25th IEEE International Symposium on
Robot and Human Interactive Communication (RO-MAN)
August 26-31, 2016. Columbia University, NY, USA

Interacting with Curious Agents: User experience with interactive sculptural systems*

Matthew T. K. Chan[1], Rob Gorbet[1], Philip Beesley[2] and Dana Kulić[1]

Abstract— **To enable long term, engaging social human-machine interaction, robots and other autonomous systems must be able to move beyond purely reactive interaction control strategies, and engage in shared-initiative interaction. In this paper, we describe an implementation of an interactive art sculpture which generates interactive behaviours using curiosity-based learning. Using its own internal motivation formulated as a curiosity drive, the system initiates interaction with and responds to human visitors, generating continuously evolving interactive behaviours. The proposed system was tested in a user study with a prototype interactive sculpture installation.**

I. Introduction

Interactive arts are a type of art form that invites viewers to become active participants in the enjoyment and creation of the artist's vision. These systems sense and respond to participant presence and actions, creating an emergent, interactive experience unique to each visitor. Recent advances and miniaturization of computers, sensors and actuators have enabled artists to create highly complex interactive artworks, such as the Hylozoic Series of kinetic sculptures built by Philip Beesley Architect Inc. (PBAI). In this system, a network of microcontrollers controls and samples hundreds of actuators and sensors [1][2]. Each node in the network can perform a simple set of interactive behaviours. While individual behaviours have previously been either prescripted or random, complex group behaviours can emerge through communication among nodes and interaction with the spectators [3]. Figure 1 shows a Hylozic Series interactive sculpture that was installed in the Museum of Modern and Contemporary Art in Seoul, South Korea.

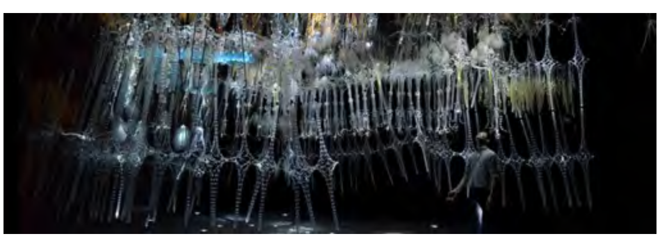

Fig. 1: *Epiphyte Chamber*, installed in the Museum of Modern and Contemporary Art in Seoul, South Korea, 2014.

Interactive arts have been a focus of increased interest and development in recent years, and a number of interaction modalities have been investigated, including sound [4], [5], visual arts [6] and interactive architecture [7]. In addition to the interaction modality, a key distinguishing feature of interactive artworks is the type of interaction. Edmonds et al. [8] differentiate between four types of interaction: static, dynamic-passive, dynamic-interactive and dynamic-interactive (varying). To date, the majority of systems developed in the literature are of the non-varying dynamic-interactive type, as their responsive behaviours do not change over time.

One of the goals of the Hylozoic Series is to invite the users to contemplate whether architecture can become alive, albeit in very primitive ways [3], [9] (Figure 2). The series is named after the ancient Greek philosophy, hylozoism, which is a belief that all matter is alive in some sense. While earlier generations of the sculptural systems exhibited primitive responsive behaviours, they were manually designed and remain unchanged over time. To take on the characteristics of a higher level living system, the sculptural system should have the ability to generate and modify its own behaviours, and adapt to changes in its external environment.

Fig. 2: Hylozoic Ground, one of LASGs first immersive environments, explored a new generation of responsive spaces. Hylozoic Ground, Venice Biennale, Italy (2010).

The evolution of behaviours of artificial systems has long been a focus of research in artificial life [10], [11] and developmental robotics [12], [13]. Oudeyer et al. [14] proposed algorithms for robot learning mimicking curiosity called Intelligent Adaptive Curiosity. In the robot, curiosity is implemented as a reinforcement learning algorithm with the objective of maximizing learning progress, defined by the change in prediction error. However, the majority of previous work has been developed for small robots with few sensors and actuators, and did not consider how such algorithms can scale to a distributed system with a large sensorimotor space.

Furthermore, as public interactive art pieces, interactive

*This work was supported in part by Canada's Natural Sciences and Engineering Research Council, and by Mitacs in partnership with PBAI.

