

Novelty-Seeking Multi-Agent Systems

Simo Linkola

Department of Computer Science and HIIT
University of Helsinki
slinkola@cs.helsinki.fi

Tapio Takala

Department of Computer Science
Aalto University School of Science
tta@cs.hut.fi

Hannu Toivonen

Department of Computer Science and HIIT
University of Helsinki
hannu.toivonen@cs.helsinki.fi

Abstract

This paper considers novelty-seeking multi-agent systems as a step towards more efficient generation of creative artifacts. We describe a simple multi-agent architecture where agents have limited resources and exercise self-criticism, *veto* power and voting to collectively regulate which artifacts are selected to the domain i.e., the cultural storage of the system. To overcome their individual resource limitations, agents have a limited access to the artifacts already in the domain which they can use to guide their search for novel artifacts.

Creating geometric images called spirographs as a case study, we show that novelty-seeking multi-agent systems can be more productive in generating novel artifacts than a single-agent or monolithic system. In particular, *veto* power is in our case an effective collaborative decision-making strategy for enhancing novelty of domain artifacts, and self-criticism of agents can significantly reduce the collaborative effort in decision making.

Introduction

Novelty is often considered a central component of creativity (e.g. Boden (1992)). Obviously, an artifact that is not novel can hardly be considered creative. This paper studies the capability of cooperative multi-agent systems to seek and produce *novel* artifacts, and the effects of social decision-making strategies on this capability. Our focus is on seeking novelty; other aspects of creativity, such as surprise and value, are left for future work.

According to the systems view of Csikszentmihalyi (1988), creative systems consist of three intertwined parts: individual agents, society and domain. A set of interacting agents forms a society. The domain is a cultural component constructed by the society by selecting artifacts worth preserving. Each part in the system is in constant interaction with other parts, e.g. individuals try to learn from the domain and bring about transformations, while it is the society that collectively decides which transformations are valued and stored in the domain.

In this work, we view the agent society as a whole, and consider the artifacts introduced to the domain as the end result of the agent population's cultural knowledge of the artifact type. From this point of view, it is important that

the agent society is capable of distributed self-regulation in controlling which artifacts are accepted to the domain.

We examine how the number of agents, the amount of their collective resources and their access to the domain amalgamate with decision-making strategies of the society. Specifically, we are interested in how self-criticism, voting and *veto* power (the ability of individual agents to reject artifacts) enhance the overall novelty of artifacts accepted to the domain. Further on, we study how much work the system has to do to produce a certain amount of domain artifacts. In our case study, we use simple agents that create spirographs.

Our main contribution is the study of overall novelty of domain artifacts produced using different social decision-making strategies, especially self-regulation and *veto* power.

This paper is structured as follows. After reviewing related work in the next section, we describe the novelty-seeking agent architecture. We then illustrate and evaluate the architecture using spirographs as the artifacts.

Related Work

Multi-agent systems are a large research area (for an overview, see, e.g., Shoham and Leyton-Brown (2009)). Within the field, our work can be characterized as a system with multiple autonomous agents, where the agents diverge in information they possess (they each have a location and some memory) but not in their interests (they all aim to generate novel artifacts). Further on, the agents are cooperative rather than competitive. The focus of this work is on creativity of agent systems and more specifically on novelty-seeking agents. Next, we briefly review related work on creative agents; a more comprehensive overview can be obtained from the review of computational social creativity by Saunders and Bown (2015).

We build our research upon existing work on creative and curious agents, especially work done by Saunders and Gero.

Saunders and Gero (2001a) present a curious agent searching for novelty in the space of geometric images produced by a spirograph. The agent learns a categorization of the produced images by showing them as input to a self-organized map, or SOM (Kohonen 1995). The novelty of a new image is computed as the pixel-wise deviation from the best matching cell's image in the SOM. The agent's curiosity is modeled as a tendency to make smaller mutations in the generating parameters when more novelty is found. This

helped the agent to concentrate on areas in the parameter space where more variability was found.

In another experiment they let a society of agents seek novelty in images produced by genetic programming (Saunders and Gero 2001b). The agents have variable degrees of curiosity, modeled as a hedonic function that gets its maximum at a certain level of novelty. The agents communicate through their creations, giving positive feedback to those artifacts that match their hedonic function. Societal formations, such as cliques, were found to emerge.

We have adopted a similar approach, simulating a society of communicating agents that try to produce novel spirographs. However, we do not utilize the hedonic function but seek only to maximize novelty. Moreover, the agents in our experiments do not learn a model, such as a SOM, of previously seen artifacts. Instead, they memorize a limited number of the encountered artifacts as they are. This is a simpler solution and also less sensitive to parameters of the model (e.g. those of SOM).

Sosa and Gero (2005) have studied design as a social phenomenon with change agents (designers) and adopter agents (consumers). They conclude that emergent social phenomena — such as gatekeepers and opinion leaders — can stem from simple social mechanisms, and that the effect of an individual on a society depends both on the individual attributes and on the social structures.

