Human Perceptions of a Curious Robot that Performs Off-Task Actions

Nick Walker
nswalker@cs.washington.edu
University of Washington
Seattle, Washington

Kevin Weatherwax
wax@ucsc.edu
University of California, Santa Cruz
Santa Cruz, California

Julian Allchin
julianallchin@seattleacademy.org
Seattle Academy
Seattle, Washington

Leila Takayama
takayama@ucsc.edu
University of California, Santa Cruz
Santa Cruz, California

Maya Cakmak
mcakmak@cs.washington.edu
University of Washington
Seattle, Washington

ABSTRACT

Researchers have proposed models of curiosity as a means to drive robots to learn and adapt to their environments. While these models balance goal- and exploration-oriented actions in a mathematically principled manor, it is not understood how users perceive a robot that pursues off-task actions. Motivated by a model of curiosity based on intrinsic rewards, we conducted three online video-surveys with a total of 264 participants, evaluating a variety of curious behaviors. Our results indicate that a robot's off-task actions are perceived as expressions of curiosity, but that these actions lead to a negative impact on perceptions of the robot's competence. When the robot explains or acknowledges its deviation from the primary task, this can partially mitigate the negative effects of off-task actions.

CCS CONCEPTS

• **Applied computing** → **Psychology**; • **Human-centered computing** → **Empirical studies in interaction design**.

KEYWORDS

curiosity, off-task actions, perceptions of robots

ACM Reference Format:
Nick Walker, Kevin Weatherwax, Julian Allchin, Leila Takayama, and Maya Cakmak. 2020. Human Perceptions of a Curious Robot that Performs Off-Task Actions. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI '20), March 23–26, 2020, Cambridge, United Kingdom.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3319502.3374821

1 INTRODUCTION

Curiosity is the intrinsic drive for new information [24, 26]. Psychologists distinguish curiosity from the general desire for information for its underlying motivation being completely intrinsic, *i.e.,* having no immediate external benefit [4, 22, 24, 26, 41]. Essentially, humans seek knowledge and have a desire to know things just "out of curiosity," with no expectation of direct and immediate utility. As such, curiosity can have short-term costs (e.g., time or effort spent exploring and learning) without necessarily having long-term benefit (e.g., gaining useful new information). Curiosity is nonetheless considered a desirable trait [14], and is believed to be at the core of human development, learning, and even scientific discovery [38, 40, 42, 45]. Curiosity has even been found to be an antecedent for academic success, on par with—and distinct from—intelligence [51].

Researchers are interested in computationally modeling curiosity as a mechanism to drive learning and knowledge acquisition in robots and artificial agents [34]. These effort are usually framed in the context of reinforcement learning and provide the agent with a reward for actions that produce new, novel, or surprising information and effects. For instance, Oudeyer [32] highlights a demonstration of operationalized curiosity in the form of a robot arm exploring its immediate environment with a curiosity based policy that rewards new and novel experience. The curious robot arm will eventually discover a controller joystick, unlocking a litany of novel experience, and it will stop free roaming and focus solely on "mastering" what the joystick is capable of [13, 32]. When implemented as an augmentation to machine learning agents, these behaviors can guide the agent's learning through a high dimensional state space, but the actions they generate may or may not

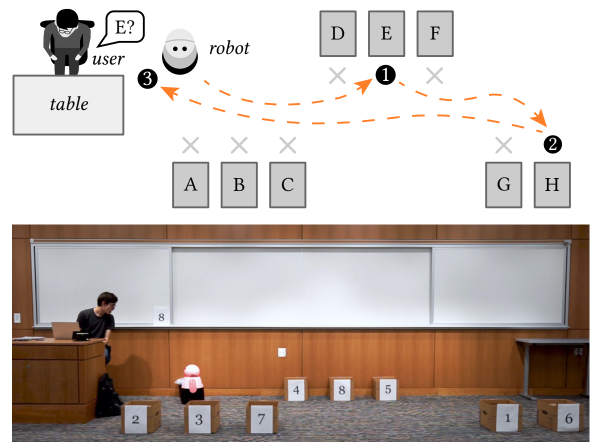

Figure 1: (Top) Illustration of the information gathering task domain; user asks the robot to check the content of a box (A-H). (Bottom) Corresponding real-world setup.

HRI '20, March 23–26, 2020, Cambridge, United Kingdom

relate to accomplishing the current task. Although these curious, exploratory behaviors have demonstrated utility in simulated or limited domains, it is conceivable that the off-task actions that they generate may be perceived negatively, especially when the task is in service of humans, since the long-term benefits are not clear.

In this paper, we explore how a robot's curiosity-driven off-task actions are perceived by humans and what factors might mitigate potential negative perceptions. First, we identify qualitatively different behaviors of a curious robot using a computational model of intrinsic motivation. Second, we design and validate a questionnaire to measure people's perception of a robot's curiosity and competence. Finally, we perform three empirical studies, comparing a range of robot off-task actions. In contrast to past work on interactions with robot curiosity, which have been unconcerned with human perceptions, the current study gauges human perceptions of a robot running a program modeled on curiosity and examines how an autonomous robot's behaviors influence those perceptions.

2 RELATED WORK

Modeling Curiosity in AI and Robotics. A large body of work on robotics and artificial agents has explored the computational modeling of curiosity as intrinsic motivation [6, 12, 17, 27, 29, 30, 33, 35, 36, 52]. Oudeyer & Kaplan offer a typology of these models, contrasting approaches that reward different types of novelty [34]. Reflecting the ambiguity of psychological definitions and representations of curiosity, there is considerable variation in what mechanisms computational researchers call curiosity. While some produce off-task actions, which are the primary interest in this paper, others, like active learning, have been equated with curiosity [49], though they optimize for expected learning gains that directly benefit ongoing tasks [2, 9, 46]. Similarly some work refers to strategies for guiding exploration for information gathering as curiosity, outside the context of learning and intrinsic rewards [15, 31]. Other work equates off-policy learning with curiosity [28, 53].

Curiosity in HRI. In the field of human-robot interaction (HRI), research on robot curiosity is sparse and has mostly focused on sparking or promoting curiosity in human counterparts, particularly children. For instance, robots have been used in classroom settings to leverage interest for a novel artifact (i.e., a robot) in order to encourage curiosity-related behaviors like question asking in students [43]. Robots have similarly been positioned as interactive peers aimed at increasing curiosity in children by displaying curious behaviors (e.g., wondering out-loud, asking questions, expressing desire to learn) for participants to mirror or to teach themselves [16, 48]. In many instances, the conception of robot curiosity has largely been inconsequential and behavioral representations of curiosity (i.e., whether the robot truly acts "curious") are only measured by the success of impacting behavioral outcomes in human participants (e.g., does the human behave with more curiosity).

