International Journal of Humanoid Robotics
Vol. 3, No. 3 (2006) 1–20

World Scientific
www.worldscientific.com

YOUNG CHILDREN'S UNDERSTANDING OF ANIMACY AND ENTERTAINMENT ROBOTS

SANDRA Y. OKITA[*,†,‡] and DANIEL L. SCHWARTZ[*,§,¶]

[*]*Cognition and Technology Laboratory,*
Psychological Studies in Education, Stanford University,
485 Lasuen Mall, Stanford, CA 94305-3096, USA
[‡]*yuudra@stanford.edu*
[¶]*danls@stanford.edu*

[†]*Psychological Studies in Education, School of Education, Stanford University,*
485 Lasuen Mall, Stanford, CA 94305-3096, USA

[§]*Learning Science, Technology & Design, School of Education, Stanford University,*
485 Lasuen Mall, Stanford, CA 94305-3096, USA

Received 4 November 2005

Complex interactions, biologically-inspired features and intelligence are increasingly seen in entertainment robots. Do these features affect how children interpret robots? Children have "animistic intuitions" that they use to attribute intelligence, biology, and agency to living things. Two studies explore whether young children also apply animistic intuitions to robotic animals, and whether attributes vary by the child's age, robot behavior and appearance. A total of ninety-three three- to five-year-olds participated in two experiments. They observed or interacted with robots that exhibited different behaviors and levels of responsiveness to their environment. They then answered simple questions that probed their attributions of biology, intelligence, and agency. The results indicated that regardless of the robots' look and behavior, younger children over-generalized their animistic intuitions about real animals and older children attributed some animistic qualities but not others. One implication is that young children's criteria and attributions do not depend on robot features that are important for older children and adults. Another implication is that children do not have a theory of aliveness, and they develop the category of robot slowly and piecemeal as they learn discrete facts about how technology differs from living things.

Keywords: Young children; human–robot interaction; robotic pets; social responses to technology; cognitive psychology; conceptual development; naïve theories.

1. Introduction

Entertainment robots are rapidly increasing the degree to which they can mimic the behavior of living beings. Robotic dogs for example, can recognize their name, they can spot a ball and approach it, and they can dance to music. If we set aside the goal of improving robot technology, a central assumption of entertainment robots appears to be that their increased realism will have important practical benefits. Many hope to see robots as social companions for elderly populations and for those with special needs.[1–4] Entertainment robots are entering the home, and more biologically inspired robots are replicating live animals.[5] These new applications raise the question of whether realism makes a practical difference. Robots are boundary objects that include some qualities of living things, but not others. It has become an important question whether realistic features make a difference in the feelings and behaviors they elicit from people[6,7] and how people integrate robots into their everyday experience.[8]

Entertainment robots are also being designed for children in the long-term hope that biologically-inspired technologies may eventually serve as pets and become sources of comfort and learning. Children, however, differ from adults in many ways, and relatively little research exists on how children interpret these "boundary objects"[9] or how such advanced technologies can assist children in their learning and development.[10,11] The purpose of the current work is to explore what young children think about entertainment robots. This has some practical importance, because the findings can tell us whether it is possible to use advanced technologies to get children to treat robots in similar ways to live pets. It also has some theoretical importance, because it can clarify how children come to understand artifacts that stretch the limits of their existing beliefs about what it means to be alive.

2. Research on Children's Attributions of Animate

In the following two studies, children evaluated the animate qualities of robotic dogs. We wanted to find out if young children confer robotic dogs with the properties of living dogs. To know whether children treat robotic dogs as living, it is necessary to understand how children view living things generally. To adults, living dogs have a cluster of animate properties that include relatively intelligent behaviors (e.g. locate a bone), biological properties (e.g. grow), plus the agency to take independent initiative (e.g. run away). Young children also have "animistic intuitions" that they use to attribute intelligence, biology, and agency. For example, research has shown that infants are quite precocious at distinguishing biological and non-biological movements[12] and can differentiate goal-directed behaviors from random movements.[13,14] It is somewhat controversial, however, how these different intuitions become integrated into a concept of animate or living. Carey,[15] for example, argues that children do not achieve an integrated concept of alive until middle school. In contrast, Inagaki and Hatano[16] found that five-year-old children make the living/nonliving distinction and believe that animals and plants are both alive. The authors proposed that children have an intuitive theory of "vital powers," whereby living things need food to sustain their powers of growth.

