Learning to Engage with Interactive Systems: A Field Study on Deep Reinforcement Learning in a Public Museum

LINGHENG MENG, DAIWEI LIN, ADAM FRANCEY, ROB GORBET, and
PHILIP BEESLEY, University of Waterloo
DANA KULIĆ, University of Waterloo and Monash University

Physical agents that can autonomously generate engaging, life-like behavior will lead to more responsive and user-friendly robots and other autonomous systems. Although many advances have been made for one-to-one interactions in well-controlled settings, physical agents should be capable of interacting with humans in natural settings, including group interaction. To generate engaging behaviors, the autonomous system must first be able to estimate its human partners' engagement level. In this article, we propose an approach for estimating engagement during group interaction by simultaneously taking into account active and passive interaction, and use the measure as the reward signal within a reinforcement learning framework to learn engaging interactive behaviors. The proposed approach is implemented in an interactive sculptural system in a museum setting. We compare the learning system to a baseline using pre-scripted interactive behaviors. Analysis based on sensory data and survey data shows that adaptable behaviors within an expert-designed action space can achieve higher engagement and likeability.

CCS Concepts: • **Human-centered computing** → **Human computer interaction (HCI)**; • **Computing methodologies** → *Artificial intelligence; Machine learning*; • **Applied computing** → Arts and humanities;

Additional Key Words and Phrases: Living architecture, human-robot interaction, reinforcement learning, engagement, voluntary engagement, social robot, interactive system, group interaction, adaptive system, natural setting interaction, open-world interaction, robotic arts, robotic sculpture

ACM Reference format:

Lingheng Meng, Daiwei Lin, Adam Francey, Rob Gorbet, Philip Beesley, and Dana Kulić. 2020. Learning to Engage with Interactive Systems: A Field Study on Deep Reinforcement Learning in a Public Museum. *ACM Trans. Hum.-Robot Interact.* 10, 1, Article 5 (October 2020), 29 pages.

<https://doi.org/10.1145/3408876>

This work was supported by the Living Architecture Systems Group, funded by the Social Sciences and Humanities Research Council of Canada.

Authors' addresses: L. Meng and D. Lin, Department of Electrical and Computer Engineering, University of Waterloo, 200 University Avenue West, Waterloo, Ontario, N2L 3G1, Canada; emails: {lingheng.meng, daiwei.lin}@uwaterloo.ca; A. Francey, Department of Psychology, University of Waterloo, 200 University Avenue West, Waterloo, Ontario, N2L 3G1, Canada; email: alzfrancey@uwaterloo.ca; R. Gorbet, Department of Knowledge Integration, University of Waterloo, 200 University Avenue West, Waterloo, Ontario, N2L 3G1, Canada; email: rbgorbet@uwaterloo.ca; P. Beesley, School of Architecture, University of Waterloo, 7 Melville Street South Cambridge, Ontario, N1S 2H4, Canada; email: pbeesley@uwaterloo.ca; D. Kulić, Department of Electrical and Computer Systems Engineering, University of Waterloo and Department of Mechanical and Aerospace Engineering, Monash University, 14 Alliance Lane, Clayton campus, VIC 3800, Australia; email: dana.kulic@monash.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Association for Computing Machinery.

2573-9522/2020/10-ART5 \$15.00

<https://doi.org/10.1145/3408876>

1 INTRODUCTION

As robots enter human environments, engaging with human in a suitable manner becomes increasingly important [1, 11, 15, 34, 40, 61, 67, 68]. To facilitate long-term interaction, interactive systems should be able to continuously adapt to generate engaging and appropriate behavior. Although for simple interactive systems it may be possible to manually design engaging behaviors, this approach is time consuming and sometimes unfeasible for complex non-anthropomorphic interactive systems with many degrees of freedom (DOF; e.g., hundreds of DOF). In addition, manually designed behaviors constrain the system to a limited set of reactions, whereas autonomous generation of actions provides the possibility of adaption and continuous behavioral evolution in a non-stationary environment. Therefore, understanding how to autonomously generate engaging behavior is necessary for long-term human-robot interaction (HRI), and learning algorithms can be advantageous.

When introducing learning into an interactive system, human input could be explicit, where human teachers aim to teach the interactive system how to interact, or implicit, where human teachers do not know the interactive system is learning from their interactions. An example of explicit input is a human taking a teacher role, giving feedback to guide the robot to achieve some goals. With implicit input, a learning signal is generated based on observations of human behaviors during interaction, without requiring active feedback. For successful long-term interaction in social and public settings, the implicit model is more appropriate than the explicit teacher model, because in the implicit model learning is not dependent on human expertise or willingness to train the system. Finally, for successful interaction in social and public settings, the system should successfully engage with and learn from both individual and group interactions.

We study the challenge of long-term autonomy with living architecture systems (LAS). LAS are interactive systems at architectural scale that emulate living environments aiming to engage occupants in long-term interaction (Figure 1 presents sample installations by the Living Architecture Systems Group (LASG) and Philip Beesley Architect Inc. (PBAI)¹). An immersive LAS can have dozens of sensors and hundreds of actuators with various modalities to enable interaction with visitors, which poses a challenge when designing effective interaction control strategies. Historically, LASG/PBAI environments have used pre-scripted, human-designed behaviors. In this work, we are interested in the potential for adaptive machine learning behavioral algorithms to generate interactive behaviors. LAS aims for long-term continuous interaction, and this necessitates the capability of LAS to evolve with and to learn in a non-stationary environment rather than using a single pre-scripted solution or assuming the existence of a single optimal solution.

Although LAS aims for long-term engagement with occupants, the architectural scale and non-anthropomorphic, immersive nature of LAS makes it distinctive compared to robots more typically studied in HRI [26, 30, 37, 64, 76, 82]. Since LAS is intended to create life-like behaviors at architectural scale, it facilitates accommodating multiple occupants and group interaction rather than the one-to-one interaction most commonly studied in HRI. In addition, unlike HRI studies with humanoid robots that can be directly inspired by human-human interactions, LAS is a non-anthropomorphic robot, making the design of interactive behavior less intuitive. Although it is possible to manually design life-like interactive behaviors to some extent, it is very time consuming, and the resulting behaviors are non-adaptive. Therefore, implementing interactive systems that can adapt and learn in dynamic crowd settings and can autonomously generate engaging behavior is critically important, not only because it is useful in many applications, including public spaces, schools, workplaces, and family residences, among others, but also because HRI in such interactive systems could extend our understanding of socially acceptable and engaging interactive behavior.

¹More LAS environments by LASG/PBAI can be found at <http://www.philipbeesleyarchitect.com/sculptures/>.

(a) *Radiant Soil*, installed in Daejeon Museum of Art in Daejeon, South Korea, 2018

(b) *Epiphyte Spring*, installed in Design Institute of Landscape and Architecture in Hangzhou, China, 2015

(c) *Hylozoic Series: Sibyl*, installed in the Industrial Precinct in Sydney, Australia, 2012

Fig. 1. Living architecture systems. (Photos courtesy of Philip Beesley Architect Inc.)

Manually designed behavior patterns can generate engaging life-like behaviors based on the designer's understanding [6, 7], but this can be time consuming and result in behaviors that become predictable during long-term interaction. To overcome this limitation, machine learning may be used to automatically generate behavior and study whether such behavior is attractive or engaging for participants. For example, in the work of Chan et al. [16, 17], the Curiosity-Based Learning algorithm (CBLA) was implemented in the LAS to automatically generate behavior based on a computational notion of curiosity [58] of the LAS. However, the action produced by the CBLA is purely intrinsically motivated by the curiosity mechanism and does not consider any extrinsic motivation, such as maximizing measures of human response, which could play a more important role when an interactive system aims to engage participants. Another issue with the work of Chan et al. [16, 17] and many works on social robots [12, 37] is that the proposed approach is only designed and tested for one-to-one interaction in a controlled setting rather than group interactions in natural settings, such as multiple people interacting with an interactive system simultaneously without instruction or guidance by the researchers.

Learning interactive behavior for an extrinsically motivated LAS in natural settings with group interactions is a difficult challenge that is not typically considered in reinforcement learning (RL) tasks in controlled settings with only one-to-one interaction. First, the LAS is an interactive system at architectural scale, with very large state and action spaces, typically having dozens of sensors and hundreds of actuators. The learning challenge is exacerbated by a complicated and non-stationary environment, where the LAS might be manipulated by occupants with different cultural backgrounds, interests, and personalities. In addition, interactions happening in the real world happen less frequently than interactions in the simulators or video games typically studied in RL. This results in fewer interactions, which poses a huge challenge for data-driven learning algorithms. Last but not least, there is no standard measurement for estimating engagement (i.e., the extrinsic motivation of LAS). If a measurement of engagement is sparse and time delayed, learning interactive behavior becomes more difficult.

In this article, we investigate extrinsically motivated learning in natural settings with group interaction. Specifically, we use sensors embedded in the interactive system to estimate the overall occupancy and engagement of the visitors in the LAS, and use this estimate as the extrinsic reward for learning. We investigate whether we can exploit the designer's pre-scripted behaviors (PBs) to bootstrap learning. We use the Deep Deterministic Policy Gradient (DDPG) algorithm to train an Actor-Critic agent, which acts in the designer-parameterized action space, which we call the *parameterized learning agent* (PLA). Both fixed pre-scripted and learning behaviors are evaluated

over 3 weeks at a public exhibition at the Royal Ontario Museum (ROM), and all data is collected from museum visitors. We evaluate how engaging the behaviors are based on a quantitative analysis of estimated engagement and the number of active interactions, and on qualitative analysis of human survey results.

2 RELATED WORK

Adaptive control. Traditional robotic systems operating in environments where the system dynamics are unknown or varying have used adaptive control [3]. In the domain of HRI, Wang et al. [80] applied adaptive control to robot navigation, where a social proxemics potential field is constructed and used to design a robot motion controller that is able to adapt its desired trajectory smoothly and at the same time comply with the proxemics constraints. Nakazawa et al. [56] proposed a potential field imposing a repulsive fin to allow adaptive control of an accompanying robot, where the robot is able to adapt its relative position to the accompanied human in the presence of obstacles. Vitiello et al. [78] proposed a neuro-fuzzy-Bayesian approach for adaptive control of a robot's proxemics behavior, where recognized human activities and human personality acquired by questionnaires are input into an adaptive neuro-fuzzy inference engine system (ANFIS) to determine a robot's stopping distance. Adaptive control has also been applied to rehabilitation robots [81], robots driven by Series Elastic Actuators [48], and many other applications. However, adaptive control relies on a known structure of the dynamics model. In our case, it is difficult to apply adaptive control because we do not have access to a good model of the environment dynamics, as the number of visitors is unpredictable and visitors' interaction style varies.