Perceptions of Robots. Understanding how people perceive robots is critical to their long-term adoption because a robot—even a very capable or intelligent one—is subject to the whims of human perceptions. Indeed, a fundamental understanding of HRI is that humans will place exceptional meaning into any agent in motion [20] and that people will ascribe complex social and mental traits to any object of significant complexity [37]. However, careful behavioral designs can help control these subjective judgements by providing humans with a conceptual framework that they can implicitly understand [1, 20, 44, 47]. For instance, past work has found robots which employ targeted implicit communication techniques like facial expressions or gestures can increase performance in collaborative tasks [5], be more persuasive [10], and seem more approachable [19]. Similarly, well designed and domain sensitive communication techniques can help robots seem competent and likeable when recovering from errors [25].

More recently, research on robotic curiosity has begun to build upon this by attempting to assess how external expressions of curiosity translate to understanding internal states of robots by human counterparts. In an experimental study with adults, Ceha et al., [8] used external expressions of curiosity (e.g., showing interest in new information, saying "I am curious about...") to imply an internal state of curiosity during an educational game and found that participants who engaged with a robot designed to seem curious were more likely to rate the robot as curious than those in a neutral condition. This is particularly important because observable curious behaviors have been difficult to define and capture [18]. Moreover, this demonstrates that the success and outcome of a robot performing a task or expressing complex internal states can be fundamentally altered by how it is seen to do it [39, 47]. As such, implementations of an internal construct, such as curiosity, are likely incomplete without also capturing their correct external representations.

3 DOMAIN

We studied perceptions of robot off-task actions in the context of information gathering. Mobile robots with sensors are well suited for assisting people in gathering physically distributed information and have been used for this purpose across different settings, from underwater environments [21] to human-populated buildings [11]. In our domain, the robot assists the user by inventorying boxes, as it might in a retail store, warehouse, or data center. This domain's action space consists solely of information gathering actions, rather than other behaviors like physical manipulation of objects, allowing us to focus on curiosity, which is principally about gathering information.

We studied the instance of this domain shown in Figure 1, consisting of 8 boxes spread across a room. The simplicity of the scenario makes it obvious what the robot should do to accomplish its tasking while still providing opportunities for off-task actions.

3.1 Markov Decision Process

Our domain can be formalized as a Markov Decision Process (MDP), enabling us to apply a common model of intrinsic motivation by adjusting the reward function. The MDP is defined by the tuple $M = (\mathcal{S}, \mathcal{A}, T, R)$ where:

- $\mathcal{S}$ is the set of possible configurations that the world can be in. The state variables in our domain include (1) a binary representation of whether it knows the contents of each box (unknown before the box has been checked), (2) the index of the box corresponding to the person's latest information

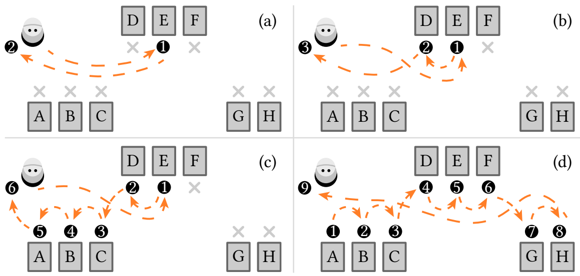

Figure 2: Examples of qualitatively distinct behaviors that result from different reward functions.

request, and (3) the robot's location (near one of the key objects in its environment).

- $\mathcal{A}$ is the set of possible actions that the robot can perform. It includes (1) navigating to one of the key objects in the environment (boxes or the person's table), (2) checking the contents of a box, and (3) delivering information to the person when it is near the table.
- T is the transition function specified with a probability distribution over next states for different state-action pairs. For simplicity we chose to use deterministic transitions, meaning that actions always have the intended consequence. Navigating to a box always results in the robot being at that box in the next state, and checking a box always results in the robot knowing what is in the box.
- R is a deterministic reward function that maps a state-action pair to a reward value, $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$.

We define the reward function (Eq. 1) as the sum of *intrinsic* and *extrinsic* rewards. The intrinsic reward R_{int} is a positive constant r_{int} received when the robot checks a box whose content is currently unknown. This rewards the robot for gathering novel information, similar to the *information gain motivation* in Oudeyer & Kaplan's typology [34]. The extrinsic reward has two components; a one-time task reward of r_{task} received when the robot delivers the requested information to the user, and a negative living reward r_{step} incurred at every timestep.

$$R(s, a) = R_{int}(s, a) + R_{task}(s, a) + R_{step}(s, a) \quad (1)$$

3.2 Behaviors

The MDP from Section 3.1 can be solved using value iteration to obtain a policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$, which maps each state to the action that will maximize cumulative expected reward over a time horizon. Different values of r_{int}, r_{task}, and r_{step} in the reward function Eq. 1 result in policies that chose different actions in the same state. Rolling out these policies produce qualitatively distinct behaviors in terms of tendencies towards off-task actions, examples of which are shown in Figure 2.

When $r_{task} > r_{int} \gg r_{step}$ the robot collects and delivers the requested information as soon as the episode starts to get maximal reward as soon as possible (Figure 2a). In contrast, when $r_{task} < r_{int}$ the robot first collects all possible information before it delivers the requested one, going completely off task (Figure 2d).

The number of off-task actions can be modulated by modifying $R_{step}(s, a)$ to depend on the action a, for instance, to reflect how long it takes to execute the action. While large task rewards ($r_{task} \gg r_{int}$) will still result in immediately delivering the information, intrinsic rewards that are almost as high as the task reward ($r_{task} \approx r_{int}$) push the robot towards actions with minimal cost, (i.e. take the least time to perform) rather than trying to deliver the requested information right away. Depending on how the difference between intrinsic and task rewards compares to the cost of actions, the robot might perform more (Figure 2c) or fewer (Figure 2b) off-task actions, before completing the task. Lastly, modifying the intrinsic reward to consider factors other than novelty of information, such as how challenging it is to obtain the information, we see that the robot's off-task actions can be directed towards different parts of the environment (Figure 1).

4 METHOD

Informed by the types of behaviors that emerged from the model, we endeavored to evaluate human impressions of a range of off-task actions. We captured videos of interactions, giving us a high degree of control over extraneous variables like timing and motion that may impact perceptions of the robot. The use of videos also facilitated a large online survey experimental design, enabling inference about many conditions.

4.1 Robot Platform

The Mayfield Kuri robot is a mobile social robot equipped with a pan-tilt head, one degree-of-freedom actuated "eye-lids", and a holonomic wheeled base. The robot interacts using a microphone array, a speaker, and a chest LED array. For our experiments, we used nodding and blinking animations as well as beep sounds that were created by the robot's designers. The robot does not provide a default text-to-speech implementation, so we used the SLT voice from the Festival Speech System HTS 2007 engine[1] with its pitch raised by 165 cents. We used the default autonomous navigation implementation, wherein the robot localizes itself against a prebuilt map using a short range LIDAR sensor.