[1]M. T. K. Chan, R. Gorbet and D. Kulić are with the University of Waterloo `dana.kulic@uwaterloo.ca`

[2]P. Beesley is with the University of Waterloo and with Philip Beesley Architect Inc.

sculpture and architecture systems should also be engaging and interesting to their visitors. We hypothesize that, if system behaviours do not change, users will find the system less interesting over time as the behaviours become predictable. This is undesirable for a permanent installation in which the same users may interact with the sculpture over an extended period of time. Moreover, in a system with a large number of sensors and actuators, programming a complex set of carefully choreographed behaviours is complicated and requires lengthy implementation and testing processes. To address the challenge of long-term, adaptive engagement and self-motivated autonomy, we proposed a new approach for behaviour generation, based on the Curiosity-Based Learning Algorithm (CBLA) [15]. The CBLA re-casts the interactive sculpture as a set of agents driven by an intrinsic desire to learn. Presented with a set of input and output variables that it can observe and control, each agent tries to understand its own mechanisms, its surrounding environment, and the occupants, by learning models relating its inputs and outputs. The CBLA algorithm is based on Oudeyer's Intelligent Adaptive Curiosity algorithm [14]. The CBLA builds on Oudeyer's work, applying it on an architectural-scale distributed interactive sculptural system with the aim of developing time-varying dynamic interactive interaction between the sculpture and visitors. The design of such systems and their evaluation with human users may also be of interest to the robotics community, to understand how users respond to robots that learn and change their behaviour over time.

To better understand users' response and perception of the learning behaviour, it is important to be able to capture their reactions as they are interacting with the sculpture. An experimental test bed running CBLA was therefore constructed and users were invited to interact with the sculpture. Through this user study, we examined how users perceive and respond to the behaviours generated by the CBLA system in comparison to behaviours that are scripted by a human designer. In particular, we investigated the relationship between the behaviours of the sculptural system and users' self-reported interest.

This paper is organized as follows: In Section II, the CBLA algorithm, first introduced in [15], is briefly described. In Section III, the physical implementation of the testbed and the algorithm configuration are described. Section IV describes the initial validation experiments, demonstrating the response of the system to simple stimuli and investigating the influence of parameters on system behaviour. Section V describes the user study, while Section VI provides a discussion of the results. Section VII concludes the paper and provides directions for future work.

II. CBLA ALGORITHM

The CBLA algorithm [15] is based on Oudeyer's Intelligent Adaptive Curiosity algorithm [14]. In this approach, the learning system consists of two learning mechanisms, the classic learner and the meta learner [14]. Based on the context (the sensors' inputs) and the action (the actuators' outputs), the classic learner computes a prediction of the consequence, i.e., the resultant sensors' inputs at the next time step. Then, it implements the action and compares the actual consequence with the prediction and modifies its model in order to reduce the error in the prediction. The meta learner predicts the error of classic learner, i.e., it estimates how accurately the classic learner is able predict the consequence. The actual prediction error is then fed back to the meta learner, which modifies its estimate of prediction error. This change in the prediction error is recorded at each time step. The expected learning progress is then estimated by calculating the change in error and used as the reward, R. As the classic learner's prediction gets better in a given area of sensorimotor space, the expected learning progress diminishes and the reward for exploring that area gets smaller. Higher rewards in other as-yet unexplored regions encourage the algorithm to move on, to satisfy its "curiosity".

A key feature of the learning algorithm is the idea of regional experts. Each region collects exemplars of similar sensorimotor context, and has an expert that is trained on the exemplars in the region. Exemplars are the training data for the prediction model and they are collected as the system selects actions and observes the consequences. The "features" are a vector of the sensory inputs $S(t)$, and a vector of selected actions $M(t)$ at time t; and the "labels" are the resultant sensory inputs $S(t+1)$ at time $t+1$. The regional experts constrain the estimate of learning progress within their respective sensorimotor contexts $SM(t)$ which is a concatenation of the vectors $S(t)$ and $M(t)$. This is important because it allows each expert to use a simpler model, as it covers only a small region of the state space.

During experiments implemented on a dog-like robot, Oudeyer [16] observed that the resulting robot behaviour had similarities to children's learning and playing behaviours. Over time, the robot would change focus to interact with objects that have increasingly complex responsive relationships. However, in the experiments done in [14], human participants were never part of the environment that the robot tries to learn. In this work, we focus on examining the interaction between the learning system and human users, and how the two influence each other. Since we expect human responses to be difficult to predict, after the robot has modelled its own mechanisms, the robot should then focus on generating behaviours that evoke reactions from the users. We hypothesize that this type of learning mechanism is well suited to interactive architecture systems, as the learning process itself will generate interesting behaviours and be an interesting feature of the installation to the visitors. In fact, it is important to note that our end goal relies on the variability of human input, which we assume will be difficult to learn completely. Unlike much of the work in robotics, our objective is not to produce a well learned model, rather, we hypothesize that the *process* of learning is interesting to observe and will generate engaging and time-varying behaviours. We believe the very complexity which makes human interaction difficult to learn and predict will result in increased visitor engagement.