Gabora and Tseng (2014) have studied a society of agents capable of inventing and imitating ideas, and of realizing the ideas as actions. In their work, each agent has a set of limbs and the agents make actions by moving the limbs. Gabora and Tseng (2014) observe that societies where agents can chain simple actions to more complex ones obtain higher average fitness and that self-regulation increases the mean diversity of the actions.

Finally, Lehman and Stanley (2008) introduce a novelty search where the main interest is not, *per se*, in satisfying certain objective goal. Instead, the aim is to find a diverse set of behaviors, i.e. behaviors that are novel enough with respect to other behaviors in the set. The search for an expanding set of novel behaviors often leads to a point where a fixed objective goal is also satisfied. Our work has a similar interest, a set of novel behaviors or artifacts, but we consider multi-agent systems without central control.

Agent Architecture

We now describe our architecture of a novelty-seeking agent system. The designs of individual agents and the society of agents have been kept as simple as possible. We make no claims of the novelty of the architecture; rather, our contribution is in the aim to maximize the diversity of artifacts created and the experimental results concerning factors behind the resulting diversity. We outline the big picture of the architecture first and then give the details.

We have a society (population) S of homogeneous agents. Each agent $S_i \in S$ has a fixed amount of resources at its disposal, in particular a constant amount of individual memory; in other respects, the agents are identical.

We model the behavior of the population via iterations: at each iteration, each agent creates a candidate artifact based

on its current position and memory. Agents then proceed to collectively decide which of the candidate artifacts to add to the domain.

In our model, the agents can be self-critical and choose not to present their own artifact as a potential candidate. They can also exercise *veto* power to reject other agents' candidates. The agents are cooperative so self-criticism and especially the *veto* power are intended to be used for the benefit of the society, not of any individual agent.

We will next more closely explain how individual agents function, and then how the multi-agent system operates as a whole.

Individual Agents

We consider agents that have a generative function producing artifacts from one or more parameters. In our model (following Saunders and Gero (2001a)), the agents live in the generative function's parameter space and can only explore different artifacts by moving in the parameter space.

Agents appreciate artifacts based on their novelty: the more novel the artifact is to the agent, the more it is appreciated. To this end, each agent has a limited memory of artifacts, and a function which can measure a distance between any two artifacts. An agent can memorize artifacts it sees during the process to its memory. If the memory is full, memorizing a new artifact will erase the oldest one.

An agent calculates the novelty of a new artifact as the minimum distance between the new artifact and any artifact currently in the agent's memory. More precisely, an agent S_i with artifact memory M_i of size m , $M_i = (A_1, A_2, \dots, A_m)$, calculates the novelty $N_i(A)$ of artifact A to be

$$N_i(A) = \min_{A' \in M_i} d(A, A'), \quad (1)$$

where $d(\cdot)$ is the distance function.

Pseudocode for the behavior of a single agent is given in Algorithm 1; details are given in the text below.

Algorithm 1 Agent behavior during a single iteration

- 1: invent a new artifact close to the agent's current location and move to the new location
 - 2: **if** the new artifact passes self-criticism **then**
 - 3: memorize the new artifact
 - 4: publish the new artifact as a candidate for the domain
 - 5: **end if**
 - 6: participate in social decision making to select which artifact, among candidates published by all agents, is added to the domain
 - 7: select and memorize artifacts from domain
-

To invent a new artifact and to move to a new location (line 1), the agent considers a fixed number of possible new locations using random walk in the parameter space (called a search beam). For each possible location, it then considers the artifact produced by the respective parameter values and chooses the one with maximum novelty with respect to the agent's own memory. It then moves to the corresponding position in the parameter space.

In order to model self-criticism, agent S_i has a novelty threshold s_i which it uses to determine if the created artifact is novel enough for its liking (line 2). If the created artifact passes the threshold, i.e. if $N_i(A) \geq s_i$, the agent memorizes the artifact and also publishes it as a potential domain artifact candidate (lines 3–4). In a single agent setting, these published artifacts will create the domain on their own.

Multi-Agent Architecture

To keep our model simple, the multi-agent system runs with minimal agent-to-agent interaction. The interactions are done solely via generated artifacts and are twofold: (1) agents use collective decision making to select artifacts to the domain D , and (2) agents can examine and memorize current domain artifacts in D to guide their own search.

In each iteration, domain artifact candidates are published by individual agents. The selection to the domain takes place in two phases (line 6).

First, agents exercise *veto* power: any agent S_i rejects any other agent’s artifact A whose calculated novelty is below a threshold v_i , in a manner similar to self-criticism. Formally, given a set C of candidate artifacts, the set

$$C^* = \{A \in C \mid \forall S_i : N_i(A) \geq v_i\} \quad (2)$$

of candidates survives to the next step.

