

Evolving Virtual Embodied Agents using External Artifact Evaluations

Lesley van Hoek, Rob Saunders, and Roy de Kleijn

Leiden Institute of Advanced Computer Science, Leiden University, Leiden, the Netherlands

{l.s.van.hoek@umail, r.saunders@liacs, kleijnrde@fsw}.leidenuniv.nl
<http://www.cs.leiden.edu>

Abstract. We present *neatures*, a computational art system exploring the potential of digitally evolving artificial organisms for generating aesthetically pleasing artifacts. Hexapedal agents act in a virtual environment, which they can sense and manipulate through painting. Their cognitive models are designed in accordance with theory of situated cognition. Two experimental setups are investigated: painting with a narrow- and wide perspective vision sensor. Populations of agents are optimized for the aesthetic quality of their work using a complexity-based fitness function that solely evaluates the artifact. We show that external evaluation of artifacts can evolve behaviors that produce fit artworks. Our results suggest that wide-perspective vision may be more suited for maximizing aesthetic fitness while narrow-perspective vision induces more behavioral complexity and artifact diversity. We recognize that both setups evolve distinct strategies with their own merits. We further discuss our findings and propose future directions for the current approach.

Keywords: aesthetic evaluation · artificial intelligence · artificial life · autonomous behavior · computational creativity · embodied agents · evolutionary art · neural networks · situated action · situated cognition

1 Introduction

Computational systems that produce artworks with high levels of autonomy have always provoked discussion about the definition of art and creativity. Researchers and artists working in the field of evolutionary and generative art cede control to autonomous systems that produce artworks, often intending to eliminate human intervention where possible [17]. Digital evolution is an established algorithmic process that has proven very capable of innovation [18]. In art and design, appropriate implementation of this technique can aid the generation of novel, valuable and surprising artifacts [4][2] that may be deemed creative by unbiased observers [8]. It has also been essential in the field of artificial life (a-life) [26] where researchers have been consistently surprised by creative solutions invented by artificial organisms evolving in computational environments [28]. Naturally, the process of digital evolution merely imitates life itself. The biological mechanism of natural selection is known to find and cause inventive adaptations that

2 L. van Hoek et al.

enhance the survival and reproduction of organisms [15]. Consequently, these may lead to the appearance of design without a designer [10]. Adaptations may include changes in behavior. We aim our attention at a particular behavior in some non-human organisms, namely the creation of artifacts.

Several species in the natural world are known to decorate and produce structures that resemble visual art in the sense that they are intended to be attractive to potential mates. This structure creation is an important behavioral characteristic of male bower birds [12] and white-spotted pufferfishes [31]. In this paper, we explore whether artificial organisms could adapt to similar, but digitally induced pressures as a consequence of constructing artifacts. The following sections briefly discuss some challenges related to building such a computational system.

1.1 Computational aesthetic evaluation

Early examples of evolutionary art include the highly influential work of Sims [41] and Latham [47], who used genetic algorithms to mutate symbolic expressions for the composition of unpredictable yet interesting visual shapes and patterns. Both adopted a top-down approach that relies on human aesthetic judgment for the evaluation of artifacts using an interactive genetic algorithm (IGA). This technique facilitates easy exploration of large parameter spaces [44] but suffers from significant limitations: (1) IGAs rely on human evaluation at every iteration and so suffer from the *fitness bottleneck* [46], and (2) human fatigue and inconsistency make it difficult to capture universal measures [44]. Attempts to overcome these limitations have included massively multi-user systems [39] and the application of machine learning to capture user preferences [33].

Challenges in IGA helped inspire the research field of computational aesthetic evaluation (CAE), where people seek computational solutions for the assessment of human aesthetics [23]. Machado and Cardoso [29] created *NEvAr*, an autonomous system that evolves Sims' symbolic expressions with an automated evaluation procedure for images that focuses exclusively on form. Here, a speculative fitness function inspired by the study of *information aesthetics* [35] was designed which favors images that are "simultaneously, visually complex and that can be processed (by our brains) easily". In the science of aesthetics, *NEvAr*'s fitness function indicates a *formalist theory* as it proposes aesthetic experience relies on the intrinsic beauty of the artifact. In contrast, a *conceptual theory* relies on other factors that may be more important for aesthetic preference like socio-cultural contexts of the work and the previous experience of the artists and observers [40]. In a more recent publication, Redies [36] proposes a model of visual aesthetic experience that unifies these two theories. Ultimately, there is currently no agreement on which paradigm offers the most effective computational framework of human aesthetics.