In previous work [15], we adapted Oudeyer's algorithm to apply it to the Hylozoic Series interactive art sculpture, which is a distributed sculptural system with a large number of sensors and actuators. A distributed approach is necessary since the sculptures often contain hundreds of sensors and actuators distributed over an area the size of a large room. Each system node runs its own instantiation of the algorithm, called the CBLA Engine. This gives the sculpture the ability to generate its own behaviours through self-motivated learning. A key feature of the CBLA system is the use of sensors that can sense both the system's internal state (i.e., are used for *proprioception*) and human interaction behaviours, such as proximity or touch. The distributed CBLA agents run asynchronously but are coupled through two mechanisms: (1) shared inputs, and (2) virtual inputs. Shared inputs couple the agents through a shared sensor, while virtual inputs couple agents by using the motor output of one agent as a virtual sensor input of another agent. The asynchronous nature of the system allows learning agents controlling actuators with different operating bandwidths to coexist. Meanwhile, the links among the different learning agents enable information to travel within the system, allowing distributed learning. Furthermore, to make the learning process visible to the user, the maximum output level of the system is scaled based on the knowledge gain potential, so that the system is less active when it believes that its learning progress is low.

III. CBLA Testbed

An experimental test bed was built to investigate how users interact with the CBLA system. This is a four-cluster system that resembles a typical interactive sculpture produced by Philip Beesley Architect Inc. (PBAI), shown in Fig. 3.

Fig. 3: Photograph of the multi-cluster test bed when actuated.

The testbed consists of three types of components: Light, Fin and Reflex units. A Light unit is made up of one high-power LED and one ambient light sensor (shown in 4a). The high-power LED is mounted on top of a flask containing coloured liquid. The ambient light sensor is mounted beside the flask under the LED (red square in Fig.4a). This allows the ambient light sensor to measure the intensity of the light emitted by the LED. A Fin unit is made up of two shape memory alloy (SMA) wires, a 3-axis accelerometer, and an IR proximity sensor. The two SMA wires pull on two levers that move a Fin, which is a mechanism made of compliant acrylic rods that curls up (shown in 4b). The IR proximity sensor and accelerometer are mounted around midway between the tip to the root of the Fin. Located below the Fin, the Reflex unit consists of a vibration motor, a pair of LEDs, and an IR proximity sensor that are mounted in the middle of two black frond-like objects (shown in 4c).

(a) Light Unit

(b) Fin Unit

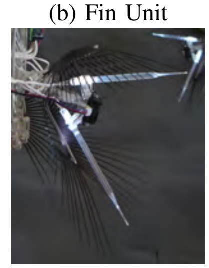

(c) Reflex Unit

Fig. 4: Photograph of components in the multi-cluster test bed.

The resulting hardware and sensing mechanisms allow the system to sense its own state and a variety of environmental and user activities. The ambient light sensors are both proprioceptive and external, they can sense changes in lighting level due to the system's own LED activations as well as changes in ambient lighting (e.g., night vs. day). Similarly, the accelerometer sensors sense the movement of the fins, which may be caused by system-initiated actuation through the SMA actuators, by environmental effects (e.g., air current in the space), or by a user touching or moving the fin. The IR proximity sensors sense the presence of nearby visitors as well as any permanent structures.

A. Configuration of CBLA Nodes

Each isolated CBLA Node is an instantiation of a CBLA engine and is associated with one actuator. A CBLA system is constructed by linking these isolated CBLA Nodes through virtual inputs and shared inputs. In this test bed, three main types of isolated CBLA Node exist: Half-Fin Node, Light Node, and Reflex Node.

Fig. 5 shows the isolated CBLA Nodes in each cluster. Two Half-Fin Nodes control the bending of a Fin through their respective SMA Controller Nodes (Fn.SMA-L and Fn.SMA-R). The pair of Half-Fin Nodes share a Fin-mounted IR proximity sensor (Fn.IR-F), and the 3 axes of the accelerometer (Fn.ACC). A Light Node controls the brightness of a high-power LED through an LED driver Node (Ln.LED) and its sensory space consists of an ambient light sensor (Ln.ALS). There are two types of Reflex Nodes, one is associated with a pair of LEDs (Fn.RFX-L), and one is associated with a vibration motor (Fn.RFX-M). In their sensory space, they share one IR proximity sensor (Fn.IR-S).

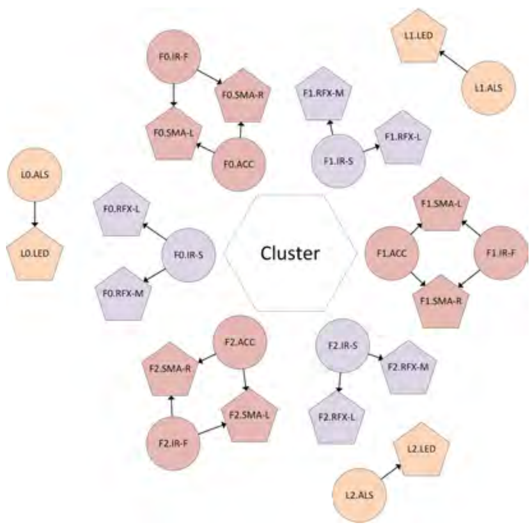

Fig. 5: Make-up of a cluster of isolated CBLA Nodes (pentagons). Half-Fin Nodes are shown in red; Light Nodes are shown in orange; and Reflex Nodes are shown in blue. Ambient Light Sensors (ALS), Accelerometers (ACC), and IR Sensors (IR) are shown as circles.