Second, agents vote on which remaining artifact in C^* to add to the domain. (If C^* is empty, none is added.) The voting procedure considers the calculated novelties of artifacts in C^* , and the winner is the artifact A^* which is considered on average most novel:

$$A^* = \arg \max_{A \in C^*} \left(\frac{1}{|S|} \sum_{S_i \in S} N_i(A) \right). \quad (3)$$

The artifact A^* is then added to the domain D .

Agents have access to the domain artifacts which they can examine and memorize (line 7). Memorizing an artifact will add it to the agent’s memory (and erase the oldest artifact from the memory if its full). In our model, agents have two means to explore domain artifacts: draw k artifacts at random or select the closest k artifacts in the parameter space. We will denote these domain artifact memorizing strategies as random_k and closest_k . In both strategies the agent memorizes the artifacts blindly in the sense that a single artifact can appear multiple times in the agent’s memory.

The domain is a set of artifacts, but for notational purposes we consider it as a temporally ordered sequence of artifacts $D = (A_1, A_2, \dots, A^*)$. This allows us later to denote all the artifacts in the domain up to the j th artifact by $D^j = (A_1, A_2, \dots, A_j)$.

Case study: Spirographs

We illustrate the novelty-seeking agent architecture by generating spirographs, a type of geometric images, like Saunders and Gero (2001a) did. While generation of a spirograph is a mechanistic process given the necessary parameters, finding parameter values that produce creative spirographs — in our case more specifically novel ones — is a non-trivial problem.

Spirograph

Spirograph is a toy used to draw epicyclic curved patterns with two interlocking gears of different sizes. A rotating gear (g) of radius r is positioned next to a fixed gear (G) of radius R such that the gear’s teeth interlock. A pen fixed to some point in g at distance ρ from the center draws a pattern when the gear is rotated. Points on the curve are given by equations

$$x = (R \pm r) \cos(\theta) + \rho \cos(\theta + t) \quad (4)$$

$$y = (R \pm r) \sin(\theta) + \rho \sin(\theta + t) \quad (5)$$

where the sign of r determines whether g is exterior or interior to G . θ is the rotation of g ’s center around G , and t is the rotation of g self, given by

$$t = \theta(R - r)/r. \quad (6)$$

The pen’s movement is cyclic, returning to the starting point when both gears have made an integer number of rotations, i.e. when $\theta = 2\pi N/R$, where N is the least common multiple of r and R . Small N gives distinguishable calligraphic patterns, whereas shaded circular bands result when r/R tends towards irrational ($N \rightarrow \infty$).

A real physical spirograph is constrained by $R > 0$ and $\rho < r$, and $r < R$ if g is inside G . In our experiment, we use an abstract computational toy, allowing any (real) values in the formula. Without loss of generality, R can be fixed and r, ρ defined relative to that. Values of $\rho > r$ (meaning that the pen is outside of g) and $\rho < 0$ are also possible, though the latter only produces mirrored equivalents of positive values (the pen is in a reversed position w.r.t. g ’s center).

Compared to Saunders and Gero (2001a) the main difference is that we also let the pen radius ρ vary, giving us two parameters to mutate while traversing the search space.

A Spirograph-Generating Agent

We will now describe in detail how a spirograph-generating agent in our experiments behaves. As described above, we run our agents in a simulation where each agent is triggered to act on every iteration. Agents follow the procedure illustrated in Algorithm 1 every time they act.

Agents live in the 2-dimensional parameter space of spirographs, where the location of an agent is determined by its values for r and ρ . Each point (r, ρ) in the parameter space corresponds to a single spirograph defined by r, ρ , and $R = 200$. Agents are initialized to start at random locations in the continuous parameter space by drawing the initial location (r, ρ) from the uniform distribution $r, \rho \sim \mathcal{U}(-199, 199)$.

Spirographs are first drawn as 500×500 greyscale images where gear G is located in the center. Because r can be negative (gear g is exterior to G), some areas of the parameter space actually produce plain white images as the whole spirograph is drawn outside the image.

To reduce the spirograph generation time, each spirograph is drawn with only 20 full rotations of gear g around gear G ’s center. This has the effect that some spirographs are only drawn partially, but as neither the completeness of the spirographs nor the generating function is in the focus here,

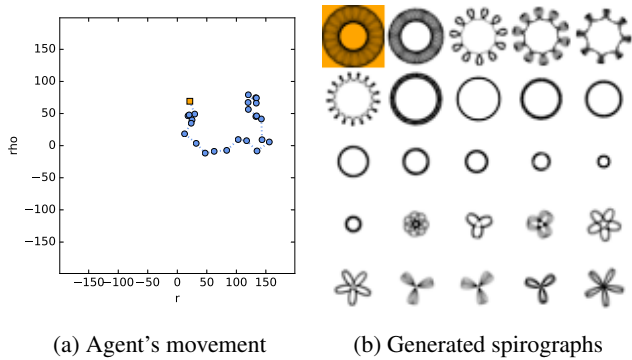


Figure 1: A single agent’s behavior, its movement in the 2-dimensional parameter space (1a) and generated spirographs (ordered left-to-right, top-to-bottom) (1b).

it does not affect the experiments. Finally, to reduce evaluation time, spirographs are rescaled to 32×32 greyscale images.