In general, there have been two stories for how children's concept of animate develops. The "theory" story proposes that children have a coherent body of knowledge with some generality and that these bodies of knowledge evolve with maturation. This approach views the child as a theory builder, and it argues that there are concepts that cannot be learned, and therefore, they must be innate (e.g. the very concept of a goal; the belief that other people have minds). Gopnick and Meltzoff,[17] for example, propose that children are equipped with innate theories that they start revising as early as birth. Children's theories mature as new theories replace or transform old ones.[18] Carey,[15] for example, argues that children's concept of animal evolves from a theory based on behavior to a theory based on biology.

The second story proposes that children's acquisition of the concept of animate is piecemeal.[19] In this approach, children's ideas do not start with the coherent texture of a theory. Instead, the children have pockets of poorly integrated facts and beliefs. diSessa,[20] for example, proposes that people develop more coherent physical knowledge by sorting through their "knowledge in pieces" and progressively selecting the pieces that explain more of the facts.

Much of this debate has taken place in three domains; intelligence, biology, and agency. A typical research paradigm involves showing children several different objects that may perform different actions. The researchers then ask the children probe questions to see what animate properties the children are willing to attribute to the object. Gelman and Gottfried,[21] for example, wanted to see if children make the critical agency distinction between self-initiated movements versus externally caused movements. Preschool children viewed animals and artifacts (e.g. a lizard, a wind up toy) that were transported by a human hand. Children attributed the movement to a person when the object was an artifact. In contrast, children said the animal's movement was self-generated, even though they saw a hand move the animal. Even three-year-old children make attributions in ways that indicate that they believe that animate things have the agency to move on their own.[22]

We adopted a similar research paradigm. We showed children different robots that behaved in different ways. We then asked them probe questions about intelligence, biology, and agency. Robots provide an interesting variation on the usual paradigm, because they imitate features of animacy by design, but they are not animate. By examining how children develop a new concept for robot, we may be able to learn something about the "theory" versus "piecemeal" nature of their knowledge of animacy.

3. Predictions about How Attributions of Animacy Develop

Our hypothesis is that young children do not have a coherent theory of animacy. Rather, the properties of animacy are loosely coupled. One way to demonstrate this hypothesis is to show that children will attribute different aspects of animacy at different frequencies. For example, children may attribute a high degree of intelligence and biology, but not agency. Or, they may attribute some biological properties but not others. Another way to demonstrate this hypothesis is to show that as children get older, they reject some aspects of animacy for robots but not other ones. In other words, they do not have a theory in the sense that their theory can be falsified if one of the necessary properties is not present. A third way to demonstrate that children lack a theory in any meaningful sense is to show that their "theoretical attributions" are not tied to data. For example, imagine a robot that shows extremely intelligent behavior by spotting a ball and then shows high agency by walking towards the ball and kicking it. Will this affect children's attributions of intelligence and agency more than their attributions of biology? Or, is it the case that children do not have a theory that relates to evidence so much as an associative network of facts that tend to activate one another?

To reiterate our position, we assume that children learn to conceptually distinguish robotic animals from living animals by developing a new category of robot (or mechanical object) in piecemeal fashion. Our proposal is that "animate" is not a monolithic category, and therefore, the child does not learn in a single moment of insight that a robotic dog is not a living dog. Rather, we propose that children learn to exclude specific animate properties from the category of robot. For example, children may learn that things with hard plastic shells do not have certain animate properties (e.g. they do not grow). Because this is a discrete fact rather than necessary feature of any coherent body of knowledge, children are likely to maintain their other animistic attributions.

4. Study 1: Attribution Based on Watching

In the first study, children from three- to five-years-old watched three robots. One robot danced, another robot spotted a ball and kicked it, and a third robot did nothing. Afterwards we asked them the animacy questions shown in Table 1. We coded which robots the children pointed to as having the property or could perform the action mentioned in the probe question.

4.1. *Method*

Participants. Thirty-two children from a university day-care program participated. The children comprised three age groups: young ($n = 10$, $M = 43.5$ months, range 35–44 months), middle ($n = 11$, $M = 52.5$ months, range 45–52 months), and old ($n = 11$, $M = 57.8$ months, range = 53–66 months).

