The Machine as Autonomous Performer

Oliver Bown, Petra Gemeinboeck and Rob Saunders

Abstract In this chapter we explore the experience of interactive system-based artworks that exhibit autonomous behaviours in an interactive context. Engaging with such autonomously behaving works opens up experiences that are more akin to conversing, performing, or negotiating. We introduce cybernetic influences and take a closer look at the performance of the participant/machine system. Following this, we discuss the ways in which artists approach working with adaptive systems and observe audiences to iteratively improve their system designs. At the core of the chapter is a discussion of five artworks that serve as our case studies: two influential works: Edward Ihnatowicz's *The Senster* and Ken Rinaldo's *Autopoiesis*, and three projects developed by the authors: *Uzume*, *Accomplice*, and *Zamyatin*. We use these case studies to explore the artists' approach to autonomy, how it shapes the audience's experience and the methods used in the development and evaluative process.

P. Gemeinboeck, College of Fine Arts, University of New South Wales
e-mail: Petra Gemeinboeck <petra@unsw.edu.au>
O. Bown and R. Saunders
Design Lab, Faculty of Architecture, Design and Planning, University of Sydney
e-mail: oliver.bown@sydney.edu.au; rob.saunders@sydney.edu.au

1 Introduction

Distinctions between the dynamic and autonomous behaviours of the natural world and the static and controllable properties of the built environment are disintegrating. Artists have long attempted to enrich our experiences with the enchanted and magically animated, and, in this context, have also experimented with concepts from cybernetics, Artificial Intelligence (AI), and Artificial Life (ALife). In this chapter, we will discuss system-based artworks that exhibit autonomous behaviours, and what may appear to be intentions, independently of being 'activated' by a human participant.

Interactive artworks that behave or perform autonomously challenge the most common interaction paradigm of reacting to what is sensed according to a pre-mapped narrative. Engaging with autonomous works opens up new experiences that are more akin to conversing, performing, or negotiating with something that has its own mind. Our exploration concerning the evaluation of these works focuses on the ways in which artists approach autonomy and the methods and criteria they develop to achieve autonomous behaviours. This is evaluation of a qualitative nature, leading to questions such as: how does the system behave, how does it involve the participant(s), and in what ways does this 'dialogue' or co-performance differ from interactive experiences with non-autonomous artworks?

2 The Art of Behaviours

We can distinguish different types of interactive artworks on the basis of the work's or the system's complexity and how it relates to its interactive capacity. A simple or weak form of interactivity can be achieved whenever any kind of input is provided to a system, e.g., a button press. In most cases, this simple input creates a singular, linear link between the 'interactor' and the system. In contrast to this, we can sketch the outline of an interactive system, which exhibits some form of autonomous behaviour and is likely to produce a more complex form of interaction between a person and the system[1] .For some, the distinction underlies the very meaning of the term 'interactivity'. As Ruairi Glynn, creator of the adaptive work *Performative Ecologies*, states, "the widespread misuse of the term 'interactivity', has trivialised its meaning to the point that it holds no more conceptual value than reactivity to most of today's artists, architects and designers" (Glynn 2008: 1).

---

[1] Such complexity, for example, may be captured by the formalisation of digital art systems by Cornock and Edmonds (1973), including the category 'dynamic interactive (varying)' in which the conditions of interaction change over time. The works described here are most likely to fit this category, but may achieve a complexity of interaction by other means than devising mechanisms for long-term variation.

For Usman Haque, the failure of much interactive art and architecture lies in its inability to enter into a conversation with a person. In Haque's view a typical interactive work "invokes a mutually reactive relationship only slightly more sophisticated that that between a person and an automated cash machine." (Haque 2007: 26). Whilst the continued novelty of digital computing technologies has allowed simple interactive systems to flourish, more complex interactive relationships remain challenging to creative practitioners and their audiences.

The re-emergence of certain strains of cybernetic thinking in the form of ALife in the 1990s, opened up the possibility for a transformative creative development in which artefacts might transcend their status as mere objects. Artists producing such behavioural artefacts often have been inspired by the interdisciplinary experimentation that occurred in cybernetic explorations (Boden and Edmonds 2009). At the core of this influence of cybernetic thinking is the idea of looking at artworks as systems and systems as artworks. This can be seen in relation to different strands of thinking at the time: on the one hand, drawing inspiration from the growing awareness of the diversity of dynamics exhibited by systems in the natural world, and, on the other, a cybernetic influence cutting across all of modern art (Ascott 2002).

In more complex approaches to interactivity, the system is addressed "as quasi-organism, in autopoietic or enactive sensorimotor loops with user(s)" (Penny 2011:80). Autopoiesis, applied to aesthetics, can be understood as a self-propelling system of aesthetics that is open to negotiation (Hall 2010). An interactive system as such expands or completely evades the bi-directional input-process-output modalities of many interactive systems. Rather, participants are more likely to engage with the work in ways akin to encountering another life form or inhabiting an alternative world. No matter how strange the artificial creature or world we encounter may be, we are already in our element as 'interactors', and don't need to learn a new language or interaction paradigm. As Penny observes, our ability to interact with digital systems is rooted in our evolved adaptation to embodied experiences in the world:

> We are first and foremost, embodied beings whose sensorimotor acuities have formed around interactions with humans, other living and non-living entities, materiality and gravity. We understand digital environments on the basis of extrapolations upon such bodily experience-based prediction. (Penny 2011:78).