For inventing a new spirograph, an agent located in a point (r, ρ) in the parameter space considers a fixed amount of new points around it. Each new point (r', ρ') is sampled from a two-dimensional normal distribution with $r' \sim \mathcal{N}(r, 8)$ and $\rho' \sim \mathcal{N}(\rho, 8)$, then both r and ρ are clamped to $-199 \leq r, \rho \leq 199$, and a spirograph corresponding to the point is created as described above.

For each new spirograph, its novelty is calculated as in Equation 1, and the spirograph considered the most novel is selected. The difference $d(\cdot)$ between two images, used in the equation, is defined as the Euclidian distance between the 1024–element vectors formed from grey-scale values of each 32×32 image’s pixels. Although this does not fully correspond to perceptual distance between images, it technically serves our purpose.

Figure 1 illustrates a sample of 25 iterations of a single agent’s behavior, its movement in the parameter space and the spirographs it has created.

Evaluation

We next report on empirical evaluation of the proposed agent architecture using spirographs as the creative artifacts.

The questions we aim to answer empirically are the following. (1) How does the number of agents affect the novelty of artifacts produced to the domain? (2) What is the effect of the beam size on the performance? (3) How does self-criticism of agents affect the novelty, and what is the effect of the *veto* power? (4) How does agents’ access to the domain affect novelty? We also study how these factors affect the rate at which artifacts are introduced to the domain.

Experimental Setup Novelty can be difficult to define in many domains, and it obviously depends a lot on the background. In the experiments of this paper, the novelty of each artifact added to the domain is measured in relation to the artifacts that the agent society has already added to the domain. Such a measure allows comparison across different

Simulation parameter	Default value
Target domain size, $ D $	200
Number of agents, $ S $	16
Self-criticism threshold, s_i	3.2
<i>Veto</i> power threshold, v_i	3.2
Total search beam width	256
Total agent memory	512
Memorization strategy	closest ₃

Table 1: Default parameter values for the experiments.

systems that aim to produce novel artifacts of the same type, whether they are single-agent or multi-agent systems.

Let A_j denote the artifact added to the domain D as its j th artifact. The novelty of A_j is measured as its distance to the nearest artifact already in the domain:

$$N^j(A_j) = \min_{A' \in D^{j-1}} d(A_j, A'), \quad (7)$$

where D^{j-1} is the set of artifacts in the domain before A_j is added to it. Further on, we define $N^1(A_1) = 0$.

Based on the novelty of individual artifacts in the domain, we define an aggregate measure as the average over all artifacts’ novelties:

$$N^*(D) = \frac{1}{|D| - 1} \sum_{2 \leq j \leq |D|} N^j(A_j), \quad (8)$$

and use $N^*(D)$ to compare performance of different system configurations.

In the experiments, we simulate the agent system until a fixed number (200) of artifacts has been accepted to the domain and compute their mean novelty N^* as the measure how novel the artifacts in the domain are on average.

The effort needed to produce a given number of artifacts varies across different settings since the exercise of self-criticism and *veto* power can result in iterations with no candidate artifacts at all. We therefore also study the number of iterations of the agent system needed to produce the artifacts.

Each agent has some resources, in particular a fixed amount of memory and a search beam (the number of locations it considers per iteration). To make comparisons fair across different numbers of agents, the total amount of these resources in the society are kept constant when the number of agents varies.

(There are other aspects that affect the computational complexity but they are ignored here. For instance, with the above division of a constant amount of memory across agents, a society consisting of a smaller number of agents makes a larger total number of comparisons between artifacts in the search beams and the memory. On the other hand, a larger society spends more efforts on mutual evaluation, vetoing, and voting on candidate artifacts produced by the society.)

The default parameter values of our experiments are listed in Table 1. The total search beam width and agent memory are divided equally to agents.

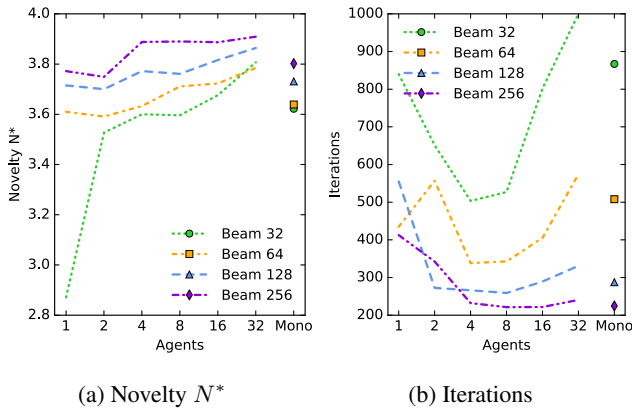


Figure 2: Effect of the number of agents on the novelty N^* (2a) and on the effort required to produce 200 novel artifacts to the domain (2b). Points at the right ends of the panels are for the baseline method Mono.