1.2 Embodiment

Theorists in situated cognition view the environment as highly significant to driving human cognitive processes. Clark and Chalmers [7] suggest that the

environment directly influences an agent’s behaviors as part of a two-way interaction between action and perception. Here, embodiment is key because it allows us to manipulate it to our needs. Biological brains have evolved to take advantage of the environment by offloading cognition to it through the body. Simultaneously, our visual systems evolved to rely on it more. This perspective supports the view of *externalism*, in which the cognitive process is considered something that occurs in- and outside of the mind [7]. In this context, embodiment is key to the creation of art and can be imagined as a feedback loop of action and perception occurring through a body. Brinck [5] states that the production (and consumption) of visual art can be accounted for by the theory of situated cognition [6]: ”Artist and canvas form a coupled system. Artistic practice starts with gaze, and then comes the gesture that accomplishes itself when the artist is in touch with the piece [they are] working on.” [5]

Experiments in the use of embodied artificial organisms and situated cognition for computational art and creativity have largely been unexplored. Thus, we present *neatures*: a prototype for an autonomous art system that simulates artificial organisms capable of producing visual art in their environment.

2 Related work

There have been several interesting art and research projects involving the use of embodied agents to create visual art. Jean Tinguely experimented with mechanical drawing machines in the 50s, exploring notions of automated artists and artificial creative processes [13]. Influences to his work can be seen in the field of swarm painting, which involves the simulation of agents supplied with some form of cognition producing emergent artworks. *Robotic Action Painter* [34] is an autonomous abstract art system based on behavioral studies of ants and other social insects. An artwork is created by employing several small wheeled robots that leave colored lines (pheromone) as they travel. A color detection sensor on each robot recognizes these lines in the environment and triggers specified behaviors for particular colors—a process analogous to *stigmergy*; a form of self-organization [14]. The result is a painting with chaotic structures that are free from preconceptions and merely represent the actions themselves. McCormack developed similar experiments using biological processes of niche construction to enhance the diversity and variation of agents’ behaviors in his art system [32].

Drawing machines that take a more anthropomorphic approach can be classified as robot painters. *eDavid* [11] is an industrial robot that simulates the human painting process using a visual feedback loop to explore painterly rendering on a real canvas. Explorations in expanding its artistic skill demonstrated the possibility of expressing a given collection of images in a different style [48]. With *neatures*, we take inspiration from the flexibility of robot painters and the emerging complexity of swarms to explore the effects of aesthetic selection pressures in an evolutionary art system.

4 L. van Hoek et al.

3 Implementation

*Neatures*¹ is a prototype computational visual art system that was developed in an attempt to employ artificial organisms for the production of visual artifacts. The current implementation is heavily inspired by the seminal work of Sims (1994), *Evolving Virtual Creatures*, in which a genetic algorithm was used to guide the evolution of specific abilities such as locomotion and jumping. *Neatures*' artificial organisms 'live' in three-dimensional space and are subject to physically plausible simulation. This is achieved using the Bullet physics engine [9]. The software comprises of a controller server which stores the population and commands the complete evolutionary process. A simulator client can connect to a controller and receive queries for queued rollouts. This component features a graphical user interface, allowing the user to observe the virtual organisms in real-time. The following sections briefly cover the system implementation.

3.1 Agent morphology

Virtual organisms situated in physically plausible environments are subject to strict laws of physics and, like real organisms, require an appropriate body to fulfill their purpose. Designing such a body is a difficult task, and perhaps best suited for an evolutionary process to solve. Sims [41] used a genotypic encoding of nodes and connections for the morphology of his creatures, and genetic operators, allowing for the evolution of morphology alongside control policy. In this system, a genotypic encoding scheme is used to generate a hexapod at the start of a simulation and remains fixed. The reason for this is that evolutionary optimization of morphology dramatically increases the complexity of the search landscape and is incompatible with fixed-topology neural network architectures.




Fig. 1: Agent morphology genotype (left) and phenotype (right).

Each element stores some information about their phenotypic transformations such as size, attachment points, and node or joint type. A phenotype generation algorithm recursively traverses the graph and builds a hierarchical

¹ *Neatures* is open-source and available at <https://github.com/lshoek/creative-evo-simulator>

structure of boxes connected to each other by joints. Fig. 1 depicts the morphology encoding and phenotype of a hexapod. The algorithm in the current work was implemented after Krčah’s example [25] with some alterations tailored to suit this work’s purpose. One notable difference, for instance, is that we use a single degree of freedom per joint for simplicity.

3.2 Agent control policy

In every simulation rollout, agents are tasked to produce an artifact in their environment. In order to achieve this in *neatures*, we chose to implement a painting system. Each agent is equipped with a single brush-type node capable of applying virtual ink drops to the canvas; a specified surface area in the environment that the agent can sense and manipulate. Four invisible walls are located at a specified distance from the canvas edges to prevent agents from moving too far away from the center. Ink is only released under the conditions that the brush node is in contact with the canvas, and the agent has decided to activate it.

