Multi-Agent Based Models of Social Creativity

Rob Saunders

Abstract This chapter provides an introduction to the computational modelling of social creativity using multi-agent systems. It reviews motivations for computationally modelling socio-cultural aspects of creativity and describes a systems view of creativity that has influenced approaches to computationally modelling social creativity. A minimal model of an 'artificial creative system' is described and the components of an individual agent are given in some detail. The Digital Clockwork Muse is presented as an implementation of an artificial creative system together with some results from some small scale investigations into the self-organisation of creative fields. Extensions of the computational model are described, including the evolution of domain specific languages, more sophisticated individuals and alternative models of inter-agent interactions.

Key words: computational creativity, social creativity, agent-based models

1 Introduction

Popular definitions of creativity maintain a distinction between personal and social creativity; Boden (1990) defines both psychological creativity (*p-creativity*) and historical creativity (*h-creativity*), while Gardner (1993) distinguishes between *little-c* (mundane) and *big-C* (eminent) creativity. These definitions maintain that creativity has two important but distinct meanings: a person's perception of their own work; and, an honorific title awarded by society. Models of creativity that attempt to reconcile these different meanings are complicated by the fact that an individual's creativity is productively connected with culture and learning (Lindqvist, 2003).

Rob Saunders
The MetaMakers Institute, Falmouth University, Cornwall, UK,
e-mail: `rob.saunders@falmouth.ac.uk`

Computational creativity has attempted to address questions around the social nature of creativity in one of two ways, either; (1) producing highly sophisticated models of individual creativity that can interact with society at large, or (2) developing computational models of artificial societies that exhibit recognisable features of social creativity. The first approach requires the development of highly capable computational systems that are not only able to produce creative works, judged by the standards of human experts, but are also capable social actors, e.g., by conjuring a persona for the computational system such that it may achieve some recognition of autonomy (Colton, 2012; Cope, 2005; Hoffman & Weinberg, 2011; McCorduck, 1991; Wiggins, 2008). The second approach requires the development of computational models of salient social and cultural aspects of creativity, e.g., by developing multi-agent based models of creative individuals (Bown & Wiggins, 2005; Gabora, 1995; Macedo & Cardoso, 2001; Saunders & Gero, 2001; Saunders & Grace, 2008; Sosa & Gero, 2005). This chapter explores the second approach.

Computational models of social creativity can be applied in different ways: as a way to understand creativity as a complex social phenomenon, similar to computational social science (Saunders & Bown, 2015); as a practical approach to developing distributed computational creativity systems, similar to other applications of multi-agent systems for distributed computing (Wooldridge, 2001); or, as a way to support human creativity as a social activity (Saunders, Chee, & Gemeinboeck, 2013). This chapter focuses on the first of these applications, although the approach and techniques can be applied to the development of distributed systems for practical applications.

Cellular automata and agent-based models are well established in the synthetic study of social phenomena (Axelrod, 1997; Epstein & Axtell, 1996; Schelling, 1969; Wooldridge, 2001). In computational creativity, the study of creativity through the development of agent-based models has spanned multiple domains and different aspects of creativity as a social phenomenon. For example, Gabora's 'Meme and Variations' (MAV) is one of the earliest multi-agent based models to examine the interactions between individuals that drive social creativity and cultural evolution through imitation and mutation of ideas (Gabora, 1995). Colton, Bundy, and Walsh (2000) developed a computational model involving multiple agents working together to explore a mathematical domain. Macedo and Cardoso (2001) explored the ability of agents to gain the attention of others through the production of 'surprising' artefacts. Sosa and Gero (2005) used multi-agent based models to examine the role of society in design. Bown (2008) developed multi-agent based models to explore cohesion, competition and maladaptation in the evolution of musical behaviour. In these agent-based models, creativity can be subjective or objective, individual or collective, direct or indirect (Sosa & Gero, 2008). This chapter presents a particular multi-agent model based approach that combines computational models of personal and social creativity to produce 'artificial creative systems'.

The next section briefly describes a systems view of creativity that provides a useful framework for developing computational models of social creativity. Section 3 presents a multi-agent approach to computationally modelling based on the systems view of creativity. Section 4 describes an implementation of an artificial

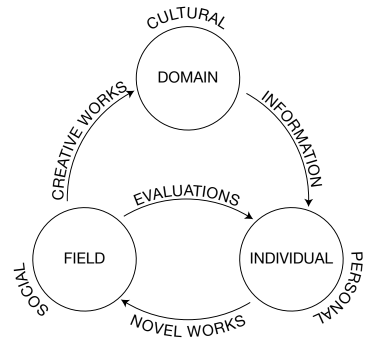

Fig. 1: Csikszentmihalyi's Systems View of Creativity. Based on an illustration from (Saunders & Grace, 2008) adapted from (Csikszentmihalyi, 1988).

creative system, The Digital Clockwork Muse, and presents some results from experiments with this implementation. Section 5 discusses possible extensions to the model presented, including the evolution of domain specific languages, alternative individuals agents, and interactions between agents.

2 A Systems View of Creativity

Vygotsky (1971) first proposed a systems theory of creativity emphasising a reciprocal relationship between individuals and their socio-cultural environment where individuals are both influenced by their understanding of their socio-cultural environment and through their actions cause it to change. An individual may determine that their work is 'creative' independent of the judgement of others but their determination is naturally informed by their experiences of their socio-cultural environment, e.g., the work of others (Martindale, Moore, & West, 1988; Tardif & Sternberg, 1988). In addition, for a work to be given the honorific title of 'creative' other members of a society must agree based on their own experiences of the socio-cultural environment. Consequently, regardless of whether creativity is personal or social, it is the result of a dynamic system of interactions between multiple individuals and their socio-cultural environment (Engeström, 1996). Csikszentmihalyi (1988) argued that creativity is the product of three shaping forces that result from the cultural, social and personal context of the creative activity. The resulting model, illustrated in Figure 1, defines creativity as the result of the interaction between three subsystems: a *domain*, an *individual*, and a *field*. Each subsystem performs a specific function; the domain transmits information to the individual, the individual produces a variation, and the field selects variations to pass on to the domain. These subsystems are described in more detail below.