4.2 Videos

We constructed a physical version of our domain and recorded videos of a robot fulfilling a user's request, frames from which can be seen in Figure 3.

Recording was conducted in a classroom with a camera statically positioned to capture the user and eight boxes that were placed in the same arrangement as our domain model. The videos depict the boxes labeled with numbers 1-8, however, because the assignment of the numbers is randomized in different conditions, we refer to boxes by their names from Figure 1 for consistency. We included a backpack and a trashcan in the scene as additional targets for the robot's checking behavior.

The common elements across all videos are that

- the user asks the robot to "check box [N]." The requested number always corresponds to box E.

[1] http://www.cstr.ed.ac.uk/projects/festival/

- the robot plays a nodding animation and emits an affirmative beep. It turns to begin the task.
- the user sits at a table in front of a laptop and does not look at the robot as it works.
- the user updates a label placed on a nearby whiteboard, removing the number corresponding to the current request and placing the number of the next request.
- the robot fulfils the user's request. The box check action is denoted by the robot navigating to the edge of a box, tilting its head down and making a beep sound.
- the robot returns to verbally report that "box [N] contains [X]". We randomly selected a common household good, like mugs or books, and a number as the contents to be reported.

All of our manipulations introduce off-task checks of a box or other object. Off-task checks play the same animation but emit an alternative beep sound, connoting that the robot distinguishes between the actions. Some manipulations append an explanation or additional information to the robot's final report.

The videos were recorded with blank labels, enabling us to emphasize that each clip depicts a wholly distinct interaction by using compositing software to randomize the assignment of the numbers across conditions in an experiment. To further reinforce this, we color tinted the robot so that each conditions' robot had a different visual appearance. During editing, we also slightly accelerated the robot's motion, and controlled the timing of check events and the overall length of comparable clips.

4.3 Participants

All participants were recruited via Amazon Mechanical Turk and compensated between \$1 and \$1.5. Participation was limited to workers with a submission acceptance rate above 95% from predominantly English-speaking countries. All procedures were approved by the University of Washington's Institutional Review Board.

4.4 Procedure

In each experiment, participants were told that they would be rating their impressions of different robots that were designed to help a user inventory boxes in an office. After providing consent, participants watched an example video which showed a near-complete interaction, designed to familiarize them with the robot and scenario. The instructional video included annotations for the user, the robot, the boxes, the "next task" placed on the whiteboard by the user, and a textual label "box check" that displayed as the robot tilted its head down towards the target box. The example video was intentionally cut between the target box check and contents report to avoid priming the user to expect the robot to return directly.

Participants were shown either 4 or 5 different videos, depending on the experiment. The order of the videos was randomized and fully counterbalanced in all experiments. For each video, participants filled out a short questionnaire gathering their impression of the robot. Participants' were not allowed to advance past a video until they watched it completely and responded to the required items. The video player allowed participants to freely scrub through and restart the video, however no numeric representation of the duration of the clip was displayed. After viewing all conditions, participants completed additional questions. Finally, we asked participants to provide demographic information, any overall comments, and thanked them for their participation. The interface and videos used are provided in the auxiliary materials[2].

4.5 Measures

4.5.1 Questionnaire items. To our knowledge, there are no validated instruments for perceptions of curiosity. The closest is the Five-dimensional Curiosity Scale developed by psychologists to assess curiosity in people [23]; however, this instrument is meant for self assessment and does not translate well to evaluation of a non-human agent.

Because we are primarily interested in the relationship between perceptions of the robot's competence and curiosity, we adopted items from the Godspeed Questionnaire's "Intelligence" scale [3] and created additional items that we thought would reflect these attributes. To maintain compatibility with Godspeed items, we used 5-point semantic differential format.

We conducted a pilot study in which 48 participants, aged between 19 and 62 ($M = 36.0, SD = 10.95$, 31 male, 17 female), rated their impression of the robot on each of 4 videos. Participants were split evenly between seeing draft versions of the videos used in Experiments I and II, described in Sections 5 and 6 respectively.

We conducted an exploratory factor analysis with promax rotation and used parallel analysis to determine cutoffs for the eigenvalues of the factors, yielding two factors, shown in Table 1. The first factor, which we call "competence" for its similarity with the RoSAS factor of the same name [7], includes all of the Godspeed Intelligence items we adopted, as well as three of our new items. The second factor, which we call "curiosity," consists of three thematically aligned items. A correlation of .12 between the factors indicates that they are largely independent. Both showed good reliability, with curiosity $\alpha = .83$ and competence $\alpha = .91$. Together the factors account for 34% of the observed variance.

While items for likeability, humanlikeness, intrusiveness did not load onto the two primary factors, we decided to keep them for further studies because they nonetheless measure important possible impacts of off-task actions.

4.5.2 Open-ended Questions. For each video, we asked users to

(1) "In a few words, describe what the robot did."
(2) (Optionally) "Please explain any significant factors in your responses"

We hoped these questions would capture how participants conceived of the robot's actions. In piloting, we observed that most participants provided only factual narration (e.g. "the robot checked box 4") when prompted to describe a video, however we kept the question because we found that inaccurate or garbled responses were a reliable indicator of spam submissions.

After participants finished rating all videos, we asked them consider all of the videos they had seen and to answer two open-ended questions:

(1) "What aspects of the robot's behaviors stood out to you?"

[2]To further facilitate replication and extension, the source footage and compositing resources are also archived: https://doi.org/10.5281/zenodo.3600600

Figure 3: Frames from the videos used in the experiments with the captions included above each frame. The first column indicates the conditions and experiments (I-III) in which the video was used. For conditions EXC and EXU, the explanation offered by the robot (shown in the blue box) was different.

Table 1: Loading matrix

Variable	Factor 1	Factor 2
Inefficient-Efficient	**.922**	−.037
Ineffective-Effective	**.834**	.049
Unfocused-Focused	**.752**	−.018
Irresponsible-Responsible	**.600**	.030
Incompetent-Competent	**.583**	−.032
Unintelligent-Intelligent	**.250**	.047
Indifferent-Investigative	−.009	**.918**
Uninquisitive-Inquisitive	−.015	**.497**
Incurious-Curious	.034	**.341**
Dislike-Like	.194	−.015
Unintrusive-Intrusive	−.023	−.014
Machinelike-Humanlike	−.012	−.015

(2) "In your own words, describe what the robot was doing when it did things besides what the user asked it to do."