An agent’s decision-making process and behavior are determined by its control policy. This is defined by a neural controller that continuously accepts sensory data as input, and based on this data, outputs a set of activation values. Agents sense their environment through two types of sensors: (1) a proprioceptive sensor, implemented by tracking the current joint angles and storing these in $a \in \mathbb{R}^j$, where j is equal to the number of joints in the agent’s morphology and (2) a vision sensor capturing a 64x64px grayscale bitmap representation of the current canvas’ content. The data of both sensors is appended to form an observation to be fed to the neural controller at regular time intervals. The physics engine and control policy are updated 60 and 20 times per second of simulated time, respectively. Fig. 2 presents the complete cognitive model of an agent.

The neural controller involves two cognitive modules; a vision model V for processing visual data inside the incoming observation, and an action model C to generate the agent’s next action. V is a *convolutional variational autoencoder* (CVAE), pre-trained to compress the canvas data to a latent vector $z \in \mathbb{R}^{32}$. C

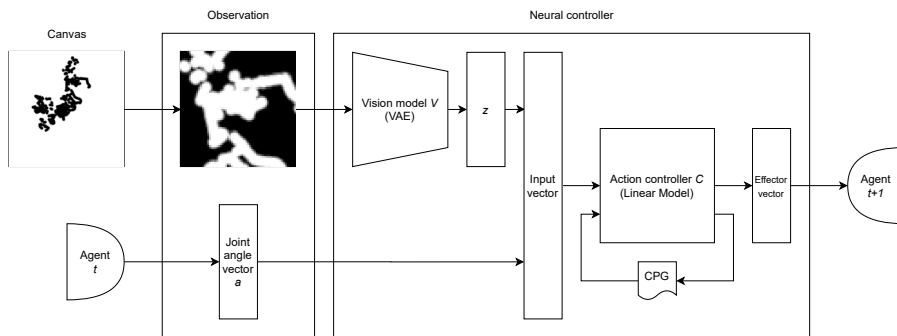


Fig. 2: Cognitive model of an agent.

6 L. van Hoek et al.

is a simple linear model that takes as input a combination of latent vector z , a joint angle vector, and an additional value to stimulate continuous movement. Compression of the visual data allows the action controller to be kept small, which alleviates the credit assignment problem in difficult reinforcement learning (RL) tasks and tends to iterate faster [19]. The output layer of C uses a \tanh activation function to output to produce a vector of effector values, including target joint angles used to update the motor parameters of the agent’s joints and a value indicating the stroke width of the brush. Finally, a stimulation output value connects to a central pattern generator (CPG) after which a feedback connection to the corresponding action model input is made for the next time the neural controller is queried [24]. This minimal recurrent network structure is set up this way to evoke changing joint angle outputs. Without it, the agent would cease to move in cases where its observations remain unchanged over multiple frames and its body incidentally has zero momentum. Additionally, as sensory input drives neural excitation, it grants C control over the agent’s movement speed, which could bring about more interesting behaviors. Section 4.3 describes the training procedure for V and C .

4 Experiment

We carry out two experiments where an artificial organism is evolved by optimizing for the aesthetic quality of its artifacts. The artistic medium of expression chosen for this task is painting. The main reason for this is that there exists a multitude of interesting theories and evaluation techniques of visual human aesthetics—suitable for two-dimensional content—that could be pursued to design an acceptable fitness function [16].



Fig. 3: The *neatures* simulator showing an agent painting.

As stated in Section 3.1, we decided to exclude morphology from evolutionary optimization, meaning we must formulate an appropriate body design for

the current experiment ourselves. We take inspiration from behavioral robotics research, where it has long been common practice to use biologically based robot designs to study artificial organisms [1]. As a matter of course, the insect-like hexapod was chosen for the current task. This design is a popular benchmark that we suppose will allow for an adequate degree of flexibility required to explore the possibilities of the virtual environment. Fig. 3 shows a screenshot of the agent as it appears in the simulator client of the system.

4.1 Setup

The following is a brief description of the realized experiments. In the first setup, the agent is supplied with a wide-perspective vision sensor. This is defined as a 64x64px grayscale bitmap representation of the environment that is equal to the size of the canvas. The orientation of this representation is at all times aligned with the facing direction of the agent and centered around the point where it last touched the canvas with its brush node. Fig. 4a shows an example of how the canvas is sensed with this perspective. The second setup supplies the agent with a narrow-perspective vision sensor, encompassing 6,5% of the canvas area as shown in Fig. 4b.



Fig. 4: The mapping from canvas (left) to visual field (right), marked in red, for wide-perspective (4a) and narrow-perspective (4b).

The vision capabilities of the agent exist in a separate conceptual space from the one it is situated in. Agents' visual capabilities exist in *artifact space*, whereas their neural controllers output actions in *effector space*. The former is a two-dimensional representation of the environment, cultivated by the agent itself. The latter relates to objects in the three-dimensional virtual environment. Other than muscle memory (the action controller parameters), an agent has no other capabilities of memorization. As a result, the environment is the only cognitive resource to the agent by which an approximate model of situated cognition is realized. The key idea to this experiment is that, under the given conditions, a mapping between these two may be learned. If successful, the creature would be able to produce an aesthetically pleasing artifact in *artifact space* by means of its motor function in *effector space*.