Inherent to such systems is their performative nature, which goes beyond the fact that there is action involved. If systems are to act autonomously then their actions clearly cannot be staged; only in their real-time enactment can we know what they do. This performance necessarily unfolds in the present, without the certainty afforded by rehearsal. Participants or 'interactors' are "caught up in a direct experience of the work's dynamics" (Tenhaaf 2008:12). We can say that both the participant and the system become performers in this process.

In embracing this performative potential of machine agents, artists have striven to shift the focus from representational issues to questions of agency and relationality. This leads to works that are not artefacts to interact with but rather, in the words of Nathaniel Stern, “relations to be performed” (2011: 233). The work is continuously constructed and composed, more akin to an event or a performance than a fixed interface or installation. An analytical approach to investigating the experience of, with, or within these interactive works is therefore best achieved by focusing on the performance of the participant/machine system, and its performative capacities.

The above provides a cursory overview of the intent and thinking behind a cybernetics-inspired approach to art making. In the following we introduce ways of asking about autonomy, as a means to gain insight into a system’s behaviour and how it relates to its aesthetic performance.

3 Approaches to Autonomy

The evaluation of a performative experience that does not reveal itself except through interaction is central to the conceptual approaches outlined throughout this book, in particular by Candy (Chapter {candy.pdf}). Our focus is on the means by which artists, and audiences, conceptualise and realise interactions with systems that exhibit autonomy. We see the qualitative, and often anthropomorphizing categorisation of behaviour as key to evaluating both the creative and the experiential process.

Anthropomorphism helps us to predict behaviours, to engage in them, as well as to maintain our interest (Horowitz and Bekoff 2007). For instance, encountering an autonomous work or system, we naturally apply anthropomorphic terms to categorise behaviour: we may ask, is it ‘alive’, ‘aware’, ‘curious’? Does it have ‘intent’? Perhaps we even wonder if it is ‘playful’, ‘mischievous’, or ‘stubborn’? Does it matter what I do and how I do it? A pivotal question for evaluating an autonomous work is: what are the perceptual or conceptual cues and rationalisations that lead people to make such judgements?

As an entry point to building autonomous systems, some artists draw on scientific theories and models, which offer a formal basis for approaching autonomy. For example, Seth (2010) has operationalized autonomy as a measure of *self-determination* in a system based on the predictability of systems in terms of different causal factors. If a system’s behaviour can be predicted entirely on the basis of its external context then it is said to be heteronomous. In contrast, we can find evidence of a system’s autonomy if its history is required to predict future behaviour. This approach to autonomy connects systematically with other key concepts such as causality and complexity of behaviour, via the mechanics of predictability.

We will take a closer look at the artists’ motivation and audiences’ perception in our case studies later. As several of the chapters in this book discuss, a necessary part of the process of developing interactive works is to investigate, anticipate or otherwise understand audience behaviour. In the artworks included here, evaluation is informal, but nevertheless involves identifiable methods. It typically involves prolonged observation and probing of the system itself, which can take a more empirical approach in terms of searching the parameter space of a computational system, setting up test cases or specific studies to develop the mappings for interaction, and an analysis of behavioural properties. The artist typically spends long periods observing the system's behaviour under different circumstances. Since the systems are behaviour generators, this often requires extended periods of waiting and watching as different behaviours are revealed. Exactly what is being sought in this process may be hard to define for the artist. Given the metacreative nature of this process the artist must develop an understanding of the range of possible behaviours, and the nature of their distribution. In this way, whilst interaction with an audience will be different from that observed in the studio, the artist is able to anticipate the likely system behaviour in the exhibition context.

A further stage of evaluation involve modes of participatory ethnographic observation, which are often integrated into a process of tinkering to improve the work as it is iteratively adapted for a presentation context, site–specific setting and in certain cases professional performers. A categorisation and discursive analysis of both system and participant behaviour is a critical part of this evaluation process. This may be informal and forms part of the artist's discourse surrounding the work.

4 Interacting with Autonomous Agents – Five Case Studies

To address our core question of how artists and audiences approach autonomy as an aesthetic and experiential concept, the following invites the reader to engage with the questions we raised above more tangibly – from inside the artworks themselves. We discuss two pioneering artworks, as well as three of our own artworks, as we can offer an intimate working knowledge of those. This selection of works is not definitive in any respect; there are a great number of significant interactive systems that could also be included here, such as Penny's *Petit Mal* (Penny 2009), Glynn's *Performative Ecologies* (Glynn 2008) or McCormack's *Eden* (McCormack 2001).

