Evolving Virtual Embodied Agents using External Artifact Evaluations

Lesley van Hoek, Rob Saunders, Roy de Kleijn

Media Technology MSc program, Leiden Institute of Advanced Computer Science

Research

We present a computational art system exploring the potential of digitally evolving artificial organisms for generating aesthetically pleasing artifacts. Populations of agents are optimized for the aesthetic quality of their work using a complexity-based fitness function that solely evaluates the artifact.

Two experimental setups are investigated: painting with a narrow- and wide-perspective vision sensor.

Background

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 [1]. This structure creation is an important behavioral characteristic that plays part in their survival. We explore whether artificial organisms could adapt to similar, but digitally induced **external selection pressures** of aesthetic nature.

Implementation

Agents act in a three-dimensional virtual environment. Their cognitive models are designed in accordance with theory of situated cognition. Agents' neural controllers are comprised of a pre-trained convolutional variational autoencoder (visual model V) and a linear model (action controller C) that continuously receives sensory data as input and outputs activation values, generating painting behavior.

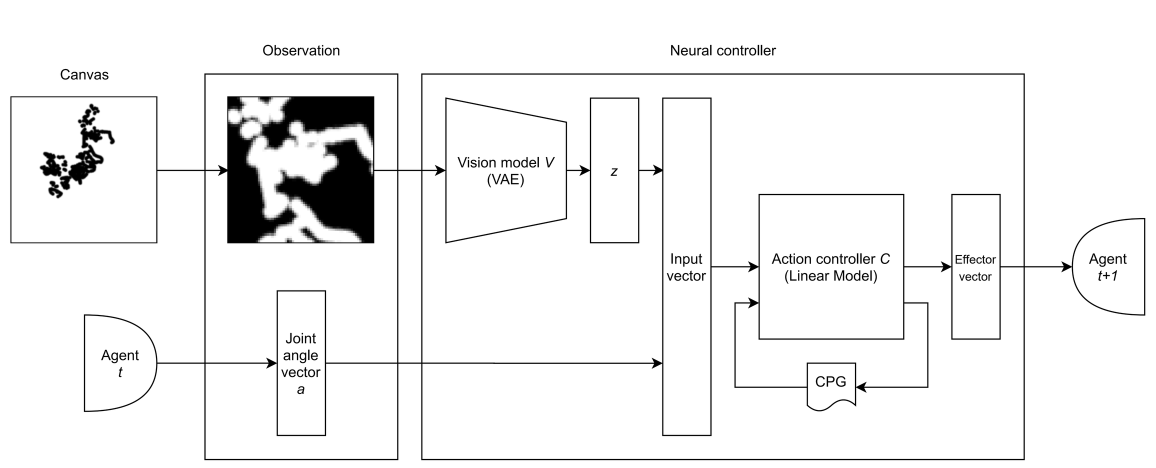

Experimental Setup

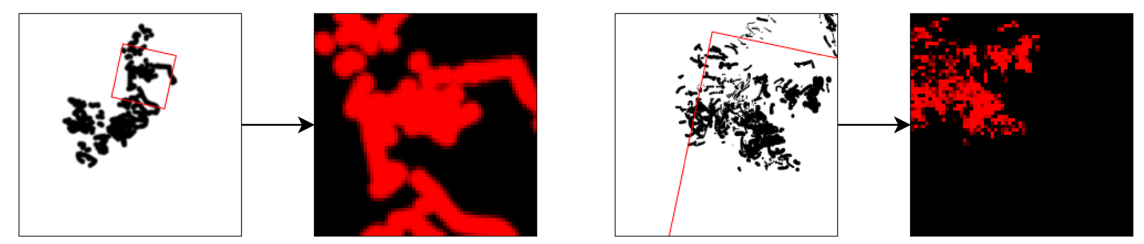

The training experiment was run with a narrow- and wide-perspective vision population. The canvas is cultivated by the agent itself and is its only cognitive resource.