8 L. van Hoek et al.

4.2 Measuring aesthetic quality

After a rollout has ended, the resulting artifact is queued for fitness evaluation. In our computational environment, the fitness function is a proxy for natural selection pressures that cause the evolution of adaptations [15]. As outsiders to this virtual world, we can design this function externally, and observe what behaviors emerge from evolutionary optimization. Taking inspiration from some animal species’ mate selection indicators that are attributed to external artifacts, we intentionally ignore any behavioral aspects of an agent’s existence. Our fitness function is designed to evaluate images in accordance with speculative visual aesthetic theory, essentially assuming the role of an art critic.

To measure the aesthetic quality of an artifact, we use a metric closely related to Birkhoff’s [3] formalist aesthetic measure, defining the formula $M = O/C$, where M is the aesthetic effectiveness, O is the degree of order and C is the degree of complexity. Birkhoff theorizes that aesthetic response to an object is stronger when the degree to which psychological effort is required to perceive it—induced by its complex features—is met with a higher degree of tension being released as the perception is realized—originating from orderly features such as symmetry and self-similarity. This formula has been disputed early and is generally regarded as inaccurate [49]. Scha & Bod [37], for instance, note that it penalizes complexity too considerably and is better suited as a measure of the degree of self-similarity. Galanter [16], however, notes that at least two aspects of Birkhoff’s work remain legitimate today; the intuitive connection between aesthetic value and order/complexity relationships, and the search for a neurological base of aesthetic behavior. These aspects are reflected in the fitness function of Machado & Cardoso [29], defined in Eq. 1. Inspired by information aesthetics [35], Machado & Cardoso speculate an image’s intrinsic aesthetic value to be equal to the ratio of image complexity IC to processing complexity PC .

$$reward_{aesthetic} = \frac{IC}{PC} \quad (1)$$

PC is measured at two temporal instances (t_0 and t_1) in the time it takes to perceive an image and provide Eq. 2. The processing complexity is maximized as PC_{t_1} and PC_{t_0} approach each other.

$$PC = (PC_{t_0}PC_{t_1})^a \left(\frac{PC_{t_1} - PC_{t_0}}{PC_{t_1}} \right)^b \quad (2)$$

In order to find PC_{t_0} and PC_{t_1} , we calculate the inverse of the root mean square error (RMSE) between the original image i , and the same image after fractal compression $Fractal(i)$, as shown in Eq. 3.

$$PC_{tn} = \frac{1}{RMSE(Fractal(i), i)} \quad (3)$$

Machado et al. [30] compared several complexity measures with human ratings across a selected set of images in five distinct stylistic categories. Among

the results of their feature extraction experiments, their JPEG-Sobel method was found to correlate the most with human ratings, especially those related to the abstract artistic category. We calculate IC following this method as shown in Eq. 4. First, the Sobel [42] edge detection operator is applied to i horizontal and vertical directions, after which the resulting gradients are averaged. Then, JPEG compression is performed on the edges. In the dividend, $size$ defines the total number of bytes required to store the image data.

$$IC = \frac{RMSE(Sobel(i), JPEG(Sobel(i)))}{size(Sobel(i))size(JPEG(Sobel(i)))} \quad (4)$$

Taylor et al. [45] note the fractal qualities of late-period action paintings by Jackson Pollock and suggests their fractal dimensions are correlated with their aesthetic qualities. Therefore, we decided to parameterize Eq. 2 using $a = 0.6$ and $b = 0.3$, increasing bias towards artifacts with more orderly features with respect to the reference implementation [29]. We argue that this suits the current experimental setup by countering excessive levels of image complexity in the artifacts due to the generally chaotic nature of agents' behaviors that generate complex and incidental painting patterns by default.

In early experiments, we found that additional encouragement to act through an easily attainable coverage reward could help agents to advance faster in early generations. This has the added benefit that a minimum specified amount of content is imposed on the artifacts. Eq. 5 defines $reward_{coverage}(x)$, where x is the mean of all normalized pixel intensities of the artifact and p is the peak coverage rate. It is essentially a smooth interpolation between x and p , ensuring a result of 1 when $x \geq p$.

$$reward_{coverage}(x) = 1 - \sin\left(\pi \frac{\frac{1}{p}x + 1}{2}\right)^4 \quad (5)$$

with initial condition

$$x = \min(x, p) \quad (6)$$

In our experiments, we use $p = 0.0625$, meaning that the maximum coverage reward is already reached when 6,25% of the canvas area is painted. Finally, the total artifact fitness is calculated as defined in Eq. 7. This shows the aesthetic reward is proportional to the coverage reward until peak reward p is reached, thus penalizing paintings that have little content. Table 1 presents a set of images and their fitness values.

$$fitness = 100 \cdot reward_{coverage} + reward_{aesthetic} \times reward_{coverage} \quad (7)$$

We find these results to be satisfactory for our purposes. Although the fitness function is arguably too generous on Gaussian noise (Table 1d), such an artifact is practically impossible for an agent to produce. The Pollock-snippet (Table 1e) is evaluated far more positively and represents a more plausible result.