Csikszentmihalyi (1999) argues that before an individual can produce a variation there must already exist a culture, with traditions and conventions in place for the individual to draw on. A *domain* is defined as the body of knowledge, the set of rules and procedures, the symbolic system, which is used by an individual to produce variations. There will be a multitude of domains in a culture (Feldman, Csikszentmihalyi, & Gardner, 1994) and domains will evolve and change over time. An individual must reference a domain to produce a contribution that a field will understand. As Boden (1990) points out: ‘To be appreciated as creative, a work of art or a scientific theory has to be understood in a specific relation to what preceded it’.

An *individual* is the producer of variation within the systems model. Csikszentmihalyi argues that a person’s background, personal traits and motivations, together with their ability to internalise domain knowledge as well as the expectations of the field, combine to enable an individual to be successfully creative within the system. This view of an individual emphasises the need for them to learn and adapt in order to gain a mastery of a domain and anticipate the response of a field to proposed variations. It also emphasises the importance of successful communication for a creative individual.

A *field* is composed of all of the individuals in a society who possess domain knowledge and have influence over its contents. Sawyer (2012) defines the field as “a complex network of experts, with varying expertise, status, and power”. Possible members of a field include creators, educators, critics, agents (marketers) and consumers. According to Csikszentmihalyi, if the members of a field judge a contribution from an individual to be creative it will be added to the domain for other individuals to reference, thus continuing the cycle.

Csikszentmihalyi argues that creativity can be found when and where these three subsystems interact. In this view, an individual is necessary but not sufficient for a creative system; all three subsystems and their interactions, are equally important. For example, Csikszentmihalyi (1988) argues that highly structured domains, e.g., mathematics, promote creativity by assisting individuals to refer to relevant knowledge and for fields to assess an individuals contribution. Similarly, social structures, e.g., the emergence of ‘gate-keepers’ (Feldman et al., 1994; Sosa & Gero, 2004), have significant impact on an individual’s ability to have variations accepted. To emphasise the importance of all of three elements and their interactions the model is often referred to as the Domain-Individual-Field-Interaction (DIFI) framework (Feldman et al., 1994).

2.1 Computationally Modelling the Systems View of Creativity

Liu (2000) first proposed an approach to computationally modelling the DIFI framework: The *Dual Generate-and-Test* model of creativity, illustrated in Figure 2, encapsulates two generate-and-test cycles: one at the level of the individual and the other at the level of society. The domain is modelled as a repository of artefacts for the individual to draw on and the field to contribute to. The individual generate-and-test cycle implements Newell et al's generate-and-test model of creative thinking (Newell, Shaw, & Simon, 1962) incorporating problem finding, artefact generation and personal creativity (p-creativity) testing. The socio-cultural generate-and-test cycle incorporates the individual as the generator and the field as a monolithic test of the creativity of the variations generated, which determines whether they are sufficiently creative to be added to the domain. The apparent simplicity of Liu's computational model masks the complexity of modelling the field as a monolithic socio-cultural creativity test. Liu suggests that such a system would likely defer to some form of oracle, most likely a human, to provide judgements of creativity. The following section explores the use of multi-agent based models to develop 'artificial creative systems' that support emergent notions of creativity from the interactions of individuals, removing the need for such oracles.

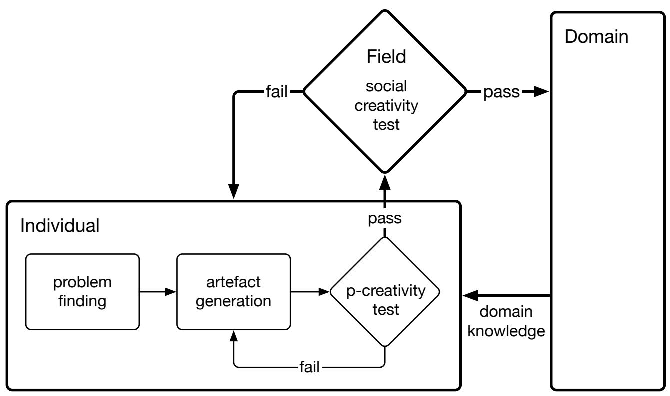

Fig. 2: The individual and socio-cultural generate-and-test loops in Liu's dual-generate-and-test model of creativity. Based on an illustration in (Saunders, 2002).

3 Artificial Creative Systems

The agent-based model presented here does not attempt to define a 1-to-1 mapping with the subsystems of the DIFI framework. In an artificial creative system, no individual agent contains a test for 'big-C' or 'h-creativity' but instead tests for 'little-c' or 'p-creativity', which may or may not be judged by the other agents within a field to be "creative". The ability of agents to make independent judgements of both novelty and value is fundamental to the model, permitting the emergence of social definitions of creativity as a collective function of many individual evaluations of creativity. Consequently, the notion of a "creative work", or a "creative individual", is honorific as it must be determined as the consequence of some form of negotiation between at least two agents.

The autonomy of agents equipped with the ability to determine what is interesting, and therefore potentially p-creative, is the key to adapting Liu's dual generate-and-test model to the study of emergent notions of creativity. This approach substitutes the monolithic socio-cultural creativity test with one based on a distributed agreement that emerges from communication between individual agents. In such an artificial creative system, the socio-cultural creativity test is modelled through the communication of artefacts and evaluations of p-creativity between individuals. The following describes how the subsystems of the DIFI framework can be modelled as agents and interactions.

3.1 Domains

Creative domains, as described by Csikszentmihalyi (1988), are dynamically maintained and contain symbolic, e.g., rules, language etc., as well archive material, e.g., previous works. Consequently, domains should be considered as being distributed across creative fields, existing within a variety of media, with each individual in a field having a partial view of the whole. The simplest computational model of a domain is a repository of artefacts that have been judged to be creative, i.e., an archive. This is the model described in Liu's model above and used in the simple implementations below. But this model lacks both the distributed and multi-faceted nature of the domain described by Csikszentmihalyi.

Other computational models of the domain are possible that better capture the distributed nature of the domain, e.g., Gabora's MAV where each agent maintains some part of the whole domain in memory. These partial views may overlap, i.e., two agents may have artefacts in common. In these models it is only by considering the intersection of these partial views, i.e., parts that are commonly held, that the domain can be understood. Saunders (2011) proposed a way to extend the model of the domain to encompass more than exemplars of previous work, by incorporating a model of the evolution of domain specific languages to capture symbolic representations, e.g., descriptions of artefacts.

