Supporting Creativity Using Curious Agents

Rob Saunders University of Sydney, Australia
rob@arch.usyd.edu.au

ABSTRACT

This paper briefly examines the evidence for a link between creativity and curiosity and argues that one way to better support creative processes is to develop computational systems that incorporate models of curiosity. Empirical evidence about curious behaviour provides some useful guidance in the development of simple computational models of explorative curiosity. Existing agent-based models of curiosity are presented and discussed with reference to models of individual creativity. A scenario for supporting creativity is presented where a curious agent models a user's preference for novelty. Limitations this approach are discussed and a ways to overcome them suggested.

Author Keywords

creativity, curiosity, computational modelling, machine learning, human-computer interaction

ACM Classification Keywords

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—intelligent agents, multiagent systems

INTRODUCTION

One potentially important way to computationally support creativity is to develop computational systems that stimulate curiosity in users and support the exploration of possibilities. The question of supporting curiosity in the service of creativity is more often considered in education and professional development. For example, intrinsic rewards, such as those that result from the discovery of new knowledge through curious behaviour, have been shown to be more important to creative individuals than extrinsic rewards that are often used to motivate them [1].

The impact of curiosity on the design of user interfaces for conventional creativity tools can be seen in the desire to facilitate exploration [13]. The computational modelling of curiosity has focussed primarily on its application to autonomous agents, e.g., robotics, machine learning, and data mining, and the simulation of curious behaviour. The question of how such models can be used to support creative activity has not been explored sufficiently.

Sternberg and Lubart [19] provide the following definition of creativity:

> Creativity is the ability to produce work that is both novel and appropriate.

Much of the research in computational creativity and creative computing has focused on supporting the production of work that is appropriate, e.g., the production of complex musical scores. In such systems the production of novelty is often left to chance, either through the use of stochastic algorithms or by relying on the user to introduce 'noise' into the system. Some computational processes are well suited to introducing novelty, e.g., evolutionary systems, and can be used to develop tools to support creative exploration, but until such systems can filter out works that are uninteresting the use of such tools will remain laborious [11].

The production of interesting novelty is a quite different proposition from the generation of something merely new, and leads to an examination of what is interesting and the study of human behaviour in the search of interesting novelty, i.e., curiosity. In the remainder of this paper, we will look at some foundations for understanding creativity and curiosity, some previous work on computationally modelling curiosity, and discuss what role computational models of curiosity can play in the support of creativity.

MODELS OF CURIOSITY

Empirical research suggests a strong connection between novelty and aesthetic preference in various creative fields including literature, art, architecture and music [10, 6]. These reports support the argument that curiosity plays an important role in creative activities. In "The Clockwork Muse", Martindale [10] presented an extensive investigation into the role that individual novelty-seeking behaviour played in literature, music, visual arts and architecture. He concluded that the search for novelty exerts a significant force on the development of styles.

Novelty and Surprise

The subjective evaluation of novelty is quite different from that of appropriateness: a creative product is likely to remain appropriate for some time but it loses its novelty as soon as it is experienced. This makes the application of fixed heuristics to determine novelty inappropriate, as was demonstrated by the continued reliance of early discovery systems, e.g., AM and EURISKO, on human assistance.

Berlyne conducted extensive research into the effects of perceiving novelty on the behaviour of humans and animals [4] and its role in the judgement of aesthetics [5]. An important distinction between different types of novelty is between novelty that is due to atypical stimuli and novelty due to a stimuli being uncommon. Atypical stimuli are unlike previous experiences, their novelty lies in the differences between it and previous experiences. Uncommon stimuli are familiar from previous experiences but are rarely experienced or have not been experienced for some time.

Some concepts are naturally related to novelty, in some cases so much so that they are commonly considered synonymous. A surprising stimulus is not just atypical or uncommon; it is a stimulus that disagrees with one or more expectation [4]. Surprise involves anticipation of an experience that is not fulfilled by the actual experience that follows. The degree of surprise depends upon the confidence put in the expectation and the degree to which the expectation is confounded. When a stimulus sets up an expectation that is not satisfied within the same experience, it is said to be incongruous rather than surprising.

Expectations obviously play an important role in the perception of surprising and incongruous stimuli. Expectations can be formed in three main ways. The most common way is through the repeated experience of combinations of stimuli or sequences of events. Classical conditioning then leads the perception of a stimulus, X, to evoke an expectation of a response that usually accompanies or follows it, Y. The strength of the expectation will reflect the reliability with which Y can be predicted given X, i.e., $p(Y|X)$. Secondly, expectations of something can be formed on the advice of a reliable information source. Finally, expectations can be formed through a reasoning process.