Results

We now report our experimental results with the above-described architecture of novelty-seeking agents.

Population size The effect of population size on the overall behavior of an agent system is of key interest. Ideally, a multi-agent system should have emergent properties that a single-agent system does not have while not introducing excessive overhead due to agent communication and coordination.

Figure 2 shows how the behavior of our multi-agent system is affected by the number of agents in the society. Different lines show different search beam widths; for now, consider the shapes of the curves, we will return to a comparison between them below.

Panel 2a shows that the overall novelty N^* of artifacts added to the domain increases with the number of agents. This is a desired effect for an agent architecture and indicates that agent collaboration, in particular the selection of artifacts to the domain works effectively. The effect is clearer with smaller beam widths (lower lines in the figure).

Panel 2b complements the picture by showing the corresponding effort, expressed in terms of the number of iterations required to produce 200 novel artifacts to the domain. Here, we observe a less trivial behavior when the number of agents increases. First, the required effort drops until about 4 agents. This is explained by the fact that a larger number of agents can search a more diverse set of options. The required effort starts to increase, however, when the number of agents grows further. When the number of agents grows, the society also becomes collectively more critical about the novelty of candidate artifacts. In our case, some 16–32 agents seem to be the critical amount, but the exact amount is of course dependent on the application.

The two panels of Figure 2 illustrate an inherent trade-off in systems like this: the more critical the society, the higher the novelty of its output is but smaller in size. Based on the

figure, in our setting some 4–16 agents seem to give a good compromise between quality and efficiency.

We next briefly compare the results of the multi-agent system to three different simple alternatives.

First, a comparison to a single-agent system with otherwise similar functionality and identical resources (Figure 2, leftmost points of the lines) shows that as a rule, a multi-agent system produces more novelty and often in less time than a single agent.

Second, an efficient and simple method to obtain 200 spirographs is to sample 200 random points uniformly from the parameter space. Artefacts produced this way have an average novelty of $N^* = 1.14$, markedly lower than the novelties obtained by agent systems with at least two participants (3.5–3.9).

Third, consider a monolithic hybrid between the two baselines above called “Mono”. Mono has no location in the parameter space and so it does not use random walk. It instead samples points uniformly from the parameter space at each iteration and, like our agents, chooses the best of them at each iteration. The Mono system also exercises self-criticism/veto with the same threshold as the agents. In contrast to our agents, Mono has a complete memory of the domain artifacts and is maximally informed in that sense.

A comparison to the novelty obtained by the Mono baseline (panel 2a, separate points at the right end of the panel) shows that from approximately four agents up, agent societies are competitive with and even outperform the monolithic system with complete memory. At the same time, the agent system is more effective in producing the 200 artifacts, up to some 16 agents (panel 2b).

Search beam width Let us now consider the different search beam widths in Figure 2.

First, a comparison of the relative performances of different search beam widths gives the expected results: a wider search finds more novel results (2a) and does it more effectively (2b). Among the different beam widths, the narrower ones tend to be more interesting because a common assumption in multi-agent systems is that the agents are relatively simple and operate under severe resource constraints. In contrast, when the beam width grows without limit, agents start to have complete information about the search space.

As already suggested above, different search beam widths behave differently when the number of agents is changed. As a rule, the number of agents has a larger effect when the search beam is narrow. This is natural, since with narrow beams the individual agents are more constrained. A larger number of agents helps overcome the limitation and find more novel results (2a). On the other hand, when the number of agents becomes large, self-criticism and especially the *veto* power hit the constrained agents harder and they need a longer time to find novel results (2b).

Selection of candidates to the domain We now move on to consider how different methods to select candidates to the domain affect the behavior of the society. This is the central social aspect of our model: we model social interaction by

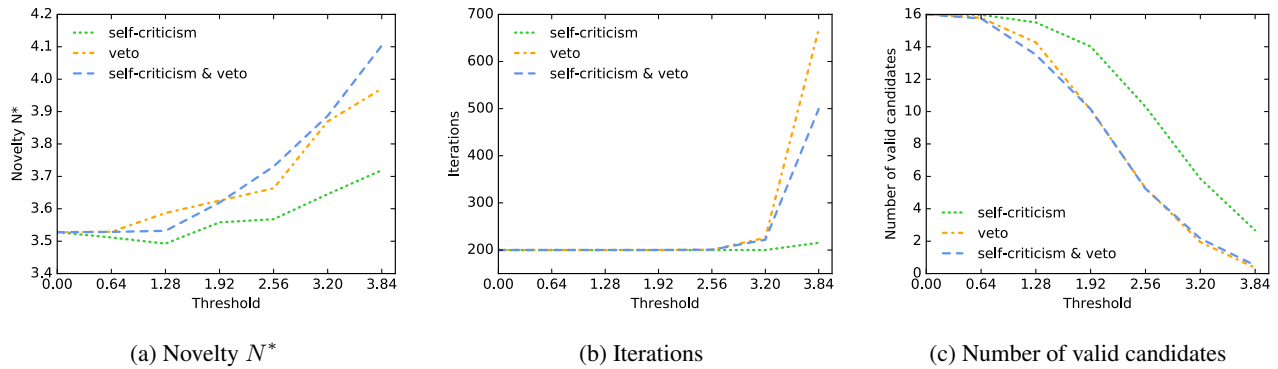


Figure 3: Effect of self-criticism and *veto* thresholds on the novelty N^* (3a), on the effort required to produce 200 novel artifacts to the domain (3b) and on the number of artifacts passing the thresholds (3c).

