Aesthetics of Adaptive Behaviors in Agent-based Art

Proceedings of A Body of Knowledge - Embodied Cognition and the Arts conference
CTSA UCI 8-10 Dec 2016

Sofian Audry

Abstract

Since the post-war era, a range of artists have used embodied, artificial agents in media installations. Their work runs in parallel with scientific research in domains associated with computer science, such as cybernetics, crtificial intelligence (AI) and artificial life (AL). Two important concepts are central to these scientific approaches: emergence – the mechanism whereby higher-order forms or processes emanate from the complex interactions of lower-order units – and adaptation – the real-time adjustment of a machine to achieve a better performance in its environment. Recent advances in AI are largely attributable to major breakthroughs in the fields of artificial neural networks and machine learning, both of which feed upon these two core ideas. But while notions of adaptation and learning are an extremely important part of that research, artists and media theorists working with agent-based systems have widely ignored them, often in favor of emergence and self-organization. Inasmuch as emergence offers a rich ground for art-making, adaptation offers an equally important, yet complementary dimension of it. To re-position adaptive systems within the theoretical and practical field of agent-based artworks, I examine (1) the historical context surrounding adaptive systems; (2) its relationship with emergence and self-organization; and (3) the aesthetic potential of machine learning algorithms by examining their intrinsic characteristics. Building upon that research, I propose an aesthetic framework for adaptive systems based on the morphological aspects of agent behaviors as they evolve through time, supported by examples from my own art practice.

Introduction

Since the 1960s, artists have been creating bodies of work using and/or inspired by computer technologies. In this paper, I am interested in a specific branch of artistic works that make use of artificial agents, that is, man-made autonomous systems who act within their environment in response to what they perceive. Examples include pioneering cybernetic artworks such as Nicholas Schöffer's *CYSP1* (1956) or Edward Ihnatowicz's *The Senster* (1970—1974); more recent works include Bill Vorn and Louis-Philippe Demers' large-scale robotic piece *La Cour des Miracles* (1997) and Ken Rinaldo's artificial life installation *Autopoiesis*. 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" [31, p. 398]. They are distinct from so-called "generative art" or "algorithmic art" which use computer algorithms to produce stabilized morphologies such as images and sound, as they are rather about the performance of a program as it unfolds temporally through an artificial body.

This paper examines a particular facet of this broader work: agent-based, *adaptive* computational artistic installations that make use of machine learning methods.Examining the cultural-social-technical repercussions that arise in the use of such techniques in artistic works, I argue for an aesthetics of adaptive agents rooted in the distinctive way their behavior evolves and stabilizes as they couple with their environment.

Machine learning (ML) is a sub-field of the computer science area of artificial intelligence (AI) that employs mathematical models able to classify and make predictions based on data or experience rather than on logical rules. Learning systems usually consist of computational structures that adjust themselves when submitted to large quantities of data. Machine learning is directly related to the biologically-rooted concept of adaptation which refers to a "process whereby a structure is progressively modified to give better performance in its environment" [20, p. 7].

Despite the increased use of machine learning in many facets of contemporary industrial and commercial culture, one site where it has not seemed to make a meaningful impact is the field of art practice – oddly enough, considering that the use of computational systems in art goes back to at least until the early 1950s. Interest in agent-based systems within the arts has mainly focused on techniques and concepts such as self-regulation, evolution and emergence, while there has been little rigorous work on machine learning and other forms of adaptive computation by artists [21, 30, 39].

This paper attemps to provide 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 Peter Cariani's categorizations of adaptive and emergent systems and Simon Penny's "aesthetics of behavior", and looking at specific considerations surrounding machine learning technologies, I put forward an aesthetic framework to understand the evolution of behaviors through time.

Historical Overview

While recent years have seen the field of machine learning grow at an unprecedented rate, the underlying idea of a computational system able to adapt to or learn from a flow of real-life data is certainly not new. On the contrary, it recurs throughout the history of computing, from early concepts of negative feedback in cybernetics to evolutionary computation. The first conceptions of adaptivity in organisms can be found in the work of early cyberneticians such as in Norbert Wiener's notion of *negative feedback*, where 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. [42] Building upon both cybernetician models of the brain [2, 27] and psychologist Donald O. Hebb's theory of self-assembling neurons [19], Frank Rosenblatt proposed in the late 1950s one of the first adaptive connectionist devices, the *perceptron* [33], a simplified model of a human neural network that could learn to recognize patterns. Yet, the excitement for such connectionist structures which was growing in the 1950s was hindered by Minksy and Papert's forceful critique of perceptrons [28]. For the two following 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) – a period referred to as the "first AI Winter". This approach espoused *computationalism* (or "strong AI"), a theory of mind based on the premise that cognition is computation [15].