10 L. van Hoek et al.

Table 1: A set of images and their fitness: (a) perfect symmetry, (b) an early-generation artifact with little variability in stroke width, (c) an early-generation artifact with high variability in stroke width, (d) Gaussian noise, (e) a contrast-enhanced snippet of No. 26A: Black and White by Jackson Pollock (1948).

	(a)	(b)	(c)	(d)	(e)
<i>Fitness</i>	101.2	116.2	145.9	399.5	871.5
<i>Coverage</i>	27.7%	15.5%	12.4%	18.2%	46.0%
<i>IC</i>	0.0697	0.4042	0.8046	44.269	48.443
<i>PC</i>	0.0561	0.0250	0.0175	0.0148	0.0063

4.3 Training procedure

Before any control policies can be evolved, visual model V must be pre-trained to discern between visual observations. First, 20,000 artifact samples (256x256 grayscale bitmaps) were collected in a preliminary run using an untrained visual model V . Then, a new dataset was generated by applying random affine transformations to each collected sample. This new dataset is more representative of an agent’s visual observations. Finally, using the updated dataset, V was trained to encode visual observations into latent vector $z \in \mathbb{R}^{32}$ for 200 epochs.

Agents’ control policies are optimized through evaluation of the quality of their work, rather than the means by which it was achieved. This indirect correspondence between goal and action may reduce credit assignment accuracy of gradient-based numerical optimization algorithms as adaptations to action controller C could have unanticipated effects on an artifact’s fitness. Therefore, gradient-free methods such as evolution strategies [38] might be best suited for solving this problem. *Neuroevolution* methods have a long history of success with evolutionary robotics and have recently increased in popularity as they have been found to perform considerably well on deep RL tasks [43]. With this in consideration, we chose *covariance matrix adaptation evolution strategy* (CMA-ES) [20] for the optimization of C ’s parameters. Evidence shows that the algorithm performs relatively well on deceptive landscapes or sparse-reward functions up to a couple of thousands of parameters [22]. We use an open-source Python implementation of the algorithm by Hansen [21].

At the start of every evolution process, the weights of every action controller C in the population are randomly initialized with $\mu = 0$ and $\sigma = 0.1$. A population size of 32 is used, where each candidate’s behavior is determined by their corresponding C , comprising 658 trainable parameters each. Every generation, one rollout is performed per agent and results in 32 artifacts. A rollout is defined as 240 seconds of simulated time an agent spends in the environment. Evaluations occur immediately after each rollout in a separate process. After all rollouts and evaluations are finished, CMA-ES uses the collected fitness values to update

each candidate’s action controller parameters for the next generation. Both experiments are performed using an evolutionary process of 350 generations.

Our training setup marks several notable limitations. Foremost, the experiments are carried out separately on two mid-range laptops (i7-7700HQ/GTX1050 and i7-8750H/GTX1070), each running a single simulator client and controller server at the same time. The most significant bottleneck comes from the fractal compression procedure required for each artifact evaluation. In the current setup, we simulate two populations of 32 candidates for 350 generations and takes about 40 hours to complete. More reliable results could be collected by increasing the population size and averaging fitness over multiple rollouts for a more representative metric of the agent’s general painting strategy. This is however outside of the scope of this research.




Fig. 5: Fitnesses of the narrow- (left) and wide-perspective populations (right).


5 Results

Fig. 5 presents the fitness results of the narrow- and wide-perspective vision experiments. Here, we see that the narrow-perspective population’s mean fitness starts with a steep positive trend and converges towards a local optimum before the 50th generation. The wide-perspective population’s mean fitness improves gradually up to around the 100th generation before a local optimum is reached. We also see that the wide-perspective population is generally about 150 points ahead of the narrow-perspective population. From these results, it is evident that the wide-perspective population performs better in terms of fitness. However, it barely shows any signs of improvement after a local optimum has been reached, until the final generation of the simulation. This is unlike the narrow-perspective population, which shows a slight upward trend around the 300th generation, and some new best-ever artifacts of the population. Table 2 presents the highest-rated artifacts of both experiments along with some key statistics. Almost every artifact shows a clear trajectory on the canvas that is telling of the strategy that was used to produce it. Fig. 6 below shows the highest-rated artifacts of the first 64 generations of both populations. We see that the sort of artifacts produced by both populations can easily be distinguished from approximately the 40th generation. From there on, we see that nearly all artifacts of the wide-vision population

12 L. van Hoek et al.

Table 2: The highest-rated artifacts of all generations of wide-perspective and narrow-perspective populations and their key statistics.