4.1 The Senster, by Edward Ihnatowicz (1970)

Edward Ihnatowicz's interest in developing kinetic sculptures stemmed from his conviction that the behaviour of something tells us far more about it than its appearance. *The Senster* (Fig. 1) is celebrated for its originality and the notable effect it had on people who experienced it. It was Ihnatowicz's largest and most ambitious work: standing 2.5m high 'at the shoulder', the body of *The Senster* was constructed from tubular steel, with no attempt to disguise its mechanical nature.

The long articulated neck of *The Senster* contained six hydraulically operated joints and two additional actuated joints in the head allowed it to be positioned much more quickly than the rest of the neck. Four microphones and two Doppler radar units were mounted on the head. The radar units were used to detect motions of visitors. The microphones were arranged in two pairs—one horizontal and the other vertical—allowing sound to be localised by cross–correlating the inputs on each pair of microphones. A Philips P9201 digital computer, with 8Kb of core memory, was used to control *The Senster* according to programs loaded from punched paper tape. The sixteen–bit servo control output of the computer was fed into racks of custom electronics that provided the interface to *The Senster*. At the heart of this interface was the *predictor*, which smoothed the output voltages so that they followed spline-like curves and made the movement of *The Senster* look natural.

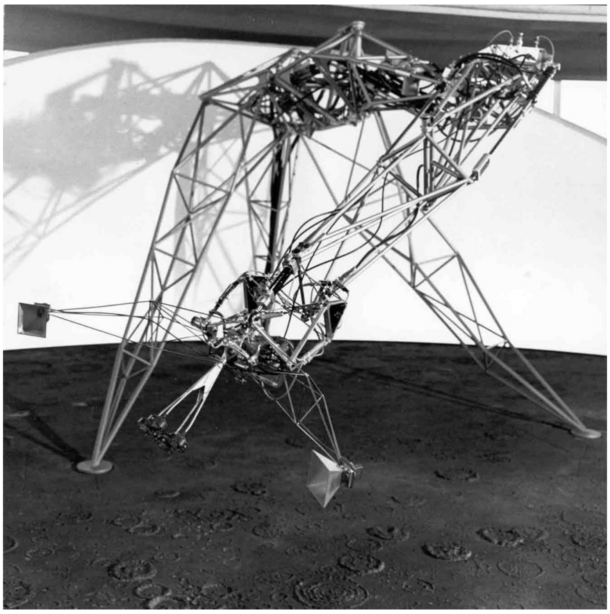

Fig. 1. Senster on display at the Philips Evoluon, Eindhoven, 1970-1974. © Edward Ihnatowicz, courtesy of Olga Inhatowicz.

As an example of behaviour–based robotics, *The Senster* was two decades ahead of its time. Using an approach that came to dominate robotics research from the late 1980s onwards, *The Senster* implemented a small set of simple behaviours that combined to produce seemingly more complex ones. Upon detecting a sound, *The Senster* would quickly turn its head in the direction of the source, with a speed proportional to the volume of the sound. If the direction of the source remained constant for some time, the rest of the body would slowly 'home–in' on the sound in stages. Sudden movements detected by the radar units, would 'frighten' *The*

Senster, causing it to withdraw. Likewise, loud noises would make it shy away from the source and if the general sound level became too loud, *The Senster* would rise up and "disdainfully ignore further sounds until the volume subsided." (Zivanovic 2005: 104).

When installed at Eindhoven, Ihnatowicz remained for three months after *The Senster* was unveiled at the Evoluon, spending much of his time in the exhibition, reprogramming *The Senster* and observing its interactions with audiences. The unpredictable behaviour of the audience, combined with the acoustic dynamics of the hall, apparently made *The Senster*'s behaviour seem more sophisticated than it actually was. In his own words, Ihnatowicz stated, "[p]eople seemed very willing to imbue it with some form of animal-like intelligence and the general atmosphere around it was very much like that in the zoo" (Ihnatowicz 1988: 6).

Ihnatowicz was disconcerted by his observation that people would refer to it as intelligent because "there wasn't an iota of intelligence in it: it was a completely pre-programmed responding system" (Reffin Smith 1984: 149). After his return to London, Ihnatowicz started to engage with artificial intelligence research, partly because, knowing how simple *The Senster*'s control software was, he "felt like a fraud and resolved that any future monster of mine would be more genuinely intelligent" (Ihnatowicz 1988: 6).

From the accounts we have of *The Senster* it is clear that Ihnatowicz was keenly interested in evaluating the nature of the interaction between his creation and audiences. One of Ihnatowicz's most interesting observations was the effect that the appearance of controlled movement and response to the environment had on audiences:

> When [Ihnatowicz] was testing [The Senster] he gave it various random patterns of motion to go through. Children who saw it operating in this mode found it very frightening, but no one was ever frightened when it was working in the museum with its proper software, responding to sounds and movement. (Michie & Johnston 1984: 153)

The difference in the observed reaction of audiences to *The Senster* moving randomly versus moving in response to environmental stimuli is entirely understandable, given Penny's observation about the embodied nature of our expectations: truly unpredictable movement is almost always a sign of something dangerous to be around. The secret as to why people were so willing to ascribe autonomy to *The Senster* may have been that, through careful observation and calibration of the work within it's complex environment, Ihnatowicz was able to strike just the right balance between unpredictability and responsiveness in its interactions with audiences.

