or not it is computational, doesn't occur in a vacuum, it is a situated, embodied activity that is connected with cultural, social, personal and physical contexts. Artificial creative systems are computational models that attempt to capture personal, social and cultural aspects of human creativity. The physical embodiment of artificial creative systems presents significant challenges and opportunities. This paper introduces the "Curious Whispers" project, an attempt to embody an artificial creative system as a collection of autonomous mobile robots that communicate through simple "songs". The challenges of developing an autonomous robotic platform suitable for constructing artificial creative systems are discussed. We conclude by examining some of the opportunities of this embodied approach to computational creativity.

1 Introduction

Human creativity is situated within cultural, social and personal contexts. From a computational perspective this suggests that the processes involved in creativity should be open to the environment, other creative agents, and a history of creative works. Physical embodiment is an important aspect of human creativity that presents significant challenges and opportunities for the development of computational creativity. Katherine Hayles argues that embodiment is always contextual, enmeshed within the specifics of place, time, physiology and culture, which together compose enactment [?]. Following Pickering [?], creativity cannot be properly understood, or modelled, without an account of how it emerges from the encounter between the world and intrinsically active, exploratory and productively playful agents. 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.

Why is embodiment important for computational creativity? The enactment described by Hayles, emphasises creativity as a situated act, e.g., in personal histories, social relations and cultural identity. The computational study of situated cognition as proposed by Clancey [?] does not require physical embodiment, but many of the more successful examples of situated computational systems are robotic in nature. Perhaps this is because, despite every effort that a developer might make to maintain a separation, there is always the sense that agent and environment are of the same ‘type’ within the simulation and consequently that the agent is not truly situated within the environment.

Physical embodiment requires that agents deal with the material nature of the creative activity that they engage in—the importance of working with an external material in creative activity was highlighted by the work of Schön studying designers and the process he termed as reflection-in-action [**?**]. Schön’s reflection-in-action illustrates the utility of ideas from distributed cognition [**?**] in understanding the creative acts of designers, providing insights into the situated nature of creative cognitive process. Distributed cognition and reflection-in-action provide useful frameworks for designing artificial creative systems because they emphasise the relationship between the agent and its environment.

The implementation of autonomous robots imposes constraints upon the hardware and software that can be incorporated. These constraints focus the development process on the most important aspects of the computational model. At the same time, embodiment provides opportunities for agents to experience the emergence of effects beyond the computational limits that they must work within. Taking advantage of properties of the physical environment that would be difficult or impossible to simulate computationally, expands the behavioural range of the agents [**?**].

Finally, embodiment allows computational agents to be creative in environments that humans can intuitively understand. As Penny [**?**] 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.

Curious Whispers is a project to investigate the nature of embodiment in an artificial creative system and explore the potential of placing this artificial society within a human physical and social environment.

2 Background

In 1738, Jacques de Vaucanson exhibited his *Flute Player* automaton. In 1769, Baron Wolfgang von Kempelen presented to the public his chess playing *Mechanical Turk*; it was not until 1834, that an article appeared in Le Magazin Pittoresque revealing its inner workings and the man hidden within [**?**]. In developing these machines, both Vaucanson and von Kempelen engaged the public in philosophical questions about the nature of creativity and the possibilities of automation [**?**]. Our apparent fascination with the prospect of building machines that can exhibit creative behaviour continues today with the development of embodied agents as robots. Following Vaucanson, many of these robotic experiments are within the domain of music.

Ja’maa is a percussion ensemble for human and robotic players, including Haile a robotic drummer that listens to the drumming of human players and responds with its own improvisations [**?**]. Eigenfeldt has developed software-based multiagent systems to emulate improvised percussion ensembles of [?] and has embodied these agents within a robotic performer, *MahaDeviBot* [?].

DrawBots [?] and *Mbots* [?] are two examples of recent attempts to develop robots capable of exhibiting creative behaviour in the production of abstract drawings. Portraitist robots have been implemented [?,?] but, while these projects have overcome significant technical challenges, they have mostly neglected to examine issues associated with embodied creativity. Cagli et al. [?] proposes the study the behaviour of realistic drawing to focus on the physical aspects of the creative process. In particular, they focus on visuomotor coordination and present a control architecture based on computational models of eye movements, and the eye-hand coordination of expert draughtsmen.