submission and evaluation of candidate artifacts and collaborative selection of which of them to add to the domain.

Self-criticism and veto power. Recall that the selection of candidates to the domain is controlled by two thresholds, the self-criticism threshold s_i and the *veto* threshold v_i , and an artifact is acceptable if its novelty is not lower than the respective threshold. For simplicity, in our experiments the thresholds are not agent-specific but rather constant across all agents.

Figure 3 illustrates the effects of self criticism and *veto* power using three curves in each panel: one where the threshold s_i for self-criticism varies over the experiments and the *veto* threshold is zero, one where the *veto* threshold v_i is varied and the self-criticism threshold is zero, and one where both are varied in synch ($s_i = v_i$).

Figure 3a shows how the novelty of artifacts selected to the domain varies as a function of the threshold. The immediate and expected observation is that a higher threshold increases the novelty of artifacts.

It is more interesting to compare the three curves. Among them, using a threshold for self-criticism has the smallest effect, while using a *veto* threshold has a much more pronounced effect. In the case of *veto* power, the effect of the threshold is multiplied when it is applied by multiple agents, even if they are on average less informed of the kind of artifacts produced by an agent than the agent itself. The result speaks for the “wisdom of the crowd”. The effect of using both thresholds is practically equal to just using *veto* with the same threshold.

Figure 3b shows the corresponding amounts of efforts required to produce 200 novel artifacts to the domain. The results are very sensitive to the *veto* threshold: the required effort grows suddenly at a certain point while the self-criticism threshold has at the same point almost no effect.

The conclusions from panels 3a and 3b are two-fold. First, the use of *veto* power and self criticism can improve the novelty of results significantly without increasing the effort needed. Second, however, an excessive *veto* threshold can have a sudden negative effect on the efforts. This is at least partially due to our application, spirographs, and how the generating function can only generate certain types of

images causing the distance between any two images to cap at ~ 4.5 .

Figure 3c provides further insight into the use of resources when the thresholds change, by showing how many artifacts on average pass the threshold(s) per iteration. Obviously, higher thresholds reduce the amount of valid candidates. In our setting, at a *veto* threshold of 3.84 the number of valid candidates drops approximately to 0.5 artifacts per iteration, causing a deep increase in the number of iterations needed to produce the required number of artifacts to the domain (panel 3b).

The most interesting result here is the effect of self-criticism: it controls the number of candidate artifacts submitted, reducing the efforts invested by the society to evaluating and selecting candidates to the domain. It turns out that self-criticism behaves nicely: its use improves novelty (3a) without increasing the number of iterations much (3b), but most importantly it can effectively reduce the collective evaluation effort of the agent society (3c).

Voting method. In addition to the ‘best mean’ voting method to choose one of the candidates from C^* to add to the domain, we also experimented with several other voting methods, namely ‘best singular’, ‘least worst’ and ‘instant run-off voting’ (IRV). In ‘best singular’ voting, an artifact with the highest single agent’s novelty calculation is chosen. ‘Least worst’ can be seen conceptually as a variant of the *veto* mechanism: it chooses an artifact which has least worst single novelty calculation. In ‘IRV’, agents first rank all candidates to a preference order, and then proceed to recursively prune candidates from the rankings based on which are not in the first place in any of the already pruned ranking lists.

Our empirical results with these alternative voting methods (not shown) indicate that ‘best mean’ clearly outperforms ‘best singular’ and ‘least worst’ methods and is on par with ‘IRV’. We use ‘best mean’ because of its simplicity.

Domain memorization In our model, agents have a limited memory of both their own experience and of artifacts in the domain. In each iteration, an agent accesses k artifacts in the domain and uses them to replace the oldest artifacts in the agent’s memory. We experiment with mem-

