Aesthetics of Adaptive Behaviors in Agent-based Art

Sofian Audry
Concordia University, Faculty of Fine Arts, Montréal, Canada
info@sofianaudry.com

Abstract

Since the post-war era, a number of artists have been exploring the use of embodied, artificial agents, in parallel to scientific research associated to Computer Science in domains such as Cybernetics, Artificial Intelligence and Artificial Life. While notions of adaptation and learning have been an extremely important component of that research, artists and media theorists seem to have focused on the concept of emergence. Whereas emergence offers a rich ground for art-making, adaptation is an equally important, yet complementary dimension of it. In an effort to re-position adaptive systems within the theoretical and practical field of agent-based artworks, an aesthetics of computationally adaptive artistic installations is proposed in this article. To do so, I examine (1) the historical context surrounding adaptive systems; (2) its relationship with the concept of emergence; and (3) the aesthetic potential of Machine Learning algorithms by examining their intrinsic characteristics. An aesthetic framework based on the morphological aspects of the temporal unfolding of agent behaviors is offered as a tool to comprehend both adaptive and non-adaptive behaviors in works of art.

Keywords

Adaptive Systems, Aesthetics of Behavior, Agent-based Art, Artificial Life, Cybernetics, Machine Learning, Media Art Installations.

Introduction

Since the 1960s, media artists have been creating bodies of work using and/or inspired by computer technologies. In this article, I am interested in a specific branch of artistic works that make use of artificial agents, which include works such as Nicholas Schöffer's cybernetics sculpture *CYSP1* (1956), Ken Rinaldo's artificial life installation *Autopoiesis* (2000) and Yves Amu Klein's living sculpture series. Artist and media theorist Simon Penny calls these kinds of work "embodied cultural agents" or "agents as artworks" and integrates them within the larger framework of an "aesthetic of behavior", a "new aesthetic field opened up by the possibility of cultural interaction with machine systems" [24]. These works are distinct from so-called generative art which uses computer algorithms to produce stabilized morphologies such as images and sound: their aesthetics is about the real-time performance of a program as it unfolds in real-time in the world through a situated artificial body.

This paper focuses on a particular facet of this broader work: agent-based *adaptive* computational artistic installations. It argues for an aesthetics of adaptive agents rooted in the distinctive way behavior morphologies unfold in time.

Among the significantly vast literature that exists in media theory, art history and STS on the topic of artificial agency and machinic life in art and science, most studies focus on related concepts such as embodiment/ situatedness [23, 6], autonomy [5] and emergence [4, 29], while very few directly address questions of adaptivity and machine learning.

I hereby wish to fill this gap by proposing an aesthetics of adaptive agent-based installations. My main objective is to provide a description of the experiential mechanisms that are made possible by adaptive behaviors in media installations by connecting the dots between the scientific perspectives over such systems and the aesthetic effects they afford. While artistic media installations cannot be separated from their visual and audio qualities, the focus of my analysis is on their processual dimension.

This paper is concerned with providing conceptual tools to support reflection and creation by artists and researchers engaging with adaptive systems. To contextualize my research, I first present an overview of the history of adaptation from the 1950s onwards, focusing mainly on cybernetics, artificial life, and machine learning, showing their impact on new forms of art. Building upon Cariani's categorizations of adaptive and emergent systems, Penny's "aesthetics of behavior" and Xenakis' theory of morphological evolution, and looking at specific considerations surrounding Machine Learning technologies, an aesthetic framework is put forth to understand the evolution of behaviors through time.

Historical Context

History is imbued with a fascination for human-fabricated life, from Al-Jazari's 13th century's moving peacocks to Jacques Vaucanson's digesting duck (1739). A change in paradigm operated in the post-war era with the advent of computers which, contrary to mechanical automata, are uniquely powerful in both their speed and programmatic capacity. But while often seen as fixed, logic-based systems, an important strand of research in computer science rather focuses on their malleable, organic properties, approaching them as adaptive, self-organized, statistics-based devices. This section offers an overview of this research while examining its role in contemporary media art.

Cybernetics, Perceptrons and Classic AI

The first conceptions of adaptivity in organisms can be found in the work of early, so-called "first-order", cyberneticians. Norbert Wiener's notion of control in Cybernetics systems is closely linked to the concept of *teleological* or *negative feedback*. In a system that displays such negative feedback control, the difference between the goal of the system and its current outputs is sent back to the inputs, allowing the system to correct its course; in other words, to constantly adapt to small changes in its environment. [32] A key and related concept is that of *homeostasis*, referring to the ability of living systems to maintain stability within an unstable milieu using self-regulation. [3]

Building upon both cybernetician models of the brain [3, 19] and psychologist Donald O. Hebb's theory of self-assembling neurons [14], Frank Rosenblatt proposed in the late 1950s one of the first adaptive connectionist devices, the *perceptron* [25], a simplified model of a human neural network that maps a set of binary data (input neurons) to a binary output (output neuron) using a layer of parametric values called weights (representing the synapses) which are initialized randomly. The training procedure allows the model to adjust its weights based on a series of example inputs for which the expected output is known, using a feedback error-correcting mechanism.