4.2 Autopoiesis, by Ken Rinaldo (2000)

Ken Rinaldo's interactive installation *Autopoiesis* consists of fifteen articulated robotic arms suspended from the ceiling (Whitelaw 2004). Each arm is made from multiple lengths of untreated grapevine connected together by four plastic joints and held in tension with steel wire. The arms are three metres in length and tapered towards a tip. Mounted at the top of each arm is the control unit, housing a microcomputer and an array of sensors, three infrared proximity sensors to detect the presence of visitors and an array or four microphones to localise sounds. At the tip of each arm is an additional infrared sensor, which is used to probe the environment, and a microphone used to sense telephone touch-tones emitted by each arm.

A microcontroller in each arm implements “a collection of co-operating real-time processes” (Rinaldo 1998: 407) that interact to produce the machine’s performance. For example, simple reactive behaviours, such as using the infrared sensors to avoid collisions with audience members as they walk around, combine with others, such as moving towards the source of sounds but recoiling if the sound is too loud, to produce the ongoing behaviours of each robotic arm. The arms use telephone touch-tones to communicate between themselves, signalling the positions of detected audience members. Upon hearing these signals, they will move towards the position given, and it is this audible communication that allows the arms to coordinate their global behaviour.

In many ways, *Autopoiesis* can be considered as a single, independent entity, sensing and responding via its network of limbs. When an audience member enters this system, they together “make a second order autopoietic system that activates through a highly complex negotiated system of organized functioning of its parts” (Hall 2010: 2). As audience members move through the installation the arms move their tips toward them without ever touching them. When a person is present between the arms, the behaviour of the sculptures is more agitated, complex and probing. When the audience observes from the outside, the installation falls into a more serene state of 'waiting'. At the tip of two of the arms, cameras capture the scene, which is projected onto the walls of the gallery space, giving audiences a sense of being observed as much as of observing (Tenhaaf 2008: 13).

Autopoiesis is a clear example of a work that has been strongly influenced by formal theories of autonomy, explicitly referencing the ‘self–making’ property of all living things first described as ‘autopoiesis’ by Francisco Varela and Humberto Maturana. The work *Autopoiesis* produces a system that both “functions as an autonomous entity, made of both the biological and mechanical parts, and as an operationally open ‘life form’, when coupled with its phenomenological environment through interactivity” (Hall 2010: 3). The ‘organic’, autopoietic mechanisms, bring to the fore the interdependence of this machine, as well as how its evolution is closely coupled with the audience and the environment. The work conceptualises, implements and exhibits autonomy through the system’s adaptive properties, pertinent group behaviour and long-term change.

4.3 Uzume, by Petra Gemeinboeck, Roland Blach and composer Nicolaj Kirisits (2000-2003)

In *Uzume* (Fig. 2), an abstract, dynamic and sensitively responsive environment immerses the visitor. Its whirly, transitory nature is based on spatial representations of the temporal behaviour of nonlinear, chaotic systems, so-called strange attractors[2]. The work was implemented for a CAVE Virtual Reality (VR) System, where participants enter a cube the size of a small room, defined by 4 to 6 projection screens, and are coupled to the VR system by means of a head sensor, mounted on a pair of LCD shutter glasses, and, commonly, one to two hand sensors.

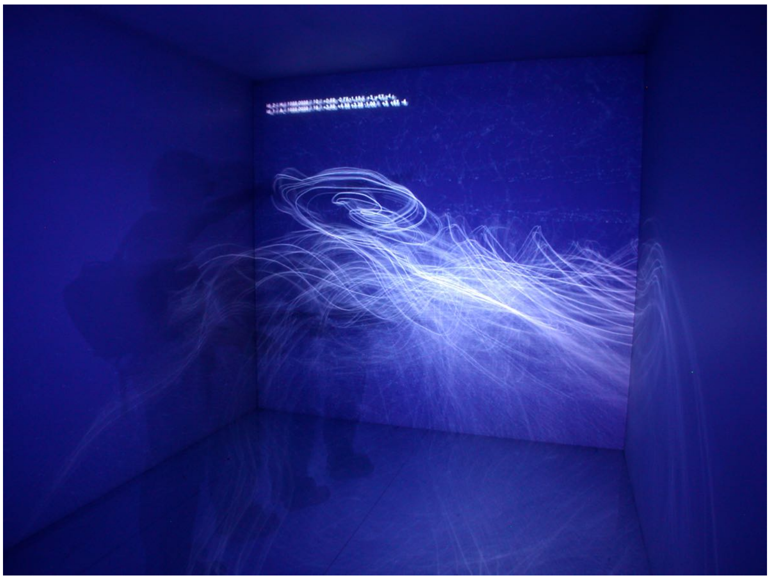

Fig. 2. Uzume, immersive virtual environment, 2003. © Petra Gemeinboeck.

