

Artificial Creative Systems and the Evolution of Language

Rob Saunders

Design Lab
University of Sydney
Sydney, NSW 2010 AUSTRALIA
rob.saunders@sydney.edu.au

Abstract

Most studies of human creativity have focused on individuals, assuming that creativity can be defined with respect to the characteristics, processes or activities of extraordinary people. Computational models of creativity have often inherited this assumption and emphasised generative processes to the exclusion of considering social or cultural aspects. This paper presents work to extend computational models of social creativity to support the evolution of domain specific languages. Artificial creative societies provide the opportunity for studying *creativity-as-it-is* in the context of *creativity-as-it-could-be*. The computational model of an artificial creative systems presented here extends previous computational models by introducing a linguistic component that supports the production and sharing of works with associated descriptions. This paper examines the potential for this extended model of social creativity to support the study of the roles that language plays in the formation, interaction and maintenance of creative domains.

Introduction

The need to define the nature of creativity has haunted attempts to develop theories of creative thinking: the difficulty is apparent from the abundance of definitions; Taylor (1988) gives some 50 examples. Expressed in the definitions of creativity are some widely different opinions about what it means for an individual to be creative, yet two broad categories of definitions can be identified: (1) creativity as a mental phenomenon; and, (2) creativity as a social construction. For example, the models of creativity proposed by Koestler (1964), Newell, Shaw, and Simon (1958), and Hofstadter (1979) go into great detail about the cognitive processes involved in creative thinking, particularly the processes involved in the generation of potentially creative ideas. Computational models of creativity are often based directly on such models, e.g., Langley et al. (1987), Hofstadter (1995), or are based on similar models of creative thinking from psychology, e.g., Partridge and Rowe (1994). Creativity as a social construction has a strong honorific sense that is as much the result of an audience's appreciation of a work as it is the creator's production. Proponents of these definitions contend that creativity cannot occur in a vacuum and must be studied in the context of

the socio-cultural environment of the creator (Gruber 1974; Simonton 1984; Martindale 1990). Attempts to combine these two views of creativity into unified theoretical frameworks often maintain the distinction between personal and socio-cultural notions of creativity, as in Bodens P-creativity and H-creativity (Boden 1990) and Gardners small-c and big-c creativity (Gardner 1993).

Dong (2009) argues that language plays a central role in creative behaviours, semantic and sentiment analysis of the use of language in design texts have been used to illustrate how the reality producing effect of language is itself an enactment of design. This insight is compatible with Clark's argument that language is the 'ultimate artefact' whose primary purpose is not to communicate ideas between individuals but to overcome cognitive limitations of the human brain through the externalisation of complex thought in a grounded symbolic form (Clark 1996).

A Systems View of Creativity

The systems view of creativity was developed by Csikszentmihalyi as a model of creativity to include interactions between individuals and the social and cultural environments they are embedded within (Csikszentmihalyi 1988). A map of the systems view of creativity is presented in Figure 1.




Figure 1: The Systems View of Creativity.

An individual's role in the systems view is to bring about

some transformation of the knowledge held in the domain. The field is a set of social institutions that selects from the variations produced by individuals those that are worth preserving. The domain is a repository of knowledge held by the culture that preserves ideas or forms selected by the field. In a typical cycle, an individual takes some knowledge provided by the culture and transforms it, if the transformation is deemed valuable by society, it will be included in the domain of knowledge held by the culture, thus providing a new starting point for the next cycle of transformation and evaluation. Using the language of Gardner, what distinguishes small-c creativity from big-c creativity is that big-c creativity affects changes to the domain whereas small-c creativity does not. In Csikszentmihalyi's view, creativity is not to be found in any one of these elements, but in the interactions between them.