<i>Pers.</i>	<i>Fitness</i>	<i>Cov.</i>	<i>IC</i>	<i>PC</i>	<i>Gen.</i>
Wide	401.05	36.05%	2.8339	0.0094	194
Narrow	251.82	21.33%	1.7714	0.0116	335



Wide Narrow

indicate a circular movement strategy, with little diversity among paintings. The fitness results and artifacts of this population show that this strategy is further exploited in subsequent generations, likely because of its effective contribution to maximizing fitness. In contrast, the narrow-perspective population struggles to escape a local optimum early on but demonstrates far more diversity among its artifacts in all generations. This suggests that potentially fit strategies are being explored rather than being exploited.




Fig. 6: The best artifacts of the first 64 generations (top-left to bottom-right) of the narrow- (left) and the wide-perspective population (right).

The discrepancy between the fitness results and the type of artifacts produced by both populations led us to believe that coverage and fitness may be strongly positively correlated. To investigate, we plotted coverage against fitness (Fig. 7) and observed that coverage is an accurate predictor of fitness in the wide-perspective population, but not necessarily for the narrow-perspective population.




Fig. 7: Mean coverage and fitness in narrow- (left) and wide-perspective populations (right).

6 Discussion

Our results of the current experiment demonstrate a notable distinction between the narrow- and wide-perspective setup. In our experiment, we observe that virtual organisms with narrow-perspective vision trigger explorative search of the fitness landscape by the evolutionary algorithm and demonstrate more complex and distinct behavior. We also see that this is not necessarily in the interest of maximizing fitness. One explanation for this could be that relatively small adaptations to a narrow-perspective controller’s weights lead to greater variations in the emerging painting strategy. In the agent’s cognitive model, perception and action are closely coupled together. Therefore, distinct actions may be more likely to be triggered when visual observations are more volatile, as is the case with the narrow-perspective agents. This is in line with Brinck’s [5] argument that art creation is a situated activity, noting that what the artist perceives is directly transformed into action. We further observe that narrow-perspective agents generally appear more sensitive to the environment in their painting strategies than wide-perspective agents. Narrow-perspective agents show more effective corrective behavior such as turning near the edge of the canvas. This is not as apparent in wide-perspective agents who barely appear to discernibly change their behavior near edges. Little response to edges is likely induced by the exploitation of circular movement patterns—evidently an effective strategy for painting highly fit artifacts. We further think that the widespread coverage of paint in the environment reinforces an agent’s behavioral pattern. This may be due to the relatively poor compression quality of global features in visual observations of developed circular patterns, leading to similar encodings of z . Incidentally, this fact may have greatly contributed to finding the circular movement strategy.

From our observations, we theorize that volatile visual information, as demonstrated by the narrow-perspective experiment, considerably complicates the shape of the fitness landscape. For instance, a consistent circular movement strategy would be much more difficult to sustain over the length of a rollout, and over multiple generations, with narrow-perspective vision than with wide-perspective vision. Even more so, this automatically concerns any potential strategy. Although volatile visual information may impede the evolution of consistent action and perception, it does have creative merit in the sense that it elicits greater

14 L. van Hoek et al.

behavioral complexity in agents. Hence, the narrow-perspective population has explored the greatest *artifact space*. This is demonstrated in Fig. 8 which presents two random selections of artifacts created in both populations.




Fig. 8: Two random selections of artifacts drawn from all of the narrow-perspective population (left) and the wide-perspective population (right).

Considering our evaluation procedure; if we, hypothetically, consider Pollock’s work as an aesthetic benchmark for this system (Table 1e), we consider the current fitness function helpful at guiding agents’ technique towards this aesthetic up to a certain point. Fig. 7 however suggests a possible perverse instantiation problem; at least one strategy exists in which coverage can be exploited to maximize fitness. However, we believe an adjustment to the fitness function would be premature. This is because, as the fitness function is based on complexity, coverage cannot be positively correlated with fitness as it approaches 50%. The highest recorded coverage of all artifacts in both populations is 36%, whereas the coverage of our Pollock example (Table 1e) is measured at 46%. We are confident that under the current time pressure of 240 seconds, it is physically not possible for agents to cover a significantly greater part of the canvas. Therefore, we believe that agents should be assigned sufficient time so that 50% coverage could be achieved. After this is explored, we believe that a worthwhile addition to the fitness function would be a novelty reward term to overcome local optima by encouraging exploration [27].

In our experiment, we see that a proxy for selection pressures based in aesthetic properties of an external artifact can evolve a virtual organism with some success. Our agents’ artificially emergent and autonomous behaviors resemble those of simple biological organisms in some ways on a superficial level and are rather interesting to observe. Whether some of the resulting artifacts are aesthetically pleasing is up to the beholder. Their chaotic patterns and compositions certainly parallel abstract expressionist action paintings to some degree. The agents’ paintings share an interesting connection to this art movement as all brushstrokes represent nothing but the actions themselves. With that, one could argue for their artistic value.