The immersive, spatial interface of the CAVE was essential for *Uzume*'s development, for it allows the 'interface room' to be simultaneously inhabited by a real-time generated data space and a participants' body, and the space to be sculpted by the body's movements. Penny describes conventional VR systems, where "the disembodied gaze had the ability to 'move' on preordained paths within a pre-structured architectonic environment" (Penny 2011: 88). *Uzume* challenges all of these: the disembodied gaze, preordained paths, and pre-structured envi-

---

[2] Dynamical systems can be highly sensitive to initial conditions, and very small differences in initial conditions can result in very different behaviours, often referred to as the 'butterfly effect'. Strange attractors are semi-stable, on the borderline between instability (where they blow up to infinity) and stability (collapsing into a singular equilibrium point) and show the unique property that they never travel through space along the same trajectory twice.

ronment. In contrast to many CAVE environments, *Uzume*'s virtual environment is bound to the physical limits of the CAVE theatre, and participants need to move around and gesture with their two hand sensors to 'negotiate' with a dynamic, ever-changing space. None of its behaviours are scripted and its dynamic nature makes it appear wilful, eluding any control, even the illusion of control.

As the participants move around inside the projection space, they traverse the attractors' parametric fields that are mapped around their body and thus affect the environment's current state. The behaviours adapt over time based on the system's history and the interplay between its internal dynamics and the constant stream of data supplying the participant's position and movements. Each strange attractor is connected to an invisible particle grid that also reacts to the participants' presence. The effect is similar to moving in a viscous medium, gently warping the whirly lines (trajectories) when moving inside them. This elastic connection also made it possible to slowly push and pull the chaotic entities. Thus, the environment responds sensitively in endless fluid variations to each individual visitor.

Communicating with *Uzume* is similar to pursuing a dialogue without knowing the language of the other: all we can do is explore the other's gestural language but the actual meaning is never revealed, never completely decoded, never fully confirmed. Visitors, at first, approach *Uzume* like a puzzle that they can learn to 'figure out' if only they find the key to how it 'works'. But each of their movements causes a myriad of changes in the whirling environment: in shape, scale, density, speed, position, and even the potential for change (computation of new trajectory points per frame). Soon participants realise that *Uzume*'s world is different; too complex to understand and impossible to control. They stop moving to measure, analyse or tame the constantly changing space and begin to dance with *Uzume*. While participants probe *Uzume* with choppy gestures like a specimen, the environment evolves based on the sudden movement data input and appears to be more chaotic, more uncontrollable. Yet, when the participants' movements become more fluid and sinuous, it loses its strangeness, responsively mirroring the visitors' expressive playfulness and, at times, appears to unfold like an extension of their bodies. Now they perform together.

The development of *Uzume* involved hundreds of hours of observation, as it was impossible to directly compose or control how *Uzume* acted. Similar to the experience of the participants, we were confronted with a non-reproducibility and complexity, whose openness also relies on the fact that its potential evolution exceeds our imagination. The only direct control we had was in the mapping between participants' movements and the system's parametric input. This became the focus of our evaluation: that is, how well we were able to respond to the participants' gestures. In the process, we often resigned ourselves to describing *Uzume*'s response in anthropomorphic terms, simply to communicate what it looked or felt like but also to develop more expressive mapping relations. We continued to expand the expressiveness of these relations after each exhibition, based on observing the participants' behaviours and affective responses.

4.4 Accomplice, by Petra Gemeinboeck and Rob Saunders (2013–ongoing)

The robotic installation *Accomplice* (Fig. 3) embeds a group of autonomous robots into the architectural fabric of a gallery. The robots appear to inhabit the wall, sandwiched between the existing wall and a temporary wall that resembles it. Each robotic agent moves along the wall and is equipped with a punch and a camera eye, which they use to interact with their surrounds. This interaction is self-motivated; they are autonomous, curious agents, driven to explore their 'world' and discover 'things' they didn't expect. With a punch 'at hand' they are able to affect their world and create new 'things' whenever it seems already too familiar and they lose interest. Moving along the wall they share, they also use their punch to develop rhythmic knocking signals to communicate their presence to each other. As a result of this ongoing piercing, sculpting and signalling activity, the wall increasingly breaks open, and configurations of cracks and hole patterns appear that mark the machines' presence and traces their autonomous agency.

Fig. 3. Accomplice, robotic installation, 2013. © Petra Gemeinboeck.

Accomplice explores a similar notion of interactivity to *Autopoiesis*, while the audience's presence and actions matter, the individual robots in *Accomplice* do not rely on input from its visitors to interact with each other, allowing the work to evolve autonomously. The audience plays a part in the work's wider ecology but *Accomplice* doesn't necessarily respond to or perform for them. This is a conception of interaction that, in Simon Penny's words, "has been expanded beyond user-machine, to larger ideas of behaviour between machines and machine systems, and between machine systems and the world" (Penny 2011: 100).

The control system of the robots combines machine vision to detect features from the camera with audio processing to detect the knocking of other robots and computational models of intrinsic motivation based on machine learning. Movements, shapes, sounds and colours are processed, learned and memorized, allowing each robotic agent to develop expectations of events in their surrounds. This adaptive model of their 'world' allows the robotic agents to expect learned behaviours and proactively intervene. To these curious machines, learning and adapting are not goal driven but evolve based on what they discover and interpret as 'interesting' (Saunders 2001).