Interest

To describe the response to arousing stimuli, Berlyne [5] coined the term "hedonic value". Interest can be considered a special case of hedonic value associated with heightened states of learning. Berlyne's model of the relationship between arousal, e.g., novelty, and hedonic value, e.g., interest, uses a non-linear function called the Wundt curve sketched in Figure 1.

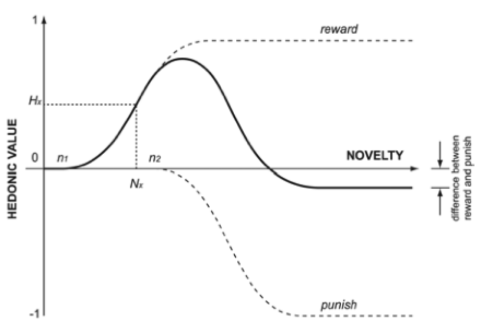

Figure 1. The Wundt Curve: a hedonic function used to calculate interest. The hedonic function is shown as a solid line, the reward and punishment sigmoidal curves summed to form the hedonic function are shown dashed.

The inverted U-shape of the Wundt curve means that the most interesting experiences are those that are similar-yet-different to those that have been experienced previously. Berlyne supported his model using empirical evidence gathered from studies of aesthetic preference and creative thinking. Berlyne proposed that his model of arousal is also the basis of behaviour commonly referred to as 'curiosity'.

Curiosity

Berlyne [5] defines curiosity as a form of motivation that promotes exploratory behaviour to learn more about a source of uncertainty, such as a novel stimulus, with the goal of acquiring sufficient knowledge to reduce the uncertainty. He presented two types of motivations for exploration, diversive and specific, in keeping with his model of hedonic reward. In diversive exploration, an organism is under-stimulated and seeks arousal from the environment. In specific exploration, an organism is over-stimulated and seeks to reduce its arousal by reducing the novelty of the situation and its associated collative variables, in particular, uncertainty. Whether motivated by diversive or specific needs the goal of such exploratory behaviour is to gain knowledge and this typifies curious behaviour.

In creative activities, we can think of curiosity as the motivation to explore possibilities to relieve the uncertainty that accompanies an incomplete understanding of the conceptual space. In divergent curiosity, creations that are similar-yet-different to those that have been experienced before will be preferred.

Curious Agents

Schmidhuber demonstrated with a number of autonomous agents engaged in self-directed learning in complex environments that curiosity can be very effective in guiding the exploration of dynamic environments [16, 17]. Schmidhuber implemented curious agents using neural controllers and reinforcement learning with intrinsic rewards generated in response to an agent improving its model of the world [16].

Marsland et al. developed robots that display orienting and locomotive exploratory behaviour motivated by curiosity that they call neotaxis [9]. The habituated mechanisms used by the robots to detect novelty respond to how recently an input was last experienced. The robots have been shown to detect novel features of an environment that aid the efficient exploration of complex environments that are initially unknown.

Gomes et al. [7] presented a function for novelty evaluation for use with case-based reasoning (CBR) design systems. Macedo and Cardosa [8] took the next logical step and developed a model of surprise and curiosity within a case-based reasoning (CBR) design system as both a search heuristic and a model of emotion. Although they describe their model as one of surprise, it is closer to Berlyne's definition of novelty.

Other models of curiosity include the curious selection mechanism based on Shannons's measure of entropy developed by Scott and Markovitch [18] and Baker et al.'s [2] scheme for novelty detection in text documents based on a hierarchical classification scheme used to track breaking news stories. Merrick [12] has developed motivated learning agents, similar to curious agents, as autonomous, adaptive non-player characters for on-line gaming environments.

Saunders and Gero [14] introduced *curious design agents* as a model of curious behaviour in creative design processes. Curious design agents are capable of autonomously exploring design spaces for "interesting" designs, based on their previous experiences. Curious design agents have been used to computationally explore the role that curiosity plays in designing and have been combined with a range of different design generators.

SUPPORTING CREATIVITY

Generative systems such as evolutionary algorithms have the potential to support creativity by supporting exploration and experimentation. Using them, however, can often be laborious because they do not filter out works that are unlikely to appeal to the user. As Takagi [20] reported, some researchers have attempted to model user preference to reduce that labour involved in the exploration process.