Computational Models of Creative Systems

Liu's dual generate-and-test model of creativity was the first attempt to produce a computational model of Csikszentmihalyi's creative systems (Liu 2000). The dual generate-and-test model encapsulates two generate-and-test loops: one at the level of the individual and the other at the level of the society. The generate-and-test loop at the individual level provides a model of creative thinking, incorporating problem finding, solution generation and evaluation of potential creativity. The outer generate-and-test loop models the field in Csikszentmihalyi's systems view of creativity; providing a model of peer evaluation and a repository of works. The limitations of Liu's computational model lay in the centralised nature of the socio-cultural test and the limited notion of the domain in the model as a repository of artefacts. The dual generate-and-test model provides a way to integrate computational models of creative thinking with models of social creativity but can say little about how fields and domains emerge as a consequence of the actions of individuals.

The artificial creativity approach proposed by Saunders and Gero (2001) provides a framework for developing computational models of individual and social creativity to support the emergence of social structures as the result of the actions of multiple individuals. Early implementations explored the role that an individual's search for novelty plays in social creative systems. Individuals who produce works that are considered interesting by other agents are rewarded. Works communicated between agents that are considered worth sharing by peers are added to the domain.

Other multi-agent models of social creativity have examined the relationship between the field and domain. Gero and Sosa (2002) explored the emergence of 'gatekeepers' in creative fields, i.e., individuals with the ability to strongly affect the contents of the domain. Bown (2008) developed multi-agent models to explore cohesion, competition and maladaptation in the evolution of musical behaviour. Colton, Bundy, and Walsh (2000) present a computational model involving multiple agents working together to explore a mathematical domain, which proved to be so successful that the agents produced new knowledge that has been accepted into the domain of number sequences.

Axelrod's model of the dissemination of culture, while

not attempting to model cultural creativity, illustrate the significance that individual acts of communication can have on the formation and multi-cultural societies (Axelrod 1997). Meme and Variations (MAV) is a computational model of cultural evolution in a society of interacting individuals (Gabora 1995) based on the premise that novel ideas are variations of existing ones. Each agent in an artificial society can acquire new ideas through innovation, by mutating a previously learned idea, or by imitation, by copying a neighbour agent. Thus, cultural evolution occurs through the collective choices of individual agents about which ideas to mutate, how to mutate them, and which ideas to copy. Miranda, Kirby, and Todd (2003) developed a model of the evolution of simple musical forms using a language game, the imitation game, similar to the one presented later and used in this study. In the society of musical agents, compositions are shared through agents performing for each other. The success of a tune is measured by the ability of another agent to reproduce it. A tune is successfully reproduced when the agent who produced the initial performance knows no tunes that are more similar to the imitators recital than the one it initially performed. Simulations show that the society of agents quickly develop coherent sets of tunes and were capable of successful recitals. The tunes, as a set of artefacts collectively agreed upon by a society of individuals, represent a simple form of domain distributed across the memories of the individuals in a similar way to Gabora's MAV.

Creative domains, as described by (Csikszentmihalyi 1988), are dynamically maintained and contain symbolic as well archive material. Domains are distributed across creative fields, existing within a variety of media, with each individual in the field having a partial view of the whole. The computational models described above share a limited notion of domains as repositories of artefacts. Some of these models use a centralised database while others, e.g., MAV, capture the distributed nature of a domain described by Csikszentmihalyi where each agent maintains some part of the whole in memory. None of the models described here however, maintain a distinct symbolic description of the knowledge stored in the domain, i.e., none of them model a domain-specific language that can be used to describe the artefacts or practices.

Computationally modelling the evolution of language in creative domains opens up the possibility of computationally investigating a range of important aspects of creativity that are outside the scope of studies focussed on individuals, including: the emergence of specialised languages that are grounded in the practices of a field; the effects of a common education on the production and evaluation of creative works; and, the emergence of subdomains as a consequence of differences in language use across a field.

The computational model below attempts to address some of the limitations of the earlier implementation of artificial creative system by Saunders and Gero (2001) by including the evolution of language as a central component in the negotiation of works between individuals and the distribution of domain knowledge across a field.