For the development of computational models of creativity one of the key advantages of embodiment with a physical and social environment may be the access it brings to a cultural context beyond the confines of the computational elements. As Penny [?] observes in relation to his embodied cultural agents "viewers (necessarily) interpret the behavior of the robot in terms of their own life experience. [...] The machine is ascribed complexities which it does not possess. This observation emphasises the culturally situated nature of the interaction. The vast amount of what is construed to be the 'knowledge of the robot' is in fact located in the cultural environment, is projected upon the robot by the viewer and is in no way contained in the robot." In Penny's works, the robots are viewed within the context of their cultural environment but this has no impact upon intrinsic behaviour of the robots, having no access to the situation that the audience brings.

2.1 Curious Agents

Martindale [?] proposes that the search for novelty is a key motivation for individuals within creative societies. Curious agents embody a computational model of curiosity based on studies of humans and other animals, where curiosity is triggered by a perceived lack of knowledge about a situation and motivates behaviour to reduce uncertainty through exploration [?]. Unlike earlier models of creative processes that try to maximise some utility function, curious agents are motivated to discover something 'interesting' based on their previous experiences using an hedonic function, the Wundt curve (see Figure **??**).

Curious agents provide a useful foundation for developing embodied agents to engage in an artificial creative system because they have been shown to be useful in modelling autonomous creative behaviour and have been used to robots to promote life-long learning in novel environments. Schmidhuber [?] presents a model interest and curiosity, based on the compressibility of information, and introduced a distributed model of curiosity based on a pair of agents competing to surprise each other [?]. Saunders focused on the role of curiosity in creativity to develop computational models of creativity to search for novelty and interest in design [?]. In these models, the computation of interest and boredom are based on novelty detection, a technology that was originally developed to detect potential faults in processes where it is critical to stay within "normal" operating

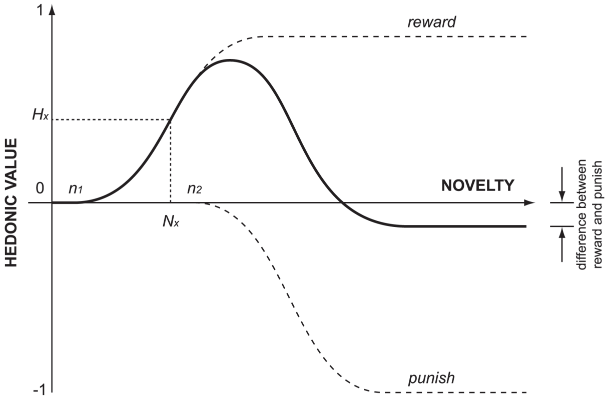

Fig. 1. The Wundt Curve: an example hedonic function for curious agents and robots.

limits. Unlike in monitoring applications, novelty is considered a desirable quality when modelling curiosity, and detected novelty is used as the basis for positive reinforcement of behaviour.

Research developing embodied curious agents has focussed on the utility of modelling curiosity as a motivation for learning about physical environments and social relations. Marsland et al. [**?**] introduced the idea of "neotaxis", movement based on perceived novelty, as a useful behaviour for autonomous robots to map physical spaces. Peters [**?**] presented the WRAITH algorithm as a layered architecture for building curious robots suitable for modelling creativity. Oudeyer and Kaplan [**?**] presents the use of curiosity to support the discovery of communication in social robots. Merrick [**?**] presents an architecture for curious, reconfigurable robots for creative play that can learn new behaviour in response to changes in their structure.

When computational models of curiosity are used as the model of motivation in intelligent environments a new kind of space emerges: a curious place [**?**]. Curious places are intelligent environments using curious agents to adapt to changing user behaviour and anticipate user demands. Curious places offer new opportunities for supporting and embodying creativity in the physical environment. In addition to supporting human activities, curious places work proactively to anticipate, identify and enact creative behaviour.