The excitement for such connectionist structures which was growing in the 1950s received a cold shower with the publication of Minksy and Papert's forceful critique of perceptrons [20]. By showing that even simple problems are unsolvable by such linear neural networks, the book put a halt to the non-symbolic and distributed approach which had great attention in the field since the 1940s. The funding switched sides and for two decades, AI research turned towards the symbolic and heuristic approach pioneered by Minksy, Papert and Simon, which would later be known as "classic AI" or Good Old Fashioned AI (GOFAI).

Systems Aesthetics

Whereas the impact of the advent of computer science on Western societies in the 1940s and 1950s has been thoroughly documented, often overlooked is how it affected the artistic world. In 1961, Roy Ascott's fascination for Cybernetics made him envision a new conception of art as embodied in interactive systems rather than in physical objects. As a replacement for "visual art" which has become too narrow to describe the new paradigm he attempts to describe, Ascott suggests the name "behavioural art" which he defines as "a retroactive process of human involvement, in which the artifact functions as both matrix and catalyst" whose "structure must be adaptive" with feedback as its core mechanism [2, p. 128].

Jack Burnham's "systems aesthetics" echoes Ascott's vision of emergent, adaptive, behavior-based art. Burnham considered how art as an institution could be understood as a hierarchical system, with artists as its basis being "similar to programs and subroutines", while at the very top a "metaprogram" constantly rearranges the long-term objectives of art. Key to Burnham's vision is the conclusion that this self-organizing, adaptive system does not produce new objects, but rather new information, embodied in the creation of works of art [8].

Artificial Life

At the beginning of the 1980s, classic approaches in AI were still dominating, showing no interest in any form of biologically-based computation such as genetic algorithms and neural computation. Nonetheless, two strands of research would come to life in that era, challenging the statu quo: Artificial Life and Machine Learning.

In the 1970s, chaos theory and complex system theory had revealed how highly non-linear systems often display *emergent* properties, that is, unpredictable behavior as the result of simple interactions between a large number of entities. Emergence directly challenges the distinction between human and machine: starting from simple rules, we can simulate complex and unpredictable behavior on the computer. This idea is core to the early 1980s apparition of Artificial Life (ALife), a synthetic approach to biology that seeks to create "life-like behaviors". This new "biology of possible life", directly influenced by Cybernetics, supplements traditional biological sciences: "By extending the empirical foundation upon which biology is based beyond the carbon-chain life that has evolved on Earth, Artificial Life can contribute to theoretical biology by locating *life-as-we-know-it* within the larger picture of *life-as-it-could-be*." [18, p. 1]

Like Cybernetics in the 1960s, the field of ALife would open up a whole new territory for artists. New media theorist Mitchell Whitelaw remarks that ALife is an area of experimental science which is less preoccupied by observation and representation than it is by intervention and action. Tracing through the interests for synthetic life in art history, he hypothesizes that "a-life art" might just be the latest addition to "a modern creative tradition that seeks to imitate not only the appearance of nature but its functional structures" by using or appealing to technology. ALife might then just be the true destiny of art and the realization of Jack Burnham's vision of a "living, cyborg art form". [31, p. 19]

Machine Leaning

In parallel, part of the people in the AI community had become interested in questions of learning systems [17, p. 275], paving the way to the institutionalization of a new research field within AI that would employ mathematical models to classify and make predictions based on data or experience rather than on logical rules.

At its beginnings, the new field of Machine Learning was still mostly based on symbolic methods such as decision trees and logic. But in mid-decade, a major breakthrough would suddenly bring connectionism back on the scene as David E. Rumelhart, Geoffrey E. Hinton and Ronald J. Williams proposed an algorithm known as *backpropagation* to train a multi-layer perceptron (MLP) [27], a model which consists in stacking several perceptrons on top of each other in interconnected layers. Each layer projects the previous layer's outputs using a non-linear threshold, circumventing the main caveat of perceptrons as pointed out by Minsky and Papert: their inability to separate non-linear data [20].

The revival of connectionist adaptive systems in the 1980s irremediably changed the field of Machine Learning, moving it away from logic towards statistics and biologically-inspired methods. In particular, the field of neural computation would gain impetus. In 1987 took place the first Conference on Neural Information Processing Systems (NIPS) took place, bringing together researchers interested in connectionist approaches from both neurosciences and computer science. It would become, over the years, the most important conference in the field of Machine Learning.