Learning with humans as the explicit teacher. Interactive reinforcement learning (interactive RL) studies RL algorithms in the context of HRI, where humans explicitly provide rewards or actions to guide an agent in RL. Isbell et al. [33] applied RL in a complex human online social environment, where a human interacts with a learning agent by providing a reward signal, and highlighted that many of the standard assumptions, such as stationary rewards, Markovian behavior, and appropriateness of average reward, for RL are clearly violated. Thomaz and Breazeal [74] and Thomaz et al. [75] proposed interactive RL for training human-centric assistive robots through natural interaction, where a human coach's feedback is used to shape the pre-defined reward. In addition, those authors introduced anticipatory cues to allow the human coach to predict the robot's action and provide timely feedback. In the work of Suay and Chernova [72], interactive RL was first studied in real-world robotic systems and showed that the positive effects of human guidance increase with state space size. Other works converting human feedback to reward include those of Knox and Stone [41, 42]. Griffith et al. [28] proposed an algorithm (Advise) to formalize the meaning of human feedback as policy feedback, which is more effective and robust than using human feedback as extra reward. Krening and Feigh [43] studied the effect of different teaching methods with the same underlying RL algorithm on human teachers' experience in terms of frustration. This work emphasizes that high transparency will decrease frustration and increase perceived intelligence. Different from most of these works where the human plays the direct role of teacher (i.e., directly providing either a reward or a policy advice in an interactive way), we aim to learn how to engage visitors without requiring them to consciously teach or train the interactive system. Therefore, in this work, the learning agent does not explicitly receive rewards from visitors but uses a measure of engagement as reward, which is calculated based on observed visitor behaviors, without any constraints on the frequency and consistency of interaction on the visitor.

Learning with humans as the implicit teacher. RL algorithms have been deployed in social robotics to enable a robot to acquire socially acceptable skills, where a pre-defined reward function is employed to implicitly infer the reward signal from sensed human behavior. In the work of

Kumagai et al. [45], emotion recognition based on videos captured during participants' interaction with a robot is exploited as a reward within an RL framework to adapt the robot's behavior toward participants' personal preferences. Leite et al. [46] proposed to model the user's affective state by combining facial expression recognition and game state in a supervised way. The predicted affective state is then used to calculate the reward in an RL framework to personalize response to users. Macharet and Florencio [51] investigated the effectiveness of using RL to learn socially acceptable approaching behavior for a mobile robot. Gordon et al. [27] studied affective personalization in an integrated intelligent tutoring system, where in a one-to-one interaction setting, facial expression based engagement and valence estimation are used as reward in a standard SARSA algorithm. These studies all work on relatively small state and action spaces without exploiting deep neural networks. Recently, deep reinforcement learning (DRL) has also been investigated in social robotics. Qureshi et al. [60] proposed the multimodal deep Q-network (MDQN) for a handshake robot, where a deep neural network was employed to approximate the Q-value function based on both image and depth information. Chen et al. [18] studied social aware collision avoidance deep reinforcement learning (SA-CADRL), where the reward function was designed to penalize actions leading to states where a robot is too close to nearby pedestrians and a deep neural network is used to approximate the Q-value function. Different from these works where either DRL is not exploited or the action space is discrete, our work applies continuous control DRL to an interactive social robot.

One-to-one HRI in natural settings without learning. Interactions in natural settings (e.g., public spaces) are too complex to be simulated in controlled laboratory settings, so to understand natural HRI it is important to conduct field studies. Many HRI systems have been tested in public spaces such as hospitals [55], train stations [31], service points [36, 38], airports [35], shopping malls [65], hotels [19], and museums [66], among others. Although robots investigated in these works are sometimes surrounded by a group of people in a public place, these robots typically interact with only one person at a specific time (i.e., one-to-one interaction). Unlike these works, the larger scale of LAS enables LAS to simultaneously accommodate interactions with multiple people. In addition, the behaviors of robots studied in most of these works are pre-designed by researchers without continuous learning and adaptability.

Interactive artworks without learning. Interactive artworks are outcomes of combining arts and engineering, and have brought a new research direction for understanding HRI, especially with non-anthropomorphic robots, not only from the robotics perspective but also from an artistic perspective. LAS in this work and previous installations [6] are examples of such immersive interactive systems that promote roboticists' and architects' understanding of life-like interactive behavior. Other examples include Levillain et al. [47], who investigated visitors' interaction and observation behaviors with mobile pianos in a museum. St. Onge et al. [70] and St. Onge and Reeves [71] studied flying cubes in multiple publicly accessible spaces where flying cubes play the role of living creatures. In the work of Cuan et al. [20], roboticists collaborated with a professional dancer to design interactive dancing robots. Wakefield et al. [79] created spatial interactions with immersive virtual reality technology. Augugliaro et al. [4] studied dancing quadcopters, without an interactive component. Close collaboration between artists and roboticists has been fostering the creation of richer modes of interaction and extending the scope of studies in HRI. Most works mentioned here rely heavily on the design of choreographers, whereas in this work we add adaptability on top of choreography by learning engaging behavior based on the action space defined by choreographers.

Engagement estimation in HRI. As the intention of HRI is to engage human participants in interactions, detecting and measuring initial and ongoing engagement in HRI is critical for

initiating and maintaining interaction [23, 25, 57]. Many approaches for measuring engagement rely on rich sensors, such as cameras, lasers, sonar, and audio sensors, and on sophisticated facial expression, gesture, gaze, body movements, physiological signals, and speech recognition technologies [1, 44, 73]. Sanghvi et al. [62] proposed to estimate engagement by extracting features of body motion from video by training an engagement prediction model in a supervised way, but the proposed method requires a lateral camera and users to wear a specific-colored coat. Keizer et al. [38] trained an engagement classifier based on a multimodal corpus in a supervised way, which classifies engagement into three levels: *NotSeekingEngagement*, *SeekingEngagement*, and *Engaged*. Michalowski et al. [54] studied a spatial model of engagement for a social robot and tested a robotic receptionist to engage visitors, and suggested that direction and speed of motion are more appropriate measures of engagement than location or distance alone. However, this work used distance to determine the engagement level and what PB to take rather than using the distance as a learning signal. Bohus and Horvitz [9, 10] studied the engagement model for a multi-party open-world dialog system, but the system can only actively engage in at most one interaction even though it could simultaneously keep track of additional, suspended interactions. Engagement in the work of Bohus and Horvitz [9, 10] is viewed as several states. Behaviors corresponding to each engagement state are pre-defined. Even though the authors claim the system is a multi-party interaction, the number of participants is limited because of the field of view of the camera and the interaction screen. Khamassi et al. [39] and Velentzas et al. [77] proposed to estimate engagement from human gaze and use it as a reward function in an RL framework to allow fast adaption to environment change, where experiments are conducted in a well-controlled setting. Most works studying engagement in HRI focus on humanoid social robots, one-to-one interactions, and discrete measures of engagement. In this article, we study a non-anthropomorphic robotic environment, group interaction, and measure engagement as a continuous value. Most importantly, we use a measure of engagement as the extrinsic reward to drive learning rather than using the measure of engagement as a contingent event to trigger a pre-defined behavior. In addition, since our environment has highly varying lighting conditions and the potential for severe occlusions, camera-based methods may be less reliable.

Interactive systems, whether humanoid robots or non-anthropomorphic robots like LAS, are human oriented and aim for continuous engagement. These characteristics intrinsically require agents to be able to autonomously generate behaviors and continuously adapt to their human partner(s). In this work, we exploit the standard RL framework and a novel reward function estimating engagement to realize autonomous agents, and study if this reward could generate behavior that is more engaging than PB. All experiments are conducted in a public museum exhibit with museum visitors.

3 THE INTERACTIVE ARCHITECTURE TESTBED

In this section, we describe the physical system used as the testbed in this article, as well as the design of PBs that drive the interactive behavior of the system. The PBs, which are designed by expert architects and interactive system designers, are the baseline we use to compare to the learning system described in Section 4.

Our testbed *Aegis Canopy* is part of the exhibition Transforming Space that consists of two sculptures—the Aegis Canopy and Noosphere—as shown in the top left of Figure 2. The exhibition was exhibited at the ROM in Toronto, Canada, from June 2 to October 8, 2018 (<https://www.rom.on.ca/en/philip>), and was publicly accessible to more than 60,000 visitors who had toured the exhibit.

Since this research mainly used the Aegis Canopy part of the installation, we will describe the design of the Aegis Canopy in detail and subsequently refer to the Aegis Canopy part of the installation as the LAS.

Fig. 2. Installation diagram and interaction types.

Fig. 3. Diagram of node in the LAS.

3.1 Physical LAS

The LAS hangs overhead within the Aegis Canopy space, with an approximate height of 1.8 m. Figure 2 shows the front view of the LAS. The active part of the system is composed of eight speakers and 24 nodes.

Each node in the LAS consists of six Fronds, one Moth, and one high-power LED as its actuated systems and one infrared (IR) sensor, as shown in Figure 3. Each Frond includes a shape memory alloy (SMA) wire that contracts when voltage is applied, pulling a cord attached to a flexible co-polyester sheet, as illustrated in Figure 3(a). The contraction generates a smooth and gentle movement, and when the applied voltage is removed, the SMA slowly relaxes to its original shape.² The Moth consists of wing-like flexible flaps attached to a small DC vibration motor that vibrates when activated, making the Moth appear to be flapping its wings. The Moth also houses two small LED lights that illuminate during vibration. The single high-power LED located beneath the central flask can be faded to illuminate the colored liquid in the flask. The IR sensor senses the proximity of visitors, and it generates a continuous reading proportional to the distance between any part of the body of a visitor and the sensor location in the Aegis Canopy.

There are eight speakers distributed throughout the LAS. These speakers play two types of sound samples. The first sound is a background sound played on a continuous loop. The second sound is triggered by the IR sensors. These speakers are independently controlled by specialized software, so here we treat them as background behaviors.

The arrangement of the speakers and nodes is illustrated in Figure 4, where the 24 nodes are highlighted by red circles. A photo of the physical LAS is shown in the bottom left of Figure 4. The 24 nodes are at varying height levels. Specifically, nodes at the left and right edges are slightly higher than those in the middle of the LAS. This spatial arrangement distinguishes three types of visitor engagement with the system. When visitors observe the LAS but are not underneath the LAS, no IR sensor is activated (i.e., visitors are observing the LAS but cannot be observed by the LAS sensors). As shown in Figure 2, when visitors walk or stand underneath the LAS, which we call *passive interaction*, the IR sensors above them are activated, but the distance between the visitor and the system is still large, corresponding to a small reading of the IR sensor. Visitors engaging in *active interaction* might also reach their hand upward to interact with the LAS, resulting in a higher activation value of the closest IR sensor.

Two web cameras (labeled Camera1 and Camera2 in Figure 2 and Figure 4) are used to record video during our experiment and to calibrate sensory data. These two web cameras are mounted on the wall in the front-right and back-left corners of the LAS space.

²A video illustrating the SMA assembly contracting and relaxing can be found at <https://youtu.be/YcreirXrRF4>.

Fig. 4. LAS: Aegis Canopy.

Fig. 5. Pre-scripted behavior.