Art historian Edward A. Shanken describes the influence of cybernetics on art in the 1960s through the work of Roy Ascott [36]. The scope of cybernetics as an encompassing theory of systems' behavior and communication, would allow Ascott to merge cybernetics and art, in an effort to "theorize the relationship between art and society in terms of the interactive flow of information and behavior through a network of interconnected processes and systems" (p. 4).

Cybernetics' conceptions of adaptivity, homeostasis and feedback loops are thus an integral component of Ascott's perspective [1]. He claimed that visual arts had entered a new era where the interactive and participative experience of the spectator in relationship with the artwork was central, hence suggesting the name "behavioural art" as a replacement for "visual art".

In 1968, artist and critic Jack Burnham published "Systems Aesthetics", where he explained how the society of his time was transiting from an "object-oriented to a systems-oriented culture" [8, p. 31]. Burnham argued how art as an institution could be understood as a cybernetic, self-organizing, adaptive system that does not produce new objects, but rather new information embodied in works of art [9].

Both Ascott and Burnham highlight an important point in their focus on *behaviors* in machine-human configurations in the artistic domain. Furthermore, their interest in cybernetics aligns with a vision of society, culture, and art, as profoundly adaptive systems, evolving through a network of self-organizing agents which adjust to one another through a myriad of feedback loops.

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 challenged the distinction between human and machine because we could now, starting from simple rules, simulate complex and unpredictable behavior on a computer. This idea is core to the early 1980s apparition of artificial life (AL or ALife), a synthetic approach to biology that seeks to create "life-like behaviors". [23, p. 1] Like cybernetics in the 1960s, the field of artificial lie would open up a whole new territory for artists in the 1980s and 1990s. [4]

The 1980s also saw the development of machine learning, in particular due to the discovery of an efficient way to train muli-layer perceptrons (MLP) [35], which consist in stacking many perceptrons on top of each other in interconnected layers of neurons. The revival of connectionist adaptive systems had a tremendous impact on the development of machine learning and neural computation as fields of research. Yet, dissatisfaction with the progress of neural-based learning by the end of the decade called once more for new approaches to AI.

Influenced by approaches in both ML and AL, as well as by the work of Maturana and Varela [26], MIT robotics scientist Rodney Brooks challenged classical AI by proposing a "New AI" at the end of the 1980s. Opposing widespread AI beliefs, Brooks proposed that the behavior displayed by living beings results from an embodied, situated interaction with their environment which does not have a need for intermediate representations of the world [6]. Brooks' subsumption architecture, which allowed him to create his first walking robot, *Genghis*, displays learning capabilities and has some close ties with reinforcement learning. [4, 7, 24, 25, 5]

As an efficient, bottom-up approach to robotics, New AI had an important influence on AL robotic art in the 1990s. Artists such as Louis-Philippe Demers, Bill Vorn, Ken Rinaldo and Simon Penny claim Brooks as a direct inspiration for their work [32, 12, 30]. In the same spirit, they give prevalence to the material embodiment of the machinic agents they build over forms of intermediate representations and computations, and to the emergence of behavioral properties through the interactions of these autonomous, self-organizing systems with the world.

Emergence and Adaptation

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 AL, cybernetics and connectionism. However, emergence also evokes an idea that somehow goes beyond the automated configuration of a system: the generation of novelty. [37] Within an artistic context, emergence promises to spawn unforeseen patterns and to surprise even its own designer. [16]

Peter A. Cariani is an interdisciplinary researcher who has developed one of the most compelling theoretical model on the role of adaptation in cybernetic and AL systems through an original and constructive critique of computationalism [11, 10]. He has contributed a uniquely stimulating taxonomy of artificial systems that establishes a clear relationship between adaptation and emergence. 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. [11, p. 132] 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. [11, p. 132]

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 [37]. 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 AL – and *generation of novelty emergence* (GNE) – which is directly connected to Cariani's emergence-relative-to-a-model.