2.2 Artificial Creative Systems

The Domain Individual Field Interaction (DIFI) framework is a unified approach to studying human creativity that provides an integrated view of individual creativity within a social and cultural context [**?**]. According to this framework, a creative system has three interactive subsystems: domain, individual and field. A domain is an organised body of knowledge, including specialised languages, rules, and technologies. An individual is the generator of new works in a creative system, based on their knowledge of the domain. A field contains all individuals who can affect the content of a domain, e.g., creators, audiences, critics, and educators. The interactions between individuals, fields and domains form the basis of the creative process in the DIFI framework: individuals acquire knowledge from domains and propose new knowledge evaluated by the field; if the field accepts a proposed addition, it becomes part of the domain and available for use by other individuals.

Inspired by the DIFI model of creativity, Saunders and Gero used curious agents to develop *artificial creative systems*, composed of curious design agents capable of independently generating, evaluating, communicating and recording works [**?**]. Other distributed approaches to computationally modelling creativity include McCormack's "ecosystemic" approach, which recognises the importance of the environment and the agent's relationship with the environment as primary concerns for modelling creative activity [**?**].

3 Implementation

Building on the previous work of Saunders [**?**], we are currently developing the *Curious Whispers* project in an attempt to develop an artificial creative systems using embodied curious agents, i.e., curious robots. Inspired by the thought experiments of Braitenburg [**?**] we have implemented the robots as simple vehicles with the addition of a loudspeaker, a pair of microphones and sufficient processing units to determine their 'interest' in the sonic environment. The robot architecture has been developed as a set of function-specific modules: audio capture and processing, song categorisation and analysis, interest and boredom calculations, sound generation and output, and servo and motor control.

The *Curious Whispers* robots have been built on top of the Ardubot bare bones mobile robot platform developed by Sparkfun Electronics[3]. The Ardubot platform was designed as a minimal, low-cost platform for developing mobile robots using the Arduino[4] interface boards. The Ardubot platform is based around an oversized expansion board for the Arduino integrating a DC motor driver integrated-circuit (IC) and a pair of mounts for motors. Arduino and Atmel ATmega168 microcontrollers are used for this application due to their relatively fast operation speed (20 MIPS), ample memory (16kb flash, 1kb SRAM), flexibility, compatibility and affordability. The Arduino acts as the primary interface between the ATmega168 microcontroller and other components attached to the robot, e.g., the DC motor driver, sound generation chips, etc. This provides simple access for programming the ATmega168 and offers expandability through the use of "shields" that can be stacked on top of the Arduino to provide additional functionality.

A custom shield has been developed for the Arduino to provide the sound generation and sound capture and processing functions to the Arduino. To produce

---

[3] http://www.sparkfun.com/

[4] http://www.arduino.cc/

the audio signal to drive the loudspeakers each robot is equipped with an FM synthesis subsystem based around a Soundgin[5] audio processor. The Soundgin processor has two independent sound engines, each with three oscillators and a mixer, providing a large variety of possible sounds.

To allow the robots to move about their environment without damaging themselves, each robot has a pair of front-facing "whiskers" attached to a touch sensor, allowing the robot to stop and back away from obstacles encountered.

3.1 Audio Capture and Processing

Two audio signals are captured by small microphones mounted on lightweight movable arms. The ATmega168 microcontroller performs a 64-point, fixed-point Fast Fourier Transform (FFT) operation on each of the audio signals. Using this 64-point FFT, a sampling rate of 16kHz and a frequency resolution of 250Hz, we are able to achieve a Nyquist frequency of 8KHz, i.e., we have 32 frequency bands at 0Hz, 250Hz, 500Hz, 750Hz, 1kHz,...8kHz. This is a sufficient frequency range for our application, since we do not generate sounds above 8kHz. The onboard 16kb of memory can hold enough samples to perform two 64 point FFT calculations. Therefore the robots can monitor a stereo pair of signals enabling left-right interest differencing, suitable for driving neotaxis. The result of the FFT calculation is passed to the Arduino board for analysis and processing.