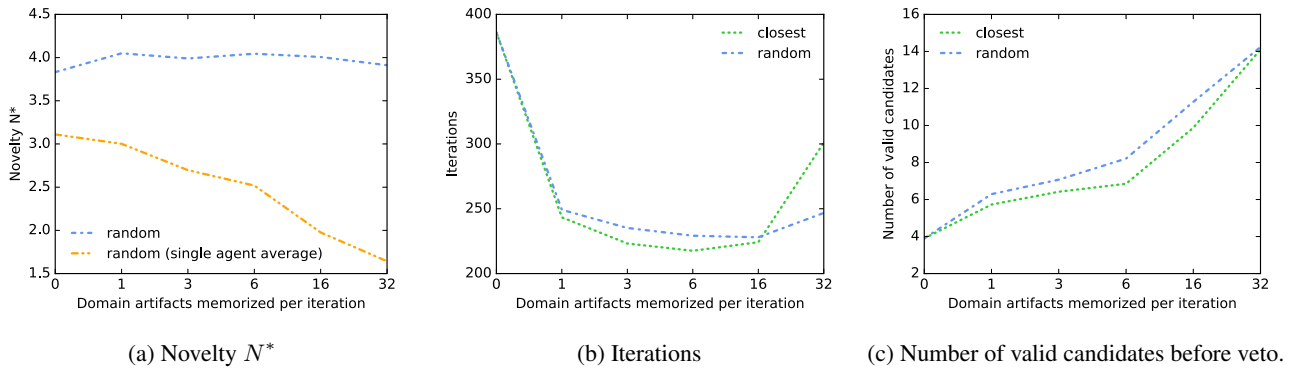


Figure 4: Effect of domain artifact memorization on novelty N^* of the domain compared to that of a single agent (4a), on the effort needed to produce 200 novel artifacts to the domain (4b) and on the number of artifacts passing the self-criticism threshold (4c).

orization techniques closest_k and random_k (as explained in section Multi-Agent Architecture) in a setting of 16 agents, each with 32 slots of memory, and the number of memorized items varying as $k \in \{0, 1, 3, 6, 16, 32\}$. Obviously, with $k = 0$ there is no memorization from the domain and the agents generate artifacts independently. The results are shown in Figure 4.

The upper line in Figure 4a shows that k has practically no effect on N^* (for clarity, we show random_k only, as the behavior of closest_k turned out to be practically identical; see Discussion). The lower line shows for different memorization settings the average novelty of a single agent, i.e. N^* computed from the candidate artifacts an agent has produced itself. In contrast to the overall novelty (the upper line), a larger value of k has a negative effect on the average performance of a single agent, which plunges to about $1/2$ when $k = 32$. This is expected: as k grows, an agent has less memory about its own products (at most one own artifact per k artifacts from the domain) and therefore is more prone to produce similar artifacts again.

Figure 4b shows the efforts needed to produce 200 domain artifacts. We observe that any amount of memorization produces the artifacts in about $2/3$ of the iterations compared to what $k = 0$ needs, but the memorization strategy does not seem to have much impact. The effort needed is at its lowest when $k \in \{3, 6\}$, and rises somewhat at $k = 32$ when the agent’s whole memory is repopulated at each iteration.

Figure 4c shows the average number of candidates that passed an agents’ self-criticism on each iteration. The curves are strictly increasing with k , suggesting that memorization of domain artifacts has a positive effect on guiding a single agent’s search.

Overall, the memorization with a conservative k (in our case $k \in \{3, 6\}$) has a positive effect on the society when comparing to $k = 0$ as the multi-agent system performs more efficiently as a whole (4b). The optimum appears to be a compromise: with very low k the society takes more time to produce the domains artifacts (4b) while high k overrides the self-criticism (4c) as the agents do not remember their own artifacts, lowering their own individual novelty (4a).

Discussion

We discuss selected technical aspects, reliability of the results we obtained, and paths towards creative multi-agent systems.

Population size With random initialization, smaller populations are clearly more prone to system-wide aberration (higher iteration counts) as all agents might be initialized into unproductive areas of the parameter space. Increasing the number of agents improves the average effectiveness of our multi-agent systems as at least some agents are more likely to be instantiated in (or at least near to) the productive areas.

Selection of candidates to the domain At a first sight, self-criticism and *veto* power seem to be surprisingly effective: self-criticism lowers the amount of collective effort needed to choose domain artifacts, and *veto* increases their novelty. However, in our setting each candidate artifact still needs to be evaluated by all agents. As a future work, it would be useful to revise the domain selection procedure to be more local in order to acquire better scalability.

The effects of the population size and social decision-making methods in our experiments are similar to what Sosa and Gero (2005) report. In small populations the effect of interaction between individuals is limited because of the low number of agents, and larger populations take more time to form a consensus. In our experiments this is reflected in how smaller populations do not reach as high overall novelty for the artifacts (Fig. 2a), and the time to reach a certain number of artifacts grows in larger populations as more agents exhibit their right to use *veto* power (Fig. 2b).

Memorization The two memorization strategies introduced, random_k and closest_k , behave nearly identically in the experiments, although one could think that the more informed closest_k would guide the agent’s search more effectively. Our initial examination suggests that the identical behavior might be influenced by two different reasons. First, the topology of the parameter space in our experiments is complex: a small change in the parameters can cause a rapid change in the artifacts. This fluctuation might inhibit closest_k from guiding the search effectively. Second, the

number of memory slots that the society collectively has is quite large compared to the amount of domain artifacts generated. This might imply that there is enough memory for random_k to continuously sample a representative set of the domain items into the society's collective memory.