Training procedure

Visual model V was pre-trained on a set 20,000 artifact samples to discern between visual observations. The agents' control policies are evolved by running the painting simulation for 350 generations of 32 candidates each. Action controller C is optimized using CMA-ES [2]. Rollouts have a duration of 240 seconds of simulated time after which the aesthetic quality of the resulting painting is evaluated.

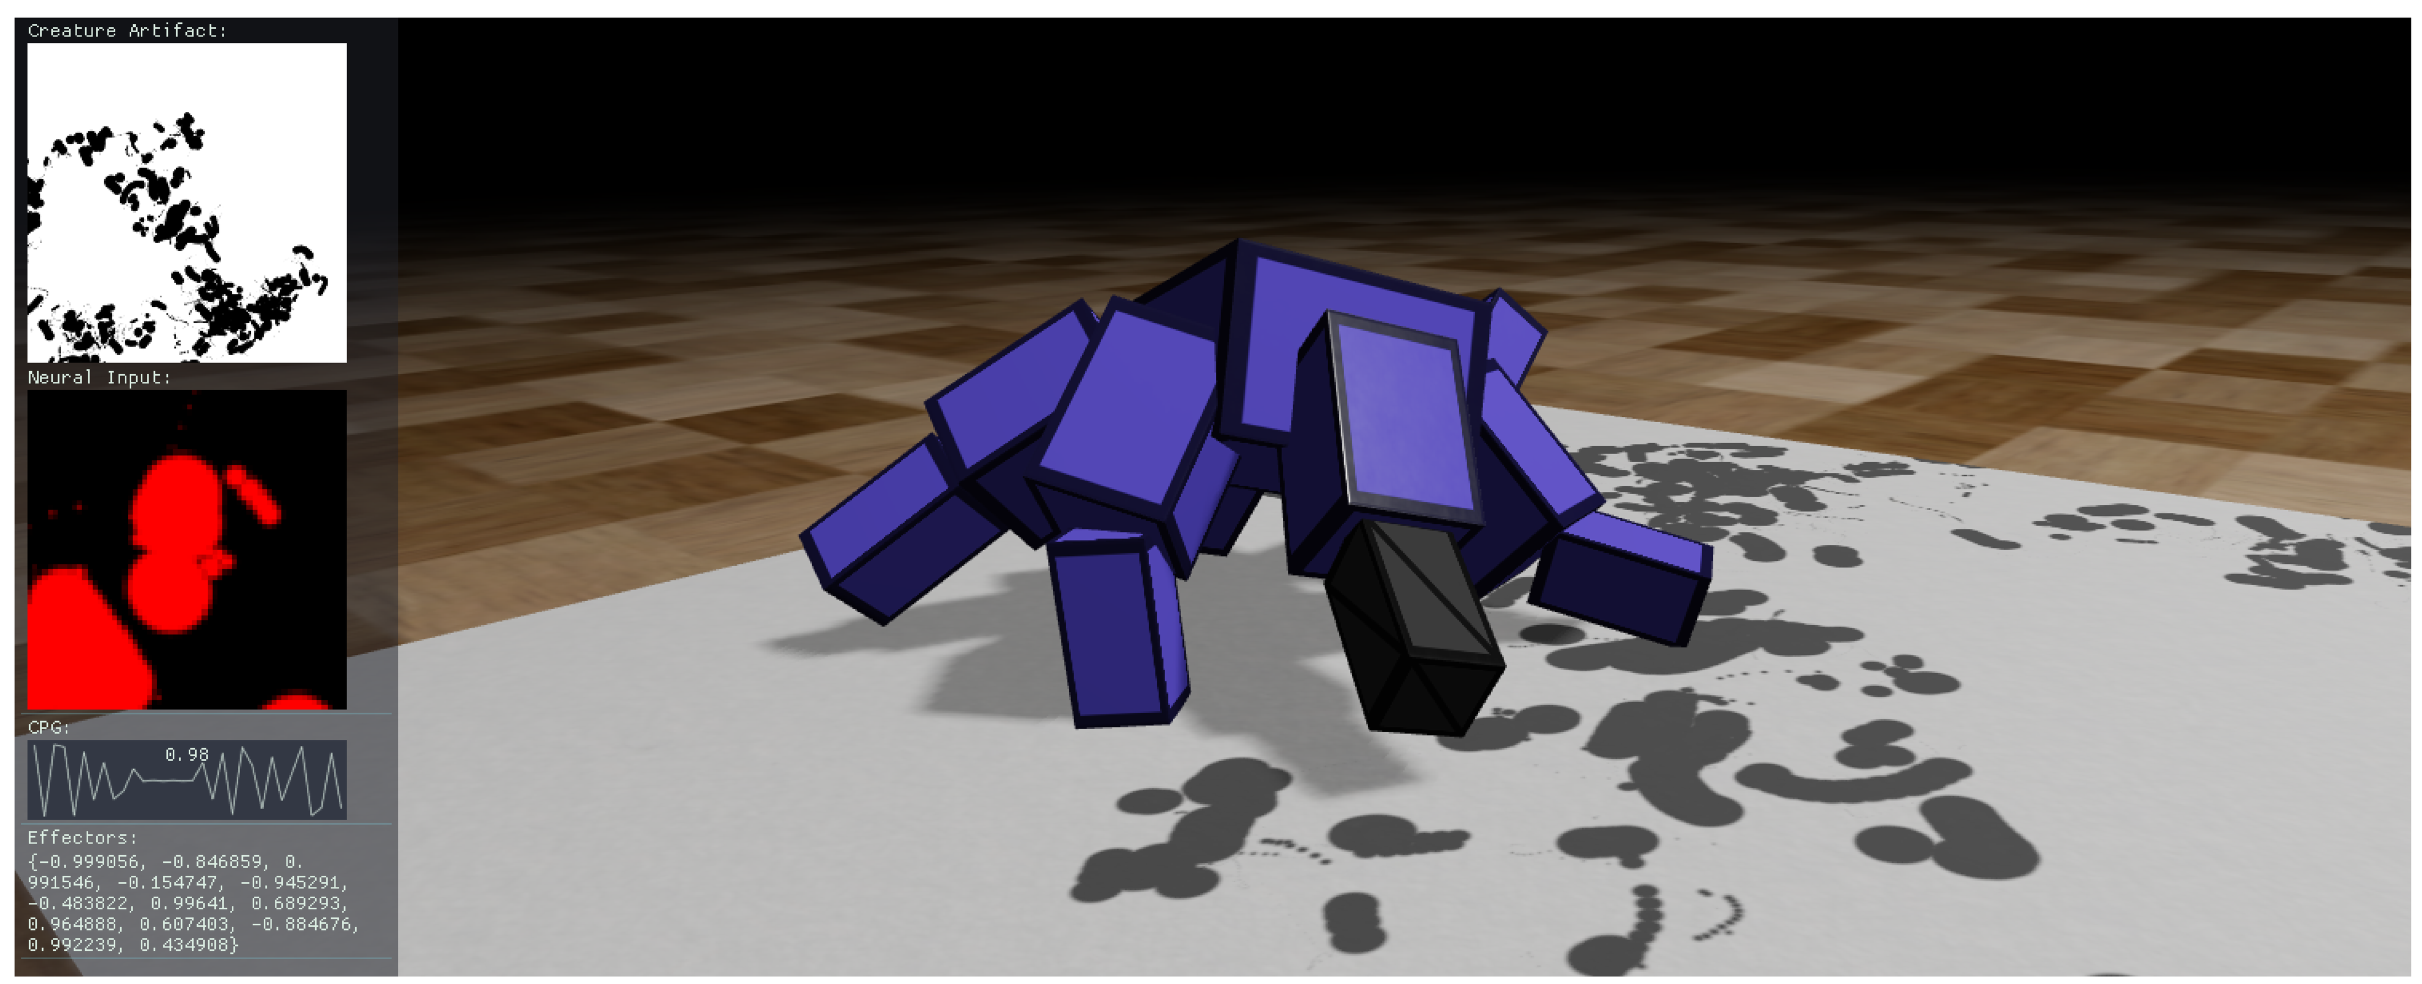

Measuring aesthetic quality

Our fitness function is designed to evaluate agents' paintings in accordance with speculative visual aesthetic theory [3]. We assume an image's intrinsic beauty to be equal to the ratio of image complexity IC to processing complexity PC (Eq. 1).