The Evolution of Language

In the extended model of artificial creative systems presented here agents continue to share works with peers in a field as before, sending 'interesting' works for evaluation to other agents; in addition, agents communicate descriptions of works as simple linguistic expressions. This extended model incorporates a model of the evolution of domain-specific languages using models of the evolution of language proposed by Steels (1996b) based on the playing of 'language games'. A language game is an abstract and simplified method of communication, first proposed by Wittgenstein to study the use of language in society. Wittgenstein (1953) describes language games in which participants can communicate to describe or learn about objects, report events, give commands or solve problems.

Steels (1995) introduced the use of language games to simulate the evolution of language in multi-agent systems. In the language game first proposed by Steels, the guessing game, one agent, the initiator, describes an object using a simple utterance to a second agent, the recipient, who attempts to identify the topic of the utterance based on their experience of the previous utterances. Steels has shown that repeated playing of such language games is capable of evolving languages grounded in shared experiences to describe, for example, other agents (Steels 1996a) and coloured shapes in a shared context (Steels 1996b; 1998). In the course of attempting to succeed at as many language games as possible, the society of agents is driven to adopt common meanings for their initially random words and as a consequence a shared lexicon emerges. Steels uses this model to support the position that language is an autonomous adaptive system and that its emergence in humans could have been the result of self-organisation rather than the acquisition of a specific language-capable area of the brain.

Other types of language games have been developed by researchers to explore the evolution of language under different conditions. For example, imitation games have been used to explore the self-organisation of vowel systems (de Boer 2000) and the evolution of simple musical forms (Miranda, Kirby, and Todd 2003) described earlier. The distinction between the evolution of musical forms through the playing of language games and the model presented here is that while the model of the evolution of tunes distributes domain knowledge of acceptable forms across the individuals in the associated field it does not support the co-evolution of a set of symbolic descriptions.

The evolution of language is distributed and self-organising; through the repeated playing of language games between pairs of agents, a shared lexicon of words and their associated meanings evolve in combination. Of particular interest, from the perspective of modelling domain-specific languages, are the ambiguities that arise in the languages evolved; a single word may have multiple meanings and multiple words may have the same meaning. Anyone who has tried to communicate across disciplinary boundaries, no matter how similar they may appear at first, will likely have experienced something similar, e.g., familiar words having unfamiliar meanings. But the resolution of tensions created when individuals from different fields communicate has the

potential for creative output as the meanings of words are negotiated (Gemeinboeck and Dong 2006).

By extending previous models of social creativity with the capacity to negotiate a commonly understood lexicon, the model presented here distributes the domain of recognised works with associated descriptions across its associated field, or fields, with each agent holding a subset of all known works and descriptions within their internal model of the domain. Such a model opens up new opportunities for simulating social institutions, e.g., education, and for studying the effects on domains when fields come into contact through the interaction of individuals.

The Computational Model

Individuals are modelled as curious design agents (Saunders 2002): each agent is capable of generating new works and assessing its novelty. If a generated work is appropriately novel, the agent produces an utterance and uses this to communicate the work and its description to another agent. To assess the novelty of new works and determine an appropriate utterance for them, each agent maintains an associative memory based on a Category Adaptive Resonance Theory (CART) network (Weenink 1997). This memory maintains vector prototypes for classes of work with an associated label, in this case the utterance. A threshold around each prototype defines a hyper-ellipsoid within which similar works will be associated with the same label.

In the following experiments, individuals explore the design space of simple, coloured shapes of varying sizes; similar to the space of coloured shapes that Steels used in "Talking Heads" (Steels 1998). Unlike the "Talking Heads" experiment, however, shapes are not selected from a relatively small finite set, but rather are generated by individual agents. The process of generation implemented for the following simulations is simple: an agent uses the prototypes of shapes that it has stored in its ontology to generate a variant.