3.2 Fields

In an artificial creative system, a field is modelled as a set of agents that interact with each other and the domain according to a communication policy, as illustrated in Figure 3. In this example, agent *i* communicates an artefact that it considers to be p-creative to agent *j*, which evaluates the artefact according to its own p-creativity test and sends its evaluation back to *i*. Each agent's evaluation of an artefact is affected by the traits of the individual, e.g., preference for novelty, prototypes held

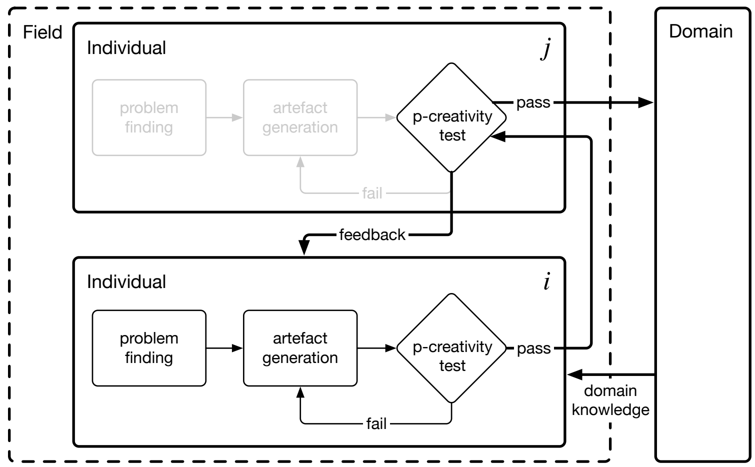

Fig. 3: A minimal social creativity test in an artificial creative system. Based on an illustration from (Saunders, 2002).

in memory. Consequently, through the communication of evaluations, *j* can affect the generation of future artefacts by *i* by rewarding *i* when it generates artefacts that *j* considers to be p-creative.

Indirectly, *i* can also affect the evaluation of p-creativity by *j* because *j*'s evaluation of p-creativity involves an evaluation of novelty, which is partly, or wholly, based on artefacts it has previously experienced. Hence, *i* affects a change in *j*'s evaluation of p-creativity every time it causes *j* to evaluate an artefact and update the prototypes held in memory. By exposing *j* to artefacts that *i* considers to be p-creative it can alter *j*'s evaluation of novelty and hence it's p-creativity test.

To implement the socio-cultural creativity test as a collective function of p-creativity tests a communication policy is required. A simple communication policy, implemented in the system described in Section 4, is for agents to communicate an artefact when their evaluation of that artefact's p-creativity is greater than some fixed threshold. In addition agents have a domain interaction policy, for example, agents will add artefacts, generated by other agents, if the p-creativity evaluation of the artefact is greater than a domain submission threshold. In this way no individual is allowed to submit their own work to the domain, thus, at least one other agent must find an individual's work creative before it is entered into the domain.

3.3 Individuals

An individual in the DIFI framework must be able to transform knowledge from the domain and produce some novelty for the field to determine whether or not it is creative. Consequently, there are three main requirements for an agent-based model of an individual; (1) it must be able to access the contents of the domain, (2) it must be able to generate some novelty, and (3) it must be able to communicate with other members of the field.

Given a simple repository-style model of a domain, the ability to access the contents of the domain, can be accomplished with a suitable interface for querying the repository, e.g., a database. For more complex models of a domain, e.g., where the some of the knowledge held in the domain may be distributed across a group of agents, then an individual agent may need to communicate with other agents to gain access to their knowledge, as in the case of Gabora's MAV (Gabora, 1995).

Common to other multi-agent based models, individual agents must be able to communicate with members of the field. Simple message passing protocols can be used to accomplish this. As a minimum, individual agents can pass artefacts to other members of the field and receive feedback in return. More complex models of communication may include meta-information about artefacts, e.g., a description using a domain-specific language (Saunders & Grace, 2008).

The ability of an individual to generate some novelty, or more precisely the ability to detect that some potentially interesting novelty has been generated, poses the greatest challenge. While the mechanics of producing novelty can be implemented in a variety of ways the ability to detect novelty and use this to implement a test for 'p-creativity' places specific requirements on the agent. The following describes the components of an agent that uses a novelty detector and an 'hedonic function' to achieve these requirements.

3.3.1 Novelty Detection

A novelty detector determines the novelty of a new input based on a model of expected inputs. Novelty detectors can be implemented in different ways depending on the type of novelty to be detected (Markou & Singh, 2003a, 2003b). One way to implement a novelty detector is to use a classifier that has been trained on a set of expected stimuli, such that when a new stimulus is presented to the classifier, the classification error is an indication of the novelty. In such an agent-based model, an agent i has a memory M_i with K learned categories of artefacts, such that, $M_i = \{m_i^1, \ldots, m_i^K\}$ where m_i^k is the kth learned category. Given an artefact a, the novelty detector calculates the novelty, $N_i(a)$ to be:

$$N_i(a) = \min_{m_i^k \in M_i} \Delta(a, m_i^k) \tag{1}$$

Where $\Delta(a, m_i^k)$ is the classification error, which measures the difference between an input a and the kth learned category m_i^k. How the classification error, Δ, is calculated will differ between implementations, e.g., it may be the Euclidean distance from a prototype or a function of the error generated by a neural network.

Simply detecting novelty is sufficient for many applications, e.g., monitoring of equipment to identify potential faults (Markou & Singh, 2003a), but in computa-

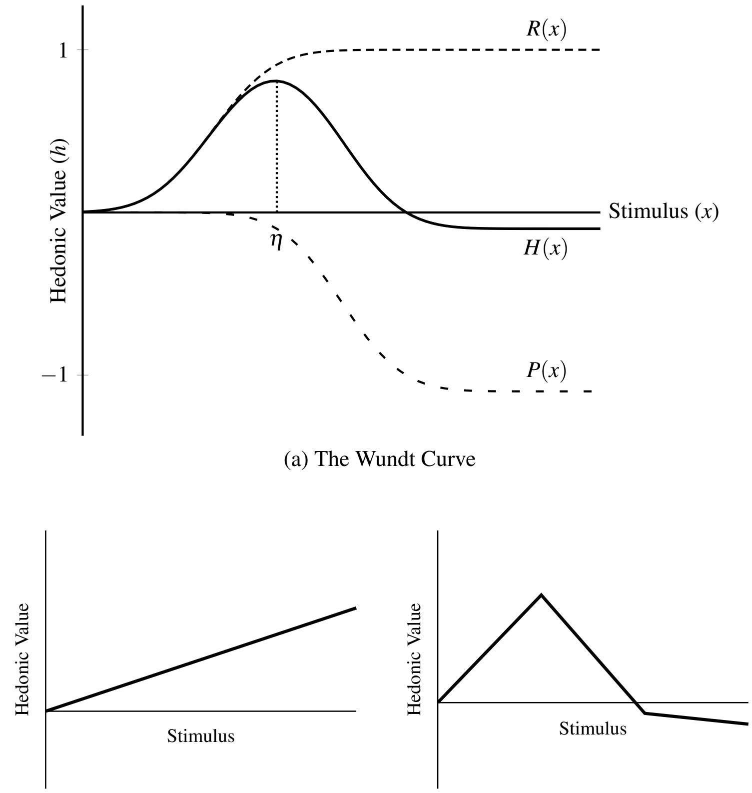

(a) The Wundt Curve

(b) Linear

(c) Piecewise

Fig. 4: Example Hedonic functions (a) Wundt Curve, (b) linear and (c) piecewise approx. of Wundt Curve. Illustration of Wundt Curve based on (Saunders, 2002) after (Berlyne, 1960).