$$reward_{aesthetic} = \frac{IC}{PC} \tag{1}$$

Eqs. 2 & 3 show calculations of IC/PC using JPEG and fractal compression of image i.

$$PC_{tn} = \frac{1}{RMSE(Fractal(i),\ i)} \tag{2}$$

$$IC = \frac{RMSE(Sobel(i),\ JPEG(Sobel(i)))}{size(Sobel(i))size(JPEG(Sobel(i)))} \tag{3}$$

Results

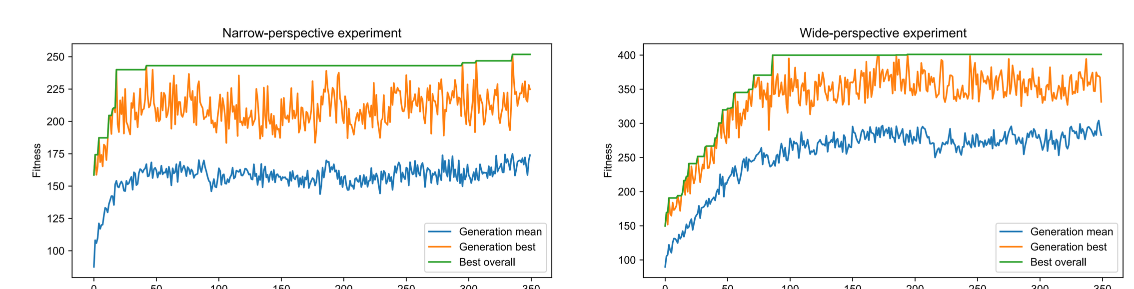

In fitness, the wide-perspective population (right) is generally about 150 points ahead of the narrow-perspective population (left). However, we note no signs of improvement after a local optimum has been reached until the final generation of the simulation, unlike the narrow-perspective population.

The artifacts themselves show that the populations' evolved strategies are easily distinguished from approximately the 40th generation.