The Isolated CBLA Nodes are connected to each other via virtual inputs, by treating the output of one CBLA node as an input variable to other Nodes. This creates an interconnected network of CBLA Nodes that spans the entire sculptural system, allowing information regarding the external environment to travel through the sculpture.

Different network configurations produce different system behaviours. Here we investigate one type of network configuration, where nodes are linked based on their spatial proximity. Fig. 6 shows the connection graph: neighbouring nodes are connected via bidirectional links, and nodes in neighbouring clusters that are in close proximity are linked via unidirectional links. The connections that daisy-chain the four clusters allow information to spread throughout the test bed.

In this configuration, nodes' responses to external disturbances or environmental changes are expected to concentrate close to the source, and spread to other nodes gradually.

B. *Prescripted Behaviours*

For the purpose of comparing between CBLA and the current practice of designing prescripted behaviours, each CBLA Node has a Prescripted Engine in addition to the CBLA Engine. This allows us to quickly switch between the two kinds of behaviours during the user study. The prescripted behaviours are implemented based on a specification by a human designer, and are similar to previous behaviours of the Hylozoic series [17].

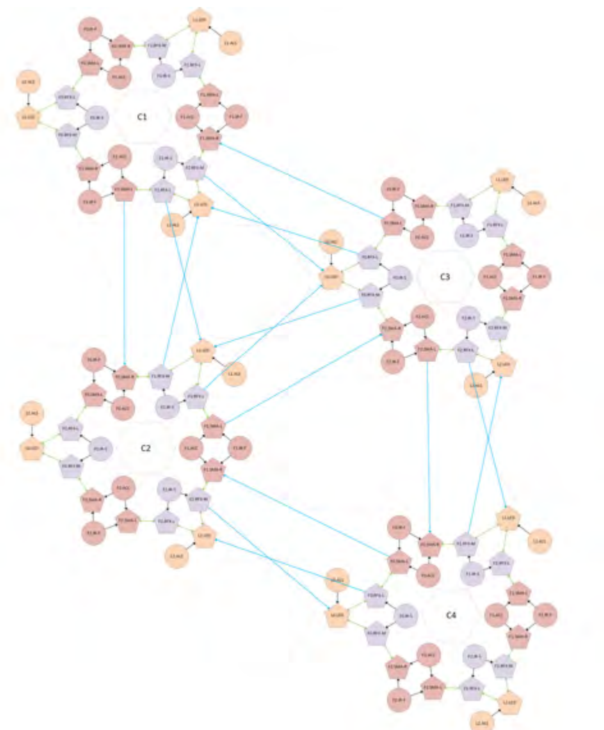

Fig. 6: Connectivity graph of the test bed. Each blue unidirectional arrow represents a virtual input link. The output of the Node at the origin of an arrow is fed into the Node that it is pointing toward as an input variable. Green bidirectional arrows represent virtual links between neighbours within a cluster.

IV. User Study

We next conducted a user study to investigate user perceptions of the interactive behaviours generated by the CBLA. The study aimed to compare the users' responses to the behaviour of the CBLA system and the prescripted behaviours. In addition, the relationships between the intensities and types of activations and users' level of interest as they interact with the sculpture were examined. The study was approved by the University of Waterloo Office of Research Ethics.

A total of 10 participants were recruited to interact with the test bed described in Section III. Participants were recruited from the researchers' contact lists. All participants were healthy, physically able adults within the age range of 18 to 65 years old, male and female. To avoid the subject-expectancy effect, participants had never encountered any of the Hylozoic Series sculptures before, and had little prior knowledge about the CBLA.

Each participant was invited to interact with the prototype described in Section III, installed in the Toronto studio of PBAI. To interact with the prototype, participants could chose to move around within the system, move/wave their hands near individual prototype elements, touch the sculpture as well as move/hold the fins. As described in Section III, participant movement was detected with the proximity IR sensors, while contact with the sculpture was detected with the accelerometers. Only one participant was interacting with the sculpture at a time. The participant was informed about the procedures of the study and asked to sign a consent form before he or she began participating in the trial. The participant was told that he or she was free to walk around the space and interact with the art sculpture in any way, and that there were no prescribed interactions. The participant was also told that the entire trial would last 20 minutes and that he or she may request to terminate the study at any time. The floor was lined with a grid, which allowed the participants to specify their locations within the sculpture. The participant was given a pen and an envelope with a stack of 8 identical business-card-sized questionnaire cards shown in Fig. 7. He or she was asked to fill out one card every time an audible tone was played. The first question on the card asks the subject to write down the current sample number as shown on a screen visible to them. The second asks for his or her subjective interest level regarding the behaviour of the sculpture at that moment. The third question asks the subject to mark his or her location at that moment by looking down at the grid on the floor.