Table 1. Questions for intelligence, biology and agency.

Questions: Do you think any of the dogs. . . ? Which ones?	
Intelligence	1.. . . can tell between a real and pretend bone?
	2.. . . will remember me when I come back tomorrow?
	3.. . . will know it's time to go for a walk if I grab a leash?
Biology	1. . . . get hungry?
	2.. . . ever grow?
	3.. . . have a heart?
Agency	1. . . would try to wake you up in a fire?
	2. . . would jump onto the forbidden couch?
	3. . . would be able to do anything without the remote control?

AIBO ERS-210 AIBO ERS-220A AIBO ERS-311

Fig. 1. Children in Experiment 1 watched robots perform different actions.

Design. There were three robotic dogs as shown in Fig. 1 (Sony AIBO ERS-210, ERS-220A, ERS-311[a]). There were also three behaviors the dogs could complete; *kick, dance, stand still*. For the *kick* behavior, the robot oriented towards a ball placed in front of it and then kicked it. For the *dance* behavior, the robot danced to music. For the *stand still* behavior, the robot was turned off and did nothing. For each child, each dog performed only one of the behaviors. To avoid confounding a particular robot with a particular behavior, the behaviors were counter-balanced across the three dogs. For example, some children saw ERS-210 dance, whereas other children saw ERS-220A dance. After the dogs performed, children heard three classes of animacy probes: Intelligence, Biology, and Agency. For each class of animacy, children heard three separate questions. Table 1 shows the full set of nine questions. Children indicated which of the robots had the feature, or ability, mentioned in the question. If a child indicated that a robot had the specific feature or ability, the robot behavior received a score of 1, and when the child did not refer to the robot, a score of 0 was given. All told, the factors created a fully crossed $3 \times 3 \times 3$ design: age by robot behavior by class of animacy probe, with three questions for each animacy class.

Procedure. Children participated in a one-to-one 10–15 min videotaped session. Figure 2 indicates that each child saw the three robotic dogs perform their respective behaviors in succession, with order of behavior counter-balanced across children. The experimenter lifted a box, and the dog performed its behavior. The box was then lowered, and the experimenter revealed the next dog. Each box showed a picture of the dog on the outside to help children remember which dog was where. After children saw all three different robots perform their respective behaviors, they heard the animacy questions one at a time. For each question, children indicated which dog(s), if any, could do what the question proposed. They pointed to the

[a]The entertainment robots (Aibo ERS-210, 220A, 311) were generously donated for this research by Sony Entertainment of America, San Diego, CA.

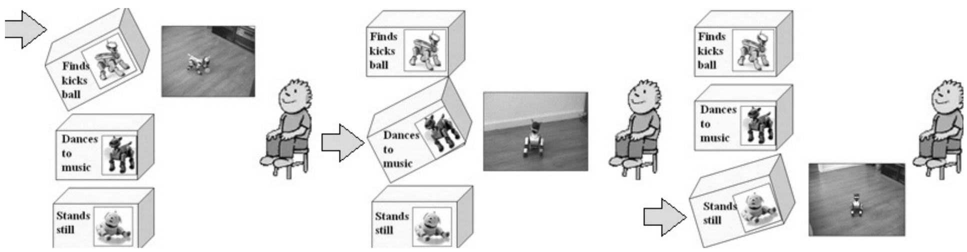

Fig. 2. Study 1 procedure for exposing children to the robots.

pictures of the robotic dogs. All children heard the questions in the same order, but they were shuffled relative to Table 1.

4.2. *Results*

We first report the broader effects and then move to the effects of specific questions. For each animacy class, children heard three questions. We found the total score for each set of three questions for each dog. For example, for biology, if a child pointed to the dancing dog for all three biology questions and pointed to the kicking dog for only two of the questions, the dancing dog would receive a biology score of 3 and the kicking dog would receive a score of 2. The behavior of the dog and scores it received for each animacy class were crossed with the children's age in a multivariate analysis of variance. In reporting the descriptive statistics, we use the more readily understood percentages (i.e. 3 out of 3 appears as 100%).