6.1 Future work

We briefly propose future directions for the current research. Foremost, the system would highly benefit from a more robust visual model, as emphasized by the poor reconstruction quality of wide-perspective visual observations. This can be achieved by using a larger dataset of intermediate visual observations. Future work could assess whether granting a virtual organism continuous agency over its visual perspective, approximating the cognitive process of *attention*, is a worthwhile approach. This feature is trainable and could explore the nuance between the benefits of the demonstrated visual perspectives.

The morphology and environmental setup we chose for the task of painting is by no means the most suitable. We recommend that future work in embodied agent art should keep exploring the evolution of morphologies. This prevents authors from making predisposed choices about the most suitable body for a given task. A significant downside to this is that it requires a flexible network structure for the action controller model that is significantly more difficult to train. A search algorithm for appropriate morphology choice is another separate topic that could be further explored in the context of art-producing artificial organisms [27]. Furthermore, agents in the current work are limited to a single type of brush, paint color, and environment to explore. Therefore, future extensions could try implementing physically based painting systems, color palettes, and varying environments, each of which could bring about interesting new artifacts and behaviors. Ultimately, painting is only one method of artistic practice, and by no means the most suitable for embodied agents to practice. Computational organisms and environments allow for other artistic modes of expression to be explored such as sculpture, dance, music, poetry, etc. The possibilities are far-reaching and may one day perhaps exceed our imagination.

7 Conclusion

We have demonstrated that virtual organisms can be evolved to make aesthetically pleasing paintings using selection pressures based on aesthetic properties of the painting. The results from our experiments show notable behavioral differences between agents employed with wide-perspective and narrow-perspective vision. The wide-perspective population achieved the best results in terms of fitness by evolving a circular movement strategy effective at maximizing fitness early on, but later showing barely any signs of improvement. The narrow-perspective population performed worse and did not evolve an exploitable strategy. Instead, it brought about a diverse set of artifacts across all generations. From this we conclude that the wide-perspective setup may be more suited for maximizing aesthetic fitness while the narrow-vision setup induces more behavioral complexity and artifact diversity. Although, the scope of this research is limited, our results provided some interesting insights and discussions which provide directions for future applications of computational art systems employing virtual organisms.

16 L. van Hoek et al.

References

1. Beer, R.D., Chiel, H.J., Sterling, L.S.: A biological perspective on autonomous agent design. *Robotics and Autonomous Systems* **6**(1-2), 169–186 (1990)
2. Bentley, P.J.: Is evolution creative. In *Proc. of the AISB* **99**, 28–34 (1999)
3. Birkhoff, G.D.: *Aesthetic measure* (p. 226). Cambridge, Mass (1933)
4. Boden, M.A.: *The Creative Mind: Myths and Mechanisms*. Psychology Press, Hove, UK (1990)
5. Brinck, I.: *Situated cognition. On artistic creativity and aesthetic experience, dynamic systems, and art* (2007)
6. Clancey, W.J.: *Situated cognition: On human knowledge and computer representations*. Cambridge Univ. Press (1997)
7. Clark, A., Chalmers, D.: The extended mind. *Analysis* **58**(1), 7–19 (1998)
8. Colton, S., Wiggins, G.A., et al.: Computational creativity: The final frontier? In: *Ecai*. vol. 12, pp. 21–26. Montpellier (2012)
9. Coumans, E., et al.: *Bullet physics library*. Open source: bulletphysics.org **15**(49), 5 (2013)
10. Dennett, D.C., Dennett, D.C.: *Darwin’s Dangerous Idea: Evolution and the Meanings of Life*. Simon and Schuster (1996)
11. Deussen, O., Lindemeier, T., Pirk, S., Tautzenberger, M.: Feedback-guided stroke placement for a painting machine. *CAe* **8** (2012)
12. Diamond, J.: Animal art: variation in bower decorating style among male bowerbirds *amblyornis inornatus*. *Proc. of the National Academy of Sciences* **83**(9), 3042–3046 (1986)
13. Dohm, K., Stahlhut, H., Hoffmann, J.: *Kunstmaschinen Maschinenkunst*. Kehrer Verlag (2007)
14. Dorigo, M., Bonabeau, E., Theraulaz, G.: Ant algorithms and stigmergy. *Future Generation Computer Systems* **16**(8), 851–871 (2000)
15. Futuyma, D.J.: Natural selection and adaptation. *Evolution* pp. 279–301 (2009)
16. Galanter, P.: Computational aesthetic evaluation: past and future. In: *Computers and creativity*, pp. 255–293. Springer (2012)
17. Galanter, P.: *Generative art theory. A Companion to Digital Art* pp. 146–180 (2016)
18. Goldberg, D.E.: The race, the hurdle, and the sweet spot. *Evolutionary design by computers* pp. 105–118 (1999)
19. Ha, D., Schmidhuber, J.: World models. *arXiv preprint arXiv:1803.10122* (2018)
20. Hansen, N.: The cma evolution strategy: A tutorial. *arxiv* (2016), preprint
21. Hansen, N., Akimoto, Y., Baudis, P.: *Cma-es/pycma* on github. *Zenodo*, doi **10** (2019)
22. Hansen, N., Auger, A., Ros, R., Finck, S., Pošík, P.: Comparing results of 31 algorithms from the black-box optimization benchmarking *bbob-2009*. In: *GECCO*. pp. 1689–1696 (2010)
23. Hoenig, F.: *Defining computational aesthetics*. The Eurographics Assoc. (2005)
24. Hülse, M., Wischmann, S., Manoonpong, P., Von Twickel, A., Pasemann, F.: Dynamical systems in the sensorimotor loop: On the interrelation between internal and external mechanisms of evolved robot behavior. In: *50 years of artificial intelligence*, pp. 186–195. Springer (2007)
25. Krčah, P.: *Evolution and Learning of Virtual Robots*. Ph.D. thesis, Univerzita Karlova (2016)
26. Langton, C.G.: *Artificial life: An overview*. Mit Press (1997)