Table 1. PB Parameters

Parameters	Meaning	Default	Range
T_{ru}^m, T_{ru}^l	Ramp up time: the time it takes for the Moths or LEDs to increase up to their maximum value	1.5	[0, 5]
T_{ho}^m, T_{ho}^l	Hold time: the time that Moths and LEDs are held at their maximum value	1	[0, 5]
T_{rd}^m, T_{rd}^l	Ramp down time: the time it takes for Moths and LEDs to fade down to 0	2.5	[0, 5]
I_{max}	Maximum percentage of duty cycle per PWM period	78	[0, 100]
T_{gap}^m	Time gap between the Moth starting to ramp up and the LED starting to ramp up	1.5	[0, 5]
T_{gap}^{sma}	Time gap between activation of each SMA arm on the nodes	0.3	[0, 5]
T_{gap}^n	Time gap between activation of each node during neighbor behavior	1.8	[0, 5]
T_{bg}^{min}	Minimum time to wait before activating background behavior	45	[15, 60]
T_{bg}^{max}	Maximum time to wait before activating background behavior	90	[60, 100]
T_w	Time to wait before trying to pick a Moth or LED	5	[0, 10]
P	Probability of successfully choosing an actuator during background behavior	0.4	[0, 1]
T_{sma}	Time between choosing SMAs to actuate	0.7	[1, 5]
T_{sw}^{min}	Minimum time to wait before performing sweep	120	[5, 200]
T_{sw}^{max}	Maximum time to wait before performing sweep	240	[200, 400]

The unit of all time parameters is seconds, except I_{max} is percentage and P is probability.

3.2 Pre-Scripted Behavior

PB is the interactive behavior manually designed by the architects, and it is also the baseline used for comparison with adaptive behaviors we will describe in Section 4. Within the PB mode, the system can be in two types of states: active and background, in which behaviors are mainly controlled by 17 parameters (shown in Table 1) specified by the architects. The values in the Default and Range columns are used in the PB and PLA modes, respectively.

The active state is entered if any of the IR sensors is triggered. In this state, the node corresponding to the triggered sensor will first activate its *local reflex behavior*. In the local reflex behavior,

the Moth, the LED, and six SMAs attached to the same node as the triggered IR sensor will be activated. When a Moth is activated, it will ramp up the vibration to its maximum I_{max} over time T_{ru}^m , hold there for a period of time T_{ho}^m , and then ramp down over (T_{rd}^m) . After a waiting period (T_{gap}^m) following the sensor trigger, the LED on the same node is activated. It ramps up over time period (T_{ru}^l) to its maximum brightness (I_{max}) , holds for a period of time (T_{ho}^l) , and then gradually dims over (T_{rd}^l) . At the same time, the SMAs are activated one after another separated by (T_{gap}^{sma}) . A step voltage is applied to contract the SMA, after which a cooling-down time is started during which this SMA will not be activated again. The activation profile of the SMA wires is fixed to protect them from overheating, so these are not included in the parameterization shown in Table 1. After the local reflex behavior is triggered, the IR-detected event will be propagated from the triggered node to neighboring nodes after a delay (T_{gap}^n) , until the edge nodes of the LAS are reached (shown in Figure 5), causing a cascade of local reflex behaviors at each node.

If no IR sensor triggering happens for a random time within $[T_{bg}^{min}, T_{bg}^{max}]$, the system goes into the background state. In this state, the LEDs and Moths will randomly activate their local reflex behaviors with probability P every amount of time T_w . The SMAs are also activated independently with the same probability P every T_{sma} .

In either state, a sweep of LEDs in either direction along the longer axis of the installation happens at random time intervals within $[T_{sw}^{min}, T_{sw}^{max}]$. During the sweep, each LED activates local reflex behavior and propagates in the direction of the sweep.

4 PROPOSED APPROACH

In this section, we describe how the PLA is designed to automatically generate interactive actions.

Figure 6 illustrates the software architecture connecting the physical hardware to the learning agents. The Middle Layer coordinates the sensor reading and action execution among the physical LAS and agents. Only one agent type controls the LAS at any given time. PB only takes and stores IR readings without issuing actions because its action policy is fixed, whereas PLA is formulated as a standard RL framework where the Observation Constructor constructs the observation, the Reward Critic generates a reward that estimates the engagement of visitors, and Actor-Critic is the learned policy and Q-value function.

4.1 PLA: Learning on Top of PB

PLA is designed to learn on top of PB (i.e., parameterized action space). The motivation for this approach is to bootstrap learning by exploiting the designer's knowledge of engaging behavior, where we hypothesize the designer already has a good idea about what types of actions might be engaging to visitors, and this can form a helpful starting point for the learner.

4.1.1 Observation and Action Space Construction. For PLA, we select 11 parameters from Table 1 as the *action space* (i.e., the dimension of the action vector is 11). This is because some parameters do not take effect until a subsequent trigger or until the current propagation finishes. This could lead to obtaining an observation that is based on both previous and updated parameters. To avoid this issue, we exclude T_{bg}^{min} , T_{bg}^{max} , T_w , P , T_{sw}^{min} , and T_{sw}^{max} from the action space. In this way, we make sure every observation is only related to the latest action. To attenuate IR sensor noise, the *observation space* for PLA is an average over 20 IR readings as defined in Equation (1):

$$obs^{(t)} = \frac{1}{20} \cdot \sum_{i=0}^{19} ir^{(t-i\Delta t)}, \quad (1)$$

where $\mathbf{ir}^{(t-i \cdot \Delta t)}$ is the value array of 24 IR sensor readings at time $(t - i \cdot \Delta t)$, $\sum$ is element-wise summation, and $\Delta t \approx 0.1$ s is the time to retrieve one set of 24 IR values from the physical LAS. Thus, the dimension of the observation vector is 24 and each observation vector represents the average IR readings over 2 seconds, which means that an agent generates an action every 2 seconds. Note that this delay between an agent's two consecutive actions is in the parameterized action space (i.e., an update to the PB parameters), which means that this will not cause a delay in the LAS's response to human visitors. The 2-second time delay between every two consecutive parameterized actions generated by PLA is chosen to allow an action to complete before initiating another action. As described in Section 5.3, IR values are converted and scaled so that 0 corresponds to the maximum distance reading (i.e., there is nothing in front of the sensor) and 1 corresponds to the minimum distance reading (i.e., there is something very close to the sensor).

4.1.2 Estimating and Using Engagement as a Reward for Learning. A key feature of our approach is the formulation of the reward function: we wish to learn and reward behaviors that foster visitor engagement. Specifically, the extrinsic reward is computed by summing over the IR observations, which can be regarded as a rough estimate of occupancy and engagement, because (1) more activated IRs means more people are standing under the LAS, thus indicating higher occupancy, and (2) closer distance between visitors and IR sensors implies more active interaction, such as looking very closely or raising hands, which are higher engagement behaviors. Therefore, higher occupancy and more active interaction will cause higher extrinsic reward. Formally, given a new observation at time $t + 1$, $\mathbf{obs}^{(t+1)} = (\mathbf{obs}_1^{(t+1)}, \mathbf{obs}_2^{(t+1)}, \dots, \mathbf{obs}_n^{(t+1)})$ ³, the reward $r^{(t)}$ for taking action $\mathbf{a}^{(t)}$ while observing $\mathbf{obs}^{(t)}$ can be expressed as Equation (2):

$$r^{(t)} = \sum_{i=1}^n \mathbf{obs}_i^{(t+1)}, \quad (2)$$

where the value of n is the number of dimensions of the observation vector (e.g. n is 24 for PLA).

Low-cost embedded IR sensors are employed to estimate engagement for three reasons. First, the varying lighting conditions in the exhibition area do not allow us to use the webcam footage for reliable pose or facial expression analysis. Second, because of occlusions and the uncontrolled nature of the space and the varying number of occupants, an array of inexpensive IR sensors that can be precisely located within the space is more effective than lower numbers of more sophisticated equipment. Third, due to privacy concerns, we could not guarantee that we would be permitted to use camera data.

4.1.3 Implementation. Given the observation and extrinsic reward, the optimal policy can be learned by an RL algorithm.⁴ In this work, learning is implemented using the DDPG algorithm [49], a variant of Deterministic Policy Gradient [69], where both the Actor and Critic are approximated with deep neural networks.

PLA uses the DDPG agent from OpenAI's Baselines package [22], following the hyper-parameters specified in Table 2 (the detailed pseudo-code can be found in Algorithm 1 in Appendix A). The structure of the neural network is shown in Figure 7. All layers are dense layers, with layer-norm applied, where the activation function and neurons for each layer are indicated above and under each layer in Figure 7, respectively. At the start of the field study, the Actor-Critic networks of PLA were randomly initialized. On all subsequent deployment days, the learned actor-critic from the previous day was loaded and training continued from the previous day's network

³In this article, we use normal lowercase for scalar and bold lowercase for vector.

⁴However, note that the interaction frequency is limited by the physical constraints of the LAS.

Table 2. Hyper-Parameters of DDPG-Based PLA

Hyper-Parameters	Value
Actor learning rate	10^{-4}
Critic learning rate	10^{-3}
Discount rate γ	0.99
Batch size N	64
Replay buffer size $ \mathcal{D} $	10^6
$train_interval$	10
$train_times$	20
$episode_length$	100
Parameter noise α	1.01
Parameter noise δ	0.1
Initial parameter noise σ	0.1

Fig. 6. Interface between LAS and PLA.

Fig. 7. Actor-Critic of PLA, where layer size and activation function are indicated under and above each layer, respectively.

parameters. There is a replay buffer $\mathcal{D}$ of size 10^6 that stores all experiences collected thus far on each day. An experience is a tuple of $(obs^{(t)}, a^{(t)}, r^{(t)}, obs^{(t+1)}, done^{(t+1)})$, where $done^{(t+1)}$ is always 0. Every 10 steps, we will sample 20 mini-batches from the replay buffer. The system does not have a terminal state, so we manually set the maximum length of an episode to 100 steps, which corresponds to about 200 seconds, and after each episode the start observation is the observation encountered at the end of the last episode. The reason for not having a terminal state for interaction with a single visitor is related to the sensor capabilities, the distributed nature of the sensors, and the fact that multiple visitors move through the space simultaneously. Since we cannot track individual visitors, we do not know when one leaves one sensor space and moves into another, or when one leaves the LAS and another arrives.

The maximum episode length is arbitrarily set to 100 steps only for logging statistic summaries (e.g., accumulated reward, estimated Q-values, etc.) regularly. It does not affect the training when the maximum episode length is reached, because there is no reset of the environment.

The Actor-Critic is trained 20 times on a randomly sampled mini-batch every 10 interactions, corresponding to approximately 20 seconds. At every step, the exploratory action is generated by the current actor perturbed with parameter noise (i.e., parameter space noise) [59].

All hyper-parameters used in this work were empirically chosen via experiments on a simplified simulator where the agent directly controls all actuators and maximizes the reward based on IR

Fig. 8. Experiment schedule.