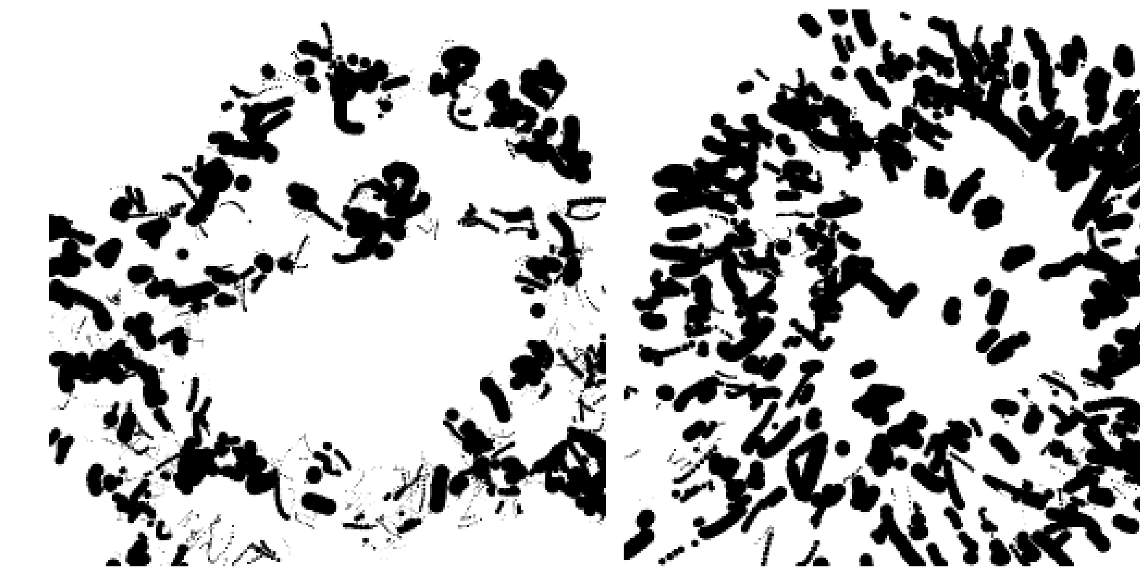

Highest-rated artifacts of the wide- (right) and narrow-perspective population (left).

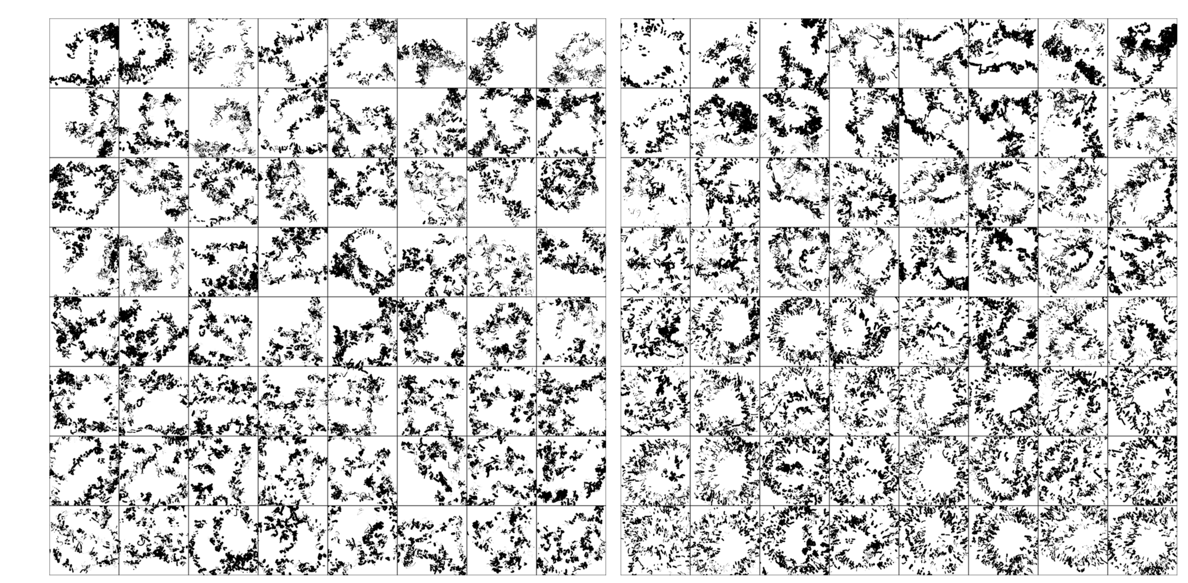

The best artifacts of the first 64 generations (top-left to bottom-right) of the wide- (right) and the narrow-perspective population (left).

Conclusion

The wide-perspective population achieved the highest fitness by exploiting a circular movement strategy.

The narrow-perspective population performed worse but brought about a diverse set of artifacts across all generations.

We conclude that virtual organisms can be evolved to make more aesthetically pleasing paintings using selection pressures based in their work, and different visual perspectives evolve distinct strategies with their own merits.

[1] J. Diamond. Animal art: variation in bower decorating style among male bowerbirds amblyornis inornatus. Proc. of the National Academy of Sciences, 83(9):3042–3046, 1986. [2] N. Hansen. The cma evolution strategy: A tutorial. arxiv.preprint, 2016. [3] P. Machado and A. Cardoso. All the truth about nevar. Applied Intelligence, 16(2):101–118, 2002.