27. Lehman, J., Stanley, K.O.: Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation* **19**(2), 189–223 (2011)
28. Lehman, J., Clune, J., Misevic, D., Adami, C., Altenberg, L., Beaulieu, J., Bentley, P.J., Bernard, S., Beslon, G., Bryson, D.M., et al.: The surprising creativity of digital evolution: A collection of anecdotes from the evolutionary computation and artificial life research communities. *Artificial Life* **26**(2), 274–306 (2020)
29. Machado, P., Cardoso, A.: All the truth about nevar. *Applied Intelligence* **16**(2), 101–118 (2002)
30. Machado, P., Romero, J., Nadal, M., Santos, A., Correia, J., Carballal, A.: Computerized measures of visual complexity. *Acta psychologica* **160**, 43–57 (2015)
31. Matsuura, K.: A new pufferfish of the genus *Tetraodon* that builds “mystery circles” on sandy bottoms in the ryukyu islands, japan (actinopterygii: Tetraodontiformes: Tetraodontidae). *Ichthyological research* **62**(2), 207–212 (2015)
32. McCormack, J.: Niche constructing drawing robots. In: *EvoMUSART*. pp. 201–216. Springer (2017)
33. McCormack, J., Lomas, A.: Understanding aesthetic evaluation using deep learning. In: *EvoMUSART*. pp. 118–133. Springer (2020)
34. Moura, L.: A new kind of art: The robotic action painter. *X Generative Art Conf., Politecnico di Milano Univ.* (2007)
35. Nake, F.: Information aesthetics: An heroic experiment. *EvoMUSART* **6**(2-3), 65–75 (2012)
36. Redies, C.: Combining universal beauty and cultural context in a unifying model of visual aesthetic experience. *Frontiers in human neuroscience* **9**, 218 (2015)
37. Scha, R., Bod, R.: *Computacionele esthetica*. *Informatie en Informatiebeleid* **11**(1), 54–63 (1993)
38. Schwefel, H.P.: *Numerical optimization of computer models*. John Wiley & Sons, Inc (1981)
39. Secretan, J., Beato, N., D’Ambrosio, D.B., Rodriguez, A., Campbell, A., Folsom-Kovarik, J.T., Stanley, K.O.: Picbreeder: A case study in collaborative evolutionary exploration of design space. *Evolutionary computation* **19**(3), 373–403 (2011)
40. Shimamura, A.P., Palmer, S.E.E.: *Aesthetic science: Connecting minds, brains, and experience*. OUP USA (2012)
41. Sims, K.: Artificial evolution for computer graphics. In: *PACMCGIT*. vol. 18, pp. 319–328 (1991)
42. Sobel, I.: An isotropic 3×3 image gradient operator, machine vision for three-dimensional scenes (h. freeman editor) (1990)
43. Such, F.P., Madhavan, V., Conti, E., Lehman, J., Stanley, K.O., Clune, J.: Deep neuroevolution: Genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning. *arxiv* (2017), preprint
44. Takagi, H.: Interactive evolutionary computation: Fusion of the capabilities of ec optimization and human evaluation. *Proc. of the IEEE* **89**(9), 1275–1296 (2001)
45. Taylor, R.P., Micolich, A.P., Jonas, D.: Fractal analysis of pollock’s drip paintings. *Nature* **399**(6735), 422–422 (1999)
46. Todd, P.M., Werner, G.M.: Frankensteinian methods for evolutionary music. *Musical networks: parallel distributed perception and performance* pp. 313–340 (1999)
47. Todd, S., Latham, W.: *Evolutionary art & computers*. Academic Press, Inc. (1994)
48. Tresset, P., Deussen, O.: Artistically skilled embodied agents. In: *AISB* (2014)
49. Wilson, D.J.: An experimental investigation of birkhoff’s aesthetic measure. *The Journal of Abnormal and Social Psychology* **34**(3), 390 (1939)