Table 3. Summary of Experiment Schedule and Data

Behavior	Days	Hours	Survey Participants
PB	14	15.5	14
PLA	13	15	15
SARA	4	4	4
ACRA	4	4	3

Video was not available on September 14.

readings, and visitors are simulated to approach LEDs with the highest intensity. Our tests on the simplified simulator showed that DDPG worked well with the chosen hyper-parameters and was able to smoothly converge to the optimal policy in 100 episodes where the maximum length of each episode is 1,000. Because the simulator is an extremely simplified model of the actual environment, we mainly used it to make sure our code is bug free rather than to pre-train a policy.

4.2 Learning in Raw Action Space

During the field study, we also tested two additional learning systems designed to act in raw action space (i.e., directly control actuators) rather than acting in parameterized action space. The *single agent raw act* (SARA) directly controls the 192 raw actuators of the LAS using a single agent, whereas the *agent community raw act* (ACRA) controls the raw actuators in a decentralized way, where a distributed multi-agent learning system replaces the single large learner. Unfortunately, we were not able to allocate enough time for examining SARA and ACRA in the field study, so their performance is not analyzed in Section 6.

5 EXPERIMENTS

5.1 Experimental Procedure

Our experiment was conducted for 2 weeks from September 14 to October 03, 2018, at the ROM. We were permitted by the ROM to collect data from 1 p.m. to 4 p.m. every day on weekdays. In addition, we conducted in-person surveys on September 18, 20, and 27. During the entire experiment period, visitors were free to visit and interact with the installation without any interference from researchers.

For each day of the experiment, the following procedure was followed:

- (1) Randomly schedule the different agent conditions into 1- or 1.5-hour time slots as shown in Figure 8. PB and PLA were scheduled on each day, whereas only one of SARA and ACRA were scheduled per day.
- (2) Automatically run scheduled behavior at each time slot, and save interaction data and learned models and videos at the end of each behavior.

During days where no visitor surveys were collected, researchers were not present in the environment. During the three survey days, researchers were present but did not provide any additional instructions to visitors. Researchers observed which visitors interacted with the LAS, passively or actively, within a specific behavior period. When visitors were finished with their visit, researchers unobtrusively approached randomly selected visitors who had interacted with the system and asked them if they were willing to participate in a survey. If a visitor agreed to do the survey, the visitor was guided to a table located around a corner and was provided with a tablet

with a questionnaire (see Section 5.2). The researchers also recorded which mode the visitor had interacted with. We only recruited visitors who had interacted with only one behavior mode.

The overall experiment schedule is shown in Figure 8, where red, blue, green, yellow, and white areas correspond to PB, PLA, SARA, ACRA, and no schedule, respectively. A summary of the experiment schedule and collected data is shown in Table 3.

5.2 Data Collection

The data collected comes in four types: sensor readings, learning agent logs, human survey data, and video data from the two webcams.

Every raw sensor reading is logged. In addition to the raw sensor data, each agent also logs its own learning algorithm data collected during the course of learning.

For human survey data, 14, 15, 4, and 3 participants completed surveys in the PB, PLA, SARA, and ACRA modes, respectively, as summarized in Table 3. The questionnaire used in our experiment is a standardized measurement tool for HRI: the Godspeed questionnaire [5]. In addition to the 24 Godspeed questions, we asked participants about their interests and background, and their general feedback and comments. The questionnaire consists of four types of questions:

- (1) Participants' interests and background (multiple-select multiple choice)
- (2) Participants' prior knowledge about interactive architecture and machine learning, including "How familiar are you with interactive architecture?" and "How familiar are you with machine learning algorithms?"
- (3) Twenty-four Godspeed questions, namely Godspeed I: Anthropomorphism; Godspeed II: Animacy; Godspeed III: Likeability; Godspeed IV: Perceived Intelligence; and Godspeed V: Perceived Safety [5]
- (4) Participants' general feedback, such as "Any additional comments regarding your experience?" and "Any overall feedback?"

Video data is collected to calibrate sensory readings and validate occupancy estimates, which will be discussed in detail in Section 5.4. Video data is available for all experiments except those on September 14.

5.3 Data Preprocessing

The IR data is scaled so that all sensory observations for the learning agent are within $[0, 1]$ and all actions that learning agent can take are within $[-1, 1]$. The IR readings are scaled into $[0, 1]$ corresponding to the nearest object being at a detected distance of 80 cm or more (no nearby humans detected) to the nearest object being 10 cm (very close human detected). The action values for all raw actuators are scaled to $[-1, 1]$. For SMAs, values in the range $[-1, 0)$ mean off and values in the range $[0, 1]$ mean on, whereas for LED and Moths the continuous values from -1 to 1 are interpreted into 0 to 255 , where 0 indicates off and 255 indicates the brightest light for LEDs, and the highest intensity vibration for the Moths. The 11 parameterized actions are scaled to $[-1, 1]$, where -1 corresponds to minimum and 1 corresponds to maximum, and their corresponding original values are shown in Table 1.

5.4 Data Analysis

The camera view includes regions outside of the LAS itself. To only focus on areas directly related to the LAS, we define three parts of the whole camera view (as shown in Figure 9(a) and Figure 9(b) for Camera1 and Camera2, respectively). In Figure 9, each camera view is divided into Camera View, Whole Interest Area, and Core Interest Area. For both IR data calibration (see Section 5.4.1) and occupancy estimation (see Section 5.4.2), we only consider the Whole Interest Area.

Fig. 9. Interest area used to estimate occupancy.

Any visitors outside this interest area will be ignored for the purposes of occupancy estimation. The Core Interest Area approximates the space directly underneath the LAS.

5.4.1 IR Data Calibration. To enable comparison between different behavior modes, the sensor data must be pre-processed to ensure consistency between conditions. Since visitors can physically interact with the system, it is possible that a visitor changes the direction of the IR sensor, thus changing its field of view and subsequent readings. To calibrate the IR data, two calibration steps are taken: (1) IR sensors, whose value is relatively constant and effectively not responding to occupants (e.g., due to obstructions), are removed, and (2) the baseline reading for each sensor is shifted to zero. Note that the calibration is only done for analysis, during the learning uncalibrated readings are used. To identify blocked IR and baseline shifts, we visually checked the videos recorded by the two webcams and selected a time period when there is no visitor within the Whole Interest Area. Then, we find the IR data corresponding to the no-visitor time. Using the no-visitor time, we determine the thresholds for noise removal and blocked IR detection for the IR data. We use these thresholds to calibrate the raw IR data.

5.4.2 Occupancy Estimation. We also use the camera data to generate a second estimate of occupancy. We estimate the number of people occupying the space during a 1-minute interval, using OpenPose⁵ [14] based on the videos recorded by one of the webcams.⁶ When estimating occupancy, we only considered the Whole Interest Area.

5.4.3 Non-Visitor Period Examination. We also used the camera data to determine whether there are significant periods when no visitors are present. To identify the time periods with no visitors in Whole Interest Area, we manually labeled the time periods when no person is under either camera in the Whole Interest Area. If a person's body is partially visible in the Whole Interest Area, we consider it as a person being in the area. The total amount of non-visitor time throughout the experiment is 1 hour, only 2.5% of total experiment time. Therefore, we use the whole time period for analysis without removing any non-visitor intervals.

5.5 Quantitative Evaluation

Two metrics—estimated engagement level and active interaction count—will be used to quantitatively evaluate the ability of PB and PLA to engage visitors. Although these two metrics both depend on IR readings, they emphasize different aspects of engagement. Specifically, estimated engagement level does not differentiate passive and active interaction (illustrated in Figure 2), whereas active interaction count focuses on measuring active interaction.

⁵<https://github.com/CMU-Perceptual-Computing-Lab/openpose>.

⁶Videos from Camera2 are highly affected by the changing light of the projector as shown in Figure 9(b), so for occupancy estimation we only used videos from Camera1.

We report the average estimated engagement and the average active interaction rather than accumulated reward commonly used in RL because of the non-episodic (i.e., no termination) and non-stationarity (i.e., transition dynamics $T(s'|s, a)$) changes with different visitors nature of the test environment. Specifically, in the natural setting of LAS, the number of visitors varies at different time periods and is highly irregular, which makes the evaluation of learned policy on a separate run unfeasible since the same environment will never be encountered twice. Similarly, comparing accumulated reward within a fixed-length episode is also unfair, because if we compare an episode from PLA during which there are no visitors with an episode from PB during which there are several visitors, we cannot say that PLA is worse than PB and vice versa because no matter what PLA does, there is no reward at all. Therefore, we regard the whole experiment as continuous learning and compare PLA and PB in terms of average estimated engagement and average active interaction.

5.5.1 Estimated Engagement Level. We use raw IR readings recorded during each behavior and Equation (3) to calculate an estimated engagement for comparison among behaviors. Specifically, given M IR readings received within 1 minute (typically sampled at 10 Hz) $\{\mathbf{ir}^{(1)}, \mathbf{ir}^{(2)}, \dots, \mathbf{ir}^{(M)}\}$ where each IR reading $\mathbf{ir}^{(i)}$ is a vector of 24 IR values, the estimated engagement level e is defined by Equation (3):

$$e = \frac{1}{M} \frac{1}{24} \sum_{m=1}^M \sum_{i=1}^{24} \mathbf{ir}_i^{(m)}, \quad (3)$$

where $\mathbf{ir}_i^{(m)}$ is the i th IR sensor in the m th IR reading. The estimated engagement is in the range $[0,1]$, where the maximum 1 corresponds to a maximally engaging state, where all IR sensors are receiving maximum readings during the entire 1-minute window, whereas the minimum 0 corresponds to a fully non-engaging state, where all IR sensors are receiving minimum readings for the duration of the 1-minute window.

5.5.2 Active Interaction Count Analysis. In addition to the estimate of engagement, we separately estimate the level of active interaction. To capture active interactions, we count the number of IR readings having a value of 0.25 or greater, which corresponds to a proximity of 35 cm or less from an IR sensor, within 1 minute. Formally, given M IR readings received within 1 minute (typically sampled at frequency $F = 10\text{Hz}$) $\{\mathbf{ir}^{(1)}, \mathbf{ir}^{(2)}, \dots, \mathbf{ir}^{(M)}\}$ where each IR reading $\mathbf{ir}^{(i)}$ is a vector of 24 IR values, the number of active interactions N_{active} is defined by Equation (4):

$$N_{active} = \frac{1}{F} \sum_{m=1}^M \sum_{i=1}^{24} \mathbb{1} \{ \mathbf{ir}_i^{(m)} \geq 0.25 \}, \quad (4)$$

where $\mathbf{ir}_i^{(m)}$ is the i th IR sensor in the m th IR reading and $\mathbb{1} \{ \cdot \}$ is an indicator function. Therefore, N_{active} is the total detected active interactions within 1 minute.

6 RESULTS

In this section, we first quantitatively compare the performance of PB and PLA based on evaluation metrics introduced in Section 5.5. After that, we analyze the human survey data.