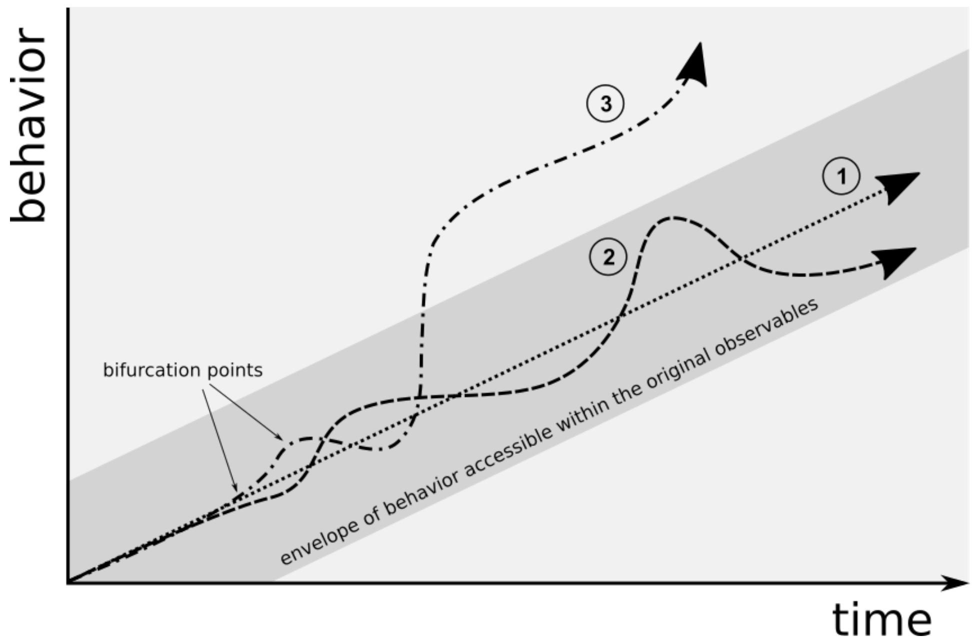

Figure 1: "Divergence of the behaviors of adaptive devices from fixed models of them". (1) Formal-computational and formal-robotic (nonemergent); (2) Adaptive device (syntactic emergent); (3) Evolutionary device (semantic emergent). Adapted from [11, p. 31].

But Soler-Adillon's aesthetics of emergence is specific to the case of interactive works and does not directly address works involving 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 does not appear to be 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. Adaptation, on the other hand, implies the existence of an autonomous, emergent/self-organizing agent whose behavior allows it to evolve in its environment. In other words, adaptive systems are emergent/self-organizing systems with agency and, in this sense, adaptation implies qualities of aliveness that go beyond those of non-adaptive emergent systems. When brought into the arts, adaptation may thus allow the generation behaviors that are more “lifelike” and perhaps, also, closer to more complex forms of life such as the brain. [4]

Behavior Morphologies

The aesthetic framework that I want to articulate in this paper resonates with the work of Simon Penny on agent-based art. Directly inspired by Rodney Brooks’ revolutionary work on situated robotics from the late 1980s that critiques representational systems in AI, Penny argues for a new “aesthetics of behavior” that contains a rejection of computationalism. [22, p. 138]

Penny hereby joins Brooks and Dreyfus in their critique of the dualistic vision of behavior and cognition that taints classical AI. Behavior, he claims, should not be understood as a purely computational, disembodied thing called “software”, but rather needs to be grasped as a situated process running through an agent’s body. Thus, both Nouvelle AI and behavior aesthetics rearticulate concepts of emergence and self-organization in AL by integrating them in a performative theory of behavior that places the agent’s body at the center of the equation. As such, Penny’s proposed artistic framework is constitutionally different from concurrent disembodied artforms such as algorithmic art and an important part of artificial life art that generates time-based simulations on the computer.

As an artist working with agent-based systems, I concur with the anti-computationalists: life and cognition are not “pure” processes that can be separated from a sensorimotor body running in the physical world. I hereby align with Harnad’s claim that cognition is at least *partly* noncomputational, though *some* computation (i.e., rule-based symbol manipulation) might be involved in it [17, 18]. I join my voice with that of Simon Penny, arguing for a new field of aesthetics opened up by computer technologies, with *behavior* as its central concept. Yet, I believe there are still important missing pieces in our understanding of the actual aesthetic qualities of such behaviors.

Gordon Pask's own definition of behaviors, which he detailed in his 1968 book on cybernetics, offers a visionary perspective over behaviors that connect well with Penny, while still allowing for a formalization in terms of their morphological evolution. In line with his view, I argue that behaviors are best defined not as algorithmic recipes, but rather as real-time material patterns as they are recognized by an observing entity. As Pask writes:

As observers we expect the environment to change and try to describe those features that remain unchanged with the passage of time. An unchanging form of events due to the activity within an assembly is called a behavior. [29, p. 18]

There are two important implications of this definition. First, while an agent's behavior involves a sequence of events that constantly change over time, its behavior has a recognizable "shape" that remains temporally invariant. Pask gives the example of a cat, which consists of "performances like eating and sleeping and, once again, it is an invariant form selected from the multitude of things a cat might possibly do" (p. 18).

Second, while a behavior is always generated by a system, it only exists through its perceptual effect on an observer. This implication is particularly appropriate to an aesthetic framework, as it focuses on the phenomenological experience generated by the agent-based performance, as it unfolds through time and space in the material world. This connects directly, in fact, to the pragmatic aesthetics of John Dewey, who claims that works of art should not be thought of as objects, but really as “refined and intensified forms of experience” [13, p. 3].