Generated shapes are perceived by the agents using a set of sensory channels similar to those used by (Vogt 2003), i.e., the agents can sense the type (square, circle, etc.), size and colour hue of shapes. All sensory channels defined for the agents in the following simulations have been normalised to fall in the range [0..1], with types mapped to specific values within this range and size and colour taking continuous values.

The perceived novelty of the generated shape is assessed as the city-block distance from the closest known prototypes. If the novelty of the new shape falls into the preferred range for the agent, the shape may be used as the topic in a guessing game. The preferred range of novelty for an agent is defined by an internal model of preference based on the Wundt curve (Saunders 2002), where similar-but-different perceptual experiences are preferred. In the simulations that follow, the centroid for preferred novelty vary between 0.025 and 0.125, and have a fixed range about the centroid of 0.05, representing suitably small distances from known prototypes in the perceptual space.

Each field in the following experiments consists of between 10 and 40 individuals. The communication policy between the individuals follows Steels (1995) and implements

either the guessing game, or a variant upon this, the education game. Through the interaction of members of a field, the development of domain-specific lexicons is modelled as a consequence of individuals generating and exchanging ‘interesting’ works with associated utterances. In the model a domain is determined to have formed when a population of agents agree upon a stable lexicon of words with agreed meaning for the associated works. In the experiments that follow, a stable lexicon is said to have formed when communicative success exceeds 80%.

Simulations and Results

This section describes the results of three simulations using the computational model. The experiments conducted so far with the extended model have focused on modelling the domain. The results below explore how domains are (1) formed under the influence of novelty seeking behaviour; (2) combined through the interaction of individuals that are members of multiple fields; and, (3) effectively maintained through the use of education.

Domain Formation

In the computational model presented here, the formation of a domain occurs when the members of a field agree upon a stable lexicon for describing a corpus of works. The model does not require a central repository of all knowledge, rather the domain is distributed amongst the members of the field, such that no individual has a complete record of the domain. Consequently, small differences in the characteristics of individuals can have a large impact on the formation of a domain. Figure 2 illustrates how individual preference for novelty affects the size of the lexicon and ontology stored in the domain as a consequence of the field’s actions. Figure 2 shows size of the active lexicon and ontology for a field of 10 individuals after playing a total of 10,000 language games. The preferred novelty reported is the mean of each individual’s preferred novelty with the range of preferred values ± 0.025 either side of the mean.

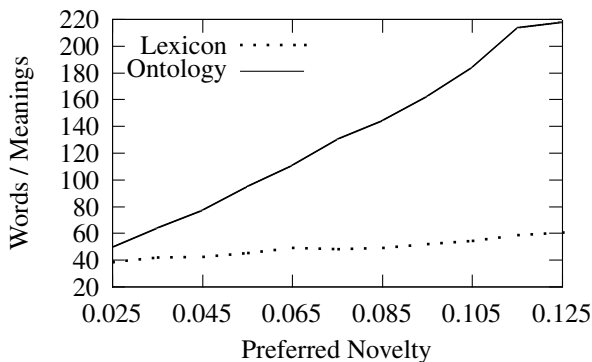


Figure 2: Domain growth as a consequence of individual preference for novelty.

The results of these simulations show that for this artificial creative system increasing the preference for novelty used

by individuals to select the topic of a language game has a modest effect on the size of the active lexicon compared to the increase in the size of the active ontology developed across the domain. In other words, the variety of meanings held by a field for a single word increases significantly as a consequence of individuals searching for novel topics.

Domain Interactions

Simulations based on the evolution of language are open; agents can be added or removed at any time. Agents that are added to a system can adapt to the lexicon in use. We use this capacity to develop models of the interaction between domains as individuals migrate between their associated fields. This type of movement allows individuals to both adapt to the lexicons used in different domains but also affect the development of language as agents transport meanings and words from one domain to another.

Figure 3 shows the results from simulating the interaction of two domains through the communication of individuals taken from two distinct fields. The results have been averaged for 10 simulation runs with 20 individuals in the combined field. The degree of overlap between fields reported is the minority percentage of the new field, i.e., where the degree of overlap is reported as 20% this means that 20% (4) of the agents have been taken from one field and 80% (16 agents) have been taken from the other field.