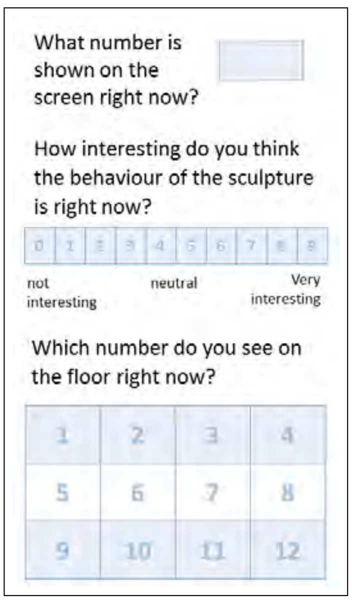

Fig. 7: Questionnaire card for the user study.

There were two versions of interactive behaviours: Prescripted Mode and CBLA Mode. The participants were not informed about the versions of the behaviours that they were interacting with nor the fact that there were two different versions of interactive behaviours. Both versions of the behaviours were run for approximately the same amount of time for each participant, with the first Mode being selected randomly. Half of the participants were exposed to Prescripted Mode first and the other half were exposed to CBLA Mode first. After the fourth tone (which was half-way through the trial), the researcher would switch to the other version manually after the participant had finished filling in the questionnaire card and resumed interacting with the sculpture.

At the end of the trial, the participant filled out an exit questionnaire. The questionnaire asked for their feedback on their overall experience, their opinion on how responsive the system was to their presence, and included a free form question allowing participants to provide any additional comments. In addition, observations of the participants' behaviours and interesting comments made in verbal conversations with the experimenter were noted to provide further insights.

A. Results

Three main sets of data were collected in each trial: the set of 8 questionnaire cards, the exit questionnaire, and the states of the system during the entire trial.

1) Qualitative Feedback: The qualitative responses obtained from the exit questionnaires indicated that some participants perceived the system to be both reactive and taking initiative. For example, one participant reported:

> *Sometimes the sculpture seemed to be guiding me. It will flash lights on more and stop when I'm within range of the sensors. Sometimes it seemed to be following me. It would actuate behind me.*

Even though the participants were not informed that the behaviour of the system would change during the trial, many participants noticed the change and commented on it in the exit questionnaire. One participant, who first interacted in the CBLA condition, followed by the Prescripted behaviour condition, reported:

> *It seemed to generate random movements at the beginning, but a pattern was repeated as the experiment goes on.*

Another participant, who observed the two conditions in the same order as above, also noted the change and commented on how the change affected their interest in the interaction.

> *Sometimes the responses I got from the sculpture were not what I was hoping for. For example, when I held up my hand to the leaf, the leaf itself was not really moving towards me as much as the two other leaves that were a bit farther away. Overall, the sculpture's behaviour did appear to be more predictable as time went on. Yet the downside of this is that it becomes less exciting. I felt as if there was not much to explore further and my curiosity level went down accordingly.*

On the other hand, another participant in a trial where Prescripted Mode was presented first described the behaviours of the system as follows:

> *...Parts that used to do something with straightforward motion/presence became unresponsive on one pillar - could not figure out the pattern. Behaviour changed a bit with time but slightly seemingly random as to how it changed.*

These responses indicate that users were aware of the differing behaviours between the two system conditions, and could identify their characteristics: the pre-scripted mode was more predictable, but to some viewers may become less interesting and engaging over time. On the other hand, in the CBLA condition, the behaviour was more complex and varying, but appeared to some viewers as random and unresponsive.

2) Quantitative Analysis: The questionnaire cards recorded each subject's levels of interest and locations at the 8 sample points. The levels of interest were used to compute the subject's overall interest level in each of the two conditions, Prescripted Mode and CBLA Mode. We further separated the data into two groups based on whether the trial started with Prescripted Mode first or CBLA Mode first. The location data was used to look for correlations between activations near the participant and his or her reported interest level.

The input, output, and internal variables of the CBLA Nodes and non-CBLA Nodes during the entire trial were collected. The activation values of each CBLA Node were calculated at 1.0s time intervals in order to analyze the correlations between activation and interest levels.

We first separated the data on levels of interest from the questionnaire cards into four groups: Prescripted Mode (first half), Prescripted Mode (second half), CBLA Mode (first half), and CBLA Mode (second half), based on the behaviour generation type and whether the participant observed pre-scripted mode in the first or second half of their session. We computed the average across all trials, resulting in four sets of data with five elements each.