I posit that different categories of system architectures allow for different kinds of behaviors, thus allowing the emergence of different aesthetic experiences. What interests me is to further analyze Penny’s artistic frame of reference by looking more closely at embodied agents with adaptive qualities. Existing taxonomies of cybernetics systems have mainly focused on relational and structural aspects of these systems [34, 11]. In this section, I propose a flexible taxonomy of embodied systems that focuses on the aesthetics of agent behaviors as their shape unfold in time.

The “zero-degree” of that categorization is the “behaviorness” 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 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”. Marc Downie heavily criticizes this hegemony of mapping in interactive arts. He argues that its apparent generality, which is seen as beneficial, makes it ineffective and sterile: precisely because its definition has “no limits” it also has “no use”. [16, p. 17]

Devoid of any kind of autonomy and agency, mapping-based devices are behaviorless, their conduct relying almost entirely upon the data that is 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 (i.e., “nonbehaviors”). Agent-based systems, which are the focus of both this dissertation as well as Downie’s, 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: their past experiences influence their present actions.

This statefulness, which in other words implies some form of structure or trace, can be found in a side variety of computer programs.Behaviors generated by these systems are thus bound within a certain domain. Hence, while an agent’s response to sensory data may change depending on context, its behavior itself does not change through time. Given enough time, it will, inexorably, come to repeat similar patterns. We will thus refer to these conducts as first-order behaviors.

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 fundamentally different from those produced by adaptive and evolutionary agents. The latter produce second-order behaviors (i.e., “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 affects the overall aesthetic effect they can 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 these ideas bring processes related to forms to the fore, they seem particularly appropriate to support an aesthetics of behavior.

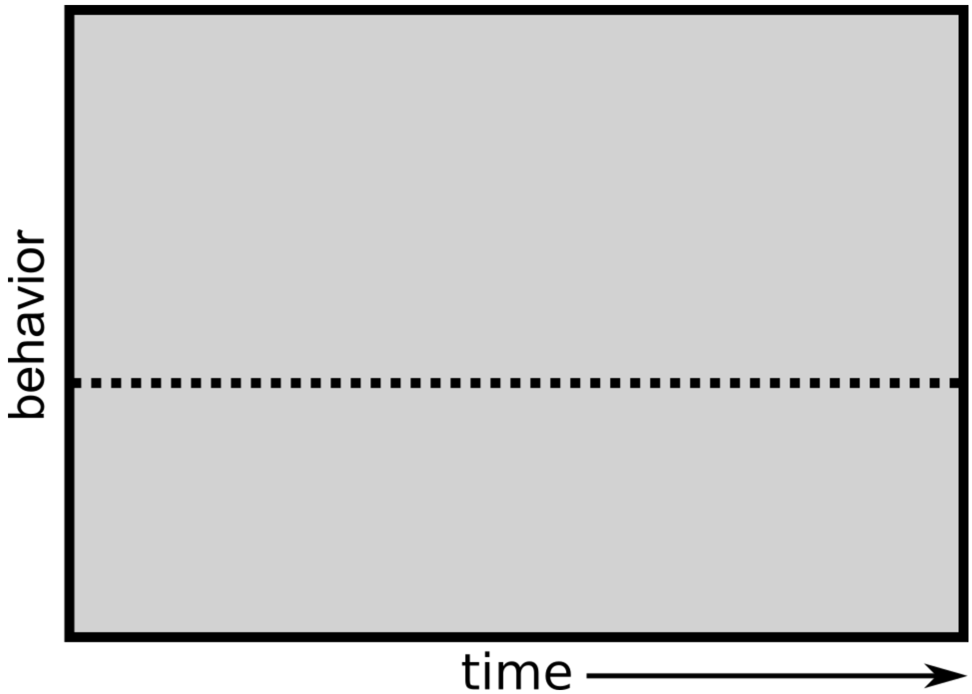

Figure 2: Example temporal evolution of a first-order behavior. The vertical axis represents the behavior of the system, understood as the temporally invariant shape of observable events the system generates through its actions. The horizontal axis represents the advance of time. The graphic shows how first-order behaviors remain temporally stable.

Morphostasis refers to the process whereby a behavior hovers around a stable state of being. While the behavioral patterns 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 behavioral 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, the term should be understood quite similarly to the way it is used in common parlance: that is, as an outstanding transformation in a living being or thing. 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.