6.1 Quantitative Comparison Between PB and PLA

In this section, we quantitatively compare the performance of the two behavior modes based on sensory data collected during the interaction between visitors and the LAS. We use two ways to quantitatively compare the two behaviors' performance: (1) comparing the estimated engagement

Fig. 10. Estimated occupancy comparison. (a) Q-Q (100-quantiles–100-quantiles) plot of estimated per-minute occupancy, using the method introduced in Section 5.4.2, for PB and PLA, where the coordinate (x, y) of the q th point from bottom left to top right corresponds to the estimated occupancy of (PB, PLA) for the q th percentile (i.e., $Q_q, q = 0, 1, \dots, 100$), and the reference line indicates a perfect match of distribution between PB and PLA. For example, the point $(4.3, 4)$ for PB vs PLA at the Q_{75} means that 75% of observations for PB and PLA are less than 4.3 and 4, respectively. (b) Average estimated per-minute occupancy and its standard error for PB and PLA.

level, as described in Section 5.5.1, and (2) comparing the number of active interactions, as introduced in Section 5.5.2.

Our experiment is run in a natural setting (i.e., a publicly accessible museum), so it is possible that there are different occupancy levels in the space due to factors not related to the behavior mode. To check whether there are different occupancy levels between conditions (which might be caused either by some behaviors being more attractive to visitors or factors not related to system behaviors), we analyze the overall occupancy level for PB and PLA, as described in Section 5.4.2. Figure 10 shows a comparison of the estimated occupancy between PB and PLA, where (a) shows that in only about 5% of data, PLA has approximately one more visitor than PB, and (b) shows that the average occupancy between PB and PLA is very similar. A Mann-Whitney U test indicates that there is no significant difference between PB and PLA in terms of occupancy level, $U = 239030.5$, $p = 0.92$ (two sided).

6.1.1 Estimated Engagement Level Comparison. Figure 11 compares the distributions of estimated engagement (see Section 5.5.1) between PB and PLA. From Figure 11, we can observe that for the first 75% of data, there is no noticeable difference between PB and PLA, whereas for the last 25% of data, PLA has larger estimated engagement than PB. Figure 12 shows the average estimated engagement, which shows that PLA achieves higher average engagement than PB.

6.1.2 Active Interaction Comparison. Figure 13 compares the active interaction count based on Equation (4) between PB and PLA. From this figure, we can see that for about 50% of the observations, PLA achieves higher active interaction than PB. Figure 14 compares PB with PLA in terms of average raw active interaction count. As shown in the figure, PLA almost doubles the PB average active interaction count.

6.1.3 Evolution of Average Estimated Engagement and Active Interaction. To analyze how performance evolved over the 3-week experiment, we plot the daily average engagement and active interaction over the whole experiment. Figure 15 shows daily average metrics of PB and PLA. In

Fig. 11. Estimated engagement distribution comparison. The Q-Q (100-quantiles–100-quantiles) plot between PB and PLA is based on average estimated engagement, where the coordinate (x, y) of the q th point from bottom left to top right corresponds to the estimated engagement level of (PB, PLA) for the q th percentile (i.e., $Q_q, q = 0, 1, \dots, 100$), and the reference line indicates a perfect match of distributions between PB and PLA.

Fig. 12. Average estimated engagement comparison, where blue bars with standard errors show the average estimated engagement and its corresponding standard error.

Fig. 13. Active interaction count comparison. The Q-Q plot is based on active interaction count per minute obtained using Equation (4), where every (x, y) corresponds to the active interaction count of PB in one percentile, and the line indicates a perfect match of distributions between PB and PLA.

Fig. 14. Average active interaction count comparison, where blue bars with standard errors show the average number of active interactions per minute for each behavior (the counting is based on 1-Hz interaction frequency).

terms of daily estimated engagement, Figure 15(a) shows that during the first 2 days, PB outperforms PLA, whereas after September 25, PLA overtakes PB for the remainder of the experiment. A similar trend can be seen in terms of daily active interaction as shown in Figure 15(b). PLA receives more active interaction than PB from the very beginning and keeps expanding the gap between PLA and PB. Even though it seems like PLA is improving, due to the uncontrolled experimental

Fig. 15. Trajectory of daily average metrics. Daily average estimated engagement (a) and daily average active interaction (b), where each data point is the corresponding average on each day, the lines are linear regressions of these data, and the translucent bands around the regression line are the 95% confidence intervals for the regression estimate.

setting, we cannot be certain whether this is caused by continuous adapting of PLA or due to factors independent from the interactive action of the LAS.

6.2 Analysis of Actions Automatically Generated by PLA

In this section, we analyze actions automatically generated by PLA. Since the dimensionality of the action space is 11 and each dimension is continuous, visualization and analysis of the action trajectory is difficult. For ease of visualization, we first cluster actions into six clusters using k -means [2], then use t-SNE [50] to embed actions generated by PLA into a 2-dimensional space, where actions that are close in high-dimensional space are modeled by nearby points and dissimilar actions are modeled by distant points with high probability. The resulting visualization allows us to compare actions over all days or between specific days. Note that the clustering is done in the 11-dimensional action space of PLA rather than the 2-dimensional embedding space, and the centroid of each cluster is an 11-dimensional data point that can be thought of as the multi-dimensional average of the points in the cluster. Figure 16(a) and (b) show embedded actions generated by PLA viewed in different colors for different days and different clusters, respectively, where PB is also embedded for comparison and is indicated by the red star. Figure 16(c) shows the actions taken by PLA separately by day, with the different clusters labeled.

Figure 17 depicts, for each PLA action cluster, the difference between each of the 11 dimensions of the cluster centroid and the corresponding default value of PB. From Figure 17, we can observe that (1) most action dimensions of the centroids of clusters 2, 3, 4, and 5 are greater than the default values of PB; (2) cluster 1 is very similar to PB on many action dimensions; and (3) the centroid of cluster 6 has more action dimensions with values lower than the default value of PB. Taking cluster 4 of PLA as a concrete example, most dimensions of the centroid are larger than the default value of PB, indicating that actions in cluster 4 are slower and smoother than PB since Moths and LEDs take more time to ramp up (T_{ru}^m, T_{ru}^l), hold on (T_{ho}^m, T_{ho}^l), and ramp down (T_{rd}^m, T_{rd}^l), and the time gaps between activating Moths and LEDs and neighbor nodes are longer ($T_{gap}^m, T_{gap}^n, T_{sma}$). Cluster 4 seems to be preferred on September 26, October 3, and October 5, as shown in Figure 16(c), as many actions in cluster 1 are taken. Cluster 1 shares more similarity with PB, as seen from the small differences in most dimensions shown in the fourth panel in Figure 17. Actions in cluster 1 are taken densely on September 17, 19, 20, and 26 and October 5, as shown in Figure 17. As

Fig. 16. Visualizing actions taken by PLA embedded in a two-dimensional space, where each axis corresponds to one embedded dimension. (a) Actions on all days with different colors for different days. (b) Actions on all days with different colors for different clusters where k -means was used to generate the six clusters. (c) Actions taken by PLA on each day, with different colors for different clusters.

Fig. 17. Difference between the cluster centroid of PLA and PB. Each panel shows the difference $a_{PLA} - a_{PB}$ between the cluster centroids of PLA a_{PLA} and the default value of PB a_{PB} for each dimension of the action space elaborated in Table 1, where a positive value means that the corresponding dimension of the cluster centroid of PLA is greater than the default value of PB, whereas a negative value means the contrary.

Table 4. Cronbach's α on Godspeed for PB and PLA

	Anthropomorphism	Animacy	Likeability	Perceived Intelligence	Perceived Safety
PB	0.74	0.77	0.85	0.89	0.52
PLA	0.64	0.80	0.93	0.85	0.27

A commonly accepted rule [21]: $0.9 \leq \alpha$: Excellent; $0.8 \leq \alpha < 0.9$: Good; $0.7 \leq \alpha < 0.8$: Acceptable; $0.6 \leq \alpha < 0.7$: Questionable; $0.5 \leq \alpha < 0.6$: Poor; $\alpha < 0.5$: Unacceptable.

an interesting and slightly contrary example to cluster 4, cluster 6 of PLA shows abrupt action where Moths and LEDs take less time than PB to ramp up (T_{ru}^m, T_{ru}^l), which seems to be preferred on September 24. Note that here we arbitrarily set the total cluster number to 6 for a relatively clear visualization. As the clusters still have considerable within-cluster variance, the analysis of centroids only provides an approximate analysis of the diversity of actions taken by PLA.

Combining Figure 16(a) and (c) and Figure 17, we do not see specific types of action that are dominant. Nevertheless, it seems that PLA continuously adapts, because actions generated by PLA on each day show slightly different coverage as shown in Figure 16(c) and demonstrated in the previous paragraph. Overall, we can observe that PLA covered a wide range of actions on each day and no obvious dominant actions were reached by the end of the experiment.

6.3 Human Survey Results

In this section, we analyze the visitor responses to the survey for PB and PLA. We first examine if there are any differences in the population characteristics between the participants who engaged with the system in PB or PLA behavior modes. Then, we compare the PB and PLA responses for each Godspeed category. Finally, we compare PB and PLA for each question in the Godspeed Likeability category individually.

We first analyze whether there are population differences between conditions. In Section 5.4.2, we confirmed that there were no significant differences between PB and PLA in terms of estimated occupancy. To test for differences in participant background and interest, we performed a χ^2 -test on participants' background and interests based on the first two questions in our questionnaire (see Section 5.2) and found no statistically significant differences between the two groups.

Cronbach's α -test was conducted on each category of Godspeed for both PB and PLA to examine the reliability of participants' responses, and results are shown in Table 4. Although α on Anthropomorphism and Perceived Safety is low, α on others is in the acceptable range, especially for Likeability $\alpha \geq 0.85$.

Figure 18 shows the Box plot and Violin plot of the calculated average scale over each Godspeed category for PB and PLA. Within the five Godspeed categories, only Likeability has a relatively large gap between the medians of PB and PLA. In addition, Likeability has a relatively small variance, whereas other categories have large variance. A normality test was conducted for the Likeability category for participants from PB and PLA, respectively. The Shapiro-Wilk test [63] indicates PB ($p = 0.12$) is normally distributed, whereas PLA ($p = 0.0008$) is not. Therefore, a non-parametric test—the Mann-Whitney U test [52]—was conducted and found that the distributions of PB (median = 4.10, M = 4.09, SD = 0.56) and PLA (median = 5.00, M = 4.48, SD = 0.64) differ significantly (Mann-Whitney U = 67, $p = 0.044 < 0.05$) in terms of Likeability, whereas for other categories there is no significant difference between PB and PLA.

In summary, PLA is rated higher than PB by the participants in terms of Likeability, whereas there are no significant differences between PB and PLA in the other Godspeed categories.

Fig. 18. Box plot and Violin plot of average scale of each Godspeed category over participants within PB or PLA.

7 DISCUSSION

In this article, we investigated how an interactive system can learn to engage with visitors in a natural setting, where no constraints are imposed on visitors and group interaction is accommodated. Relying on the standard RL framework and a novel measure of engagement as reward, an adaptive behavior (PLA) was compared to a PB. Our results show that PLA outperforms PB in terms of estimated engagement level, active interaction count, and likeability in human survey data. We hypothesize that the PLA configuration outperforms PB because it benefits from both human expert input, such as parameterized action space and manual reward function, and learning. During the design process, architects design PB based on their expertise. We exploit this human expertise to effectively restrict the parameterized action space of PLA into a region where we know good solutions can be found, and at the same time limit the dimension of the PLA action space. Therefore, PB and PLA both incorporate human intelligence, but compared with PB, PLA is endowed with adaptability by applying RL to its parameterized action space.