We then performed Welch's T-Tests for unequal variance on these four sets of data. We observed a large variance between the participant responses, resulting in no statistically significant differences between the two conditions, except for the comparison between levels of interest between CBLA Mode and Prescripted Mode when they were on during the second half of the trial. This test was able to reject the null hypothesis with a two-tailed P-value of 0.0684. The average interest level of Prescripted Mode and CBLA Mode were 6.0 and 4.375 respectively. This means that Prescripted Mode was more interesting than CBLA Mode when it was on during the second half of the trial. We did not find any significant differences in reported levels of interest between CBLA Mode and Prescripted Mode in any other cases.

We next investigated if there were any correlations between the participants' reported levels of interest and the activation values of the sculpture. In this analysis, we did not consider which version of the behaviours was running at the sample point and focused on the actual activation values of the sculpture.

The extracted activation values for each CBLA Node were the output levels of each node averaged over 1.0s windows. For each trial, there were 8 sample points which correspond to the times when the participant filled out questionnaire cards and reported his or her levels of interest and location. For each sample point, the time interval containing it and 30 time intervals (which translates to 30 seconds) preceding it were considered as activations related to the sample point. A time interval of 30s was chosen in order to include prior activations that were likely to be associated with the participant's response at a sample point. We considered both overall average and peak activations, and activations for each device type.

For the overall activation, there was a weak positive correlation between participant interest level and both average activation value ($r = 0.37$) and peak activation value ($r = 0.32$).

For device type correlation, there was a higher correlation for Light Nodes than the overall correlation when considering the average activation levels, while higher correlation was found for Half-Fin Nodes when considering the overall peak activation levels. Interestingly, we did not find any significant correlation between the activation levels of the Reflex Node to the participant's interest level.

Although the overall correlations across all ten participants were relatively weak, there were in fact large variations among different participants. For instance, the overall peak activation for participant 10 showed a strong correlation with the subject's interest level, with $r = 0.92$. On the other hand, the same metric for participant 3 showed no significant correlation.

This shows that different people responded very differently to the behaviours of the sculpture. For instance, some people might indeed be attracted by the activations and were engaged to interact with the sculpture, while some might be attracted to the sculpture due to other reasons, such as the aesthetic of the design.

In the exit questionnaire, we asked the participants to rate their overall interest level and how responsive they thought the behaviour was. There was a moderate correlation ($r = 0.68$), between responsiveness and overall interest level. Similarly, the r between responsiveness and the average interest level reported on the questionnaire cards was 0.57. This shows that the perceived responsiveness of the behaviours did have a moderate positive correlation with participant interest.

V. Discussion

The observations from the user study raise a number of questions about factors affecting visitor interest and engagement with interactive sculptures. First, can autonomously generated behaviour simulating life-like motivations be as interesting and engaging as behaviours that are designed specifically for interaction? Second, what are the factors that influence visitor perceptions of the system?

The long term aim of the Hylozoic Series installation is to generate life-like behaviours automatically. However, we don't know if a life-like system will actually be interesting to all visitors. Gaver et al. [18] suggested that ambiguity can be used as a tool to create interactive arts that are more engaging and thought-provoking by posing questions without providing solutions. However, without the balance between ambiguity and consistency, the work can also appear to be confusing and meaningless. Indeed, research has shown that people find events that are new and comprehensible interesting [19]. Novelty and complexity alone are not enough to make a behaviour interesting; it may also need to be understandable and at least somewhat predictable.

In prescripted mode, behaviours are designed to react immediately and consistently to each visitor action. Unlike the Prescripted Mode, the responsive behaviours of the CBLA Mode are much less consistent. For instance, the type, the intensity, and the timing of the actuations are often different when given the same trigger. In addition, there are far more spontaneous activations as well. This unpredictability of the response makes it difficult for the users to recognize the causal relationship between the response of the sculpture and their actions. Thus, this factor might have contributed to the lower responsiveness ratings.

The preference of some users for more predictability might be of interest also when designing socially interactive robots which incorporate learning mechanisms [20].

In addition to the types of behaviour generated by the sculpture, visitor background knowledge and interests, as well as the context and duration of their visit, may also have a significant influence on the visitors' experience. The participants recruited in this study may be different from the typical visitors to art exhibitions, where these sculptures are typically displayed. People who frequent art exhibitions may not be "naive users", having had motivation and opportunity to read about the work prior to their interactions with the sculpture. Indeed, Silvia showed that an abstract poem would appear more interesting when hints about its meaning were given [21] and that concepts that are confusing to novices can be interesting to experts [19]. In this user study, the participants didn't necessarily have any interest in the arts and were not given any information about the behaviours of the sculpture. Knowledge of the conceptual and technological background of the Hylozoic Series sculptures might enable the users to appreciate the complex and subtle patterns of the CBLA system that aren't immediately obvious. Thus, the simplicity and comprehensibility of Prescripted behaviours might seem more interesting to a naive visitor.