tional models of p-creativity more is required. An agent in an artificial creative system needs to be able to identify potentially interesting novelty by modelling a preference for novelty, which can be achieved with the use of an hedonic function.

3.3.2 Hedonic Functions

An hedonic function defines a transformation from a perceived stimulus to a response signal, which can be used to guide learning and action, e.g., in intrinsically-motivated reinforcement learning (Singh, Barto, & Chentanez, 2004). Studies of human preference suggest an inverted U-shape relationship between stimuli and interest (Heckhausen & Heckhausen, 2008; Wundt, 1910), known as the Wundt Curve, illustrated in Figure 4(a). For an agent i the Wundt Curve may be implemented as a function, $H_i(x)$, which takes a stimulus, x, and calculates a response signal as the difference of two cumulative Gaussian functions, a reward function $R_i(x)$ and a punishment function $P_i(x)$, which according to Berlyne (1960) represent a positive response to small amounts of stimuli and a negative response to large amounts of stimuli:

$$
\begin{aligned}
H_i(x) &= R_i(x) - \alpha P_i(x) \\
R_i(x) &= F(x \mid \mu_r, \sigma_r) \\
P_i(x) &= F(x \mid \mu_p, \sigma_p)
\end{aligned} \tag{2}
$$

Where μ_r and σ_r are the mean and standard deviation that define the underlying normal distribution for rewarding smaller amounts of stimulus and μ_p and σ_p define the underlying normal distribution for punishing larger stimuli and α defines the degree to which large values of the stimulus are punished, i.e., for values greater than 1 a negative reward value will be generated for values of x when $P_i(x) \geq R_i(x)$. The cumulative Gaussian function $F(x \mid \mu, \sigma)$ is calculated by:

$$
F(x \mid \mu, \sigma) = \frac{1}{2}\left[1 + \operatorname{erf}\left(\frac{x-\mu}{\sigma\sqrt{2}}\right)\right] \tag{3}
$$

Where $\operatorname{erf}(y)$ is the Gauss error function, which can be approximated for efficient calculation. It is also useful to define η, as the value of x that generates the peak response:

$$
\eta = \arg\max_x H_i(x) = \{x \mid \forall x' : H_i(x') \leq H_i(x)\} \tag{4}
$$

Berlyne (1960) identified the Wundt Curve as a model for typical reactions that animals and humans display to the presence of novel situations. That is, the most interesting experiences are those that are similar-yet-different to those that have been experienced before, or might be expected given previous experiences. By defining the stimulus, $x = N_i(a)$, the Wundt Curve can be used to calculate a reward based on novelty Where x is due to novelty, η represents the preferred novelty for an individual agent. By altering parameters controlling the reward and punishment functions the value of η can be altered to control how novel an artefact must be for it to be considered ‘interesting’.

Other hedonic functions are possible and may be desirable for certain domains. For example, if the expected novelty of any artefact can be reasonably assumed to lie within the range of values close to the origin, then a simple linear response function may be sufficient, such as that illustrated in Figure 4(b). Alternatively, if the interest values determined by an hedonic function are only used for comparison, e.g. to compare the relative interest due to different artefacts, such that the absolute value of the interest is not important, then a piecewise linear approximation of the Wundt Curve, Figure 4(c), may be a approximation.

3.3.3 Interest, Boredom and Curiosity

Given an hedonic function, an agent can determine a measure of interest and determine what action to take as a consequence. For example, given a communication threshold, τ_C, an agent i may determine to send an artefact a to another agent if $H_i(N_i(a)) > \tau_C$. Alternatively, given a domain submission threshold, τ_D, an agent may decide to submit an artefact a to the domain if $H_i(N_i(a)) > \tau_D$. The rules governing when these decisions may be acted upon form a policy for how the field is structured. A measure of interest for each artefact allows an agent to monitor the frequency with which it encounters 'interesting' artefacts. By keeping an accumulated measure of interest over time, it is possible to develop a simple computational model of 'boredom'. Accumulating interest over time can be achieved simply, e.g., $S_i = \alpha S_i + (1-\alpha)H_i(N_i(a))$, where S_i is the accumulated interest for agent i and α is a suitable decay rate. A state of 'boredom' can then be modelled whenever $S_i < \tau_B$, where τ_B is a suitable threshold for a desired minimum interest level that the agent attempts to maintain.

Given a model of boredom, it is possible to model a type of curiosity identified by Berlyne (1960) as 'diversive curiosity'. In diversive curiosity, a lack of novel stimuli produces a change in behaviour to increase potential exposure to new experiences. An agent in an artificial creative system may implement this type of curiosity very simply by retrieving an artefact from the domain. Alternatively, an agent could adjust the parameters of its generative system, such that a more diverse range of artefacts are produced. This simple model of curiosity is based on an assumption that the memory of an agent contains an implicit model of the expectations of future experiences. More sophisticated explicit models of curiosity have been developed and models with explicit expectations have been developed to model surprise (Baranès & Oudeyer, 2009; Merrick & Maher, 2006; Schmidhuber, 1991). Other types of curiosity have also been identified, e.g., 'specific curiosity', may also be computationally modelled. In addition, other forms of intrinsic motivation, e.g., competence, can be computationally modelled (Merrick & Maher, 2006).

This section has described how an multiple agents, interacting with other agents and a repository, can be used to model a creative system. The next section provides a concrete example of implementing this approach.