There was no effect of age. Descriptively, the younger children were more inclined to attribute animacy across all the questions (young = 74%, middle = 62%, old = 64%), but the difference was not significant; $F(2, 29) = 0.82$, MSE = 4.6, $p > 0.4$. Figure 3 shows the mean percentages for the factors of behavior and animacy class. There was a main effect of behavior on children's overall level of attribution; $F(2, 28) = 4.2$, Roy's Root = 0.298, $p < 0.05$. There was also a strong effect of the class of animacy on children's attributions; $F(2, 28) = 16.5$, Roy's Root = 1.18, $p < 0.001$. There were no interactions.

Post-hoc analyses (using a Bonferroni adjustment for multiple tests) confirm the patterns in Fig. 3. Children were significantly more likely to attribute animacy to the *dance* behavior than *stand still* ($p < 0.05$), whereas the *kick* behavior was not significantly different from the other two behaviors. Children were significantly more likely to attribute intelligence to the robots than biology or agency (ps < 0.01).

Based on these analyses, the behavior of the robot influenced children's attributions of animacy, particularly the comparison of a motionless robot versus a dancing robot. At the same time, the results make it clear that animacy attributions are not monolithic, because the children were much more likely to attribute intelligence than biology or agency (though children were still willing to attribute the latter two properties over 50% of the time as well). Interestingly, the type of behavior

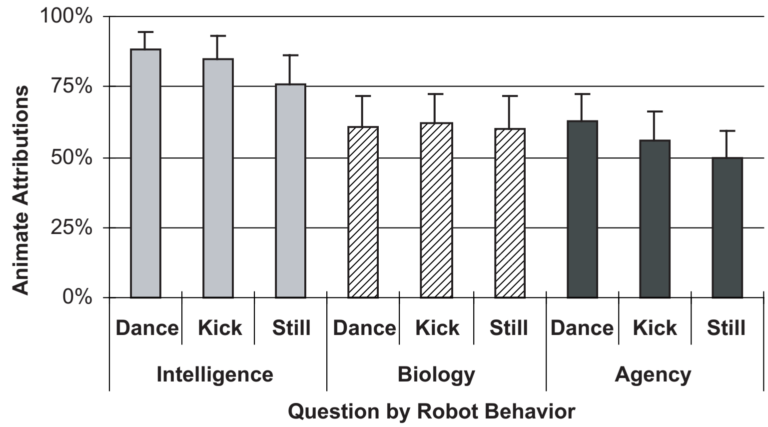

Fig. 3. Mean percentages for behavior and class of animacy.

did not influence the type of attribution. Children were not more likely to attribute intelligence to a robot that tracked down the ball and kicked it. Though the children differentiated types of animacy, these distinctions were not tied to particular behaviors.

The preceding analyses did not indicate any reliable effects of age. A more refined analysis shows an age effect. The following analyses examine how children responded to each of the questions within each animacy class. To simplify matters, the results for each animacy class are described in turn.

To examine intelligence attributions, the factors of age (3) by behavior (3) by intelligence question (3) were crossed in a multivariate analysis of variance.[b] Figure 4 shows the percentages of children who indicated the robots for each question broken down by age and animacy class. The figure does not indicate behavior because the behavior of the robots did not interact with age or question type.

The top panel of Fig. 4 shows the results for the intelligence questions. There were no effects of age, question type, or any interactions; $ps > 0.4$. There was a moderate effect of behavior; $F(2, 28) = 3.2$, Roy's Root $= 0.23$, $p < 0.1$. As before, *post-hoc* pairwise comparisons indicated that children were more likely to attribute intelligence to the *dance* behavior compared to *stand still* ($p < 0.05$), and *kick* was not different from the other two behaviors (ps > 0.1).

The same style of analysis was conducted for the three biology questions. The middle panel of Fig. 4 shows these percentages. The children showed very different levels of attribution for the three biology questions; $F(2, 28) = 6.06$, Roy's

[b]Strictly speaking, a Manova is not ideal for this analysis, because the data are restricted to 1's and 0's (did the child indicate the dog for the given question?). The previous analysis used the summed scores across the three intelligence questions, which ranged from 0 to 3. However, the results are more interpretable for most readers, and the statistical patterns are consistent with more arcane analyses.