From the written responses on the exit questionnaires and informal conversations after the studies, we observed that people had very different interests and expectations about the behaviours of the sculpture. Some expected much more coordinated, fast-pace movements, while some found the slow and organic-looking movements appealing. Some expected the sculpture to be very responsive, while some did not even know at first that there were sensors that could detect their presence. Some participants took figuring out the exact mechanism of sculpture's behaviours as a challenge; some were enjoying it and some were frustrated by their inability to figure out the patterns. Moreover, there were also participants who enjoyed looking at the design of the sculpture, and some who were interested in the design of the circuit boards and actuators. They spent a great deal of time examining the details of the sculpture itself rather than interacting with it. Furthermore, the ways in which the participants interacted with the sculpture varied greatly. Some mainly stood back and observed, while some walked around the space rapidly and touched many parts of the sculpture at great frequency. In fact, one participant was taking apart the sculpture in order to better examine the parts and how those alterations change the activation patterns of the sculpture. Perhaps by informing the users about the behaviours of the sculpture, the users might better understand its behaviours, avoiding a negative feeling of disappointment when what they expect is absent.

Moreover, we hypothesize that, over a long period of time, prescripted behaviour would become repetitive and less interesting. In this user study, the participants were only interacting with the prescripted behaviours for 10 minutes and perhaps that was insufficient for them to realize that it was repetitive and lose interest. In fact, one participant (quoted in Section V) thought it was responsive to sound, and was making noise and still trying to figure out its (non-existent) response to sound until the end. Indeed, the set up of this interactive art sculpture was also very different from a typical set-up. Typically, this kind of sculpture is set up in a public space such as a department store, office building, or museum. Visitors are free to enter or leave the space as they like. In cases of a permanent installation, the sculpture is placed in the background to some other daily activities. This is very different from the one-on-one, timed interactions that we tested in this user study. In addition, typically, visitors would be accompanied by other people, both acquaintances and strangers, and the effect of multiple occupants in the same space was not considered.

The initial pilot study also yielded useful insights about the experimental design, which should be improved in future studies to elicit further meaningful insights about user experiences. Many interesting observations were obtained from informal participant feedback following the formal questionnaire; in future studies, participants should be given the more open ended questions, or allowed to talk aloud through their experience. Methods for collecting open ended feedback as the participants are interacting with the sculpture should be devised, so that impressions of different behaviours and sculpture actions could be better temporally and spatially localized. These mechanisms are particularly important for longer duration experiments, which are needed to evaluate how user perceptions may evolve during longer interaction sessions, or over multiple sessions. To better understand how participant expectations influence their experience, a more detailed pre-study questionnaire should also be used, focusing particularly on participants' interest, knowledge and expectations of interactive artworks and art as a whole.

VI. Conclusions and Future Work

This paper reports on a user study investigating user responses to a interactive system embodied in a kinetic sculpture. The system was tested under two types of behaviours, autonomously generating behaviours based on an internal curiosity-based motivation (CBLA), and generating prescripted behaviours. Users were able to discern the differences between the two sets of behaviour generation modes, and there was a correlation between system activation and user interest level. While both increasing and decreasing levels of interest were reported when the CBLA mode was shown following prescripted mode, users' interest levels tended to increase when CBLA mode preceded prescripted mode. Different participants in the study interacted with the sculpture in very different ways and they had very different expectations about their experience. This suggests that more meaningful correlations may be revealed by categorizing the different types of users. Alternatively, to align participants' expectations, the expected behaviours and meanings behind the concepts of the sculpture can be explained to the participants prior to the study. This is similar to how informed visitors may read and learn about the sculptures prior to interacting with them in a museum or art gallery setting.

Fig. 8: An audience member interacts with interactive installation Sentient Chamber (National Academy of Sciences, Washington D.C., 2015).

In addition, though the test bed resembled a typical interactive art sculpture, there were also significant differences. For instance, the small size of the sculpture means that the proximity of the activation became less of a factor since all nodes were relatively close to the visitor. We are currently implementing the CBLA functionality in a larger installation, illustrated in Figure 8. In addition, in a public exhibition, visitors' interaction with the sculpture might be secondary to their primary activities such as socializing with friends, or simply passing by to get from one location to another. Users may find the experience more interesting as an augmentation to their primary activities in comparison to the more focused, one-on-one experience tested in this study. Their reactions to the testbed, and to interactive art more generally, may also be strongly coupled with their own general subjective perceptions of art. Therefore, in future work, we are planning to conduct user studies in more naturalistic settings to further examine the users' experience in interacting with the sculpture, and include questions to explore the users' predisposition to interactive art.