The Arduino board monitors a stream of serial data from the FFT calculation on the left and right audio channels. Each sample is represented as an integer value between 0 and 63 representing the most active frequency detected by the FFT calculation: values close to 0 represent bass sounds. When the dominant frequency detected by the FFT changes, in either the left or right audio channel, the values for both channels are appended to short-term memory. The values in the short-term memory represent the "song" the robot is hearing.

3.2 Novelty, Interest and Boredom

When short-term memory contains a total of eight frequencies, the values are packaged as a vector and presented to a small Self Organising Map (SOM) that serves as the robot's long-term memory [**?**]. Due to the limitations of the hardware platform, the SOM contains just 16 neurons, but this has proved sufficient for the task of categorising the eight-note songs that the robots are capable of producing. In contrast with typical applications of categorisation systems, the robots do not attempt to maintain a complete map of the space of all possible songs, rather each robot constructs a local map of recently experienced songs.

The novelty of a song is calculated as the shortest Euclidean distance between the vector representation of the song and all of the prototypes held in the SOM. To calculate the interest that the robot has in the current song a non-linear function, which approximates the Wundt curve as the sum of two sigmoids, is used to transform the novelty value. Consequently, a song stored in short-term

---

[5] http://oopic.com/soundgin/

memory that exactly matches an existing song in the SOM is not particularly interesting, and a song that is radically different to anything which the robot has previously experienced is also not very interesting. The most interesting songs for these robots will be songs which are similar but different to the songs recently experienced by the robot and held in the SOM.

Interest values calculated for the audio signals received from the left and right microphones are translated into movement such that interest value for the left channel will be converted into a speed for the right wheel, and vice versa. Figure **??** illustrates a scenario for neotaxis as implemented in our robots, where one of the robots, having analysed the songs of two other robots (A and B), moves in the direction of the robot that has produced the more interesting song.

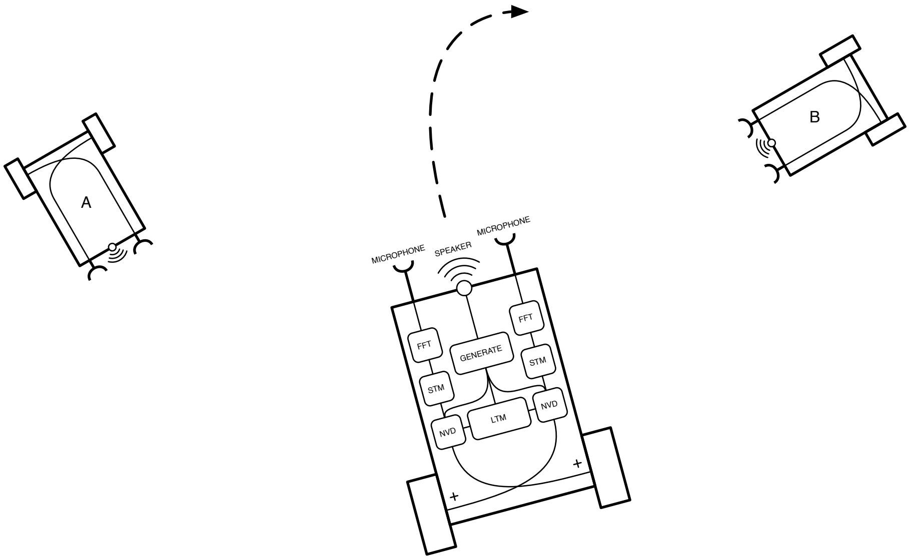

Fig. 2. The robots in the *Curious Whispers* project implement neotaxis, driving in the direction of the most interesting novelty. The architecture includes: audio analysis (FFT); short-term memory (STM); long-term memory (LTM); novelty detection (NVD); and, song generation (GEN).

In the absence of interesting songs a robot will become ‘bored’. Boredom is computationally modelled as a threshold on the long-term level of interest that the robot has had in recent songs. If the robot becomes bored, it changes from a listening mode to a generating mode.