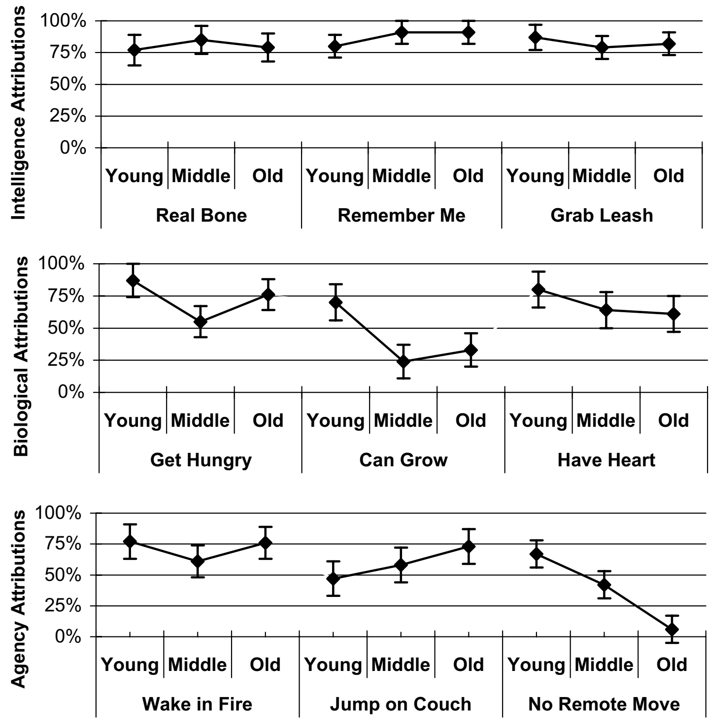

Fig. 4. Children's attributions broken down by question and age.

Root = 0.43, $p < 0.01$. Pairwise comparisons indicated that children were much less likely to accept that the robots could grow compared to having a heart or getting hungry (ps < 0.01). There were no other effects.

Finally, the analysis was repeated for the agency questions. For agency, there is the first evidence of a significant effect of age. There was a question by age interaction; $F(2, 29) = 5.7$, Roy's Root = 0.39, $p < 0.01$. The bottom panel of Fig. 4 shows opposite effects of age for the jump-on-couch and move-without-remote-control questions. The young children frequently thought the robots could move without a remote control, whereas the old children rarely thought the robots could move without a remote control. Inversely, the young children were less likely to think the robots would jump on the couch, whereas the older children thought it might. There was also a main effect of behavior; $F(2, 28) = 4.4$, Roy's Root = 0.31, $p < 0.05$. Children were more likely to attribute agency to the dance behavior but not the stand still behavior, $p < 0.05$.

4.3. *Discussion*

We proposed three predictions that would indicate that children have a piecemeal understanding of animacy, at least when they need to apply it to the boundary object of a robotic dog. The first hypothesis was that children would not confer animistic properties evenly. The results supported this hypothesis. Children attributed intelligent behavior more than biology and agency. Moreover, within classes of animacy, the children varied in their responses. For biological properties, the older children believed the robots get hungry and have a heart, but they did not believe they grow. For agency, almost no older children thought the robot could move without a remote control, but over three-quarters thought the robot might jump on the couch when left alone. Thus, if these children had a theory of animacy, it would have to be highly idiosyncratic, because it did not appear to entail a set of necessary features derived from an integrated belief system.

The second hypothesis was that children would show select changes in their attributions of animacy with age. So, rather than replacing one theory with a new one, they would simply change discrete beliefs, based on facts they may have acquired during development. The results supported this hypothesis. The younger children attributed animate properties relatively consistently across the different questions. To view this as an indication of a theory, it would be necessary for the children to reject attributions of animacy with equal consistency across categories. The older children, who presumably would have a more mature theory, did not reject animacy properties whole cloth. Instead, the older children accepted some biological attributions, but rejected the ability to grow. They also accepted some agency attributions, but they rejected the ability to move without a remote. One explanation for the spotty beliefs for the older children is that they had learned that growing things need "soft" exteriors not plastic shells and that toys need remotes. Knowledge of these select facts, however, did not prevent them from believing that the robots ate food or that the robots could choose to be bad dogs and jump on the couch.