ACKNOWLEDGMENT

The authors wish to thank all the study participants for volunteering their time for the study and providing valuable feedback. We also thank M. Memarian and M. Borland for their assistance with the electrical design of the prototype, M. Josef and S. Saines for their help on the production and testing of the circuit boards and electronics assemblies; and A. Schwartzentruber, A. Ling, D. Watson, J. Prosser, L. Li, M. Correa, M. Hu, P. Bogias, S. Taylor, and W. Fu for designing and constructing the sculptural components and the experimental test bed. Figures 1, 2 and 8 are courtesy of Philip Beesley Architect.

REFERENCES

[1] P. Beesley and C. Macy, *Kinetic Architectures & Geotextile Installations*. Toronto, Ontario, Canada: Riverside Architectural Press, 2010.
[2] P. Beesley, R. Gorbet, P. Ohrstedt, and H. Isaacs, *Hylozoic Ground: Liminal Responsive Architecture*. Toronto, Ontario, Canada: Riverside Architectural Press, 2010.
[3] P. Beesley, *Sibyl: New Responsive Immersive Architecture Projects 2010-12*, 1st ed. Kitchener, Ontario: Riverside Architectural Press, 2012.
[4] J. Drummond, "Understanding Interactive Systems," *Organised Sound*, vol. 14, no. 02, p. 124, 2009.
[5] B. Wands, "Variations : An Interactive Musical Sculpture," in *C&C '05 Proceedings of the 5th conference on Creativity & cognition*. ACM, 2005, pp. 306–309.
[6] B. Costello, L. Muller, S. Amitani, and E. Edmonds, "Understanding the experience of interactive art: Iamascope in Beta_space," *IE '05 Proceedings of the second Australasian conference on Interactive entertainment*, pp. 49–56, 2005.
[7] M. Fox and M. Kemp, *Interactive Architecture*. Princeton Architectural Press, 2009.
[8] E. Edmonds, G. Turner, and L. Candy, "Approaches to interactive art systems," in *Proceedings of the 2nd international conference on Computer graphics and interactive techniques in Austalasia and Southe East Asia GRAPHITE 04*, 2004, pp. 113–117.
[9] P. Beesley, *Near-Living Architecture: Work in Progress from the Hylozoic Ground Collaboration 2011-2014*. Toronto, Ontario, Canada: Riverside Architectural Press, 2014.
[10] M. A. Bedau, J. S. McCaskill, N. H. Packard, S. Rasmussen, C. Adami, D. G. Green, T. Ikegami, K. Kaneko, and T. S. Ray, "Open problems in artificial life," *Artificial life*, vol. 6, no. 4, pp. 363–376, 2000.
[11] T. Ikegami, "A design for living technology: Experiments with the mind time machine," *Artificial Life*, vol. 19, no. 3/4, pp. 387–400, 2013.
[12] M. Lungarella, G. Metta, R. Pfeifer, and G. Sandini, "Developmental robotics: a survey," *Connection Science*, vol. 15, no. 4, pp. 151–190, 2003.
[13] M. Asada, K. Hosoda, Y. Kuniyoshi, H. Ishiguro, T. Inui, Y. Yoshikawa, M. Ogino, and C. Yoshida, "Cognitive Developmental Robotics : A Survey," *IEEE Transactions on Autonomous Mental Development*, vol. 1, no. 1, pp. 12 – 34, 2009.
[14] P.-Y. Oudeyer, F. Kaplan, and V. V. Hafner, "Intrinsic Motivation Systems for Autonomous Mental Development," *IEEE Transactions on Evolutionary Computation*, vol. 11, no. 2, pp. 265–286, 2007.
[15] M. T. K. Chan, R. Gorbet, P. Beesley, and D. Kulić, "Curiosity-Based Learning Algorithm for Distributed Interactive Sculptural Systems," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Hamburg Germany, 2015, pp. 3435–3441.
[16] P.-Y. Oudeyer, F. Kaplan, V. V. Hafner, and A. Whyte, "The playground experiment: Task-independent development of a curious robot," in *In Proceedings of the AAAI Spring Symposium on Developmental Robotics*, Stanford, California, 2005, pp. 42–47.
[17] R. Gorbet, M. Memarian, M. Chan, D. Kulic, and P. Beesley, "Evolving Systems within Immersive Architectural Environments: New Research by the Living Architecture Systems Group," *Next Generation Building*, 2015.
[18] W. W. Gaver, J. Beaver, and S. Benford, "Ambiguity as a resource for design," in *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 2003, pp. 233–240.
[19] P. J. Silvia, "Interest - The curious emotion," *Current Directions in Psychological Science*, vol. 17, no. 1, pp. 57–60, 2008.
[20] C. Breazeal, D. Buchsbaum, J. Gray, D. Gatenby, and B. Blumberg, "Learning from and about others: towards using imitation to bootstrap the social understanding of others by robots," *Articifical Life*, vol. 11, no. 1–2, pp. 1–32, 2005.
[21] P. J. Silvia, "What Is Interesting? Exploring the Appraisal Structure of Interest," *Emotion*, vol. 5, no. 1, pp. 89–102, 2005.