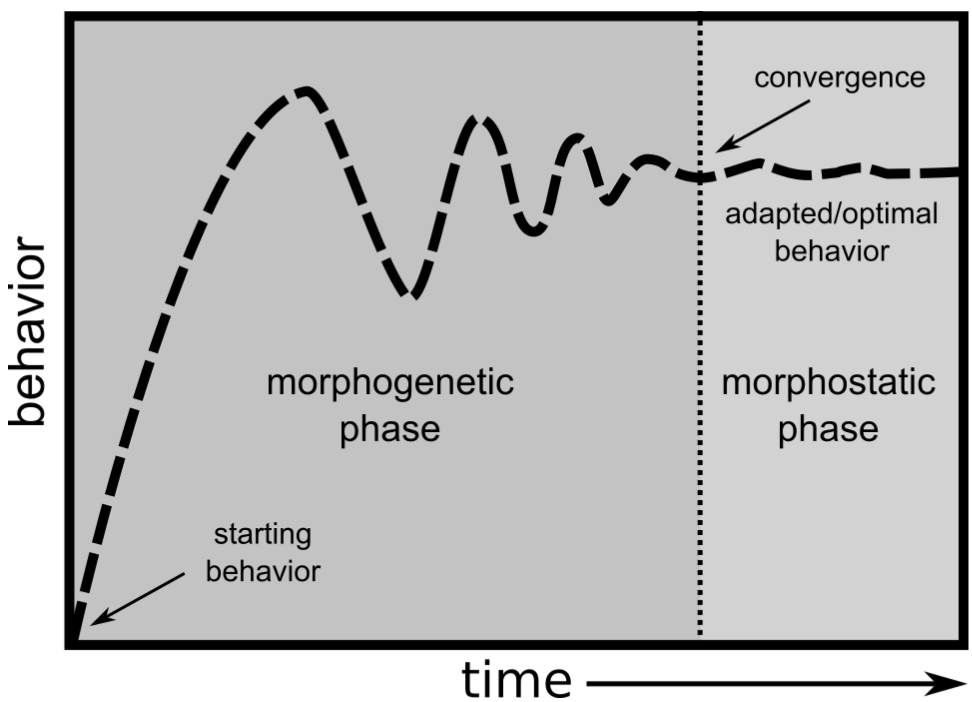

Figure 4: Example temporal evolution of an adaptive behavior. Distance along the vertical axis represent difference in the form of observable events produced by the agent. The graphic shows how second-order, adaptive behaviors iteratively change through time through a process of morphogenesis, until they stabilize into an optimal first-order behavior, thus entering a phase of morphostasis.

These aspects of an agent's performance should be seen less as hard-set categories, but rather as conceptual tools for describing processes of behavior formation. These qualifiers complete concepts such as those previously discussed (i.e., enactivism, coupling, autonomy, and emergence) by bringing attention to the morphology of behaviors and their evolution.

From this perspective, both formal systems as well as self-regulated devices such as Rodney Brook's robots or pre-trained machine learning algorithms, produce purely morphostatic behaviors. However, they are distinct from each other in the kinds of first-order, repetitive patterns they produce, which are related to their different structural and behavioral properties, as highlighted above.

At the opposite end of the spectrum, some morphogenetic systems freely move from one behavioral embodiment into another, living in a constant state of metamorphose, as if never fully coming into being. These systems are often referred to as "generative": they evolve behaviors regardless of their fitness or value [3].

Adaptive systems, on the other hand, evolve their morphologies in relationship to a usually indeterminate "ideal" (i.e., optimal in regards to whatever the evaluation function is) behavior, which they try to approach and match. In this, they differ from *nonadaptive* second-order behaviors. Adaptation, like intentionality, requires an object: systems do not simply adapt, they adapt *to* something. Adaptive systems are relational devices by definition: they are governed by their coupling with another behavior, which in turn can be of zeroth-, first-, or second-order. Their experiences effects their inner structure so as to improve their prospective performances. In other words, their past feeds their future.

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 their evaluation 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.