The third hypothesis, and most relevant to the practical design of robots, was that the children would not tie specific animacy attributes to specific behaviors. The children did attribute more animacy properties to the dancing robot than the still robot. So, this means that a moving robot is more likely to elicit the belief that it is animate. Similarly, the children were unlikely to attribute growth to any of these robots, which may be a result of their hard shell. So, it appears that the children do connect their beliefs of animacy to specific empirical features. However, this connection is not theory-like, because the children did not specify classes of attributes associated with indicative behaviors; the children did not attribute more intelligence or agency to the kicking robot, which showed the most cleverness by noticing and locating a ball. In fact, they attributed more intelligence to a dog that was turned off than they attributed properties of biology or agency to a moving dog.

This study indicated that children bring an eclectic set of attributions to robots and that these attributions are only mildly connected to the range of behavior available to the robot. A limitation of the study, however, is that the children only saw the robots briefly, and the robots did not show a large range of behaviors. Moreover, the children did not get to interact with the robots. Contemporary entertainment robots can produce behavior that is highly contingent on what children do with them. Interacting with robots with high contingency may alter the patterns of attribution.

A second limitation of the study is that some of the children may have thought they were supposed to treat the robots *as though they were alive*. For example, the older children may have been playing along that the robots have a heart. But, when it came to the remote control question, the older children decided that they were being asked if they thought the robot was really alive. A useful second study would examine whether children actually think the robot is alive and whether this affects their attributions.

5. Study 2: Attributions through Physical Interaction

The second study attempted to address some of questions raised by the first study. One question was whether an interactive experience with a more or less responsive robot influences children's attributions. Therefore, in this study, half of the children interacted with a robot that had highly contingent behavior, and half the children interacted with a robot that was not responsive. (In this study, children only interacted with one robot.)

A second question came from the interesting finding that older children thought the robots could not grow, even though they could get hungry and have a heart. Perhaps, this pattern of inconsistent biological attribution was simply the result of the children believing that not all things grow. After all, their parents had not grown taller. Alternatively, this failure to attribute growth may be just one instance of a larger pattern of inconsistent biological attribution. To address this question, each child answered six questions (shown in Table 2) that probed their biology and sensory attributions towards their robot.

A third question was whether children really thought the robot was alive, or whether they were just playing along that it was alive. There is a large body of research on children's abilities to pretend,[23,24] and therefore, it is important to find out how their attributions differ when they believe the dog is alive or not. To address this question, we asked the children if they thought the robot was alive. This permits us to analyze how biological attributions differ when children believe the robot is alive or not. It also allows us to see if a more interactive robot leads children to believe the robot is alive.

Table 2. Questions for biology and intentions.

	Questions: Do you think "Chai" . . . ? How can you tell?
Biology	1. . . . needs to have food?
	2. . . . pees and poops like we do?
	3. . . . breathe?
	4. . . . can see the toy?
	5. . . . can hear me?
	6. (Experimenter hits the robot) Do you think it **hurts** "Chai" ?
Alive/Real	1. Do you think "Chai" is **alive** or not alive? How can you tell?

NeCoRo Cat | BN-1 Cat | FurReal Cat | Paro Seal Robot

Fig. 5. Children in Experiment 2 interacted with robotic animals.

Our predictions for this experiment are similar to the previous one. Children will not make consistent attributions across biological properties. Older children will show piecemeal changes relative to younger ones; the older children will not change from accepting all biological attributions to rejecting all biological attributions. We also thought that the contingency level of the robot would make very little difference. The first study seemed to imply that the key feature for the children was whether the robot moved or not. In this study, all the robots moved. The difference between conditions was simply whether the robot moved in response to the child. Our assumption is that coherent movement is the key for children, and that their beliefs about animacy are not tied to the responsiveness of an organism.

We did not have strong predictions about the frequency children would say the robot is living. We also did not know what would happen to children's attributions when they thought the robot was alive or not. One possibility is that children do not have a good concept of alive to start with, so whether or not they think the robot is alive would be somewhat arbitrary and without influence. An alternative possibility is that all the children would say the robot is not living, in which case, our results would simply indicate the ways children are willing to pretend with entertainment robots.