Accomplice's robotic agents physically inscribe their computational processes into our built environment by turning the wall into a playful stage for creating and learning, similar to a sandpit. Such an autonomous, proactive machine performance challenges common interaction paradigms of primarily reacting to what is sensed. As the agents are intrinsically motivated to explore their environment, the audience comes into play once they have created sufficiently large openings in the wall for them to detect and study the audience members as part of their environment. The appearance and behaviours of audience members are perceived by the system as changes in their environment. In line with the work's coupling with the built environment, the way in which it involves the audience pursues an expanded, ecological perspective. Thus, it is not only the robots that 'perform' for the audience, but also the audience that provokes, entertains and rewards the machines' curiosity.

Rather than being invited to control the course of events, the audience is implicitly implicated in the material interventions of *Accomplice*; they become an accomplice in the work's ongoing transformations. Initially, it is the physical impact of the work, the loud banging, expelled bits of wall, and dust accumulating on the floor, that draws them in intrigued or confronts them with a strange feeling of discomfort. As soon as they realise that there are active machines behind the wall, they often get close to the wall, moving along slowly and peeking into the holes to catch a glimpse of these strange trespassers. This often is the moment that captures them, and they begin to listen to the rhythmic knocking signals, follow their movements, and patiently wait in front of a hole for one of the robots to peek out, curiously sweeping its camera eye, and suddenly look back at them. It is interesting to observe how keen visitors are to be 'seen' by the robots, for them to acknowledge their presence. Yet the machines will soon lose interest and move on to continue chatting with the other robots or piercing along the raggedy edges of a hole. Similar to Ihnatowicz's observation, the encounter between human and non-human agents in *Accomplice* is reminiscent of those we have in the zoo.

Accomplice is the product of an iterative experimental and evaluative process, which started with an earlier work, called *Zwischenräume* (2010-2). In the first version, we took a more anthropomorphic approach to the robots' behaviours, which we then challenged in the next version by developing more machinistic and expressionless behaviours. Based on our observations of the audience response, in connection with our artistic intent, for the third version of *Zwischenräume* we strived for a middle ground, a machinic design that had some capacity to express its curiosity. *Accomplice* builds on and expands on this hybrid approach by conceiving the robots as social actors that share their wall territory with each other and use their tool to develop rhythmic communication signals.

4.5 Zamyatin, by Oliver Bown (2009-ongoing)

Zamyatin is a system for live improvisation with a human musician, developed over several years by the first author, Oliver Bown. It fits into a class of creative research known as 'live algorithms', which aspires to develop systems that can engage in meaningful musical interaction with a performer. Whilst consistent with the focus of this chapter, this is a context with clear differences to those discussed above. It concerns the domain of musical performance, which has a specific cognitive and perceptual nature (see Cross 2007). It also involves a different presentation format, with a hierarchy of participation, distinguishing system and musician from audience.

In *Zamyatin*, the goal, inspired by behavioural robotics such as the work of Beer (1996), was to take a two-layered approach, integrating a subsystem that might lay claim to behavioural autonomy, feeding into a 'composed' system designed by the author. This approach sees the creation of live algorithms as a design problem, by asking how composers can write creative decisions into a system at the same time as allowing it an operation of its own. Here, the approach to autonomy is somewhat philosophical, based on the idea that neither copycat learning nor the expression of rules, devised by the composer, facilitates autonomy. 'Merely writing musical rules' is considered undesirable, whereas iteratively developing the system behaviour along with the musical parameters the system operates, is seen as a viable creative process leading to a system capable of meaningful musical interaction.

Zamyatin has existed in two major manifestations: first as a continuous-time recurrent neural network, taking directly from Beer (1996); and, second as a decision tree, customised by the author to incorporate internal feedback pathways. In both cases, a fixed set of low-level audio features are extracted in real-time by a 'listening' system to be passed into this decision-making unit. The decision-making unit itself is not designed by the programmer, but is shaped using evolutionary optimisation to achieve abstractly stated behaviours, such as to remain silent until sound is heard at the input, to tend to produce repetitive patterns, or to exhibit behavioural variation over long time scales. For an explanation of the above terms and implementation details see Bown (2011).

Evaluation of *Zamyatin*'s behavioural character has taken the form of informal responses from musicians and fellow practitioners of musical meta-creation (MuMe) through 'meet-ups' in which collective performances are presented, for example at the MuMe Weekend at the International Symposium on Electronic Arts in Sydney in 2013. Observation and discussion with the musicians playing with *Zamyatin* reveal interesting differences in the level of expectation, reception, tolerance and compatibility amongst musicians, and understanding how to conceptually frame musician's responses is an important preliminary step to being able to successfully describe the interplay of concerns amongst which the autonomy of the system may be buried. A novel ethnomusicological study by Banerji (2012) of his own software system Maxine turns a traditional question: how well does the system perform on its head, by asking instead how different performers deal with the system in their own playing. Banerji asks, how well did the musicians respond to the system? How does Maxine make them play? As our introductory discussion illustrates, turning to study the behaviour of participants may be the more relevant focus. Banerji approaches the discussion of interaction as if the discourse involved the interaction between two musicians. This, combined with more conventional forms of performer and audience analysis, such as the recent survey-based studies of Eigenfeldt et al. (2012) and Brown et al. (2013), offers a working methodology for analysing the perceived autonomy of the system.