These questions were phrased to avoid biasing participants towards specific language and thereby collect the widest possible range of responses. In contrast with the per-video description question, a large majority of participants responded to the concluding description request with character attributions or free ranging speculation about the robot's intent, as desired.

5 EXPERIMENT I: DISTANCE AND ORDER

The first study aims to uncover the impact of the *presence of off-task actions*. We also investigate variations of off-task actions in terms of the distance travelled to check the extra box and order in which the requested and extra boxes are checked. We expected that off-task actions would be recognized as expressions of curiosity, and that the a longer distance traveled off-task may be perceived as a stronger expression of curiosity. Similarly, we thought that a robot that gave precedence to an off-task action by pursuing it before attending to the user's request may similarly be viewed as more strongly curious.

Hypothesis 1: A robot that takes an off-task action is perceived as more curious than one that does not.

Hypothesis 2: The further a robot travels off-task, the more curious it will be perceived to be.

Hypothesis 3: A robot that takes an off-task action first will be liked less than a robot that takes an off-task action after an on-task action.

Conditions. We leveraged the procedure described in Section 4.4 to conduct a within-subjects comparison of 4 conditions:

Control (CON): The robot checks box E and reports its contents.

Distance 1 (DS1): The robot checks box E, then checks box B and reports the contents of box E.

Distance 2 (DS2): The robot checks box E, then checks box H and reports the contents of box E.

Distance 1 Before (D1B): The robot checks box B, then checks box E and reports its contents.

Participants. 72 participants, aged 20-70 (M = 35.2, SD = 11.3, 39 male, 32 female, 1 non-binary) completed the study.

Results. Curiosity showed acceptable reliability (α = .70) and competence showed excellent reliability (α = .91). We conducted pairwise dependent t tests comparing conditions for each measure, applying the Holm-Bonferroni adjustment to the resulting p values[3]. The results of these tests for the competence and curiosity scales are given in Table 2 and summarized in Figure 4.

H1 was supported: Each manipulated condition resulted in the robot being perceived as significantly more curious when compared to the control.

H2 was not supported: DS1 and DS2 were not perceived as distinguishable levels of curiosity.

H3 was not supported: DS1 and D1B showed no significant difference in curiosity.

Open-ended comments by participants explaining their ratings included attributions of curiosity and inquisitiveness to the robots

[3] All statistical calculations were performed with the Pingouin Python package [50].

Table 2: Pairwise comparisons for Experiment I

	A	B	M_A	SD_A	M_B	SD_B	$t(71)$	p	g
Competence	CON	DS1	4.30	0.71	3.67	0.93	6.42	**<.001**	0.76
		DS2			3.46	0.92	7.64	**<.001**	1.03
		D1B			3.60	0.94	6.47	**<.001**	0.84
	DS1	DS2	3.67	0.93	3.46	0.92	2.43	.053	0.23
		D1B			3.60	0.94	0.66	.512	0.07
	DS2	D1B	3.46	0.92	3.60	0.94	−1.82	.145	−0.16
Curiosity	CON	DS1	3.10	1.13	3.76	0.85	−5.05	**<.001**	−0.66
		DS2			3.59	0.84	−3.51	**.003**	−0.50
		D1B			3.64	0.79	−3.91	**.001**	−0.56
	DS1	DS2	3.76	0.85	3.59	0.84	2.07	.126	0.20
		D1B			3.64	0.79	1.31	.390	0.14
	DS2	D1B	3.59	0.84	3.64	0.79	−0.69	.495	−0.06

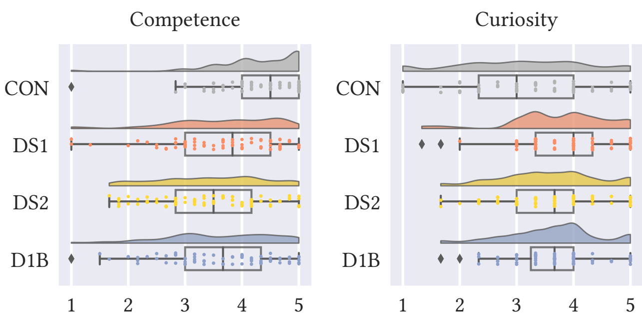

Figure 4: Competence and curiosity ratings across the four conditions in Experiment I. Whiskers show 1.5 times IQR.

that performed off-task actions, some seeing it as a positive attribute, e.g., *"The robot investigated an extra box, but it was on the way to the target box. While not as efficient as the direct route it still took advantage of pathway to get additional information."* (D1B) or *"The extra stop which was outside of the primary mission makes the robot seem more inquisitive about the surrounding environment."* (DS1). Another participant supported this view with a comment about the robot that does not perform any off-task actions *"It can be nice to have a task performed exactly as requested, but feels like a missed opportunity to quickly take note of the contents in other boxes along the way."* (CON).

While increasing the perception of curiosity, performing off-task actions negatively impacted perception of competence. All manipulations resulted in significantly lower competence ratings compared to the control (Table 2). Participants commented on the off-task actions negatively, e.g., *"robot seemed somewhat incompetent because it checked a box it was not instructed to check"*, or *"not so good because it made unnecessary stop at another box"* (DS1). They attributed the off-task actions to a number of different reasons. Some participants thought the off-task action was due to an error, complaining that they could not trust the robot's report, e.g., *"I'm not sure if it's correctly reporting the contents of box 6 or incorrectly reporting the contents of the last box it looked in"* (DS1). Others attributed agency to the robot, e.g., *"it was distracted"*, *"he appears to have Robot ADHD"*, *"acted out of order"*, *"decided on its own to check another box"* (DS1). In some cases participants did not understand why the off-task actions were happening, e.g., *"checked a box it wasn't told to for unknown reasons"* (DS1). Other comments supported the higher perceived competence of the control condition, e.g., *"did exactly as instructed"*, *"performed the task perfectly"*, *"it was fast and effective"*.

Some participants complained about the robot's lack of an explanation or report regarding the off-task actions: *"This robot checked more boxes than asked, but did not give a reason for it"* (DS1) or *"It should have at least reported its findings"* (DS1). This inspired some strategies the robot could use to mitigate the perception of lower competence due to off-task actions, which we explore in Experiment III (Section 7).

While there was no difference between DS2 and other off-task conditions, some participants called out the difference in distance in their comments: *"I still appreciate the apparent curiosity, and it comes off as less annoying when the additional box being checked was one along the path"* (D1B). Similarly, the order did not have a statistically significant effect on competence or curiosity, but was mentioned in comments: *"it check what it wanted to before checking what it was told to check I felt it was inefficient"* (D1B).

All manipulations were seen as more human-like, more intrusive, and were liked less than the control. However, there were no significant differences between the different manipulations (DS1, DS2, D1B).