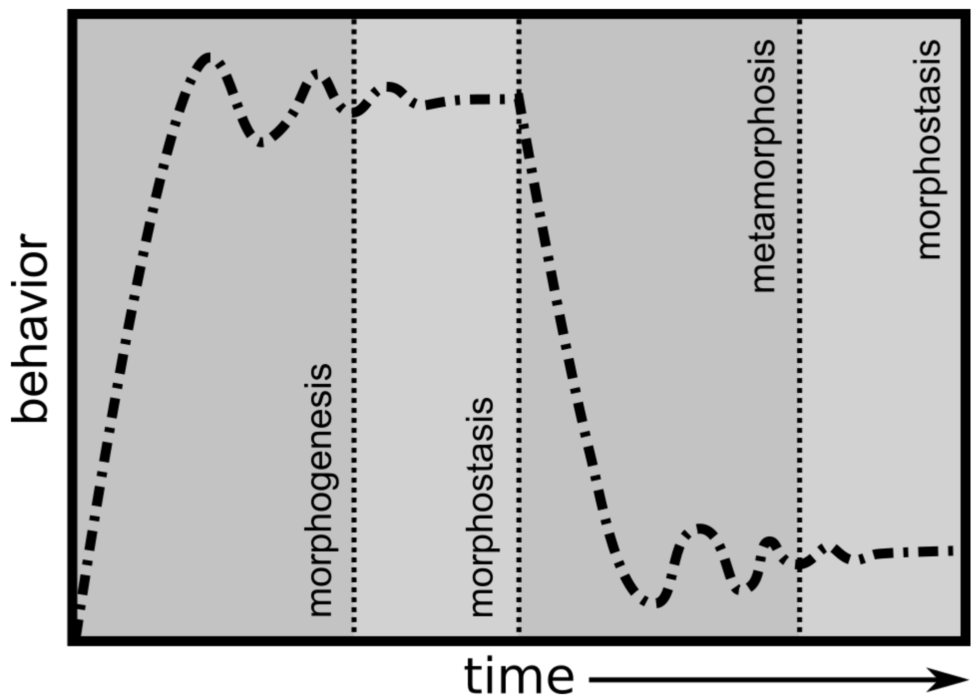

IFigure 5: Example of the temporal evolution of an adaptive behavior going through multiple phases of learning. Starting from a random behavior, it runs through a morphogenetic phase until it converges to an optimal behavior, stabilizing into morphostasis. Then, subjected to environmental changes, it needs to readjust itself, metamorphosing into another shape that performs better in the new conditions.

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 changes can bring a sense of astonishment or angst in the viewer that artists working with interactive media have learned to exploit. In contrast, longer yet steady and noticeable changes can evoke curiosity, anxiety, and uncaninness.

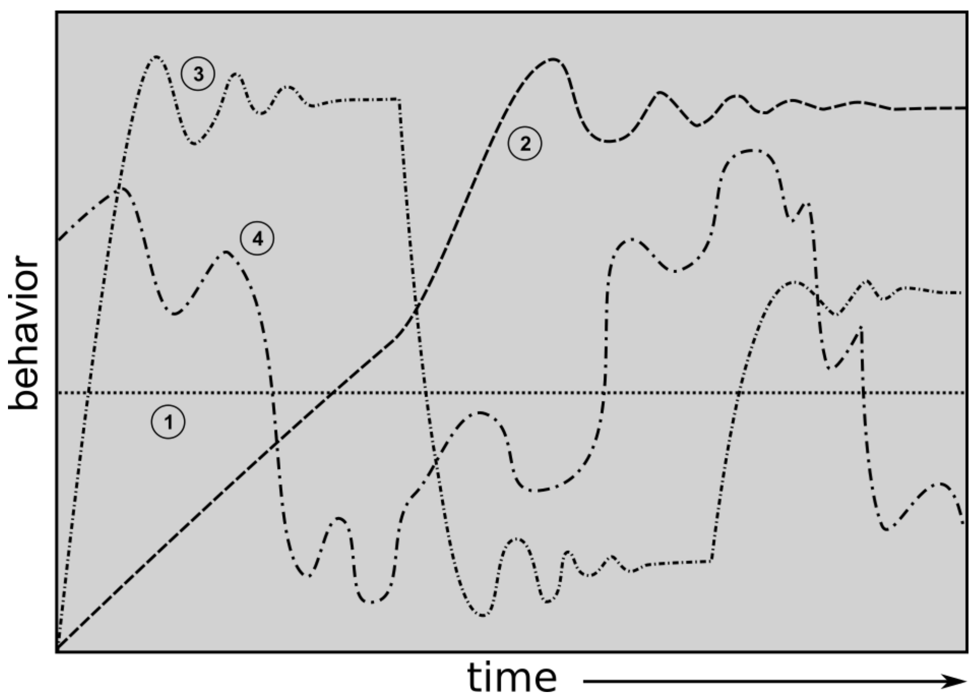

Figure 6: Example of the temporal evolution of different kinds of behaviors: (1) first-order behavior; (2) adaptive behavior converging into morphostasis; (3) adaptive behavior running through different phases of metamorphosis and morphostasis; (4) nonadaptive second-order behavior (generative).

Finally, adaptive behaviors convey a certain narrative. Unfolding before our eyes, we perceive fluctuating stories of trials and errors, of successes and failures, that evoke our own experiences of learning as fallible and imperfect entities. For instance, as I watch an agent such as a cart-pole system for an extended period of time, as it balance an inverted pendulum, discovering its environment, reaching plateaus of apprenticeship, I might start perceiving the desires of the agent, what it “wants”. I start to look ahead with apprehension, projecting both myself and the agent into the future.

I want to end with a few disclaimers. First, the “orders” of behaviors that I described should not be read as a hierarchy. Second-order behaviors are not in any way better or worse than behaviors of lower order: they are just different, they come with their own strengths and weaknesses, and can each be used efficiently (or badly) in artmaking.

Second, these categories are porous. For example, some mapping functions, such as moving averages or delays, have a short “memory” and can thus be said to have a “state”; some self-organizing adaptive systems have very limited structures which do not allow them to adjust significantly in the face of changing environments.