Adaptation and Emergence

Emergence has been widely studied by scholars interested in questions of artificial cognition and living systems. It is often associated with self-organization, such as in ALife, Cybernetics and Connectionism. However, emergence also evokes an idea that somehow goes beyond the automated configuration of a system: the generation of novelty. [29] Within an artistic context, emergence promises to spawn unforeseen patterns and to surprise even its own designer. [12]

Peter A. Cariani has contributed a stimulating ontology of artificial systems that establishes a clear relationship between adaptation and emergence in Cybernetics systems [9]. Differentiating Cybernetics devices on the basis of their adaptive qualities, he identifies three kinds of such systems: formal, adaptive and evolutionary. *Formal* devices are purely (formal-computational) or partly (formal-robotic) symbolic apparatus that respond to a fixed set of instructions and are thus non-adaptive. *Adaptive* systems are capable of adapting their computational structure based on experience but are limited by their fixed semantical components (sensors and effectors). Machine Learning systems and even adaptive robotic agents are part of that category.

Finally, *evolutionary* devices are those that are able to adaptively construct their own sets of sensors and effectors. [9, p. 132] This last category can be refined by considering systems that have adaptive semantics but a non-adaptive syntactic part, such as the immune system. *General evolutionary* devices are those that are both *adaptive* and *evolutionary*, in other words that display both semantics and syntactic adaptiveness, and there are plenty of examples of such systems in the biological world.The main "advantage" of evolutionary devices as compared to adaptive or formal systems lies in their ability to generate novelty, which Cariani directly associates with emergence (p. 148).

Cariani uses *emergence-relative-to-a-model* (or "observer-centric emergence") to integrate adaptation and emergence in a comprehensive framework. First developed by theoretical biologist Robert Rosen, it defines an emergent event as "a deviation of the behavior of the physical system under observation from its predicted behavior" (p. 30). In other words, emergence comes from the fact that since we dispose of only a finite number of observable dimensions whereas the universe contains a potentially infinite number of attributes, it follows that our models of the world are always incomplete accounts of it. (p. 157)

The taxonomy of adaptivity at the core of Cariani's theory can now be attached to the emergent qualities of a system's behavior:

> When the behavior of the physical system, in this case the device itself, bifurcates from the behavior of the model, another model will have to be constructed which will capture subsequent behavior of the physical system/device. (p. 158)

This "bifurcation" from the model's behavior is thus, according to Cariani and Rosen, the locus of novelty emergence in the agent's behavior. That emergence is realized by the agent through its adaptive capabilities, either syntactic, semantic, or both. Thus, one could say that adaptivity is the means by which emergence is realized. In that context, adaptivity is seen not just as a way for systems to self-organize but as a necessary condition for creativity.

Cariani's framework provides useful tools to think about adaptation and emergence. However, his perspective is that of a cognitive scientist, not an artist, hence it is limited when applied to works of art. Expanding upon Cariani's work, Joan Soler-Adillon has developed an extensive aesthetic framework to understand interactive artworks that make use of emergent systems [29]. This analytical tool is rooted in the distinction between two forms of emergence: self-organization emergence (SOE) – which is related to works in Cybernetics and ALife – and Generation of Novelty Emergence (GNE) – which is directly connected to Cariani's emergence-relative-to-a-model.

But Soler-Adillon's aesthetics of emergence is specific to the case of interactive works and is thus, at best, imperfectly applicable to the case of adaptive systems. Adaptation presupposes a form of self-organizing emergence through which the agent will adapt the structure underpinning its behavior. As we have seen, it also provides an anchor point for understanding novelty generation, being the process by which emerging-relative-to-one model is achieved.

Furthermore, whereas emergence is often associated with living systems, it is certainly not a sufficient condition to life, as there are many non-living systems that can be described as emerging, such as weather or cosmic phenomena. In other words, emergence can happen independently of any kind of agency. But as it presupposes an agent whose behavior allows it to evolve in the environment, I posit that adaptation brings us one step closer to *aliveness*. When brought into the arts, it promises to generate behaviors that are more "lifelike" and perhaps, also, closer to more complex forms of life such as the brain, in line with the life-imitating artistic tradition underscored by Whitelaw. [31]

With that in mind, in the coming sections, I examine the aesthetic potential of Machine Learning techniques. Combining Simon Penny's "aesthetics of behavior", Cariani's ontology of adaptive and emergent systems and Xenakis and Di Paolo's study of morphogenesis, I suggest a framework to think aesthetically about adaptive and emergent behaviors rooted in their evolution in time.

Machine Learning Aesthetics

Cybernetics-style adaptive systems evolved through the 1980s onwards into the science of Machine Learning, bringing together a vast multitude of approaches ranging from statistics, stochastics and Bayesian logic to neural and genetic computing, under a common research program within AI. Machine Learning explores algorithms that are able to make inferences and predictions about the world by looking at large quantities of data.

These techniques were, of course, never intended to be used for artmaking. Still, anyone looking at an agent tentatively trying to balance an articulated pole [30], achieving acrobatic stunts with a toy helicopter [22] or finding new ways to play Pong [21] can assess to the aesthetic qualities of these performances. But what are the dimensions of Machine Learning algorithms that can be exploited for artistic expression, and how can they be utilized as such? To approach that question, let us look at the fundamental characteristics of learning methods and explore ways they can be harnessed for art creation.