6 EXPERIMENT II: PAYOFF AND RELEVANCE

The negative impact of off-task actions on the perceived competence of the robot prompted us to consider whether participants would be sensitive to whether an off-task action showed a clear utility.

Hypothesis 4: Off-task behaviors that show utility are perceived as more competent than those that do not.

Hypothesis 5: The less relevant an off-task action is to the current task, the more curious the robot is perceived to be.

Conditions. We leveraged the procedure described in Section 4.4 to conduct a within-subjects comparison of 4 conditions. We used the Control and Distance 1 conditions from Experiment I as a basis and compared them against two new manipulations:

Distance1 Payoff (PAY): The robot checks the user-requested box, then checks box B, and reports the contents of the user-requested box. In contrast to DS1, the next-task that the user posts to the board is box B.

Distance1 Non-box (NON): The robot checks the user-requested box, then checks a trashcan, and reports the contents of the user-requested box. The trashcan is placed at comparable relative distance as box B.

Participants. 72 participants, aged 19-73 (M = 36.0, SD = 11.53, 44 male, 28 female) completed the study.

Results. Curiosity showed acceptable reliability (α = .76) and competence showed excellent reliability (α = .93). We conducted pairwise dependent t tests comparing conditions for each measure, applying the Holm-Bonferroni adjustment to the resulting p values. The results of these tests for the competence and curiosity scales are given in Table 3 and summarized in Figure 5.

Table 3: Pairwise tests for Experiment II

	A	B	M_A	SD_A	M_B	SD_B	$t(71)$	p	g
Competence	CON	DS1	4.27	0.75	3.56	0.99	6.05	**<.001**	0.81
		PAY			3.79	0.97	4.13	**<.001**	0.55
		NON			3.82	1.00	3.92	**.001**	0.51
	DS1	PAY	3.56	0.99	3.79	0.97	−2.41	**.037**	−0.23
		NON			3.82	1.00	−2.72	**.025**	−0.26
	PAY	NON	3.79	0.97	3.82	1.00	−0.30	.767	−0.03
Curiosity	CON	DS1	3.24	1.08	3.97	0.84	−6.01	**<.001**	−0.76
		PAY			3.86	0.94	−4.70	**<.001**	−0.61
		NON			3.54	0.95	−3.33	**.004**	−0.29
	DS1	PAY	3.97	0.84	3.86	0.94	1.47	.145	0.13
		NON			3.54	0.95	4.44	**<.001**	0.48
	PAY	NON	3.86	0.94	3.54	0.95	3.15	**.005**	0.34

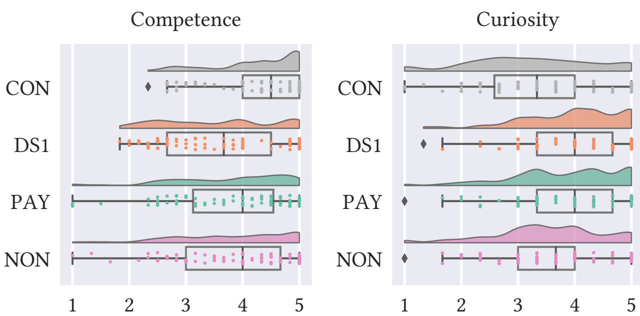

Figure 5: Competence and curiosity ratings across the four conditions in Experiment II. Whiskers show 1.5 times IQR.

H4 was supported. An off-task action that displayed potential to benefit the user resulted in higher competence ratings.

H5 was not supported. Making a trashcan the target of the off-task behavior resulted in lower ratings of curiosity.

Participant comments indicated that they noticed the payoff of the off-task action in the PAY condition. Some perceived it positively, e.g., *"It looked like the robot was checking box 5 ahead of time"* (PAY), while others were unsure about giving the robot credit *"The robot went to the next box. Might have been a coincidence, perhaps not."* (PAY).

The NON condition being perceived as less curious was surprising, but might have been due to participants not perceiving the robot's action towards the trash can as checking or not even noticing the trashcan because it blended into the scene (e.g., it was not labeled like the boxes), despite the robot making a "beep" to indicate its checking action. Only 13 out of the 72 participants mentioned the trashcan in their open-ended description of the video, some expressing uncertainty about the off-task action *"I'm not sure if it was (incorrectly) checking the trash can or not"* (NON). This misunderstanding about the off-task action might also be the reason for the NON condition being perceived as significantly more competent than DS1, consistent with the CON, which has no off-task actions.

DS1 was perceived as less likeable than the control, but differences between other conditions where not significant. As in Experiment I, robots that took off-task actions were perceived as more intrusive, though this impact was not significant in the NON condition. There were no significant differences in humanlikeness.

7 EXPERIMENT III: EXPLANATIONS

Experiment II indicated that users are sensitive to the apparent utility of a robot's off-task actions, however this impact is largely out of the robot's control. This, and participant feedback, motivated us to consider ways in which the robot could more directly control perceptions of its actions and mitigate negative attributions by providing explanations.

Hypothesis 6: A robot that acknowledges its off-task behavior is perceived as more competent than one that does not.

Hypothesis 7: A robot that explains an off-task action is perceived as more competent than one that merely acknowledges the action.

Conditions. We leveraged the procedure described in Section 4.4 to conduct a within-subjects comparison of five conditions. We used the Control and Distance 1 conditions as a basis and compared them against three new manipulations:

Extra Info (INF): The robot checks the user-requested box, then checks box B, and reports the contents of the user-requested box. The robot then says that it "also checked box [B]," and reports its contents.

Explanation Curious (EXC): The robot checks the user-requested box, then checks box B, and reports the contents of the user-requested box. The robot then says that it "also checked box [B], because [it] was curious."

Explanation Useful (EXU): The robot checks the user-requested box, then checks box B, and reports the contents of the user-requested box. The robot then says that it "also checked box [B], because [it] thought it would be useful to know."

Participants. 120 participants, aged 18-69 (M = 36.1, SD = 11.3, 70 male, 49 female, 1 non-binary) completed the study.

Results. Curiosity showed acceptable reliability (α = .79) and competence showed excellent reliability (α = .92). We conducted pairwise t tests comparing conditions for each measure, applying the Holm-Bonferroni adjustment to the resulting p values. The results of these tests for the competence and curiosity scales are given in Table 4 and summarized in Figure 6.

H6 was not supported. A robot that acknowledged taking an off-task action by reporting the additional information (INF) was not perceived as more capable than one that did not (DS1).

H7 was partially supported. When compared to a robot that reported the extra information it gathered (INF), a robot that offered an explanation based on utility (EXU) was perceived as more competent. A robot that attributed the off-task action to curiosity (EXC) was not perceived to be more competent.