Finally, these categories can (and should, when appropriate) be mixed together. Most agent-based adaptive installations actually 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. The use of lower-order behaviors gives the artist more direct control over the outcomes, which is often crucial for the success of a work.

This categorization is not meant as a systematic classification scheme, but rather as a frame of reference, a flexible analysis tool for artists and theorists. It gives an angle, a way to think and discuss about agent-based systems in art practice, that I hope can contribute to the language of new media as practitioners attempt to imagine new experiences and communicate their views with their peers.

Conclusion

The aesthetic framework deployed in this paper aims to assist artists and theorists in thinking aesthetically about behaviors. It does not, however, offer an extensive argument to the affective qualities that such behaviors might foster. In particular, one should not forget that the human observers experiencing these systems and trying to make sense of them are, themselves, adaptive agents. The process by which human observers encounter behaviors is one by which they tentatively and iteratively adjust their expectations so as to better predict the movements of the agents, becoming familiar with them, getting to know them. To illustrate this point, consider scientist Adrian Thompson's use of genetic algorithms to evolve a programmable electroniccircuit to discriminate between a 1kHz and a 10kHz tone [40]. After the circuit evolved into an optimal discriminator, in order to simplify it, Thompson tried to prune out the parts of the circuits that were not contributing to the output. Contrary to Thompson's assumptions, some parts of the circuits that were completely disconnected from any path that could leading to the output, were in fact crucial to the discrimination process – likely through some forms of local magnetic interactions. In other words, the adaptive agent that controlled the evolutionary process had learned a solution to the problem that made use of the intrinsic, embodied, physical properties of the circuit and that no human could possibly have come up with.

Similarly, human observers of adaptive or evolutionary works cannot understand their behaviors rationally, because the underlying processes that govern them follow non-logical rules. Works that are based on mappings and first-order behaviors can be rationally explained and understood: for example, this photocell triggers that sound effect, that microphone activates that video sequence, that gestures causes the agent to start running in circle for a minute, etc. But in order to experience second-behaviors in all their richness, one needs to "get to know them" phenomenologically, through her own sensorimotor body.

Acknowledgments

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

References

[1] Roy Ascott. Behaviourist Art and the Cybernetic Vision. In *Telematic Embrace: Visionary Theories of Art, Technology, and Consciousness*, pages 109–156. University of California Press, Berkeley, 2003.

[2] William Ross Ashby. *Design for a Brain*. Wiley, New York, 1954.

[3] Oliver Bown. Generative and Adaptive Creativity: A Unified Approach to Creativity in Nature, Humans and Machines (preprint). In Jon McCormack and Mark d'Inverno, editors, *Computers and Creativity*, pages 361–381. Springer Berlin Heidelberg, 2012.

[4] R. Brooks. A robust layered control system for a mobile robot. *IEEE Journal on Robotics and Automation*, 2(1):14–23, March 1986.

[5] Rodney Brooks. *Flesh and Machines: How Robots Will Change Us*. Random House LLC, April 2002.

[6] Rodney A. Brooks. Intelligence Without Representation. *Artificial Intelligence*, 47:139–159, 1987.

[7] Rodney A. Brooks. A Robot that Walks: Emergent Behaviors from a Carefully Evolved Network. Technical Report AI MEMO 1091, Massachusets Institute of Technology, 1989.

[8] Jack Burnham. Systems Esthetics. *Artforum*, 7(1):30–35, September 1968.

[9] Jack Burnham. Real Time Systems. *Artforum*, 8(1):49–55, September 1969.

[10] Peter Cariani. Adaptivity and emergence in organisms and devices. *World Futures*, 32(2-3):111–132, September 1991.

[11] Peter A. Cariani. *On the Design of Devices with Emergent Semantic Functions*. PhD thesis, State University of New York at Binghamton, Binghamton, 1989.

[12] Louis-Philippe Demers and Bill Vorn. Real Artificial Life as an Immersive Media. In *Convergence: Proceedings of the 5th Biennial Symposium for Arts and Technology*, pages 190–203, New London, CT, 1995. Center for Arts and Technology at Connecticut College.

[13] John Dewey. *Art as Experience*. A Perigee book. Putnam, New York, NY, 1959. OCLC: 247836059.

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

[15] Eric Dietrich. Computationalism. *Social Epistemology*, 4(2):135–154, 1990.

[16] Marc Downie. *Choreographing the Extended Agent: Performance Graphics for Dance Theater*. Ph. d. thesis, Massachusetts Institute of Technology, Cambridge, MA, August 2005.

[17] Stevan Harnad. What's Wrong and Right About Searle's Chinese Room Argument? In John Preston and Mark Bishop, editors, *Essays on Searle's Chinese Room Argument*. Oxford University Press, 2001.

[18] Stevan Harnad. The Annotation Game: On Turing (1950) on Computing, Machinery and Intelligence. In Grace Epstein, Robert & Peters, editor, *Parsing the Turing Test: Philosophical and Methodological Issues in the Quest for the Thinking Computer*. Springer, 2008.