Performers' responses to *Zamyatin* have varied. In three notable cases so far, performers have reported being deeply engaged by the behaviour of the system, leading to a sense of mutual musical exploration. As discussed by Bown *et al.*, (2013) and originally raised as an idea by Pachet (2003), the experience of creative flow during performance with the system is the more immediately pragmatic goal of the performer, rather than any specific sense of autonomy. This pragmatism is reflected in comments by saxophonist Ben Carey on performing with *Zamyatin*:

> "I felt there was a responsibility for me to play a leading role, to not get stuck into trying to make the system react to me." – Ben Carey

From this, and similar comments and observations, comes the suggestion that naïve responses that ease this conflict for the musician, may be more crucial to establishing a working musical partnership than more advanced and complex behaviours, even if it means the antithesis of autonomy, i.e., something that the musician can reliably manipulate. This poses the interesting problem of how a rich autonomy is even manifest in the face of the potentially pragmatic concerns of the performer, which can be examined through further iterations of system prototype development and observation in action.

5 Evaluation

Working in the tradition of experimental arts, where the nature of the outcome of the creative process cannot be predetermined, evaluation is an important part of the iteration of observation and adjustment within the creative process, much of which occurs in the studio, before the work has been finally placed in its interactive context. Reflecting on our own and others' approaches towards evaluating autonomy requires us to consider how artists conceptualise autonomy in their work in order to build a successful set of high-level concepts that can guide their work's development, and how they apply that conceptualisation in the iterative process of tweaking and observing that is required to achieve desired behavioural outcomes.

We suggest that approaches to autonomy can be understood in terms of three rationalisations for why a system is autonomous. The obvious goal is to create a system that *appears* autonomous, demanding of the participant what Dennett (2009) calls an 'intentional stance', that is, a point of view with which one simplifies the understanding of the behaviour of a system by attributing intentions to it. *Autopoiesis* and *The Senster* are examples where the perception of autonomy is a clear goal, even if, as in the case of Ihnatowicz, the creator is surprised by a deeper attribution of intentionality than aimed for.

But there are additional approaches to the concept of autonomy that impact the way in which evaluation of the work is understood. The first approach is formal. Given a formal understanding of autonomy, such as that of Seth (2010) described previously, or Ashby's notion of homeostasis (Ashby 1952), one can design a system to that specification. Formal approaches provide a way to understand how autonomy might be implemented, and might also be understood by the participants experiencing an artwork. *Uzume* and *Zamyatin* both draw on such formal notions of autonomy, *Uzume* by drawing on the theoretical basis of strange attractors, and *Zamyatin* by targeting simple, formal behavioural targets. As artists, our evaluation of these works engages with issues around how participants experience interaction in the context of these simple formal properties.

The second approach is explanatory. A system may be considered autonomous if there is a reasonable explanation for why it should have acquired autonomous traits, even if this autonomy cannot be immediately measured or observed. Thus a system evolved through simulated natural selection may exhibit autonomy because it acted in such a way as to stay 'alive'. Glynn's *Performative Ecologies* and McCormack's *Eden* are well-known examples of works that make use of audience interaction as the basis for an evolutionary process, which gives a conceptual credibility to the notion that the interacting agents are autonomous. *Zamyatin*'s evolution has a slightly different significance: to establish behaviour in the agent that is not the direct design of its maker. *Accomplice* involves an explanatory element in the curious nature of the agents, which contributes to the agents' perceived autonomy but may also enhance it conceptually.

The above do not constitute evaluative methods in themselves, but act as elements that may guide the evaluation of works that exhibit and explore autonomy, as understood both by the creator of the work and its audience.

6 Conclusion

Interactive experiences with autonomous works are qualitatively different from experiences of reactive or responsive works. The locus of control is shifted from the audience to a shared space in-between and the interaction is more akin to a negotiation or an unscripted dialogue in which the changing character of the work plays a key role in shaping the 'conversation'. In some instances, interaction may be actively sought by the work, for instance, when *The Senster*, *Autopoiesis* and *Uzume* appear attentive, continuously monitoring their domains for change. In other instances, the interaction will appear peripheral, for example, when the social knocking of *Accomplice* or the dynamic swirling of *Uzume* continues to evolve in the absence of an audience. When caught in loops of co-production the interaction becomes an intimate dialogue, as in the playful dance between *Uzume* and a person or when a musician gains fluency with *Zamyatin*. Interacting with these works becomes a co-performance, sensitive to the contingencies of a moment and place, such that each encounter is unrepeatable, a uniquely singular event.

In this chapter we have discussed several works related to the notion of building autonomy. Through reflection on our own practice, and a review of others' approaches to working with notions of autonomy, we have discussed how this forms a critical basis for the artist's evaluation and iterative development of their work, and how audiences may also take on or independently apply these concepts to the works they experience.