For example, Baluja et al. [3] attempted to model user preferences by training neural networks on user evaluations of genetic artworks. In their system, the user evolved genetic artworks using an interactive evolutionary algorithm. The aim of the study was to train neural networks to evaluate artworks similarly to the user, based on features within each genetic artwork. Several different types of networks ranging in complexity were tried but found that for the most part the networks could only predict which images were likely to be uninteresting with any accuracy. It was concluded that the disappointing performance of the networks was due to a lack of sophistication in the image processing and learning systems, however, as Baluja et al. observed: "users often will choose an image because it is different than the other images on the screen." This simple observation suggests that a system based on novelty detection and curious selection might better predict the preferences of a user.

Prediction of Novelty Preference

Curious design assistants are a form of curious design agent to support creativity either by filtering the designs produced by a generative system for interestingness before presenting them to a user, or by autonomously exploring a design space using a generative design system and then prepare a report on potentially interesting solutions at the end of the exploration. In either case, curious design assistants will have to model a user's preferences for novelty as they filter or explore a space of possibilities.

As a proof-of-concept, a system has been developed that allows a curious design agent and a person to use the same tool to evolve 2D structures. The structures evolved by the agent are called "horns", name after the similar 3D structures evolved by Todd and Latham [21]. Horns are constructed by applying a sequence of morphological processes to simple graphical elements to produce complex structures, see Figure 2 for examples. The curious agent uses a model of curiosity to rank every horn in each generation with respect to its experiences of other horns within an evolutionary run.

A preliminary study compared an agent's interestingness rankings with a user's selections. The results in Figure 3 shows

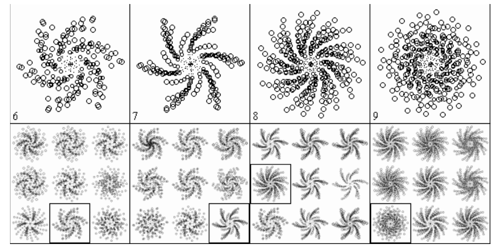

Figure 2. Example of horns as part of a report on an exploration produced using a curious design agent. Below each horn the report shows the selection of the horn from the population of other members of its generation.

that the curious design agent could predict the most interesting structure in a population (i.e. assign it a rank of 1) with up to 50% accuracy; taking the top 3 rankings as likely candidates for selection improves this score to between 60% and 72% accuracy. Unlike Baluja et al.'s study, the agent in this system is not designed to learn a user's preference, rather it is a model of user preference based on the empirical findings. The results of the initial study suggest that up to 72% of selections can be explained as a preference for novelty. The system also has the ability to function as an "auto-pilot", guiding the evolution of new horns along the most interesting paths and presenting them to the user as a report on the most interesting works found.

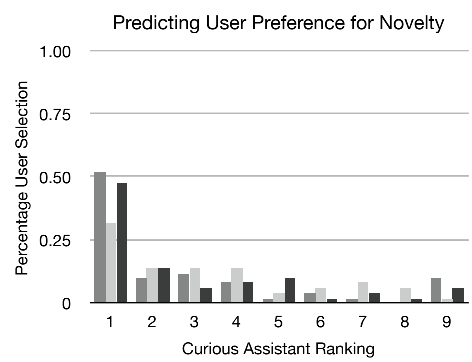

Figure 3. Results of pilot study to model a user's preferences using a curious design agent. Chart shows the number of selections made by a user against the preference judgements of a curious design assistant ranked from 1–9, where a rank of 1 is given to the most interesting structures. Three trials of 50 selections in each trial are shown.

DISCUSSION

Creativity and curiosity are linked by the exploratory behaviour typical of the early stages of a creative activities as a user learns about the possibilities within a space. This paper has attempted to argue that supporting this explorative behaviour requires more than the addition of stochastic elements to a generative process. Existing frameworks for understanding human curiosity have informed the development of computational models of curiosity for autonomous agents applied to a wide range of applications. Curious design agents have applied these models to the autonomous exploration of design spaces. Curious design assistants have shown some potential for applying them in support of human creativity by acting as filters and guides in complex design spaces.

The current model of novelty employed in the pilot study modelling user preference for novelty does not adapt it's preference for novelty based on user selections, and so only improves over time through the sharing of experiences with a user. This does not provide sufficient accuracy to be useful in a creative tool as yet, and more work needs to be done to develop the computational model further.