A Machine Learning algorithm comprises four components: (1) the *category of task* one is trying to solve; (2) the *model* used to address it; (3) the *loss function* against which the model is trained; and (4) the *search or optimization* procedure. These items represent dimensions of a learning system which all influence its outcomes in terms of efficiency, but also – and this is, of course, what should interest us the most here – relatively to the aesthetic effects it affords.

To these elements, we need to add a key component of Machine Learning: the *data* that is fed into the system. This comprises the choice of inputs and outputs and the distribution from which the examples are taken. In a real-life setting, all these elements can be understood as the context in which the learning agent is situated, including both the observables and hidden variables that form its environment.

Tasks

The field of Machine Learning is divided in three sub-fields, corresponding to three different classes of problem: supervised learning, unsupervised learning and reinforcement learning.

Supervised learning (SL) – the most common category – concerns the problem of predicting an output associated with a certain input data, based on a dataset containing examples of data points with expected target response (typically hand-labelled by humans). *Unsupervised learning* (UL) refers to classes of problems where there are no precise outputs that need to be predicted, typically referred to as “unlabeled data”. Rather, the algorithm needs to learn “something about the data distribution”.

Reinforcement learning (RL) tries to address problems involving an agent that tries to take actions in an environment in order to maximize its reward over time. [30] The agent learns by taking actions and receiving positive or negative feedback from the world through rewards. A reward is a single-value information given to the agent in response to his state or actions. In line with Holland’s definition of adaptation, the goal of a reinforcement learning agent is to modify its inner structure in order to maximize its performance – represented as the rewards it collects over time – as it evolves in the environment. [15]

While RL seems to be the most adapted to agent-based installations, it has been scarcely used in practice. One of the few examples is the installation/performance work *N-Polytope: Behaviors in Light and Sound After Iannis Xenakis* created by Chris Salter in collaboration with Sofian Audry, Marije Baalman, Adam Basanta, Elio Bidinost and architect Thomas Spier, in which RL is used to simulate bursts of light “chasing” one another along the cables that form the structure.

One of the most widespread ML techniques used in such pieces involve Genetic Algorithms (GAs) which are a form of Supervised Learning. In Karl Sims’ 1997 installation *Galápagos*, a series of twelve computers each show a single virtual 3D organism whose morphology and movements are the phenotypic outcomes of a digital genotype. Visitors supervise the evolution of the organisms by selecting the ones they prefer, directly impacting the next generation of artificial beings [28].

Self-Organizing Maps (SOM) are a kind of neural network used in Unsupervised Learning tasks which have also been abundantly used in works of art. Many of Yves Amu Klein’s *Living Sculpture* works make use of them, such as *Scorpibot* and *Octofungi* [16]. Nicolas Baginsky uses a similar approach in *The Three Sirens*, a robotic music band who plays improvisational rock music: the guitarist and the bassist use SOMs to direct their actions, playing live music in response to the sound environment they generate in real-time.Finally, artist Ben Bogart’s *Context Machines* use SOM as part of image-based installations that play on the questions of memory association and dreaming [6].

All three approaches can be mixed together, as well as with non-adaptive components. *Zwischenräume* is an installation by artist Petra Gemeinboeck and computer scientist Rob Saunders that features robotic agents that are “sandwiched” between the gallery wall and a temporary wall. Each one of them is equipped with a motorized system that allows it to move vertically and horizontally, covering a specific region of the wall, a puncturing device that allows it to make holes through the surface, as well as a camera and a microphone for sensing the environment. The system allows the robots to extract features from the camera and from the audio signal, combining these informations into a system that mixes SOMs to detect similarities between images and a RL program that tries to “maximise an internally generated reward for capturing ‘interesting’ images and to develop a policy for generating rewards through action” where the level of interest is based on a measure of “novelty and surprise” where “‘novelty’ is defined as a difference between an image and all previous images taken by the robot” and “‘surprise’ is defined as the unexpectedness of an image within a known situation” [13, p. 217].

Components of a Machine Learning Algorithm

Machine Learning algorithm can be qualified by the interoperability of three constituents: the model, the optimization procedure and the evaluation function. The optimization process gradually improves the model based on its performance over the evaluation function [1, pp. 35–36].

Models in Machine Learning refer to the computational structure that gets modified through learning. The best way to think of a model is as a function that tries to approximate as close as possible a distribution of data, based on a sample of that distribution (the dataset). The model contains free parameters that will be adjusted by the training algorithm, such as the "weights" or "synapses" in a neural network.