Our approach uses low-cost IR sensors for engagement estimation. Compared to other social HRI work with rich sensing such as cameras and microphones, we were still able to estimate engagement with limited sensing and generate engaging behaviors accordingly. This might be helpful to other large-scale interactive systems where having sophisticated measurement may be unfeasible. In summary, this work provides two useful generalizable guidelines for designing engaging behavior for long-term interaction: (1) in a group interaction setting, group engagement from low-cost ambient sensors can be used either stand-alone or as a complement of individual engagement measures, and (2) such a measure of engagement can be used as a reward signal to generate customized and evolving behavior.

Creating engaging behaviours for the LAS requires learning algorithms that continue to adapt rather than optimizing to a single best policy. First, the Markov decision process and stationary environment dynamics assumptions are broken because of the complicated interaction environment. In addition, interaction time cannot be assumed to have a constant length, and the interaction speed is bound to physical interaction and cannot be sped up. In addition, since the LAS is an architectural scale interactive system, perception of the environment is more complex and may need to include both proprioception and exteroception. Therefore, although we exploit an RL framework in our work, the role it plays is different from that of standard testbeds such as OpenAI Gym [13] or the Arcade Learning Environment (ALE) [8]. In this work, RL is used to introduce adaptability, but there is no guarantee that the learning leads to optimal policy. This is illustrated by the observa-

(a) A visitor learns how to interact from another visitor

(b) A parent lifts a child

(c) A group visit led by a guide

(d) Visitors taught by a security guard but misunderstood the touch IR sensors

Fig. 19. Sample interesting scenario. (a) How visitors can affect each other. Before the woman in blue arrives, the man in black had spent a while just looking around and did not know how to interact with the LAS. Then coincidentally he saw how the woman raised her hand and how the LAS responded. After that, he copied her action to interact with the LAS. (b) A parent helping his child experience the interaction after he explored how to interact. Since the child could not reach the IR sensor, the parent lifted his child. (c) A group visit where a group of visitors is led by a guide. (d) A more complex scenario about misunderstanding shared information. In this case, three girls were taking photos without interacting. Then a security guard taught one of the three girls how to interact. At the same time, the second girl saw and joined them. After that, the third girl learned this from her friends. However, there was a misunderstanding of the security guard's instruction, so they directly touched the IR sensors rather than just closely waving a hand.

tion that compared with PB, PLA shows very flexible action patterns, and some are very different from PB. Specifically, one observed behavior generated by PLA is LEDs turned on and propagated quickly from one node to another back and forth multiple times accompanying activated SMAs and Moths (see <https://youtu.be/2tICanYEpo0>) for the video comparing PB and PLA, which makes the LAS look like a thunderstorm. This novel behavior illustrates that the sculpture has taken the primitives composed by the designers and evolved engaging and interesting behavior from those.

The group setting presents a challenging environment for learning. During the entire experiment, we found some scenarios that highlight the complexity of using RL in LAS, such as interactions between visitors and the possibility that the LAS could be physically changed by touching as shown in Figure 19. In addition, there are many other examples of complicated environment dynamics that present challenges to a learning algorithm. Examples include visitors who use alternate interaction strategies, change their interaction strategy over time, or are influenced by other visitors, and visitors who raise their hands for reasons other than interaction. These observations illustrate the non-stationarity of the environment and the influence of

human-human interaction in group scenarios during HRI. They also emphasize the importance of developing and testing these algorithms “in the wild.”

Selecting RL algorithms and hyper-parameters with real experiments is challenging, as we do not have a “validation set” that allows us to do multiple runs and we cannot accelerate learning. Due to the non-stationary nature of the investigated environment in our case, this is also impossible to do with a simulator, because unlike some fields in robotics where a realistic simulator can be devised, in our case the diversity of visitors’ interaction styles, social influence of human-human interaction, and variations in visitor numbers during different time periods, and so forth, make it impossible to devise such a good simulator. Therefore, in this work, we only used a simplified simulator to make an initial hyper-parameter selection and confirm our code is bug free rather than pre-train a policy or fine tune hyper-parameters. Finally, although DDPG was selected as the learning algorithm because it is easy to implement and works on continuous action space, the best choice of RL algorithm requires further investigation. DDPG may not be sample efficient and can be susceptible to over-estimation and sensitive to hyper-parameters [32]. More advanced continuous control algorithms, such as Soft Actor-Critic (SAC) [29], Twin Delayed DDPG (TD3) [24], and Multi-Step DDPG (MDDPG) [53], should be investigated in the future.

8 CONCLUSION AND FUTURE WORK

In this article, we developed and evaluated algorithms for generating interactive behaviors in group environments. Specifically, we provided a way to estimate engagement during group interaction based on multiple IR sensors, where both individual engagement (i.e., passive and active interaction) and group engagement (i.e., occupancy) are taken into account. PB and PLA were examined to evaluate how the use of human knowledge influences interaction. By analyzing interaction and human survey data, we found that learned interactive behaviors (i.e., PLA) result in higher engagement and perceived likeability than PB.

Even though PLA received higher average engagement and perceived likeability than PB, we cannot be certain about the cause of this difference. Therefore, a baseline with random policy can be tested to see if there is a difference between this baseline and PLA to confirm that the learning agent is indeed learning from and adapting to its interaction experience. Other advanced continuous control DRL algorithms such as SAC, TD3, and MDDPG are also worth investigating. Upcoming installations of even larger LAS are planned, and this decentralization will become necessary as the number of actuated elements increases beyond 1,000 in a single installation. In addition, hierarchical RL with PB bootstrapping could be a promising extension, where we could design a pool of PBs and various levels of reward functions, and see how complicated action patterns could emerge. We also plan to introduce intrinsic motivation and a learning algorithm driven both intrinsically and extrinsically for the LAS. To tackle the low pace of interaction in the LAS and high sample requirement of RL, we also plan to investigate how to transfer learned models from simulation to a physical LAS.

APPENDIX

A PSEUDO-CODE FOR DDPG-BASED PLA

ALGORITHM 1: DDPG-Based PLA

Input: *first_day*

- 1 Initialize: *train_interval*, *train_times*, *episode_length*, replay buffer $\mathcal{D}$, mini-batch size N
- 2 **if** *first_day* **then**
- 3 Initialize Critic $Q(obs, a|\theta^Q)$ and Actor $\mu(obs|\theta^\mu)$ with random parameters θ^Q, θ^μ
- 4 **else**
- 5 Load pre-trained Critic $Q(obs, a|\theta^Q)$ and Actor $\mu(obs|\theta^\mu)$
- 6 **end**
- 7 Target networks $\theta^{Q-} \leftarrow \theta^Q, \theta^{\mu-} \leftarrow \theta^\mu$
- 8 *start_of_day* $\leftarrow$ *True*
- 9 **while** *True* **do**
- 10 **if** *start_of_day* **then**
- 11 Receive initial observation $obs^{(1)}$
- 12 Initialize parameter noise σ
- 13 *start_of_day* $\leftarrow$ *False*
- 14 **else**
- 15 $obs^{(1)} \leftarrow obs^{(t+1)}, \sigma \leftarrow \sigma_{j+1}$
- 16 **end**
- 17 **for** $t \leftarrow 1$ **to** *episode_length* **do**
- 18 Select action $a^{(t)} = \mu(obs^{(t)}|\theta^\mu + \mathcal{N}(0, \sigma))$ according to the current policy and exploration noise.
- 19 Execute action $a^{(t)}$, observe reward $r^{(t)}$ and observe new observation $obs^{(t+1)}$
- 20 $\mathcal{D} \leftarrow \mathcal{D} \cup (obs^{(t)}, a^{(t)}, r^{(t)}, obs^{(t+1)}, done^{(t+1)} = 0)$
- 21 **if** $t \% train_interval == 0$ **and** $|\mathcal{D}| \geq N$ **then**
- 22 **for** $j \leftarrow 1$ **to** *train_times* **do**
- 23 Sample a random mini-batch of N transitions $(obs^{(i)}, a^{(i)}, r^{(i)}, obs^{(i+1)})$ from $\mathcal{D}$
- 24 /* Update adaptive parameter noise scale */
- 25
$$\sigma_{j+1} \leftarrow \begin{cases} \alpha \sigma_j & \text{if } d(\mu_{\theta^\mu}, \mu_{\theta^\mu + \mathcal{N}(0, \sigma_j)}) \leq \delta \\ \frac{1}{\alpha} \sigma_j & \text{otherwise.} \end{cases}$$
- 26 where $d(\mu_{\theta^\mu}, \mu_{\theta^\mu + \mathcal{N}(0, \sigma_j)}) = \sqrt{\frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\left(\mu(s|\theta^\mu) - \mu(s|\theta^\mu + \mathcal{N}(0, \sigma_j)) \right)^2 \right]}$.
- 27 /* Update Actor-Critic */
- 28 Set $y^{(i)} = r^{(i)} + \gamma Q'(obs^{(i+1)}, \mu(obs^{(i+1)}|\theta^{\mu-})|\theta^{Q-})$
- 29 Update critic by minimizing the loss: $L = \frac{1}{N} \sum_i \left(y^{(i)} - Q(s^{(i)}, a^{(i)}|\theta^Q) \right)^2$
- 30 Update actor using the sampled policy gradient:
- 31
$$\nabla_{\theta^\mu} J \approx \frac{1}{N} \sum_i \nabla_{\mu(obs^{(i)}|\theta^\mu)} Q(obs^{(i)}, \mu(obs^{(i)}|\theta^\mu)|\theta^Q) \nabla_{\theta^\mu} \mu(obs^{(i)}|\theta^\mu)$$
- 32 Update the target networks:
- 33
$$\theta^{Q-} \leftarrow \tau \theta^Q + (1 - \tau) \theta^{Q-}$$
- 34
$$\theta^{\mu-} \leftarrow \tau \theta^\mu + (1 - \tau) \theta^{\mu-}$$
- 35 **end**
- 36 **end**
- 37 $obs^{(t)} \leftarrow obs^{(t+1)}, \sigma_1 \leftarrow \sigma_{j+1}$

ACKNOWLEDGMENTS

The authors would like to thank Philip Beesley Architect Inc. and the Royal Ontario Museum for allowing us to conduct research with the exhibit, Salvador Breed and 4DSOUND for developing the sound system, and Joslin Goh and Pavel Shuldiner from the Statistical Consulting and Collaborative Research Unit (SCCR) at the University of Waterloo for insightful discussions on statistical analysis.