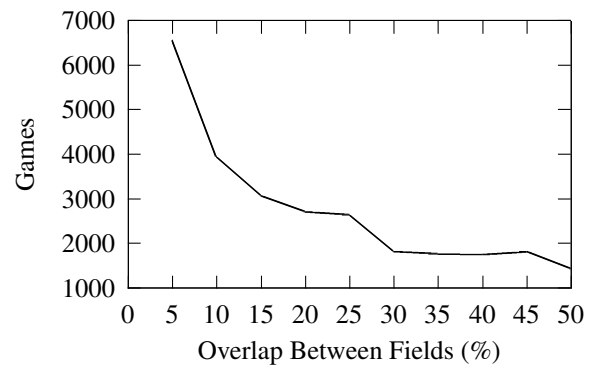


Figure 3: The number of games for a domain to re-form, i.e., reach a communication success rate of 80%, as a function of the percentage of overlap between two existing fields.

The results from the simulations shown in Figure 3 indicate that in this artificial creative system the time taken for the contents of two pre-existing domains to be combined reduces as the number of individuals combined from each existing field approaches 50%. The degree of disruption caused by a small minority of agents is perhaps surprising but it can be easily understood with few opportunities for interaction with agents from the minority percentage many more language games are required for the combined field to reach agreement.

Education

Education, whether through self-study or a more formal education process, plays an essential role in an individual's mastery of a domain: it is only by learning the history of valued works and the language used to describe them that an individual can hope to contribute something new and describe it in such a way as to have it accepted by the 'gatekeepers' of a domain (Csikszentmihalyi 1988). The guessing game can be used to model informal education through exposure to domain knowledge through interactions with members of a field. To model formal education with institutional frameworks, a modified form of the guessing game, called the education game, can be played where the initiator of a game is assumed to be an expert in the domain. In this modified version the initiator takes on the role of teacher, and selects topics for the game with which it has a high confidence based on previous communicative success. The recipient agent takes on the role of student and must choose the most likely object in the shared context that best matches the teacher's utterance. When the identity of the topic is revealed only the student updates its mappings between words and meanings.

To test the efficacy of the education game versus the guessing game at initiating new individuals to a field, a series of simulations were performed to compare how quickly new individuals achieved communicative success rate of 80% with instructors drawn from a pre-existing field using a stable domain. In each run of the simulation a population of 10 individuals engage in language games with an existing field, also containing 10 individuals, where the initiator (teacher) is always chosen from the pre-existing field and the recipient (student) is always chosen from the population of introduced individuals. Figure 4 compares the communicative success for simulations using guessing games and education games.

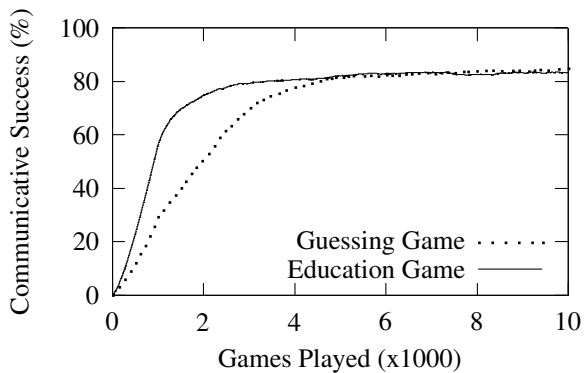


Figure 4: A comparison of communicative success rates of guessing games versus education games with the introduction of individuals to an existing domain.

The results suggest that, in these simulations, the use of the model of formal education significantly decreases the number of language games required for an individual to be able to effectively communicate with a field. To reach a communicative success rate of 80% between the existing field of initiators (teachers) and recipients (students) the

number of games required is reduced by 40%, i.e., from an average of 8,920 to 5,143 language games.