References

Ascott R (2002) Behaviourist Art and the Cybernetic Vision. In Packer, R. and Jordan, K. (eds) Multimedia: From Wagner to Virtual Reality. W. W. Norton & Company, New York, p 104-120

Ashby, W. R, (1954). *Design for a brain*, Wiley, New York.

Banerji R. (2012). Maxine's Turing Test–A Player-Program as Co-Ethnographer of Socio-Aesthetic Interaction in Improvised Music. In *Eighth Artificial Intelligence and Interactive Digital Entertainment Conference*.

Beer R (1996) Toward the evolution of dynamical neural networks for minimally cognitive behavior. In From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior, MIT Press, p 421-429

Boden, M, & Edmonds, E (2009). What is generative art? Digital Creativity, 20(1-2), p 21-46

Bown O, (2011) Experiments in Modular Design for the Creative Composition of Live Algorithms. Computer Music Journal, 35(3), MIT Press, p 73-85

Bown O., Eigenfeldt E, Martin A, Carey B, Pasquier P, (2013) The Musical Metacreation Weekend: Challenges arising from the live presentation of musically metacreative systems. In Proceedings of the Musical Metacreation Workshop, AIIDE Conference, Boston.

Brown A, Gifford T, & Voltz B (2013). Factors affecting audience perceptions of agency in human computer musical partnerships. In *Proceedings of the 9th ACM Conference on Creativity & Cognition* (p 296-299). ACM.

Cornock S, Edmonds E (1973), The Creative Process Where the Artist is Amplified or Superseded by the Computer, Leonardo 6(1), p 11-16

Cross I (2007) Music and cognitive evolution. In Dunbar R, and Barrett L (eds.) Oxford Handbook of Evolutionary Psychology, Oxford University Press, Oxford, p 649-667

Dennett, D. Intentional Systems Theory. In McLaughlin, B., Beckermann, A., Walter S. (eds), Oxford Handbook of the Philosophy of Mind, OUP, (2009), 339-350.

Eigenfeldt A. Burnett A, & Pasquier P (2012). Evaluating musical metacreation in a live performance context. In *Proceedings of the Third International Conference on Computational Creativity* (pp. 140-144).

Glynn R, (2008) Conversational Environments Revisited. In Proceedings of the 19th Meeting of Cybernetics and Systems Research, Graz, Austria

Hall J (2010) An Autopoietic Aesthetic for Interactive Robotic Installation. Available online at: www.jenhall.org/pdf_files/Autopoietic_Aesthetic_JHall.pdf

Haque U (2007) The Architectural Relevance of Gordon Pask. Architectural Design 77(4), 54–61

Horowitz A C, Bekoff M (2007) Naturalizing Anthropomorphism: Behavioral Prompts to Our Humanizing of Animals. In ANTHROZOÖS 20:1, 23-35.

Ihnatowicz E (1988) Portrait of the Artist as an Engineer, unpublished book proposal, available online at: http://www.senster.com/ihnatowicz/articles/artist_as_engineer.pdf (accessed on 15/08/2013)

McCormack, J. (2001), Eden: an evolutionary sonic ecosystem, in J. Kelemen and P. Sosik (eds), Lecture Notes in Artificial Intelligence, Vol 2159, *Advances in Artificial Life, Proceedings of the 6th European Conference (ECAL)*, Springer-Verlag, Berlin, Germany. p 133-142.

Michie D, Johnston R (1984) The Creative Computer: Machine Intelligence and Human Knowledge, Penguin Books, London

Pachet F (2003) The Continuator: Musical interaction with style. Journal of New Music Research 32 (3):333-341.

Penny S (2009) Art and Artificial Life - A Primer. Proceedings of the 2009 Digital Arts and Culture Conference, UC Irvine.

Penny S (2011) Towards a Performative Aesthetics of Interactivity. *Fibreculture Journal*, FCJ-132:72–108.

Reffin Smith B (1984) *Soft Computing: Art and Design*, Addison-Wesley, 147–155

Rinaldo K (1998) The Flock. Leonardo 31(5):407

Saunders R (2001) Curious Design Agents and Artificial Creativity. PhD dissertation. The University of Sydney

Seth A (2010) Measuring autonomy and emergence via granger causality. Artificial Life 16:179–196

Stern N (2011) The Implicit Body as Performance: Analyzing Interactive Art. Leonardo 44(3). The MIT Press, 233-238

Tenhaaf N (2008) Art Embodies A-Life: The VIDA Competition. Leonardo 41:1:6–15

Turing A (1948) Machine Intelligence. In Copeland B (ed) The Essential Turing: The ideas that gave birth to the computer age, Oxford University Press, Oxford

Whitelaw M (2004) Metacreation: Art and Artificial Life. MIT Press, Cambridge, MA

Zivanovic A (2005) The Development of a Cybernetic Sculptor: Edward Ihnatowicz and The Senster. In Proceedings of Creativity & Cognition '05, ACM, p 102–108