3.3 Song Generation

Generative systems are often computationally expensive, both in terms of the process of generation and analysis. The limited computational resources available in the autonomous robots has required that generation of new works be handled differently from previous artificial creative systems. Firstly, a simple generative system has been implemented, which takes advantage of the long-term memory of stored patterns to generate similar-but-different songs. To generate a new song the agent either mutates a pattern randomly chosen from the prototypes stored in the SOM, or chooses two prototypes and combines them using an operation similar to crossover used in genetic algorithms. Secondly, the analysis of generated songs takes advantage of the embodied nature of the robots to reuse the analysis systems already present. In the generative mode the robot changes its physical configuration by moving its left and right microphones closer to its speaker and reduces the volume of the speaker. This reconfiguration allows the robot to listen exclusively to its own songs.

To bootstrap the system, all robots begin in this generative mode. Using the random vectors assigned to the prototypes held in the robot's long-term memory, each robot generates songs until it discovers one that is interesting enough to communicate to others.

4 Planned Experiments

Three robots are in the final stages of construction and a series of experiments are planned to evaluate the utility of our approach. In particular, the experiments will examine:

1. whether embodiment has significant benefits over simulation for the study of artificial creative systems; and,
2. how humans interacting with an artificial creative system construe the agency of the robots.

Comparing the behaviour of artificial creative systems is a difficult task. The behaviour of the system cannot be validated using the principles that underlie the approach, yet these principles are important indicators of creative behaviour. Behavioural diversity is a key factor in attaining creative behaviour, and one approach to evaluating creative behaviour is to quantify behavioural diversity. We will quantify the behavioural diversity of our embodied agents and of the artificial creative system as a whole, and compare these to our simulations of the same agents to gain insights into the effects of embodiment on the creative processes. The simulation of the artificial creative system uses as much of the code running on the robots as possible, interacting within a simulated environment.

Human audiences will encounter the artificial society and its evolving tunes within a the context of a gallery environment. This will allow them to share the same space with robots and to engage with their activities and relations from within. To study how embodiment affects the way humans construe the agency of the robots, visitors will have the opportunity to interact with the robotic system using an FM synthesiser, similar to the ones used by the robots. The goal is to encourage visitors to engage with the social creative process at work in the community of robots by playing simple tunes, allowing visitors to inject elements from their human cultural context into the artificial creative system.

5 Conclusion

This paper has described the design of *Curious Whispers*, a proof-of-concept implementation for an embodied artificial creative system. Unlike typical human-robot interactions, this project does not place the human in a privileged position, able to dictate what the robots should play. Instead the human enters the artificial creative system as an equal to the robots, who is required to produce songs of interest to the robots for them to be picked up and reworked within the system. *Curious Whispers* is a system open to human engagement, potentially allowing the agents to take advantage of the social and cultural contexts that visitors bring.

References

1. Hayles, N.K.: How we became posthuman: virtual bodies in cybernetics, literature, and informatics. University of Chicago Press, Chicago, Illinois (1999)
2. Pickering, J.: Embodiment, constraint and the creative use of technology. In: Freedom and Constraint in the Creative Process in Digital Fine Art: An AHRB Invited Workshop in Association with the 5th Conference on Creativity and Cognition, Goldsmiths College, London, United Kingdom (April 12–15 2005)
3. Clancey, W.J.: Situated Cognition: On Human Knowledge and Computer Representations. Cambridge University Press, Cambridge, England (1997)
4. Schön, D.A.: The reflective practitioner : how professionals think in action. Basic Books, New York, NY, USA (1983)
5. Hutchins, E.: Cognition in the Wild. MIT Press, Cambridge, MA (1995)
6. Brooks, R.: Elephants don't play chess. Robotics and Autonomous Systems **6** (1990) 3–15
7. Penny, S.: Embodied cultural agents: At the intersection of art, robotics, and cognitive science. In: Socially Intelligent Agents: Papers from the AAAI Fall Symposium, AAAI Press (1997) 103–105
8. Reichle, I.: Art in the Age of Technoscience: Genetic Engineering, Robotics, and Artificial Life in Contemporary Art. Springer Verlag (2008)
9. Wood, G.: Living Dolls: A Magical History of the Quest for Mechanical Life. Faber and Faber (2003)
10. Weinberg, G., Driscoll, S., Thatcher, T.: Ja'maa – a middle eastern percussion ensemble for human and robotic players. In: International Computer Music Conference, New Orleans (2006) 464–467
11. Eigenfeldt, A.: Drum circle: Intelligent agents in max/msp. In: Proceedings of the International Computer Music Conference, Copenhagen (2007)
12. Eigenfeldt, A., Kapur, A.: An agent-based system for robotic musical performance. In: Proceedings of New Interfaces for Musical Expression (NIME'08), Genova, Italy (2008)