4 The Digital Clockwork Muse

The Digital Clockwork Muse is an implementation of an artificial creative system inspired by the work of Martindale (Martindale, 1990). In "The Clockwork Muse" Martindale presented an investigation into the role that individual novelty-seeking behaviour plays in literature, music, visual arts and architecture. Martindale concluded that the search for novelty exerts a significant force on the development of styles. The Digital Clockwork Muse is an attempt to computationally model a creative system to investigate some of the features of creative societies, driven by the search for novelty, described by Martindale.

Algorithm 1: The Digital Clockwork Muse

```
while t < total simulation time do
    foreach agent i in field F do
        update interest h_i for artefact a_i, h_i = H_i(N_i(a_i))
        while message queue Q_i is not empty do
            remove artefact a^n from Q_i, sent by agent n
            calculate the hedonic value h_i^n = H_i(N_i(a^n))
            update memory M_i to include a^n
            send feedback including h_i^n to sender agent n
            if h_i^n > domain submission threshold (τ_D) then
                submit artefact a^n to domain D
            end
            if h_i^n > h_i then
                adopt received artefact, a_i ← a^n, h_i ← h_i^n
            end
        end
        generate new artefact a'_i from a_i
        calculate the hedonic value h'_i = H_i(N_i(a'_i))
        update memory M_i to include a'_i
        if h'_i > communication threshold (τ_C) then
            select an agent m from F, where m ≠ i
            send artefact a'_i to agent m
        end
        if h'_i > h_i then
            adopt generated artefact, a_i ← a'_i, h_i ← h'_i
        end
        update interest level S_i = αS_i + (1 − α)h_i
        if S_i < boredom threshold (τ_B) then
            retrieve artefact a^d from domain D
            calculate the hedonic value h_i^d = H_i(N_i(a^d))
            update memory M_i to include a^d
            if h_i^d > h_i then
                adopt retrieved artefact a_i ← a^d, h_i ← h_i^d
            end
        end
    end
end
```

The operation of The Digital Clockwork Muse is expressed in Algorithm 1, every agent, i, in a field maintains a current artefact, a_i, with an associated interest value, h_i. At every step in the simulation, each agent implements up to three phases in order to process artefacts; (1) received from members of the field, (2) generated by the individual, and (3) retrieved from the domain.

In the first phase, each agent, i, evaluates every artefact, a^n in its message queue, shared by another agent n, to calculate an hedonic value h_i^n. Agent i sends an evaluation of its interest in the artefact, h_i^n, to the sending agent, n. If the agent calculates that its interest in an artefact exceeds the domain submission threshold, τ_D, then the agent adds the artefact to the domain, D. If the agent calculates that its interest in an artefact from the queue exceeds its interest in its current artefact, $h_i^n > h_i$, it will adopt the received artefact, $a_i \leftarrow a^n$.

In the second phase, each agent generates a new artefact, a_i', based on its current one, a_i, and evaluates its interest in the generated artefact, h_i'. If the agent's interest in the generated artefact exceeds the communication threshold, τ_C, then the agent will choose another agent, m, from the field and send the generated artefacts to it. If the agent's interest in the generated artefact is greater than its interest in the current artefact, $h_i' > h_i$, it will adopt the generated artefact, $a_i \leftarrow a_i'$.

In the third phase, an agent updates its internal state of accumulated interest, S_i, based on the hedonic value of the current artefact. If the level of accumulated interest falls below the boredom threshold, τ_B, then the agent will retrieve an artefact, a^d, from the domain D. If the agent's interest in the retrieved artefact, h_i^d, exceeds the agent's interest in the current artefact, then the agent will adopt the retrieved artefact.

4.1 Experiments

Martindale (1990) illustrated the influence of the search for novelty by individuals in a thought experiment "The Law of Novelty". The Law of Novelty forbids the repetition of word or deed and punishes offenders by ostracising them. Martindale argued that The Law of Novelty was merely a magnification of the reality in creative fields. Some of the consequences of the search for novelty are that individuals that do not innovate appropriately will be ignored in the long run and that the complexity of any one style will increase over time to support the increasing need for novelty.

The following experiments were designed to study the effects of the search for novelty in artificial creative societies modelled as agents that have hedonic functions with different preferred novelty value, i.e., η, as defined in Equation 4. In this implementation, η ranges from 0 to 32, equal to the range of the potential classification error generated by the novelty detectors used. More detailed accounts of these experiments can be found in (Saunders, 2002).

4.1.1 The Law of Novelty

In the first experiment a group of 12 agents were created. Ten of the agents, agents 0–9, shared the same hedonic function, i.e. the same preference for novelty ($\eta = 11$). Agent 10 was given a preference for low amounts of novelty ($\eta = 3$) and agent 11 was given a preference for high amounts of novelty ($\eta = 19$). Figure 5(b) illustrates the network of communication links developed between agents that communicate artefacts and evaluations on a regular basis.

Agent ID	Preferred Novelty(η)	Attributed Creativity
0	11	5.43
1	11	4.49
2	11	4.50
3	11	3.60
4	11	4.48
5	11	1.82
6	11	6.32
7	11	8.93
8	11	10.72
9	11	5.39
10	3	0.00
11	19	0.00

(a) The attributed creativity between agents.

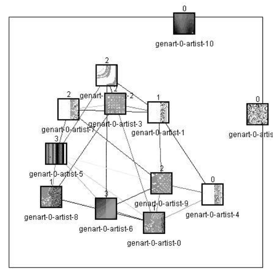

(b) Screenshot of the running simulation.

Fig. 5: The Law of Novelty simulated within a single field of agents with different preferences for novelty.

The results of the simulation are presented in Figure 5(a). The results indicate that the agents with the same preference for novelty to be somewhat 'creative' according to their peers, with an average attributed creativity of 5.57. Neither agent 10, with a preference for low amounts of novelty, nor agent 11, with a preference for high degrees of novelty, received any credit for their artefacts. Consequently none of the artefacts produced by these agents were saved in the domain.

The results illustrate the potential for the simulation of the Law of Novelty in artificial creative systems. Agents with a lower novelty preference tend to innovate at a slower rate than those with a higher hedonic preference and while an agent must produce novelty to be considered creative, it must do so at a pace that matches its audience.

4.1.2 The Formation of Cliques

In the second experiment, the behaviour of groups of agents with different hedonic functions is investigated. To do this a group of 10 agents was created, five of them had a hedonic function that favoured novelty close to $\eta = 6$ and the other five agents favoured novelty values close to $\eta = 15$. Figure 6(a) illustrates the network of communication of high evaluations between the agents for interesting artefacts.