Models are the object of important debates in the field of Machine Learning, being the defining flagships of different research strands. However, when it comes to artistic works, they are possibly the least explored dimension, as most adaptive artworks involve either GAs or SOMs. This is likely because scientists and artists have different goals and expectations: an apparently small improvement in the performance of a model can be seen as revolutionary from a scientist's perspective but will not change much how artwork is experienced.

Nonetheless, there are at least three ways in which models can affect artistic outcomes. First, the nature of the model is often an important part of the concept of a piece: the imaginary space deployed by neural nets differ from that of evolutionary computation or decision trees. The second way is the kind of artistic "hijackings" a model can allow because of its nature. For example, in the outdoor installation *Vessels* (2015), a GA is used as ways for a community of water-dwelling robots to share their personalities, evolving a kind of family resemblance. The third process by which models can impact artistic works is more subtle and has not been the object of much analysis. It has to do with the fact that, indeed, different models will yield, or afford, different kinds of behaviors. The types and variety of behavioral strategies that the model allows, and the "smoothness" – or "abruptness" – in the evolution of these strategies during learning, are examples of how models can affect agent aesthetics.

The *optimization procedure* – also called *search* or *training algorithm* depending on the context – changes the parameters of the model in an attempt to improve its responses over time. Different kinds of such procedures exist, each with their own advantages and domain of application, although most are deeply tied to specific models. For example, there is a vast amount of research on training algorithms for neural networks, using different optimization approaches such as gradient descent (backpropagation), genetic algorithm, and simulated annealing.

The *evaluation function* measures the performance of the model on its task. In both Supervised and Unsupervised learning, it is usually referred to as the *loss function* or *cost function*. In a classification task, for instance, the category predicted by the model given an example is compared to the expected target category: the more the model misses the target, the larger the loss. In RL, the evaluation function is called the *reward function*, while in GAs, it corresponds to the *fitness function*.

Among the three dimensions of a Machine Learning algorithm, the evaluation function is probably the one that is the most readily useable for authoring. This is because it has been designed specifically for the purpose of bringing a human input into the equation. Models and optimization procedures are meant to be rather agnostic: the evaluation function determines the kind of "problem" one tries to solve.

Artists can play with evaluation functions and look at how the agent responds. An evaluation function can also be learned or attributed by another agent. Finally, evaluation functions can be interactive, with either the artist or the audience replacing the function by directly giving an evaluation of the system's performance. In the case of evolutionary computation, this technique is known as Interactive Genetic Algorithm (IGA), an approach first proposed by Richard Dawkins [10]. Sims' *Galápagos* is a most famous example of using IGAs in an interactive installation [28].

Another example is the "chasers" algorithm in *n-Polytope*, which simulates agents moving across the installation's cables using a RL algorithm combined with an MLP. Each cable represents a one-dimensional "world" with twelve (12) discrete locations/cells. At any specific moment in time, an agent occupies one and only one of the twelve cells and can choose to either stay in place or move to one of the adjacent cells. The only information (observation) the agent receives is the distance (in number of cells) between itself and the next agent, in both directions. The agents' positions are represented by lightening the corresponding LED on the cable.

The reward function is the sum of three different components: (1) *touch* rewards (or penalizes) the agent for being on the same spot as another agent; (2) *move* rewards the agent for moving in a given direction (and punishes it for going the opposite way); and (3) *stay* rewards the agent for staying put (and punishes it for moving). By playing with these parameters, different behaviors can be promoted, such as prey, predation, movement, and collision avoidance.

Data

Data is an often overlooked, yet crucial dimension when thinking about adaptive behaviors, especially in an artistic context. There are practical concerns when dealing with data encoding and challenging issues that arise when dealing with high dimensional spaces, such as is the case with image or speech recognition, which are largely beyond the scope of this dissertation.

The first things to consider are the kinds of inputs and outputs that will be fed into the system, in other words, what the agent will be able to observe and how it will be able to respond to these observations. In order to be effective, these inputs and outputs need to be carefully chosen to afford the kind of experience the artist has in mind. Moreover, there needs to be a way for the agent to make inferences, otherwise no learning will happen. For example, a system that can only detect light cannot be asked to learn about the sounds made by visitors.

Second, it is self-evident that the data distribution from which the examples are picked has an important influence on the reactions and establishment of the system's behavior. One of the most dreaded issues in Machine Learning is *overfitting*, a problem that arises when a system estimates "too perfectly" a specific dataset, thus becoming less efficient at making predictions on unseen samples (ie., taken outside of the training dataset). While overfitting is a plague for data scientists, it might actually be exploited creatively by artists, by hand-picking data (such as by creating a constrained environment) to encourage a specific response in the system.

Morphologies of Behavior

The temporal dimension of agent behaviors is charged with aesthetic potential. Existing taxonomies of cybernetics systems have mainly focused on their relational and structural aspects [26, 9]. In this section, I propose an ontological frame that focuses on the aesthetics of agent behaviors as they unfold in time.