[19] Donald Olding Hebb. *The Organization of Behavior*. Wiley & Sons, New York, 1949.

[20] John H. Holland. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. A Bradford Book, Cambridge, MA, new edition edition edition, April 1992.

[21] Eduardo Kac. Foundation and Development of Robotic Art. *Art Journal*, 56:60–67, 1997.

[22] Jihoon Felix Kim and Kristen Galvin. An interview with Simon Penny: Techno-Utopianism, Embodied Interaction and the Aesthetics of Behavior. *Leonardo Electronic Almanac (DAC09: After Media: Embodiment and Context)*, 17(2):136–145, 2012.

[23] Christopher G. Langton, editor. *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*. Addison-Wesley, Redwood City, Calif, 1989.

[24] P. Maes, R. Brooks, Pattie Maes, and Rodney A. Brooks. Learning to Coordinate Behaviors. In *Association for the Advancement of Artificial Intelligence*, volume 2, pages 796–802, Palo Alto, CA, 1990. AAAI Press.

[25] Pattie Maes. Modeling Adaptive Autonomous Agents. *Artificial Life*, 1:135–162, 1994.

[26] Humberto R. Maturana and Francisco J. Varela. *Autopoiesis and Cognition: The Realization of the Living*. D. Reidel, Dordrecht, 1980.

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

[28] Marvin Lee Minsky and Seymour Papert. *Perceptrons: An Introduction to Computational Geometry*. MIT Press, Cambridge, MA, 1969.

[29] Gordon Pask. *An Approach to Cybernetics*. Hutchinson, London, 1968.

[30] Simon Penny. Embodied Cultural Agents: At the intersection of Robotics, Cognitive Science and Interactive Art. In Kerstin Dautenhahn, editor, *AAAI Socially Intelligent Agents: Papers from the 1997 Fall Symposium*, pages 103 – 105, Menlo Park, 1997. AAAI Press.

[31] Simon Penny. Agents as Artworks and Agent Design as Artistic Practice. In Kerstin Dautenhahn, editor, *Advances in Consciousness Research*, volume 19, pages 395–414. John Benjamins Publishing Company, Amsterdam, 2000.

[32] Kenneth E. Rinaldo. Technology Recapitulates Phylogeny: Artificial Life Art. *Leonardo*, 31(5):pp. 371–376, 1998.

[33] Frank Rosenblatt. The perceptron – a perceiving and recognizing automaton. Technical Report 85-460-1, Cornell Aeronautical Laboratory, Buffalo, NY, 1957.

[34] Arturo Rosenblueth, Norbert Wiener, and Julian Bigelow. Behavior, Purpose and Teleology. *Philosophy of Science*, 10(1):pp. 18–24, 1943.

[35] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning Representations by Back-propagating Errors. *Nature*, 323:533–536, 1986.

[36] Edward A. Shanken. Cybernetics and Art: Cultural Convergence in the 1960s. In Bruce Clarke and Linda Darlymple, editors, *From Energy to Information*, pages 155–177. Stanford University Press, Palo Alto, NM, 2002.

[37] Joan Soler-Adillon. *Emergence as Self-Organization and as Generation of Novelty: A Framework for Understanding Emergence in the Context of Interactive Art*. PhD thesis, Universitat Pompeu Fabra, Barcelona, Spain, 2015.

[38] Makis Solomos. Cellular automata in Xenakis's music. Theory and Practice. In Solomos Makis, Anastasia Georgaki, and Giorgos Zervos, editors, *Definitive Proceedings of the International Symposium Iannis Xenakis*, page 11 p., Greece, 2006.

[39] Nell Tenhaaf. Perceptions of self in art and intelligent agents. In Kerstin Dautenhahn, editor, *Proceedings of the 2000 AAAI Fall Symposium on Socially Intelligent Agents: The Human in the Loop*, North Falmouth, 2000. AAAI Press.

[40] Adrian Thompson. An Evolved Circuit, Intrinsic in Silicon, Entwined with Physics. In *Proceedings of the First International Conference on Evolvable Systems: From Biology to Hardware*, ICES '96, pages 390–405, London, UK, UK, 1996. Springer-Verlag.

[41] Mitchell Whitelaw. *Metacreation: Art and Artificial Life*. The MIT Press, Cambridge, MA, March 2004.

[42] Norbert Wiener. *Cybernetics: Or Control and Communication in the Animal and the Machine*. M.I.T. Technology Press publications. Technology press, Cambridge, Mass., 1948.

[43] Iannis Xenakis. Les chemins de la composition musicale. In Marc Battier, editor, *Le Compositeur et l'ordinateur*, pages 13–32. Ircam - Centre Pompidou, Paris, 1981.