Two areas of frequent communication can be seen in the matrix of communication shown in Figure 6(a). The agents with the same hedonic function frequently send high evaluations for interesting artefacts amongst themselves but rarely send them to agents with a different hedonic function, i.e., there are a large number of high evaluation messages between agents 0–4 and agents 5–9, but only one between the two groups, agent 4 sends a high evaluation to agent 5.

Sender

Receiver	0	1	2	3	4	5	6	7	8	9
0		2	8	1	2					
1										
2	2	1		1	3					
3	4	5	2		5					
4	2	3	3	2						
5					1		6	1	3	5
6						3		4	5	1
7						3	5		1	4
8						4	3	2		4
9						1	4	4	4	

(a) A matrix of the number of positive creative evaluations sent between agents.

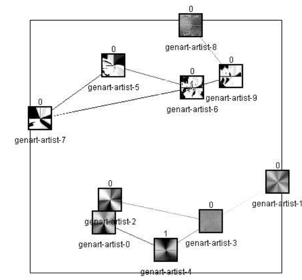

(b) A screenshot of a simulation showing two non-communicating cliques.

Fig. 6: The formation of cliques between agents with different hedonic functions.

The result of putting collections of agents with different hedonic functions in the same group appears to be the formation of cliques: groups of agents that communicate credit frequently amongst themselves but rarely acknowledge the creativity of agents outside the clique. As a consequence of the lack of communication between the groups the style of artworks produced by the two cliques also remains distinct. The different styles of the two groups can be seen in Figure 6(b), with agents 0–4 producing smooth radial images and agents 5–9 producing fractured images with clearly defined edges.

The results of this experiment suggest that when a population of agents contains subgroups with different hedonic functions, the agents in those subgroups form cliques. The agents within a clique communicate high evaluations frequently amongst themselves but rarely to outsiders. The stability of these cliques will depend upon how similar the individuals in different subgroups are and how often the agents in one subgroup are exposed to the artefacts of another subgroup.

Communication between cliques is rare but it is an important aspect of creative social behaviour. Communication between cliques occurs when two individuals in the different cliques explore design subspaces that are perceptually similar. Each of the individuals is then able to evaluate the other's artefacts highly because they have constructed appropriate perceptual categories. The transfer of artefacts from one clique to another permits new variables into the creative processes of the destination clique, the two cliques can then explore in different directions. Cliques can therefore act as "super-individuals", exploring a design space as a collective and communicating interesting artefacts within and between cliques.

5 Extensions

The agent-based model of social creativity provided by artificial creative systems provides a flexible framework for experimentation, which can be extended in a number of ways to explore different aspects of social creativity. This section explores some examples of these extensions to the domain, individual, field and interactions between these components.

5.1 Domains

The computational model of the domain presented above is lacking in many ways compared to the dynamically evolving source of cultural knowledge that Csikszentmihalyi describes. Saunders (2011) has incorporated 'language games' in artificial creative systems to explore the possibility of computationally modelling more complex knowledge structures through the evolution of domain-specific languages. 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.

Wittgenstein (1953) proposed language games as a thought experiment to explore the production of language as a consequence of action and interaction. An example of a language game requires a listener to attempt to identify the topic of an utterance within a given context and for a speaker to provide feedback on the success or failure of the listener. Computational model of language games have demonstrated the ability for agents to evolve languages as a consequence of repeated plays (Steels, 1995). In the extended artificial creative system proposed by Saunders (2011), agents produce utterances to describe artefacts when communicating with other agents. This extended model has be used to explore the impact of the preference for novelty on the formation of creative domains. For the purposes of the computational 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. A stable lexicon is said to have formed when communicative success exceeds 80%.

Figure 7 illustrates how individual preference for novelty affects the size of the lexicon and ontology, of artefacts, stored in the domain as a consequence of the field's actions. 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. The presence of

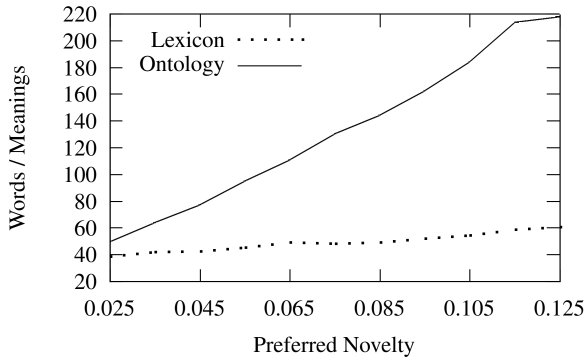

Fig. 7: Lexicon versus ontology growth as a consequence of individual preferences for novelty within an artificial creative system.

ambiguous words in the lexicon of an evolved languages has the potential to support the computational modelling at the level of domain interactions as a consequence of individuals actions (Saunders, 2011). This has implications for the modelling of creative processes; the ability to produce and evaluate novel descriptions opens up the possibility for modelling grounded forms of specific curiosity (Berlyne, 1960).

The evolution of domain-specific languages also presents opportunities for domains to differentiate within a single culture, as they present barriers to the flow of information between domains. Consequently, it is possible to computationally model interactions between domains as a result of the actions of individuals (Saunders, 2011).

5.2 Individuals

One of the obvious limitations of the computational model of individuals presented here is the lack of an explicit test for the appropriateness of artefacts. Similarly, integrating alternative generative processes, including analogy-making (Falkenhainer, Forbus, & Gentner, 1989), could provide a useful framework for evaluating the effectiveness of such creative processes within a social and cultural context. Curiosity is not the only intrinsic motivation for creative individuals, although it is one of the most persistent (Martindale, 1990). Other motivations for exploring a design space can be computationally modelled, e.g. competency (Merrick & Maher, 2006).

Building on the development of computational models of intrinsic motivation in robotics (Baranès & Oudeyer, 2009; Marsland, Nehmzow, & Shapiro, 2000), Curious Whispers 2.0 is an example of embodied artificial creative system with three robots exchanging simple tunes in much the same way as the software agents in the computational models described above (Saunders et al., 2013). The use of robots opens up new possibilities of also engaging humans in creative activities. In the case of Curious Whispers 2.0 the robots exchange tunes ‘in the open’ by performing them and listening for tunes being played by other robots. This openness allowed human participants in the creative system to intervene by playing tunes using a custom synthesiser that could play the same three notes as the robots. The Curious Whispers 2.0 platform has been used to explore interactions between humans and robots when the locus of the creative activity is in the interactions between all of the agents, rather than the human having a privileged role.