The "zero-degree" of that categorization is the "behavioreness" of the system, that is, whether it should be considered to have a behavior or not. The initial differentiation criterion, I argue, lies in the structural capacities of the system, more precisely in the existence of an internal state. Stateless devices are akin to mathematical functions: their outputs/actions only depend on their inputs/observations. By design, they are incapable of accumulating experience.

Such systems are known in the field of digital media art as *mappings*. Their widespread popularity is evidenced by the prevalence of data-flow softwares such as Max/MSP or PureData, often appearing under names such as "visualisation" or "sonification". In his critique of the hegemony of mapping in interactive arts, artist Marc Downie argues that its apparent generality, which is often seen as an advantage, actually makes it ineffective and sterile, acting as "a normative idea" of how "numbers get transformed into numbers". [12, p. 17]

Devoid of any kind of autonomy and agency, mapping-based devices are *behavioreless*, their conduct relying almost entirely upon the data fed into them. Whatever sense of aliveness associated with them truly lies in the system that generates this data, be it a human performer or a natural phenomenon. Their statelessness imprisons their "performance" into the instant: their world, if they have any, is a *succession of independent moments*. They are, in other words, *zero-order behaviors* (ie., "nonbehaviors").

Agent-based systems, which are the focus of both this and Downie's dissertations, are *behaviorful* in their ability to extend their world into the past through the use of some kind of inner structure. These stateful devices possess some sort of "memory" (whether it is discrete, continuous, long or short) which is modified by their interactions with the environment. In other words: *their past experiences influences their present actions* (at least within a certain time window).

This statefulness, which implies some kind of structure or trace, can be found in a vast variety of computer programs. For instance, formal devices as defined by Cariani can possess states, typically recognizable in computer code as named variables of different types (ie., booleans, integers, floats), however these syntactic components are fixed. Behaviors generated by these systems are thus bound within a certain domain. Hence, while its response to sensory data may change depending on context, the agent's behavior *itself* does not change through time: it will, inexorably, come to repeat similar patterns after a while. We will thus refer to these conducts as *first-order behaviors*.

To understand this idea better, consider how a behavior can have a certain, recognizable *morphology* that exists in a domain different from other forms of non-computational, "stabilized" media, so to speak, such as image of video, or even, as I explained earlier, real-time mappings such as sonifications or visualizations. The shape of a behavior is parameterized by the sensors, effectors and processing capacities of the system that generates it, and evolves within a certain space-time territory. Morphology and morphological processes have been used to describe time-based behaviors in the writings of contemporary music composers such as Iannis Xenakis and Agostino Di Scipio [33, 11].

Because of their inability to generate new forms and/or to transform their own form, I argue that the behavioral morphologies produced by formal, rule-based systems, are viscerally different from those produced by adaptive and evolutionary agents. The latter produce *second-order behaviors* or "metabehaviors", which involves the coming-into-being, and possibly transformation, of their own (first-order) behavior. They therefore exist in a "different time" than their formal/fixed counterparts, which in turn affects the overall aesthetic effect they are allowed to engender.

I propose to use the concepts of *morphogenesis*, *morphostasis* and *metamorphosis* to further characterize the different processes by which behavioral morphologies exist, emerge and/or change over time. These notions are related, each in their own way, to ideas of emergence, self-organization, self-regulation, novelty and autonomy. As the focus on processes related to forms, they seem particularly appropriate to support an aesthetics of behaviors.

Morphostasis refers to the process whereby a behavior hovers around a stable state of being. While they might look like they are changing when considered over a certain period of time, morphostatic behaviors quickly exhaust the space of dynamic patterns they can generate and start appearing repetitive. These behaviors are immutable: they stay constant through time.

Morphogenesis is the mechanism by which emergent behaviors develop their form in a continuous manner. Only adaptive and evolutionary devices, which are capable of self-organization, are able to support morphogenetic behaviors. The category implies the production of new processual morphologies through a system's interaction with the world.

Metamorphosis is intimately related to morphogenesis and refers to the process by which behaviors change *from one shape into another*. In essence, it should be understood much like its everyday usage, that is, as an outstanding transformation in the demeanor of a person or other living being. The two main dimensions of metamorphosis are (1) the *metaboly*, that is, the magnitude of the transformation undergone by the behavior; and (2) the *speed* at which the behavior transits from one form into the other.

Consider the evolution of robotic behaviors in *Vessels*. When the robots are first started, the DNA in which their behavior is rooted is initialized randomly, resulting in very diverse "personalities". Within the scope of an hour, they evolve towards a collective behavior (morphogenesis) into which they stabilize (morphostasis). But as time goes and as the environment changes (weather, ambient light, distribution of robots accross the space) they also modify their behaviors, slowly adjusting to the new conditions (metamorphosis). Change happens in two different ways. First, in a very specific, personalized, yet slow way, where each robot adjust its personality based on its own reading of the environment. Second, at a regular rythm (about one per 2-3 minutes), a random robot will become a "hub" and call the other robots, promoting rapid adaptation to its own personality. These two mechanisms interoperate in the making of the global aesthetic experience of the piece.