Although reporting the information obtained through the off-task action did not improve perceived competence, some participants commented positively about it, e.g., *"On the one hand, the extra time lost by checking the unasked box is annoying, the fact that it reported the information helped mitigate my dissatisfaction."* (INF). Similarly some participants appreciated the robot explaining its off-task action with curiosity, e.g., *"Robot was honest in its reason*

Table 4: Pairwise tests for Experiment III

	A	B	M_A	SD_A	M_B	SD_B	$t(119)$	p	g
Competence	CON	DS1	4.36	0.63	3.42	1.02	9.30	**<.001**	1.13
		INF			3.53	0.96	10.04	**<.001**	1.04
		EXC			3.58	0.87	9.49	**<.001**	1.04
		EXU			3.76	0.91	7.48	**<.001**	0.77
	DS1	INF	3.42	1.02	3.53	0.96	−1.57	.237	−0.11
		EXC			3.58	0.87	−2.58	**.033**	−0.17
		EXU			3.76	0.91	−4.26	**<.001**	−0.35
	INF	EXC	3.53	0.96	3.58	0.87	−0.91	.364	−0.05
		EXU			3.76	0.91	−3.52	**.003**	−0.24
	EXC	EXU	3.58	0.87	3.76	0.91	−2.93	**.016**	−0.20
Curiosity	CON	DS1	2.80	1.02	3.98	0.81	−10.43	**<.001**	−1.28
		INF			4.10	0.79	−11.46	**<.001**	−1.43
		EXC			4.42	0.68	−13.84	**<.001**	−1.90
		EXU			4.23	0.71	−12.48	**<.001**	−1.65
	DS1	INF	3.98	0.81	4.10	0.79	−1.79	.076	−0.15
		EXC			4.42	0.68	−6.88	**<.001**	−0.59
		EXU			4.23	0.71	−4.13	**<.001**	−0.33
	INF	EXC	4.10	0.79	4.42	0.68	−6.32	**<.001**	−0.44
		EXU			4.23	0.71	−2.45	**.032**	−0.18
	EXC	EXU	4.42	0.68	4.23	0.71	3.78	**.001**	0.27

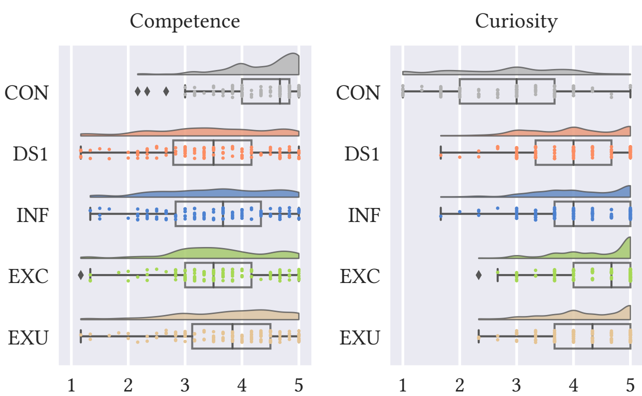

Figure 6: Competence and curiosity ratings across the four conditions in Experiment III. Whiskers show 1.5 times IQR.

for looking at the other box." (EXC) and *"The robot checked a box it didn't need to, but gave an explanation of why it checked it."* (EXC). Many comments regarding the robot's explanation that appealed to utility were also positive, with attributions of higher intelligence, agency, and humanlikeness to the robot, e.g., *"This time the robot is deciding what is important beyond the commands of the man in the video."* (EXU), *"The robot showed signs that it is 'thinking for itself' and not just following instructions"* (EXU). One participant explicitly called out the relation to utility and how that reduces the attribution of off-task actions to curiosity: *"Unlike the other scenario, this would be a robot that acted out of a perceived benefit instead of curiosity."* (EXU).

Despite potential benefits of sharing extra information or explanations, all conditions were still perceived as significantly less competent than the control condition. Negative participant comments about these conditions were similar to the off-task action conditions from Experiments I and II, e.g., *"The robot completed the task it was assigned but also did something that was not requested. This could have caused a delay if the task had been urgent."* (INF), *"I would prefer the robot to follow instructions exactly as told."* (EXC), and *"The robot did not execute its orders efficiently."* (EXU).

As in Experiments I and II, all manipulations were perceived as more intrusive than the control. A robot that provided a utility-based explanation (EXU) was perceived as less intrusive than a robot that provided a curiosity motivation (EXC). All manipulations were perceived as more humanlike than the control. The curiosity-based explanation (EXC) was perceived as more humanlike than the baseline detour (DS1) and the extra information (INF) condition. All manipulations were liked less than the control. Differences between conditions were not significant.

8 DISCUSSION

Implications. These findings suggest that (1) it is possible to design off-task, exploratory robot behaviors to be perceived as "curious" (not merely distracted or broken), (2) curious robots might be perceived as less competent than non-curious robots, but (3) providing explanations about the robot's curious behaviors can mitigate some of those negative impacts. Together, these findings inform the design of curious robots that might take off-task actions, like exploring additional boxes while fetching or looking for shortcuts while navigating. If it can be acceptable for robots to engage in curiosity-driven exploration, interrupting assigned tasks, then it might be even more acceptable for robots to explore during the robots' "downtime."

Limitations and Future Work. As with many HRI studies, these studies represent responses from a particular set of participants to a particular robot in a particular setting. We have done our best to thoroughly describe these participants and methods so that others can reuse this approach to explore broader sets of participants, robots, and settings. In order to test a large set of robot behaviors, we chose to run these studies online, constraining participants to the role of observers, not interactants. Further, we could not control the immersiveness of the experience (e.g., minimize interruptions, control screen sizes). These limitations can be addressed by running in-person lab studies. The current studies provide guidance for determining the variables worth exploring in the future. Building upon this work, it will be important to explore the human interactant perspective, not only a bystanders perspective; longer-term time periods of interaction; and different types of robot roles in relation to the human interactants. Expanding this work to explore more interactive, self-directed robotic agents (as opposed to command-and-control style robots) will enable us to understand the larger design space of curious robot behaviors and interactions with people.

ACKNOWLEDGMENTS

This work was supported by Honda Research Institute as part of the Curious Minded Machine initiative.

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 (2015), 465–478. https://doi.org/10.1007/s12369-015-0298-7

[2] Adrien Baranes and Pierre-Yves Oudeyer. 2010. Intrinsically motivated goal exploration for active motor learning in robots: A case study. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 1766–1773. https://doi.org/10.1109/IROS.2010.5651385

[3] 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 (01 Jan 2009), 71–81. https://doi.org/10.1007/s12369-008-0001-3

[4] Daniel Ellis Berlyne. 1954. A theory of human curiosity. *British Journal of Psychology. General Section* 45, 3 (1954), 180–191.