REFERENCES

- [1] Salvatore M. Anzalone, Sofiane Boucenna, Serena Ivaldi, and Mohamed Chetouani. 2015. Evaluating the engagement with social robots. *International Journal of Social Robotics* 7, 4 (Aug. 2015), 465–478. DOI : <https://doi.org/10.1007/s12369-015-0298-7>
- [2] David Arthur and Sergei Vassilvitskii. 2007. k-means++: The advantages of careful seeding. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms*. 1027–1035.
- [3] Karl J. Åström and Björn Wittenmark. 2013. *Adaptive Control*. Courier Corporation.
- [4] Federico Augugliaro, Angela P. Schoellig, and Raffaello D’Andrea. 2013. Dance of the flying machines: Methods for designing and executing an aerial dance choreography. *IEEE Robotics Automation Magazine* 20, 4 (Dec. 2013), 96–104. DOI : <https://doi.org/10.1109/MRA.2013.2275693>
- [5] Christoph Bartneck, Dana Kulić, Elizabeth Croft, and Susana Zoghbi. 2009. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International Journal of Social Robotics* 1, 1 (Jan. 2009), 71–81. DOI : <https://doi.org/10.1007/s12369-008-0001-3>
- [6] Philip Beesley, Matthew Chan, Rob Gorbet, Dana Kulić, and Mo Memarian. 2015. Evolving systems within immersive architectural environments: New research by the Living Architecture Systems Group. *Next Generation Building* 2 (2015), 31–56. DOI : <https://doi.org/10.7564/15-NGBJ20>
- [7] Philip Beesley, Pernilla Ohrstedt, and Rob Gorbet. 2010. *Hylozoic Ground: Liminal Responsive Architecture: Philip Beesley*. Riverside Architectural Press. <https://books.google.ca/books?id=Ad6gAQAACAJ>.
- [8] Marc G. Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. 2012. The arcade learning environment: An evaluation platform for general agents. arXiv:1207.4708. <http://arxiv.org/abs/1207.4708>.
- [9] Dan Bohus and Eric Horvitz. 2009. Models for multiparty engagement in open-world dialog. In *Proceedings of the SIGDIAL 2009 Conference*. 225–234. <http://aclweb.org/anthology/W09-3933>.
- [10] Dan Bohus and Eric Horvitz. 2009. *Open-World Dialog: Challenges, Directions, and a Prototype*. Technical Report. Microsoft. <https://www.microsoft.com/en-us/research/publication/open-world-dialog-challenges-directions-and-a-prototype/>.
- [11] Dan Bohus and Eric Horvitz. 2014. Managing human-robot engagement with forecasts and... um... hesitations. In *Proceedings of the 16th International Conference on Multimodal Interaction (ICMI’14)*. ACM, New York, NY, 2–9. DOI : <https://doi.org/10.1145/2663204.2663241>
- [12] Cynthia Breazeal, Kerstin Dautenhahn, and Takayuki Kanda. 2016. *Social Robotics*. Springer International, Cham, Switzerland 1935–1972. DOI : https://doi.org/10.1007/978-3-319-32552-1_72
- [13] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. OpenAI Gym. arXiv:1606.01540. <http://arxiv.org/abs/1606.01540>.
- [14] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. 2016. Realtime multi-person 2D pose estimation using part affinity fields. arXiv:1611.08050 <http://arxiv.org/abs/1611.08050>.
- [15] Ginevra Castellano, André Pereira, Iolanda Leite, Ana Paiva, and Peter W. McOwan. 2009. Detecting user engagement with a robot companion using task and social interaction-based features. In *Proceedings of the 2009 International Conference on Multimodal Interfaces (ICMI-MLMI’09)*. ACM, New York, NY, 119–126. DOI : <https://doi.org/10.1145/1647314.1647336>
- [16] Matthew T. K. Chan, Rob Gorbet, Philip Beesley, and Dana Kulić. 2015. Curiosity-based learning algorithm for distributed interactive sculptural systems. In *Proceedings of the 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS’15)*. IEEE, Los Alamitos, CA, 3435–3441. DOI : <https://doi.org/10.1109/IROS.2015.7353856>
- [17] Matthew T. K. Chan, Rob Gorbet, Philip Beesley, and Dana Kulić. 2016. Interacting with curious agents: User experience with interactive sculptural systems. In *Proceedings of the 2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN’16)*. IEEE, Los Alamitos, CA, 151–158. DOI : <https://doi.org/10.1109/ROMAN.2016.7745104>
- [18] Yu Fan Chen, Michael Everett, Miao Liu, and Jonathan P. How. 2017. Socially aware motion planning with deep reinforcement learning. In *Proceedings of the 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS’17)*. IEEE, Los Alamitos, CA, 1343–1350.

- [19] Michael Jae-Yoon Chung and Maya Cakmak. 2018. “How was your stay?”: Exploring the use of robots for gathering customer feedback in the hospitality industry. In *Proceedings of the 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN’18)*. 947–954. DOI: <https://doi.org/10.1109/ROMAN.2018.8525604>
- [20] Catie Cuan, Ishaan Pakrasi, and Amy LaViers. 2018. Time to compile. In *Proceedings of the 5th International Conference on Movement and Computing (MOCO’18)*. ACM, New York, NY, Article 53, 4 pages. DOI: <https://doi.org/10.1145/3212721.3212888>
- [21] Robert F. DeVellis. 2016. *Scale Development: Theory and Applications*. Vol. 26. Sage Publications.
- [22] Prafulla Dhariwal, Christopher Hesse, Oleg Klimov, Alex Nichol, Matthias Plappert, Alec Radford, John Schulman, Szymon Sidor, Yuhuai Wu, and Peter Zhokhov. 2017. OpenAI Baselines. Retrieved August 1, 2020 from <https://github.com/openai/baselines>.
- [23] Kevin Doherty and Gavin Doherty. 2018. Engagement in HCI: Conception, theory and measurement. *ACM Computing Surveys* 51, 5 (Nov. 2018), Article 99, 39 pages. DOI: <https://doi.org/10.1145/3234149>
- [24] Scott Fujimoto, Herke van Hoof, and David Meger. 2018. Addressing function approximation error in actor-critic methods. arXiv:1802.09477.
- [25] Nadine Glas and Catherine Pelachaud. 2015. Definitions of engagement in human-agent interaction. In *Proceedings of the 2015 International Conference on Affective Computing and Intelligent Interaction (ACII’15)*. 944–949. DOI: <https://doi.org/10.1109/ACII.2015.7344688>
- [26] Michael A. Goodrich and Alan C. Schultz. 2007. Human-robot interaction: A survey. *Foundations and Trends in Human-Computer Interaction* 1, 3 (Jan. 2007), 203–275. DOI: <https://doi.org/10.1561/1100000005>
- [27] Goren Gordon, Samuel Spaulding, Jacqueline Kory Westlund, Jin Joo Lee, Luke Plummer, Marayna Martinez, Madhurima Das, and Cynthia Breazeal. 2016. Affective personalization of a social robot tutor for children’s second language skills. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*.
- [28] Shane Griffith, Kaushik Subramanian, Jonathan Scholz, Charles L. Isbell, and Andrea Thomaz. 2013. Policy shaping: Integrating human feedback with reinforcement learning. In *Proceedings of the 26th International Conference on Neural Information Processing Systems—Volume 2 (NIPS’13)*. 2625–2633. <http://dl.acm.org/citation.cfm?id=2999792.2999905>.
- [29] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. arXiv:1801.01290.
- [30] Sami Haddadin and Elizabeth Croft. 2016. *Physical Human–Robot Interaction*. Springer International, Cham, Switzerland, 1835–1874. DOI: https://doi.org/10.1007/978-3-319-32552-1_69
- [31] Kotaro Hayashi, Daisuke Sakamoto, Takayuki Kanda, Masahiro Shiomi, Satoshi Koizumi, Hiroshi Ishiguro, Tsukasa Ogasawara, and Norihiro Hagita. 2007. Humanoid robots as a passive-social medium—A field experiment at a train station. In *Proceedings of the 2007 2nd ACM/IEEE International Conference on Human-Robot Interaction (HRI’07)*. IEEE, Los Alamitos, CA, 137–144. DOI: <https://doi.org/10.1145/1228716.1228735>
- [32] Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. 2018. Deep reinforcement learning that matters. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*.
- [33] Charles Isbell, Christian R. Shelton, Michael Kearns, Satinder Singh, and Peter Stone. 2001. A social reinforcement learning agent. In *Proceedings of the 5th International Conference on Autonomous Agents (AGENTS’01)*. ACM, New York, NY, 377–384. DOI: <https://doi.org/10.1145/375735.376334>
- [34] Serena Ivaldi, Sébastien Lefort, Jan Peters, Mohamed Chetouani, Joelle Provasi, and Elisabetta Zibetti. 2015. Towards engagement models that consider individual factors in HRI: on the relation of extroversion and negative attitude towards robots to gaze and speech during a human-robot assembly task. arXiv:1508.04603. <http://arxiv.org/abs/1508.04603>
- [35] Michiel Jooisse and Vanessa Evers. 2017. A guide robot at the airport: First impressions. In *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction (HRI’17)*. ACM, New York, NY, 149–150. DOI: <https://doi.org/10.1145/3029798.3038389>
- [36] Kirsikka Kaipainen, Aino Ahtinen, and Aleksi Hiltunen. 2018. Nice surprise, more present than a machine: Experiences evoked by a social robot for guidance and edutainment at a city service point. In *Proceedings of the 22nd International Academic Mindtrek Conference (Mindtrek’18)*. ACM, New York, NY, 163–171. DOI: <https://doi.org/10.1145/3275116.3275137>
- [37] Takayuki Kanda and Hiroshi Ishiguro. 2012. *Human-Robotic Interaction in Social Robotics*. Boca Raton, FL.
- [38] Simon Keizer, Mary Ellen Foster, Zhuoran Wang, and Oliver Lemon. 2014. Machine learning for social multiparty human–robot interaction. *ACM Transactions on Interactive Intelligent Systems* 4, 3 (2014), 14.
- [39] Mehdi Khamassi, George Velentzas, Theodore Tsitsimis, and Costas Tzafestas. 2018. Robot fast adaptation to changes in human engagement during simulated dynamic social interaction with active exploration in parameterized reinforcement learning. *IEEE Transactions on Cognitive and Developmental Systems* 10, 4 (2018), 881–893.