5.3 Fields

A significant shortcoming of the simulations described above is the small size of the simulated fields. The ability to simulate larger creative societies will permit the study of the spread of innovations and styles. It may also facilitate the emergence of new fields as cliques attain a critical size. Spatial and topological relationships will become more important issues in large population models.

There are several other important players in creative societies besides the producers of innovations (Policastro & Gardner, 1999) including, e.g. consumers, distributors, critics, etc. Each has their own role to play in creative societies; consumers evaluate artefacts, distributors distribute artefacts widely, and critics distribute their evaluations widely. Convincing other people that you’ve had a creative idea is often harder than having the idea in the first place.

Building on the extended model of domains described above, Saunders and Grace (2008) introduced ‘generation games’ as a type of language game where a speaker agent takes the place of a client and an utterance represents a ‘brief’, such that listener agents, acting as designers can attempt to satisfy the brief through the production of artefacts. Saunders (2011) also examined the possibility of computationally model of educators within an artificial creative system.

In non-homogenous societies of agents, the selection of which agents to communicate with becomes an important strategy for agents seeking recognition as a creative individual. Other computational models based on Csikszentmihalyi’s system view of creativity have also been developed that demonstrate the important role that authority figures, or ‘gatekeepers’, play in creative fields (Sosa & Gero, 2005).

5.4 Interactions

Artificial societies can have many different policies that control the interactions and decision making activities of agents. For example, simulations of technological innovation in industry show that the consideration of the costs of innovation in decision-making can lead to complex behaviour (Haag & Liedl, 2001). Simulating similar costs in the design process may provide a better understanding of the economics of creative design in creative societies and the strategies needed to manage creativity with limited resources.

Linkola, Takala, and Toivonen (2016) have implemented artificial creative systems with more complex interactions between the individuals within a field in order to select artefacts that may be added to a domain. In Linkola et al.'s model all artefacts that pass an individual's *self-criticism* test, similar to the p-creativity tests described above, are first published and every agent engages in a two stage process to vote on which, if any, artefact gets added to the domain. The first stage of voting allows any agent to veto the addition of an artefact if they assess its novelty to be below a threshold. If any artefacts remain in the set of published artefacts, the second stage of voting selects the one with the highest average novelty assessment from all of the agents to be added to the domain. Linkola et al. explored the effects of varying the self-criticism and veto thresholds on the collective effort required by a field to achieve domains of a given size and concluded that raising the self-criticism threshold reduces the collective effort, while raising the veto threshold maintains the novelty of the artefacts in a domain.

6 Conclusion

The computational modelling of creative societies opens up new opportunities for computational creativity that go beyond the modelling of the romantic figure of the lone creative genius, or the utilitarian assistant to the human creative. The aim of this chapter has been to present an approach to computationally modelling creativity using multi-agent systems and show how this can be used to explore aspects of social and cultural creativity. By using agents as models of individuals within creative fields, the framework provides a flexible basis for developing multi-agent systems that can be used to study the interaction between personal and social judgements of creativity. This chapter has also attempted to show that this type of model is open to extension to include other aspects of the social and cultural context for creative individuals, e.g., domain specific languages.

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 chapter seeks to accomplish, however, is to show that it is possible to develop relatively simple computational models that accommodate both personal and socio-cultural aspects of creativity.

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. Retrieved from http://dx.doi.org/10.2307/174371 doi: 10.2307/174371

Baranès, A., & Oudeyer, P.-Y. (2009). R-iac: Robust intrinsically motivated exploration and active learning. *IEEE Transactions on Autonomous Mental Development*, *1*(3).

Berlyne, D. E. (1960). *Conflict, arousal and curiosity*. New York, NY.: McGraw-Hill.

Boden, M. A. (1990). *The creative mind: Myths and mechanisms*. London, UK: Cardinal.

Bown, O. (2008). *Theoretical and computational models of cohesion, competition and maladaptation in the evolution of human musical behaviour* (Unpublished doctoral dissertation). University of London (Goldsmiths), London, UK.

Bown, O., & Wiggins, G. A. (2005). Modelling musical behaviour in a cultural-evolutionary system. In *Computational creativity workshop.*

Colton, S. (2012). The painting fool: Stories from building an automated painter. In J. McCormack & M. d'Inverno (Eds.), *Computers and creativity* (pp. 3–38). Berlin, Heidelberg: Springer Berlin Heidelberg.

Colton, S., Bundy, A., & Walsh, T. (2000). Agent based cooperative theory formation in pure mathematics. In G. Wiggins (Ed.), *Proceedings of aisb 2000 symposium on creative and cultural aspects and applications of ai and cognitive science* (p. 11-18). Birmingham, UK.

Cope, D. (2005). *Computer models of musical creativity*. Cambridge, MA: MIT Press.

Csikszentmihalyi, M. (1988). Society, culture and person: a systems view of creativity. In R. J. Sternberg (Ed.), *The nature of creativity: Contemporary psychological perspectives* (pp. 325–339). Cambridge, UK: Cambridge University Press.

Csikszentmihalyi, M. (1999). Implications of a systems perspective for the study of creativity. In R. J. Sternberg (Ed.), *Handbook of creativity* (pp. 313–335). Cambridge, UK: Cambridge University Press.

Engeström, Y. (1996). Development as breaking away and opening up: A challenge to vygotsky and piaget. *Swiss Journal of Psychology*, *55*, 126-132.

Epstein, J., & Axtell, R. (1996). *Growing artificial societies: Social science from the bottom up*. Brookings Institution Press. Retrieved from https://books.google.com.au/books?id=xXvelSs2caQC

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, *41*, 1–63.

Feldman, D. H., Csikszentmihalyi, M., & Gardner, H. (1994). *Changing the world, a framework for the study of creativity*. Westport: Praeger.

Gabora, L. (1995). Meme and variations: A computer model of cultural evolution. In L. Nadel & D. Stein (Eds.), *1993 lectures in complex systems.* Reading: Addison-Wesley.

Gardner, H. (1993). *Creating minds: An anatomy of creativity seen through the lives* *of freud, einstein, picasso, stravinsky, eliot, graham and gandhi*. New York, NY.: Basic Books.

Haag, G., & Liedl, P. (2001). Modelling and simulating innovation behaviour within micro-based correlated decision processes. *Journal of Artificial Societies and Social Simulation*, *4*(3). Retrieved from `http://www.soc.surrey.ac.uk/JASSS/4/3/3.html`

Heckhausen, J., & Heckhausen, H. (2008). *Motivation and action*. New York: Cambridge University Press.