These aspects of an agent's performance should be seen less as hard-cut categories but rather as conceptual tools for describing second-order processes of behavior formation. For example, from that perspective, both formal systems as well as self-regulated devices such as early Cybernetics systems [3], or pre-trained machine learning algorithms, produce purely morphostatic behaviors. However, they are distinctive in the kinds of first-order, repetitive patterns they produce, which are related to their different structural and processual properties, as highlighted before. Importantly, most agent-based adaptive installations bring together a mixture of different systems, staging different kinds of zero-, first- and second-order behaviors, intertwining phases of morphological stasis, genesis and transformation intervening at different rates.

At the opposite end of the spectrum, some morphogenetic systems freely move from one behavioral embodiment into another, constantly metamorphosing, never fully coming into being. These systems are often referred to as "generative": they evolve behaviors regardless to their fitness or value. [7]

Adaptive processes, on the other hand, transform their morphologies in relationship to a usually indeterminate "optimal" behavior which they try to approach and match. Adaptation, like intentionality, requires an object: systems do not just adapt, they adapt *to* something. Their experiences affects their inner structure so as to improve their prospective performances. In other words, *their past feeds their futurity*.

Typically starting from a state of pure randomness, adaptive agents run through a learning process of morphogenesis where they progressively and asymptotically modify the shape of their behavior to better perform in relationship to the cost/fitness/reward function. When they reach their final form, they enter a state of morphostasis, exploiting the stabilized, learned behavior which they converged to. Some adaptive systems have the ability to depart from this crystalized demeanor, either as a result of an internal intentionality, or as a response to environmental changes that require drastic adjustments to their performance. Figure 1 compares the temporal evolution of different kinds of behaviors.

The aesthetic experience of these behaviors is dependent on a number of factors. The ratio between the magnitude of change and the time period necessary to perform it during metamorphosis (which in the case of machine learning systems is directly related to the learning rate) can be used as a measure of intensity. Abrupt, fast transformations can bring a sense of astonishment or angst in the recipient, while noticeable changes that are more spread-out might activate feelings of curiosity or anxiety. An important challenge in designing interactive media experiences is to learn how to play with these parameters to generate the desired effects in the audience.

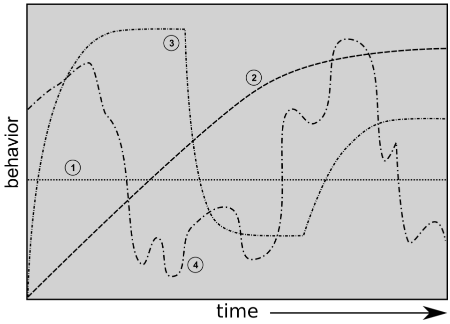

Fig. 1: Example temporal evolution of different kinds of behaviors: (1) first-order behavior; (2) adaptive behavior converging into morphostasis; (3) adaptive behavior running through different phases of metamorphosis and morphostasis; (4) non-adaptive second-order behavior (generative).

Conclusion

In this article, I have examined adaptive behaviors in embodied agent-based artworks from an aesthetics perspective. After contextualizing adaptive systems historically in both science and art, I examined Cariani's theoretical framing of emergent and adaptive systems as an approach to aesthetic understanding. An insider's view over Machine Learning algorithms was used to outline intrinsic features of such systems that can be harnessed for art practice. The temporal aspect of adaptive behaviors was finally examined through the elaboration of an ontology based on morphologies.

Marc Downie explains that emergence is problematic in artistic creation when contrasted with the question of *authorship*. He argues that emergence-based approaches try to avoid the question of authorship altogether by trying to create processes that work by themselves, without human intervention. However, despite decades of efforts, we are still waiting for the advent of higher-order emergent artificial life structures. [12, p. 29]

As I argue, the process by which emergence-relative-to-a-model realizes itself is adaptive in nature. Hence, Machine Learning offers a set of tools to navigate in the development of emergent behaviors by providing concrete ways to achieve authorship, such as playing with models, datasets, and evaluation functions. I hope that the conceptual tools that I developed here might be useful to artists and researchers who are working with self-organizing, agent-based systems.

Acknowledgments

I would like to thank the Fonds de Recherche du Québec - Société et Culture, Concordia University and Dr Christopher Salter for their support.

References

1. Alpaydin, E. 2004. *Introduction to machine learning*. Adaptive computation and machine learning. Cambridge, Mass: MIT Press.

2. Ascott, R. 2003. Behaviourist Art and the Cybernetic Vision. In *Telematic embrace: visionary theories of art, technology, and consciousness*. Berkeley: University of California Press. 109–156.