The computational models of curiosity briefly discussed in this paper have mostly concentrated on the modelling of diversive curiosity, i.e., curiosity due to a lack of stimulation. To model specific curiosity, i.e., curiosity due to an interest in an imagined concept that has yet to be realised, we will have to develop more sophisticated models of curiosity that can support the evaluation of unrealised works. To this end, models of curious agents incorporating a model of the evolution of language are being developed [15]. The potential for these models is for agents to have the ability to combine concepts as linguistic elements, assess how interesting these linguistic constructions are and decide whether to search for works that realise the implied concept.

The potential for curious agents to assist in the exploration of creative works is significant. Computational models of curiosity, with their inherent models of interest and novelty, may provide an effective way to track the changing preferences of a user as they learn about the potential of a computational creative support tool.

REFERENCES

1. T. M. Amabile. Motivation and creativity: Effects of motivational orientation on creative writing. *Journal of Personality and Social Psychology*, (48):393–399, 1985.
2. D. L. Baker, T. Hofmann, A. K. McCallum, and Y. Yang. A hierarchical probabilistic model for novelty detection in text. Technical report, Just Research, 1999.
3. S. Baluja, D. A. Pomerleau, and T. Jochem. Towards automated artificial evolution for computer generated images. *Connection Science*, 6(2&3):325–354, 1994.
4. D. E. Berlyne. *Conflict, Arousal and Curiosity*. McGraw-Hill, New York, 1960.
5. D. E. Berlyne. *Aesthetics and Psychobiology*. Appleton-Century-Crofts, New York, 1971.
6. W. W. Gaver and G. Mandler. Play it again, sam: On liking music. *Cognition and Emotion*, (1):259–282, 1987.
7. P. Gomes, C. Bento, and O. Gago. A function for novelty evaluation in design. In J. S. Gero and M. L. Maher, editors, *Computational Models of Creative Design IV*, pages 19–31, Sydney, Australia, 1999. Key Centre of Design Computing and Cognition, University of Sydney.
8. L. Macedo and A. Cardosa. Creativity and surprise. In G. Wiggins, editor, *AISB'01 Symposium on AI and Creativity in Arts and Science*, York, UK, 2001. SSAISB.
9. S. Marsland, U. Nehmzow, and J. Shapiro. Novelty detection for robot neotaxis. In *International Symposium On Neural Computation (NC'2000)*, pages 554—559, 2000.
10. C. Martindale. *The Clockwork Muse*. Basic Books, New York, NY., 1990.
11. J. McCormack. Open problems in evolutionary music and art. In F. Rothlauf, J. Branke, S. Cagnoni, D. W. Corne, R. Drechsler, Y. Jin, P. Machado, E. Marchiori, J. Romero, G. D. Smith, and G. Squillero, editors, *Proceedings of Applications of Evolutionary Computing, (EvoMUSART 2005), Lausanne, Switzerland, 30 March - 1 April 2005*, number 3449 in Lecture Notes in Computer Science, pages 428–436. Springer-Verlag, Berlin, Germany, 2005.
12. K. Merrick. Modelling motivation for adaptive non-player characters in dynamic computer game worlds. In N. Lee, editor, *ACM Computers in Entertainment*, 2008.
13. J. Preece, Y. Rogers, and H. Sharp. *Interaction Design: Beyond human-computer interaction*. John Wiley & Sons, New York, 2002.
14. R. Saunders and J. S. Gero. Designing for interest and novelty: Motivating design agents. In *Proceedings of CAAD Futures 2001*, pages 725–738, 2001.
15. R. Saunders and K. Grace. Towards a computational model of creative cultures. In *AAAI Spring Symposium on Creative Intelligent Systems*, Stanford University, 26–28 March 2008.
16. J. Schmidhuber. Curious model-building control systems. In *Proceedings of the International Joint Conference on Neural Networks*, volume 2, pages 1458–1463. IEEE, 1991.
17. J. Schmidhuber. What's interesting? Technical Report TR-35-97, IDSIA, Lugano, Switzerland, 1997.
18. P. D. Scott and S. Markovitch. Learning novel domains through curiosity and conjecture. In *Proceedings of International Joint Conference for Artificial Intelligence*, pages 669–674, Detroit, Michigan, 1989.
19. R. J. Sternberg and T. I. Lubart. The concept of creativity: Prospects and paradigms. In R. J. Sternberg, editor, *The Handbook of Creativity*, pages 3–15. Cambridge University Press, Cambridge, UK, 1999.
20. H. Takagi. Interactive evolutionary computation: Fusion of the capabilities of ec optimization and human evaluation. *Proceedings of the IEEE*, 89(9):1275–1296, 2001.
21. S. Todd and W. Latham. *Evolutionary Art and Computers*. Academic Press, London, UK, 1992.