Discussion

There is no doubt that computational modelling will continue to focus on developing analogs for creative cognition and individual creative behaviour. After all, the promise of developing computer programs able to solve problems in ways that are obviously "creative" is so tantalising that we cannot help ourselves. What this paper seeks to accomplish, however, is to show that the potential exists for developing computational models that capture how creativity works within a cultural environment.

The model presented here represents a first attempt to implement a computational model of creative activity within a field that involves language. There are different kinds of creative individual (Policastro and Gardner 1999) and each kind may take part in different types of language games as they interact. For example, Saunders and Grace (2008) proposed the use of the *generation game*, where a speaking agent takes on the role of the client and multiple listener agents take on the role of designers attempting to satisfy the design brief encapsulated in the client's utterance. Unlike Steels' guessing game and the education game presented here, there may be many possible designs that satisfy a single design requirement. This opens the possibility for judging success or failure on more than just the ability of a design to satisfy a set of required features, but to have an implicit requirement for all designs to be 'interesting', according to some function of interest that does not contradict the intended meaning of words within a lexicon. The generation game highlights the role that clients often play in the creative process.

The computational model presented here advances the modelling of the artificial creative systems by introducing a way for domain specific languages to develop from the interactions of individuals within a creative system. The simulations have shown that it is possible to integrate language games with models of individual and social creativity without undermining the grounding of words for describing works within an evolving language. Future work in this area will need to incorporate similar mechanisms for the evolution of languages to describe processes, policies and rules.

The artificial creative system supporting the evolution of language that has been presented in this paper is limited in a number of ways that will need to be addressed. In particular, the language implemented here is holistic, i.e., the words evolved cannot be decomposed into components that describe properties of the shape. To address this limitation a computational model that supports the evolution of compositional languages, similar to that described by Vogt (2003), is being investigated. Computational models of the evolution of compositional languages support the emergence of words that play particular roles in a linguistic construct, e.g., adjective, noun, etc. The meaning of utterances is then formed by the composition of words.

The use of compositional languages in computational models of cultural creativity opens up new and interesting possibilities for modelling the role that language plays in the creative process, e.g., using a compositional language it

is possible for an agent to form a sentence such that all of the words have familiar and agreed upon meanings, but that the combination of words is novel. This has implications for the modelling of creative processes; the ability to produce and evaluate novel descriptions as hypothesised experiences opens up the possibility for modelling grounded forms of specific curiosity (Berlyne 1960).