Hoffman, G., & Weinberg, G. (2011, October). Interactive improvisation with a robotic marimba player. *Autonomous Robots*, *31*(2-3), 133–153. Retrieved from `http://dx.doi.org/10.1007/s10514-011-9237-0` doi: 10.1007/s10514-011-9237-0

Lindqvist, G. (2003). Vygotsky's theory of creativity. *Creativity Research Journal*, *15*(2-3), 245-251. Retrieved from `http://www.tandfonline.com/doi/abs/10.1080/10400419.2003.9651416` doi: 10.1080/10400419.2003.9651416

Linkola, S., Takala, T., & Toivonen, H. (2016). Novelty-seeking multi-agent systems. In *Proceedings of the seventh international conference on computational creativity.*

Liu, Y. T. (2000). Creativity or novelty? *Design Studies*, *21*(3), 261—276.

Macedo, L., & Cardoso, A. (2001). Using surprise to create products that get the attention of other agents. In L. Canamero (Ed.), *Aaai fall symposium* (pp. 79–84). Menlo Park, CA: The AAAI Press.

Markou, M., & Singh, S. (2003a). Novelty detection: A review—part 1: Statistical approaches. *Signal processing*, *83*(12), 2481–2497.

Markou, M., & Singh, S. (2003b, December). Novelty detection: A review—part 2: Neural network based approaches. *Signal Process.*, *83*(12), 2499–2521. Retrieved from `http://dx.doi.org/10.1016/j.sigpro.2003.07.019` doi: 10.1016/j.sigpro.2003.07.019

Marsland, S., Nehmzow, U., & Shapiro, J. (2000). Novelty detection for robot neotaxis. In *International symposium on neural computation (nc'2000)* (pp. 554—559).

Martindale, C. (1990). *The clockwork muse*. New York, NY.: Basic Books.

Martindale, C., Moore, K., & West, A. (1988). Relationship of preference judgements to typicality, novelty, and mere exposure. *Empirical Studies of the Arts*, *6*(1).

McCorduck, P. (1991). *Aaron's code: Meta-art, artificial intelligence, and the work of harold cohen*. W.H. Freeman. Retrieved from `https://books.google.com.au/books?id=r3UyBgAAQBAJ`

Merrick, K., & Maher, M.-L. (2006). Motivated reinforcement learning for non-player characters in persistent computer game worlds. In *Acm sigchi international conference on advances in computer entertainment technology, (ace 2006).* Los Angeles, USA.

Newell, A., Shaw, J. C., & Simon, H. A. (1962). The process of creative thinking. In H. Gruber, G. Terrell, & M. Wertheimer (Eds.), *Contemporary approaches* *to creative thinking* (pp. 63—119). New York, NY.: Atherton Press.

Policastro, E., & Gardner, H. (1999). From case studies to robust generalizations: An approach to the study of creativity. In R. J. Sternberg (Ed.), *Handbook of creativity* (pp. 213–225). Cambridge: Cambridge University Press.

Saunders, R. (2002). *Curious design agents and artificial creativity* (Unpublished doctoral dissertation). University of Sydney.

Saunders, R. (2011, April). Artificial creative systems and the evolution of language. In D. Ventura, P. Gervás, D. F. Harrell, M. L. Maher, A. Pease, & G. Wiggins (Eds.), *Proceedings of the second international conference on computational creativity* (p. 36-41). México City, México.

Saunders, R., & Bown, O. (2015, August). Computational social creativity. *Artif. Life*, *21*(3), 366–378.

Saunders, R., Chee, E., & Gemeinboeck, P. (2013). Human–robot interaction with embodied creative systems. In M. L. Maher, T. Veale, R. Saunders, & O. Bown (Eds.), *Proceedings of the fourth international conference on computational creativity* (pp. 205–209). Sydney, Australia.

Saunders, R., & Gero, J. S. (2001). The digital clockwork muse: A computational model of aesthetic evolution. In G. A. Wiggins (Ed.), *Proceedings of the aisb'01 symposium on ai and creativity in arts and science* (pp. 12–21). SSAISB.

Saunders, R., & Grace, K. (2008, 26–28 March 2008). Towards a computational model of creative cultures. In *Proceedings of aaai spring symposium on creative intelligent systems.* Stanford University, CA..

Sawyer, K. (2012). *Explaining creativity: The science of human innovation* (2nd Edition ed.). Oxford University Press.

Schelling, T. C. (1969). Models of segregation. *The American Economic Review*, *59*(2), 488–493.

Schmidhuber, J. (1991). Curious model-building control systems. In *Proceedings of the international joint conference on neural networks* (Vol. 2, pp. 1458–1463). Singapore: IEEE.

Singh, S., Barto, A. G., & Chentanez, N. (2004). Intrinsically motivated reinforcement learning. In *Neural information processing 18.*

Sosa, R., & Gero, J. S. (2004). Diffusion of design ideas: Gatekeeping effects. In H. S. Lee & J. W. Choi (Eds.), *Caadria 2004.* Seoul: Yonsei University Press.

Sosa, R., & Gero, J. S. (2005). A computational study of creativity in design: the role of society. *AIEDAM*.

Sosa, R., & Gero, J. S. (2008). Creative social systems. In *Aaai spring symposium: Creative intelligent systems* (pp. 90–94).

Steels, L. (1995). A self-organizing spatial vocabulary. *Artificial Life*, *2*(3), 319-332.

Tardif, T. Z., & Sternberg, R. J. (1988). What do we know about creativity? In R. J. Sternberg (Ed.), *The nature of creativity* (pp. 429—440). Cambridge, UK.: Cambridge University Press.

Vygotsky, L. S. (1971). *The psychology of art*. Cambridge, MA: MIT Press. (Original work published 1930.)

Wiggins, G. A. (2008). Computer models of musical creativity: A review of computer models of musical creativity by david cope. *Literary and Linguistic Computing*, *23*(1), 109-116. Retrieved from http://llc.oxfordjournals.org/content/23/1/109.short doi: 10.1093/llc/fqm025

Wittgenstein, L. (1953). *Philosophical investigations*. Blackwell Publishing.

Wooldridge, M. J. (2001). *An introduction to multiagent systems*. John Wiley & Sons.

Wundt, W. (1910). *Principles of physiological psychology*. New York: Macmillan.