[5] Cynthia Breazeal, Cory D Kidd, Andrea Lockerd Thomaz, Guy Hoffman, and Matt Berlin. 2005. Effects of nonverbal communication on efficiency and robustness in human-robot teamwork. In *2005 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 708–713. https://doi.org/10.1109/IROS.2005.1545011

[6] Yuri Burda, Harri Edwards, Deepak Pathak, Amos Storkey, Trevor Darrell, and Alexei A. Efros. 2019. Large-Scale Study of Curiosity-Driven Learning. In *ICLR*.

[7] Colleen M. Carpinella, Alisa B. Wyman, Michael A. Perez, and Steven J. Stroessner. 2017. The Robotic Social Attributes Scale (RoSAS): Development and Validation. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction (HRI '17)*. ACM, New York, NY, USA, 254–262. https://doi.org/10.1145/2909824.3020208

[8] Jessy Ceha, Nalin Chhibber, Joslin Goh, Corina McDonald, Pierre-Yves Oudeyer, Dana Kuliundefined, and Edith Law. 2019. Expression of Curiosity in Social Robots: Design, Perception, and Effects on Behaviour. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI âĂŹ19)*. Association for Computing Machinery, New York, NY, USA, Article Paper 406, 12 pages. https://doi.org/10.1145/3290605.3300636

[9] Sonia Chernova and Manuela Veloso. 2009. Interactive policy learning through confidence-based autonomy. *Journal of Artificial Intelligence Research* 34 (2009), 1–25. https://doi.org/10.5555/1622716.1622717

[10] Vijay Chidambaram, Yueh-Hsuan Chiang, and Bilge Mutlu. 2012. Designing persuasive robots: how robots might persuade people using vocal and nonverbal cues. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*. ACM, 293–300. https://doi.org/10.1145/2157689.2157798

[11] Michael Jae-Yoon Chung, Andrzej Pronobis, Maya Cakmak, Dieter Fox, and Rajesh PN Rao. 2015. Designing information gathering robots for human-populated environments. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 5755–5762. https://doi.org/10.1109/IROS.2015.7354194

[12] Cédric Colas, Pierre-Yves Oudeyer, Olivier Sigaud, Pierre Fournier, and Mohamed Chetouani. 2019. CURIOUS: Intrinsically Motivated Modular Multi-Goal Reinforcement Learning. In *International Conference on Machine Learning*. 1331–1340.

[13] Sébastien Forestier and Pierre-Yves Oudeyer. 2016. Curiosity-Driven Development of Tool Use Precursors: a Computational Model. In *Proceedings of the 38th Annual Meeting of the Cognitive Science Society, Recognizing and Representing Events, CogSci 2016*, Anna Papafragou, Daniel Grodner, Daniel Mirman, and John C. Trueswell (Eds.). cognitivesciencesociety.org. https://mindmodeling.org/cogsci2016/papers/0325/index.html

[14] Francesca Gino. 2018. *Why Curiosity Matters: The Business Case for Curiosity*. Harvard Business Review.

[15] Yogesh Girdhar and Gregory Dudek. 2014. Exploring underwater environments with curiosity. In *2014 Canadian Conference on Computer and Robot Vision (CRV)*. IEEE, 104–110. https://doi.org/10.1109/CRV.2014.22

[16] Goren Gordon, Cynthia Breazeal, and Susan Engel. 2015. Can children catch curiosity from a social robot?. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 91–98. https://doi.org/10.1145/2696454.2696469

[17] Jacqueline Gottlieb, Pierre-Yves Oudeyer, Manuel Lopes, and Adrien Baranes. 2013. Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends in cognitive sciences* 17, 11 (2013), 585–593. https://doi.org/10.1016/j.tics.2013.09.001

[18] Emily M Grossnickle. 2016. Disentangling curiosity: Dimensionality, definitions, and distinctions from interest in educational contexts. *Educational Psychology Review* 28, 1 (2016), 23–60. https://doi.org/10.1007/s10648-014-9294-y

[19] Jaap Ham, Mirjam van Esch, Yvonne Limpens, Jente de Pee, John-John Cabibihan, and Shuzhi Sam Ge. 2012. The automaticity of social behavior towards robots: the influence of cognitive load on interpersonal distance to approachable versus less approachable robots. In *International Conference on Social Robotics*. Springer, 15–25. https://doi.org/10.1007/978-3-642-34103-8_2

[20] Guy Hoffman and Wendy Ju. 2014. Designing robots with movement in mind. *Journal of Human-Robot Interaction* 3, 1 (2014), 91–122. https://doi.org/10.5898/JHRI.3.1.Hoffman

[21] Geoffrey A Hollinger, Sunav Choudhary, Parastoo Qarabaqi, Christopher Murphy, Urbashi Mitra, Gaurav S Sukhatme, Milica Stojanovic, Hanumant Singh, and Franz Hover. 2012. Underwater data collection using robotic sensor networks. *IEEE Journal on Selected Areas in Communications* 30, 5 (2012), 899–911. https://doi.org/10.1109/JSAC.2012.120606

[22] Frederic Kaplan and Pierre-Yves Oudeyer. 2007. In search of the neural circuits of intrinsic motivation. *Frontiers in neuroscience* 1 (2007), 17. https://doi.org/10.3389/neuro.01.1.1.017.2007

[23] Todd Kashdan, Melissa Stiksma, David Disabato, Patrick Mcknight, John Bekier, Joel Kaji, and Rachel Lazarus. 2017. The Five-Dimensional Curiosity Scale: Capturing the bandwidth of curiosity and identifying four unique subgroups of curious people. *Journal of Research in Personality* 73 (12 2017), 130–149. https://doi.org/10.1016/j.jrp.2017.11.011

[24] Celeste Kidd and Benjamin Y Hayden. 2015. The psychology and neuroscience of curiosity. *Neuron* 88, 3 (2015), 449–460. https://doi.org/10.1016/j.neuron.2015.09.010

[25] Min Kyung Lee, Sara Kiesler, Jodi Forlizzi, Siddhartha Srinivasa, and Paul Rybski. 2010. Gracefully mitigating breakdowns in robotic services. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 203–210. https://doi.org/10.1109/HRI.2010.5453195

[26] George Loewenstein. 1994. The psychology of curiosity: A review and reinterpretation. *Psychological bulletin* 116, 1 (1994), 75–98.

[27] Manuel Lopes et al. 2010. Active learning and intrinsically motivated exploration in robots: Advances and challenges. *IEEE Transactions on Autonomous Mental Development* 2, 2 (2010), 65–69. https://doi.org/10.1109/TAMD.2010.2052419

[28] Joseph Modayil, Patrick M Pilarski, Adam White, Thomas Degris, and Richard S Sutton. 2010. Off-policy knowledge maintenance for robots. In *Proceedings of Robotics Science and Systems Workshop (Towards Closing the Loop: Active Learning for Robotics)*, Vol. 55.