- [40] Rajiv Khosla, Khanh Nguyen, and Mei-Tai Chu. 2017. Human robot engagement and acceptability in residential aged care. *International Journal of Human-Computer Interaction* 33, 6 (2017), 510–522. DOI: <https://doi.org/10.1080/10447318.2016.1275435>
- [41] W. Bradley Knox and Peter Stone. 2010. Combining manual feedback with subsequent MDP reward signals for reinforcement learning. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'10)*.
- [42] W. Bradley Knox and Peter Stone. 2012. Reinforcement learning from simultaneous human and MDP reward. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems—Volume 1 (AAMAS'12)*. 475–482. <http://dl.acm.org/citation.cfm?id=2343576.2343644>.
- [43] Samantha Krening and Karen M. Feigh. 2018. Interaction algorithm effect on human experience with reinforcement learning. *ACM Transactions on Human-Robot Interaction* 7, 2 (Oct. 2018), Article 16, 22 pages. DOI: <https://doi.org/10.1145/3277904>
- [44] Dana Kulic and Elizabeth A. Croft. 2007. Affective state estimation for human-robot interaction. *IEEE Transactions on Robotics* 23, 5 (2007), 991–1000.
- [45] Kazumi Kumagai, Daiwei Lin, Lingheng Meng, Alexandru Blidaru, Philip Beesley, Dana Kulić, and Ikuro Mizuuchi. 2018. Towards individualized affective human-machine interaction. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, Los Alamitos, CA, 678–685.
- [46] Iolanda Leite, André Pereira, Ginevra Castellano, Samuel Mascarenhas, Carlos Martinho, and Ana Paiva. 2011. Modelling empathy in social robotic companions. In *Proceedings of the International Conference on User Modeling, Adaptation, and Personalization*. 135–147.
- [47] Florent Levillain, Elisabetta Zibetti, and Sébastien Lefort. 2017. Interacting with non-anthropomorphic robotic artworks and interpreting their behaviour. *International Journal of Social Robotics* 9, 1 (Jan. 2017), 141–161. DOI: <https://doi.org/10.1007/s12369-016-0381-8>
- [48] Xiang Li, Yongping Pan, Gong Chen, and Haoyong Yu. 2016. Adaptive human-robot interaction control for robots driven by series elastic actuators. *IEEE Transactions on Robotics* 33, 1 (2016), 169–182.
- [49] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2015. Continuous control with deep reinforcement learning. arXiv:1509.02971. <http://arxiv.org/abs/1509.02971>.
- [50] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research* 9 (Nov. 2008), 2579–2605.
- [51] Douglas G. Macharet and Dinei A. Florencio. 2013. Learning how to increase the chance of human-robot engagement. In *Proceedings of the 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2173–2179. DOI: <https://doi.org/10.1109/IROS.2013.6696660>
- [52] Henry B. Mann and Donald R. Whitney. 1947. On a test of whether one of two random variables is stochastically larger than the other. *Annals of Mathematical Statistics* 18, 1 (1947), 50–60.
- [53] Lingheng Meng, Rob Gorbet, and Dana Kulić. 2020. The effect of multi-step methods on overestimation in deep reinforcement learning. arXiv:2006.12692.
- [54] Marek P. Michalowski, Selma Sabanovic, and Reid Simmons. 2006. A spatial model of engagement for a social robot. In *Proceedings of the 9th IEEE International Workshop on Advanced Motion Control*. 762–767. DOI: <https://doi.org/10.1109/AMC.2006.1631755>
- [55] Bilge Mutlu and Jodi Forlizzi. 2008. Robots in organizations: The role of workflow, social, and environmental factors in human-robot interaction. In *Proceedings of the 2008 3rd ACM/IEEE International Conference on Human-Robot Interaction (HRI'08)*. 287–294. DOI: <https://doi.org/10.1145/1349822.1349860>
- [56] Kazushi Nakazawa, Keita Takahashi, and Masahide Kaneko. 2013. Unified environment-adaptive control of accompanying robots using artificial potential field. In *Proceedings of the 2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI'13)*. IEEE, Los Alamitos, CA, 199–200.
- [57] Heather L O'Brien and Elaine G. Toms. 2008. What is user engagement? A conceptual framework for defining user engagement with technology. *Journal of the American Society for Information Science and Technology* 59, 6 (2008), 938–955. DOI: <https://doi.org/10.14288/1.0107445>
- [58] Pierre-Yves Oudeyer and Frederic Kaplan. 2009. What is intrinsic motivation? A typology of computational approaches. *Frontiers in Neurobotics* 1 (2009), 6. DOI: <https://doi.org/10.3389/neuro.12.006.2007>
- [59] Matthias Plappert, Rein Houthoofd, Prafulla Dhariwal, Szymon Sidor, Richard Y. Chen, Xi Chen, Tamim Asfour, Pieter Abbeel, and Marcin Andrychowicz. 2017. Parameter space noise for exploration. arXiv:1706.01905. <http://arxiv.org/abs/1706.01905>.
- [60] Ahmed Hussain Qureshi, Yutaka Nakamura, Yuichiro Yoshikawa, and Hiroshi Ishiguro. 2016. Robot gains social intelligence through multimodal deep reinforcement learning. In *Proceedings of the 2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids'16)*. IEEE, Los Alamitos, CA, 745–751.

- [61] Charles Rich, Brett Ponsler, Aaron Holroyd, and Candace L. Sidner. 2010. Recognizing engagement in human-robot interaction. In *Proceedings of the 2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI'10)*. IEEE, Los Alamitos, CA, 375–382. DOI: <https://doi.org/10.1109/HRI.2010.5453163>
- [62] Jyotirmay Sanghvi, Ginevra Castellano, Iolanda Leite, André Pereira, Peter W. McOwan, and Ana Paiva. 2011. Automatic analysis of affective postures and body motion to detect engagement with a game companion. In *Proceedings of the 6th International Conference on Human-Robot Interaction*. ACM, New York, NY, 305–312.
- [63] Samuel Sanford Shapiro and Martin B. Wilk. 1965. An analysis of variance test for normality (complete samples). *Biometrika* 52, 3–4 (1965), 591–611.
- [64] Thomas B. Sheridan. 2016. Human–robot interaction: Status and challenges. *Human Factors* 58, 4 (2016), 525–532. DOI: <https://doi.org/10.1177/0018720816644364>
- [65] Chao Shi, Satoru Satake, Takayuki Kanda, and Hiroshi Ishiguro. 2018. A robot that distributes flyers to pedestrians in a shopping mall. *International Journal of Social Robotics* 10, 4 (Sept. 2018), 421–437. DOI: <https://doi.org/10.1007/s12369-017-0442-7>
- [66] Masahiro Shiomi, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. 2006. Interactive humanoid robots for a science museum. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction (HRI'06)*. ACM, New York, NY, 305–312. DOI: <https://doi.org/10.1145/1121241.1121293>
- [67] Candace L. Sidner, Cory D. Kidd, Christopher Lee, and Neal Lesh. 2004. Where to look: A study of human-robot engagement. In *Proceedings of the 9th International Conference on Intelligent User Interfaces (IUI'04)*. ACM, New York, NY, 78–84. DOI: <https://doi.org/10.1145/964442.964458>
- [68] Candace L. Sidner and Christopher Lee. 2003. Engagement rules for human-robot collaborative interactions. In *Proceedings of the 2003 IEEE International Conference on Systems, Man, and Cybernetics: System Security and Assurance—Volume 4*. 3957–3962. DOI: <https://doi.org/10.1109/ICSMC.2003.1244506>
- [69] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. 2014. Deterministic policy gradient algorithms. In *Proceedings of the 31st International Conference on Machine Learning (Proceedings of Machine Learning Research)*, E. P. Xing and T. Jebara (Eds.), Vol. 32. PMLR, Beijing, China, 387–395. <http://proceedings.mlr.press/v32/silver14.html>.
- [70] David St-Onge, Pierre-Yves Brches, Inna Sharf, Nicolas Reeves, Ioannis Rekleitis, Patrick Abouzakhm, Yogesh Girdhar, Adam Harmat, Gregory Dudek, and Philippe Gigure. 2017. Control, localization and human interaction with an autonomous lighter-than-air performer. *Robotics and Autonomous Systems* 88, C (Feb. 2017), 165–186. DOI: <https://doi.org/10.1016/j.robot.2016.10.013>
- [71] David St-Onge and Nicolas Reeves. 2010. Human interaction with flying cubic automata. In *Proceedings of the 2010 IEEE/ACM International Conference on Human Robots Interaction*.
- [72] Halit Bener Suay and Sonia Chernova. 2011. Effect of human guidance and state space size on interactive reinforcement learning. In *Proceedings of the 20th IEEE International Symposium on Robot and Human Interaction Communication (RO-MAN'11)*. IEEE, Los Alamitos, CA, 1–6. DOI: <https://doi.org/10.1109/ROMAN.2011.6005223>
- [73] Andrea Thomaz, Guy Hoffman, and Maya Cakmak. 2016. Computational human-robot interaction. *Foundations and Trends in Robotics* 4, 2–3 (Dec. 2016), 105–223. DOI: <https://doi.org/10.1561/23000000049>
- [74] Andrea L. Thomaz and Cynthia Breazeal. 2006. Reinforcement learning with human teachers: Evidence of feedback and guidance with implications for learning performance. In *Proceedings of the 21st National Conference on Artificial Intelligence—Volume 1 (AAAI'06)*. 1000–1005. <http://dl.acm.org/citation.cfm?id=1597538.1597696>.
- [75] Andrea Lockerd Thomaz, Guy Hoffman, and Cynthia Breazeal. 2005. Real-time interactive reinforcement learning for robots. In *Proceedings of the AAAI 2005 Workshop on Human Comprehensible Machine Learning*.
- [76] Panagiota Tsarouchi, Sotiris Makris, and George Chryssolouris. 2016. Human–robot interaction review and challenges on task planning and programming. *International Journal of Computer Integrated Manufacturing* 29, 8 (2016), 916–931. DOI: <https://doi.org/10.1080/0951192X.2015.1130251>
- [77] George Velentzas, Theodore Tsitsimis, Iñaki Rañó, Costas Tzafestas, and Mehdi Khamassi. 2018. Adaptive reinforcement learning with active state-specific exploration for engagement maximization during simulated child-robot interaction. *Paladyn, Journal of Behavioral Robotics* 9, 1 (2018), 235–253.
- [78] Autilia Vitiello, Giovanni Acampora, Mariacarla Staffa, Bruno Siciliano, and Silvia Rossi. 2017. A neuro-fuzzy-Bayesian approach for the adaptive control of robot proxemics behavior. In *Proceedings of the 2017 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE'17)*. IEEE, Los Alamitos, CA, 1–6.
- [79] Graham Wakefield, Tobias Hollerer, JoAnn Kuchera-Morin, Charles Roberts, and Matthew Wright. 2013. Spatial interaction in a multiuser immersive instrument. *IEEE Computer Graphics and Applications* 33, 6 (Nov. 2013), 14–20. DOI: <https://doi.org/10.1109/MCG.2013.99>
- [80] Chen Wang, Yanan Li, Shuzhi Sam Ge, and Tong Heng Lee. 2016. Adaptive control for robot navigation in human environments based on social force model. In *Proceedings of the 2016 IEEE International Conference on Robotics and Automation (ICRA'16)*. IEEE, Los Alamitos, CA, 5690–5695.

- [81] Juanjuan Zhang and Chien Chern Cheah. 2015. Passivity and stability of human–robot interaction control for upper-limb rehabilitation robots. *IEEE Transactions on Robotics* 31, 2 (2015), 233–245.
- [82] Jakub Zlotowski, Diane Proudfoot, Kumar Yogeeswaran, and Christoph Bartneck. 2015. Anthropomorphism: Opportunities and challenges in human–robot interaction. *International Journal of Social Robotics* 7, 3 (Jan. 2015), 347–360. DOI: <https://doi.org/10.1007/s12369-014-0267-6>

Received April 2019; revised March 2020; accepted June 2020