References

- Axelrod, R. 1997. The dissemination of culture: A model with local convergence and global polarization. *The Journal of Conflict Resolution* 41(2):203–226.
- Berlyne, D. E. 1960. *Conflict, Arousal and Curiosity*. New York: McGraw-Hill.
- Boden, M. A. 1990. *The Creative Mind: Myths and Mechanisms*. London: Cardinal.
- Clark, A. 1996. *Being There: Putting Brain, Body, and World Together Again*. Cambridge, MA, USA: MIT Press, 1st edition.
- Colton, S.; Bundy, A.; and Walsh, T. 2000. Agent based cooperative theory formation in pure mathematics. In *Proceedings of the AISB'00 Symposium on Creative and Cultural Aspects and Applications of AI and Cognitive Science*.
- Csikszentmihalyi, M. 1988. Society, culture and person: a systems view of creativity. In Sternberg, R. J., ed., *The Nature of Creativity: Contemporary Psychological Perspectives*. Cambridge: Cambridge University Press. 325–339.
- de Boer, B. 2000. Emergence of vowel systems through self-organisation. *AI Communications* 13:27–39.
- Dong, A. 2009. *The Language of Design*. London: Springer.
- Gabora, L. 1995. Meme and variations: A computer model of cultural evolution. In Nadel, L., and Stein, D. L., eds., *1993 Lectures in Complex Systems*. Addison Wesley. 471–486.
- Gardner, H. 1993. *Creating Minds*. New York: Basic Books.
- Gemeinboeck, P., and Dong, A. 2006. Discourses of intervention: A language for art & science collaboration. In *Proceedings of International Conference: New Constellations: Art, Science and Society*.
- Gero, J. S., and Sosa, R. 2002. Creative design situations: Artificial creativity in communities of design agents. In *CAADRIA*.
- Gruber, H. E. 1974. *Darwin on Man*. Chicago: University of Chicago Press.
- Hofstadter, D. 1979. *Godel, Escher, Bach: An Eternal Golden Braid*. New York, NY.: Basic Books.
- Hofstadter, D. R. 1995. *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*. New York, NY.: Basic Books.
- Koestler, A. 1964. *The Act of Creation*. New York, NY.: Macmillan.
- Langley, P.; Simon, H. A.; Bradshaw, G. L.; and Zytkow, J. M. 1987. *Scientific Discovery: Computational Explorations of the Creative Processes*. Cambridge, MA.: MIT Press.
- Liu, Y. T. 2000. Creativity or novelty? *Design Studies* 21(3):261–276.
- Martindale, C. 1990. *The Clockwork Muse*. New York, NY.: Basic Books.
- Miranda, E. R.; Kirby, S.; and Todd, P. M. 2003. On computational models of the evolution of music: From the origins of musical taste to the emergence of grammars. *Contemporary Music Review* 22(3):91–111.
- Newell, A.; Shaw, J. C.; and Simon, H. A. 1958. Elements of a theory of human problem solving. *Psychological Review* 65(3):151–166.
- Partridge, D., and Rowe, J. 1994. *Computers and Creativity*. Oxford, UK: Intellect Books.
- Policastro, E., and Gardner, H. 1999. From case studies to robust generalizations: An approach to the study of creativity. In Sternberg, R. J., ed., *Handbook of Creativity*. Cambridge: Cambridge University Press. chapter 11, 213–225.
- Saunders, R., and Gero, J. S. 2001. Artificial creativity: A synthetic approach to the study of creative behaviour. In Gero, J. S., ed., *Computational and Cognitive Models of Creative Design V*, 113–139. Sydney, Australia: University of Sydney.
- Saunders, R., and Grace, K. 2008. Towards a computational model of creative cultures. In *Proceedings of AAAI Spring Symposium on Creative Intelligent Systems*.
- Saunders, R. 2002. *Curious Design Agents and Artificial Creativity*. Ph.D. Dissertation, University of Sydney.
- Simonton, D. K. 1984. *Genius, Creativity, and Leadership: Historiometric Inquiries*. Cambridge, MA.: Harvard University Press.
- Steels, L. 1995. A self-organizing spatial vocabulary. *Artificial Life* 2(3):319–332.
- Steels, L. 1996a. Perceptually grounded meaning creation. In Tokoro, M., ed., *ICMAS96*. AAAI Press.
- Steels, L. 1996b. The spontaneous self-organization of an adaptive language. In Muggleton, S., ed., *Machine Intelligence 15*. Oxford University Press, Oxford.
- Steels, L. 1998. The origins of syntax in visually grounded robotic agents. *Artificial Intelligence* 103(12):133–156.
- Taylor, C. W. 1988. Various approaches to the definition of creativity. In Sternberg, R. J., ed., *The Nature of Creativity: Contemporary Psychological Perspectives*. Cambridge, UK: Cambridge University Press. 99–124.
- Vogt, P. 2003. Thsim v3.2: The talking heads simulation tool. In Banzhaf, W.; Christaller, T.; Dittrich, P.; Kim, J. T.; and Ziegler, J., eds., *Proceedings of the 7th European Conference on Artificial Life, ECAL 2003*.
- Weenink, D. 1997. Category art: A variation on adaptive resonance theory neural networks. In *Proceedings of the Institute of Phonetic Sciences*, volume 21, 117–129.
- Wittgenstein, L. 1953. *Philosophical Investigations*. Blackwell Publishing.