3. Ashby, W. R. 1954. *Design for a brain*. New York: Wiley.

4. Baljko, M., and Tenhaaf, N. 2008. The Aesthetics of Emergence: Co-constructed Interactions. *ACM Trans. Comput.-Hum. Interact.* 15(3):11:1–11:27.

5. Boden, M. A. 1996. Autonomy and Artificiality. In Boden, M. A., ed., *The Philosophy of Artificial Life*. Oxford University Press, USA, first edition. 95–108.

6. Bogart, B. D. R., and Pasquier, P. 2013. Context Machines: A Series of Situated and Self-Organizing Artworks. *Leonardo* 46(2):114–143.

7. Bown, O. 2012. Generative and Adaptive Creativity: A Unified Approach to Creativity in Nature, Humans and Machines (preprint). In McCormack, J., and d'Inverno, M., eds., *Computers and Creativity*. Springer Berlin Heidelberg. 361–381.

8. Burnham, J. 1969. Real Time Systems. *Artforum*.

9. Cariani, P. A. 1989. *On the design of devices with emergent semantic functions*. Ph.D. Dissertation, State University of New York at Binghamton, Binghamton.

10. Dawkins, R. 1986. *The blind watchmaker*. New York: Norton, 1st american ed edition.

11. Di Scipio, A. 1994. Formal Processes of Timbre Composition: Challenging the Dualistic Paradigm of Computer Music. In *Proceedings of the 1994 International Computer Music Conference*, 202–208. San Francisco: : International Computer Music Association.

12. Downie, M. 2005. *Choreographing the Extended Agent: performance graphics for dance theater*. Ph. d. thesis, Massachusetts Institute of Technology, Cambridge, MA.

13. Gemeinboeck, P., and Saunders, R. 2013. Creative Machine Performance: Computational Creativity and Robotic Art.

14. Hebb, D. O. 1949. *The Organization of Behavior*. New York: Wiley & Sons.

15. Holland, J. H. 1992. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. Cambridge, Mass: A Bradford Book, new edition edition edition.

16. Klein, Y. A. 1998. Living Sculpture: The Art and Science of Creating Robotic Life. *Leonardo* 31(5):393.

17. Langley, P. 2011. The changing science of machine learning. *Machine Learning* 82(3):275–279.

18. Langton, C. G., ed. 1989. *Artificial Life: Proceedings Of An Interdisciplinary Workshop On The Synthesis And Simulation Of Living Systems*, volume 6 of *Santa Fe Institute studies in the sciences of complexity*. Redwood City, Calif: Addison-Wesley.

19. McCulloch, W. S., and Pitts, W. 1943. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics* 5(4):115–133.

20. Minsky, M. L., and Papert, S. 1969. *Perceptrons: an introduction to computational geometry*. Cambridge, Mass.: MIT Press.

21. Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. 2013. Playing Atari with Deep Reinforcement Learning. *arXiv:1312.5602 [cs]*.

22. Ng, A. Y.; Kim, H. J.; Jordan, M. I.; and Sastry, S. 2004. Autonomous Helicopter Flight via Reinforcement Learning. 799–806.

23. Penny, S. 1997. Embodied Cultural Agents: at the intersection of Robotics, Cognitive Science and Interactive Art.

24. Penny, S. 2000. Agents as Artworks and Agent Design as Artistic Practice. In Dautenhahn, K., ed., *Advances in Consciousness Research*, volume 19. Amsterdam: John Benjamins Publishing Company. 395–414.

25. Rosenblatt, F. 1957. The perceptron – a perceiving and recognizing automaton. Technical Report 85-460-1, Cornell Aeronautical Laboratory, Buffalo, NY.

26. Rosenblueth, A.; Wiener, N.; and Bigelow, J. 1943. Behavior, Purpose and Teleology. *Philosophy of Science* 10(1):pp. 18–24.

27. Rumelhart, D. E.; Hinton, G. E.; and Williams, R. J. 1986. Learning Representations by Back-propagating Errors. *Nature* 323:533–536.

28. Sims, K. 1991. Artificial Evolution for Computer Graphics. In *Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '91, 319–328. New York, NY, USA: ACM.

29. Soler-Adillon, J. 2015. *Emergence as Self-Organization and as Generation of Novelty: A Framework for Understanding Emergence in the Context of Interactive Art*. Ph.D. Dissertation, Universitat Pompeu Fabra, Barcelona, Spain.

30. Sutton, R. S., and Barto, A. G. 1998. *Reinforcement learning: an introduction*. Adaptive computation and machine learning. Cambridge, Mass: MIT Press.

31. Whitelaw, M. 2004. *Metacreation: Art and Artificial Life*. The MIT Press.

32. Wiener, N. 1948. *Cybernetics: or Control and communication in the animal and the machine*. M.I.T. Technology Press publications. Cambridge, Mass.: Technology press.

33. Xenakis, I. 1981. Les chemins de la composition musicale. In *Le Compositeur et l'ordinateur*. Paris: Ircam - Centre Pompidou. 13–32.