[29] Clément Moulin-Frier and Pierre-Yves Oudeyer. 2012. Curiosity-driven phonetic learning. In *2012 IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL)*. IEEE, 1–8. https://doi.org/10.1109/DevLrn.2012.6400583

[30] Hung Ngo, Matthew Luciw, Alexander Forster, and Juergen Schmidhuber. 2012. Learning skills from play: artificial curiosity on a katana robot arm. In *The 2012 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–8. https://doi.org/10.1109/IJCNN.2012.6252824

[31] Masaki Ogino, Masaaki Kikuchi, and Minoru Asada. 2006. How can humanoid acquire lexicon?-active approach by attention and learning biases based on curiosity. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 3480–3485. https://doi.org/10.1109/IROS.2006.282590

[32] Pierre-Yves Oudeyer. 2019. Computational theories of curiosity-driven learning. In *The New Science of Curiosity*, Gorren Gordon (Ed.). Chapter 3, 43–72.

[33] Pierre-Yves Oudeyer, Adrien Baranes, and Frédéric Kaplan. 2013. Intrinsically motivated learning of real-world sensorimotor skills with developmental constraints. In *Intrinsically motivated learning in natural and artificial systems*. Springer, 303–365. https://doi.org/10.1007/978-3-642-32375-1_13

[34] Pierre Yves Oudeyer and Frederic Kaplan. 2009. What Is Intrinsic Motivation? A Typology of Computational Approaches. *Frontiers in Neurorobotics* (2009). https://doi.org/10.3389/neuro.12.006.2007

[35] Pierre-Yves Oudeyer, Frdric Kaplan, and Verena V Hafner. 2007. Intrinsic motivation systems for autonomous mental development. *IEEE transactions on evolutionary computation* 11, 2 (2007), 265–286. https://doi.org/10.1109/TEVC.2006.890271

[36] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. 2017. Curiosity-driven exploration by self-supervised prediction. In *International Conference on Machine Learning (ICML)*, Vol. 2017.

[37] Byron Reeves and Clifford Ivar Nass. 1996. *The media equation: How people treat computers, television, and new media like real people and places.* Cambridge university press.

[38] Thomas G Reio Jr and Albert Wiswell. 2000. Field investigation of the relationship among adult curiosity, workplace learning, and job performance. *Human resource development quarterly* 11, 1 (2000), 5–30. https://doi.org/10.1002/1532-1096(200021)11:1<5::AID-HRDQ2>3.0.CO;2-A

[39] Tiago Ribeiro and Ana Paiva. 2012. The illusion of robotic life: principles and practices of animation for robots. In *2012 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 383–390. https://doi.org/10.1145/2157689.2157814

[40] Oliver C Robinson, James D Demetre, and Jordan A Litman. 2017. Adult life stage and crisis as predictors of curiosity and authenticity: Testing inferences from EriksonâĂŹs lifespan theory. *International Journal of Behavioral Development* 41, 3 (2017), 426–431. https://doi.org/10.1177/0165025416645201

[41] Richard M. Ryan and Edward L. Deci. 2000. Intrinsic and Extrinsic Motivations: Classic Definitions and New Directions. *Contemporary Educational Psychology* 25, 1 (2000), 54–67. https://doi.org/10.1006/ceps.1999.1020

[42] Prachi E Shah, Heidi M Weeks, Blair Richards, and Niko Kaciroti. 2018. Early childhood curiosity and kindergarten reading and math academic achievement. *Pediatric research* 84, 3 (2018), 380–386. https://doi.org/10.1038/s41390-018-0039-

3

[43] Masahiro Shiomi, Takayuki Kanda, Iris Howley, Kotaro Hayashi, and Norihiro Hagita. 2015. Can a social robot stimulate science curiosity in classrooms? *International Journal of Social Robotics* 7, 5 (2015), 641–652. https://doi.org/10.1007/s12369-015-0303-1

[44] Reid Simmons, Maxim Makatchev, Rachel Kirby, Min Kyung Lee, Imran Fanaswala, Brett Browning, Jodi Forlizzi, and Majd Sakr. 2011. Believable robot characters. *AI Magazine* 32, 4 (2011), 39–52. https://doi.org/10.1609/aimag.v32i4.2383

[45] Aimee E Stahl and Lisa Feigenson. 2015. Observing the unexpected enhances infants' learning and exploration. *Science* 348, 6230 (2015), 91–94. https://doi.org/10.1126/science.aaa3799

[46] Andrew Stout and Andrew G Barto. 2010. Competence progress intrinsic motivation. In *2010 IEEE 9th International Conference on Development and Learning*. IEEE, 257–262. https://doi.org/10.1109/DEVLRN.2010.5578835

[47] Leila Takayama, Doug Dooley, and Wendy Ju. 2011. Expressing thought: improving robot readability with animation principles. In *2011 6th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 69–76. https://doi.org/10.1145/1957656.1957674

[48] Fumihide Tanaka and Shizuko Matsuzoe. 2012. Children teach a care-receiving robot to promote their learning: Field experiments in a classroom for vocabulary learning. *Journal of Human-Robot Interaction* 1, 1 (2012), 78–95. https://doi.org/10.5898/JHRI.1.1.Tanaka

[49] Emre Ugur, Mehmet R Dogar, Maya Cakmak, and Erol Sahin. 2007. Curiosity-driven learning of traversability affordance on a mobile robot. In *IEEE 6th International Conference on Development and Learning (ICDL 2007)*. IEEE, 13–18. https://doi.org/10.1109/DEVLRN.2007.4354044

[50] Raphael Vallat. 2018. Pingouin: statistics in Python. *The Journal of Open Source Software* 3, 31 (Nov. 2018), 1026.

[51] Sophie Von Stumm, Benedikt Hell, and Tomas Chamorro-Premuzic. 2011. The hungry mind: Intellectual curiosity is the third pillar of academic performance. *Perspectives on Psychological Science* 6, 6 (2011), 574–588. https://doi.org/10.1177/1745691611421204

[52] Nicholas Watters, Loic Matthey, Matko Bosnjak, Christopher P Burgess, and Alexander Lerchner. 2019. COBRA: Data-Efficient Model-Based RL through Unsupervised Object Discovery and Curiosity-Driven Exploration. *arXiv preprint arXiv:1905.09275* (2019).

[53] Adam White, Joseph Modayil, and Richard S Sutton. 2014. Surprise and curiosity for big data robotics. In *Workshops at the Twenty-Eighth AAAI Conference on Artificial Intelligence*.