13. Boden, M.A., Brown, P., Husbands, P., Gere, C.: Computational intelligence, creativity and cognition: a multidisciplinary investigation. http://www.informatics.sussex.ac.uk/research/groups/ccnr/research/creativity.html, accessed 2009 (2005)
14. Moura, L., Pereira, H.G.: Symbiotic art. IAC, Villeurbanne (2004)
15. Calinon, S., Epiney, J., Billard, A.: A humanoid robot drawing human portraits. In: Proceedings of 2005 5th IEEE-RAS International Conference on Humanoid Robots. (2005)
16. Robotlab: autoportrait. http://www.robotlab.de/auto/portrait_engl.htm, accessed 2009 (2002)
17. Cagli, R.C., Coraggio, P., Napoletano, P.: Drawbot: a bio-inspired robotic portraitist. Digital Creativity **18**(1) (2007) 24–33
18. Martindale, C.: The Clockwork Muse. Basic Books, New York, NY. (1990)
19. Berlyne, D.: Aesthetics and Psychobiology. Appleton-Century-Crofts, New York, NY (1971)
20. Schmidhuber, J.: A possibility for implementing curiosity and boredom in model-building neural controllers. In Meyer, J.A., Wilson, S.W., eds.: Proceedings of the International Conference on Simulation of Adaptive Behaviour: From Animals to Animats, Cambridge, MA, MIT Press (1991) 222–227
21. Schmidhuber, J.: Adaptive confidence and adaptive curiosity. Technical Report FKI-14991, Technische Universitat München, München, Germany (1991)
22. Saunders, R.: Curious Design Agents and Artificial Creativity. PhD thesis, University of Sydney, Sydney (2002)
23. Marsland, S., Nehmzow, U., Shapiro, J.: Novelty detection for robot neotaxis. In: International Symposium On Neural Computation (NC'2000). (2000) 554—559
24. Peters, M.W.: Towards artificial forms of intelligence, creativity and surprise. In: Proceedings of the Twentieth Annual Conference of the Cognitive Science Society, Erlbaum (1998) 836–841
25. Oudeyer, P.Y., Kaplan, F.: Discovering communication. Connection Science **18**(2) (2006) 189–206
26. Merrick, K.: Designing toys that come alive: Curious robots for creative play. In: The Seventh International Conference on Entertainment Computing (ICEC 2008), Carnegie Mellon University, Springer (September 25-27 2008) 149–154
27. Maher, M.L., Merrick, K., Macindoe, O.: Intrinsically motivated intelligent sensed environments. In: The Thirteenth International Workshop of the European Group for Intelligent Computing in Engineering, Ascona, Switzerland (2006) 455–475
28. Feldman, D.H., Csikszentmihalyi, M., Gardner, H.: Changing the World: A Framework for the Study of Creativity. Praeger Publishers, Westport, CT (1994)
29. Saunders, R., Gero, J.S.: How to study artificial creativity. In: Proceedings of Creativity and Cognition 4. (2002)
30. McCormack, J.: Artificial ecosystems for creative discovery. In et. al., T.D., ed.: Genetic and Evolutionary Computation Conference (GECCO 2007). Volume 1., London, ACM, New York (7-11 July 2007 2007) 301–307
31. Braitenberg, V.: Vehicles: Experiments in synthetic psychology. MIT Press, Cambridge, MA (1984)
32. Kohonen, T.: Self-Organizing Maps. Springer–Verlag, Berlin (1995)