Reliability of the results Our results have been obtained through simulations that involve randomness. While randomness certainly has a high role in the suggested system, the behavior between different runs with same system settings is stable enough to make conclusions from the results.

A more important issue is how specific the results are to spirographs. Spirographs are a good test case in their complexity: sometimes even small movements in the parameter space can cause big changes in the resulting spirographs, while there also are large areas producing essentially the same result.

To test if our results hold in other domains, we experimented with agents that searched for different colors in an image, and found qualitatively similar results. In particular, the dependency of novelty and iterations on the threshold for criticism had a similar form as in Figures 3a and 3b. There appeared to be a turning point in the threshold, above which novelty is higher and the number of iterations turns into steep increase. The reason for this effect may be that the domain becomes 'saturated' in the sense that the probability of finding novel enough artifacts rapidly decreases.

Creativity vs. novelty Saunders and Gero (2001a) propose agents that have a bell-shaped hedonistic curve as a function of novelty. Such a curve can be motivated by the value related to novelty (very familiar artifacts are of no new value) and of utility of that novelty (very strange artifacts cannot be utilized). Our novelty-seeking agents just look at one side of this, since our goal has been specifically to create novel artifacts. Adding aspects of value will change the model, possibly resulting in something similar to the hedonistic curve.

The ultimate goal is to develop creative agent systems. While we have only been dealing with novelty here. Formally, a minimal addition to the current system to make the agents more creative is that each agent also has function $E(A)$ which calculates the value or aesthetics of the artifact. We could then use both the novelty and aesthetics in the voting process. They both might have their own thresholds, but aesthetics probably should not be so heavily vetoed as aesthetics is much more subjective than novelty.

Conclusions

Novelty is a key criterion for creativity (Boden 1992). We have described and evaluated a novelty-seeking multi-agent architecture as a step towards creative multi-agent systems.

Our evaluation shows that a society of novelty-seeking agents can be more productive in generating novel artifacts than a single-agent or monolithic system. Obviously, a larger number of agents can be more effective in exploring the search space.

We found out that self-criticism and *veto* power can be powerful features in novelty-seeking agent systems. Self-criticism of agents can reduce the collaborative effort in evaluating candidate artifacts, while *veto* is an effective way to collaboratively reject candidates that are not novel.

Future work for developing the novelty-seeking agent architecture has numerous possible directions. First, agents could interact in numerous ways, in particular exchanging coordinates, artifacts and their evaluations. Second, agents could be adaptive to their own experience as well as to the society, e.g. by adjusting their random walk step size, self-criticism, and use of *veto* power. Third, emergence of social phenomena like community structure would be interesting to study, and also to apply in making candidate artifact selection more local and thereby more scalable. Fourth, experiments in more domains are needed.

In our efforts to study and understand creative agent systems, the next big question will be to consider seeking both *novel* and *valuable* artifacts.

Acknowledgments

This work has been supported by the European Commission under the FET grant 611733 (ConCreTe) and by the Academy of Finland under grant 276897 (CLiC).

References

- Boden, M. 1992. *The Creative Mind*. London: Abacus.
- 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 University Press. 325–339.
- Gabora, L., and Tseng, S. 2014. The social impact of self-regulated creativity on the evolution of simple versus complex creative ideas. In *Proceedings of the Fifth International Conference on Computational Creativity*, 8–15. Ljubljana, Slovenia: Josef Stefan Institute, Ljubljana, Slovenia.
- Kohonen, T. 1995. *Self-Organizing Map*. Berlin: Springer.
- Lehman, J., and Stanley, K. O. 2008. Exploiting open-endedness to solve problems through the search for novelty. In *Proceedings of the Eleventh International Conference on Artificial Life (ALIFE XI)*, 329–336. Cambridge, MA: MIT Press.
- Saunders, R., and Bown, O. 2015. Computational social creativity. *Artificial Life* 21(3):366–378.
- Saunders, R., and Gero, J. S. 2001a. A curious design agent: A computational model of novelty-seeking behaviour in design. In *Proceedings of the Sixth Conference on Computer Aided Architectural Design Research in Asia (CAADRIA 2001)*, volume 1, 345–350. Sydney, Australia: CAADRIA.
- Saunders, R., and Gero, J. S. 2001b. The digital clockwork muse: A computational model of aesthetic evolution. In *The AISB'01 Symposium on AI and Creativity in Arts and Science, SSAISB*, 12–21. York, UK: AISB Press.
- Shoham, Y., and Leyton-Brown, K. 2009. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. New York, NY, USA: Cambridge University Press.
- Sosa, R., and Gero, J. S. 2005. Social models of creativity. In *Proceedings of the International Conference of Computational and Cognitive Models of Creative Design VI*, 19–44. Heron Island, Australia: Key Centre of Design Computing and Cognition, University of Sydney, Australia.