5.1. *Method*

Participants. Sixty-one children from the same university day-care program participated. The children comprised three age groups: young ($n = 20$, $M = 42.8$ months, range 35–48 months), middle ($n = 21$, $M = 53$ months, range 49–56 months), and old ($n = 20$, $M = 60.3$ months, range = 57–66 months).

Design. Children from the three age levels interacted with either a high- or low-contingency robot. There were four robotic animals (shown in Fig. 5). Two were for the high-contingency condition, and children saw and interacted with only one of them (Omron's NeCoRo Robot cat[c] or AIST's robotic seal Paro[d]). There were also two robots used for the low-contingency condition (Bandai's BN-1 cat or Tiger Electronics' Furreal cat). The high contingency robots had a relatively realistic

[c]The communication robot NeCoRo cat was generously donated by Omron Corporation, Japan.
[d]The robotic seal Paro was generously loaned from the Advanced Industrial Science and Technology (AIST), Japan.

appearance and were capable of responding to different types of input (e.g. pushing back lightly when held, responding to name). The low contingency robots were less realistic in appearance (more like a toy) and had a smaller repertoire of contingent responses. Unlike Study 1, there were no specific behaviors the robots performed; however, the robots did generate sounds and move (e.g. meow, squeak, move body parts), and in the high contingency condition their behaviors were in response to the child. During the interaction, children answered six questions that mapped into attributions of biology and sensation. Table 2 shows the listing of the questions. Children that answered "yes," received a score of 1 for that question, and received a score of 0 for "no." At the end of the session, children were asked "Do you think Chai is alive or not alive?"

Procedure. Children participated separately in 10–15 minute videotaped sessions. On entering the room, children were asked to wait quietly as the experimenter went into the back. The experimenter returned with one of the robots, already turned on. The children physically interacted with their robot through actions such as holding, waving their hands in front of it, petting, and tugging. After children had a couple of minutes to interact with the robot, they heard the biology questions one at a time in the same order. For each question, the children indicated with a "yes" or "no" if they thought the robot could do what the question proposed. Afterwards they answered whether they thought Chai (the robot) was alive.

5.2. *Results*

Older children were less likely than the younger children to attribute biological properties to the robots. Across the six biology questions, a percentage of positive attributions was calculated for each child. This aggregate score was the dependent measure in an age by contingency level univariate analysis. As reflected in Fig. 6, there was a significant effect of age; $F(2, 55) = 4.17$, MSE $= 2.99$, $p < 0.05$. Older children were less likely to attribute biological properties to the robots than young children. So, unlike Study 1, where there was a descriptive affect of age on biology attributions, the effect reached significance here. There was not a significant effect of the robot's contingency and no interaction of age and contingency; ps > 0.1. However, Fig. 6 indicates that as children get older, they are less likely to attribute biological properties to the less interactive robot. A subsequent study that used a larger sample size or more interactive robots might bring this descriptive interaction to the level of significance. Regardless, the data strongly suggest that the contingency of the robot has no effect on three-year-old children's attributions of biological properties.

A second question is whether the age of the children or the robot's contingency had an influence on whether children believed the dog was alive. Table 3 shows the percent of children at each age and contingency level who said their dog was alive. To determine if the differences were significant, we used a logistic regression of age by contingency on whether the child though the robot was alive.

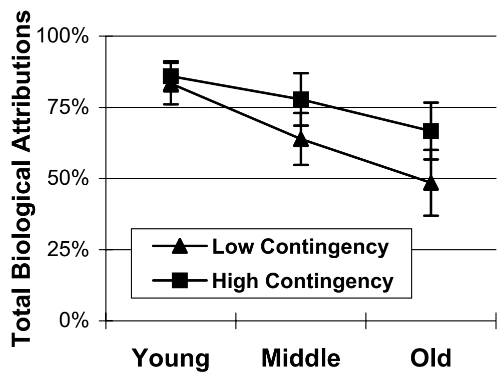

Fig. 6. Average percentage of biological attribution (out of six questions) by robot contingency and age.

Table 3. Percent of children that said the robot was alive.

Age	Low Contingency(%)	High Contingency(%)
Young	43	31
Middle	50	67
Old	45	56

There were no significant effects; $ps > 0.2$. Thus, the behavior of the dog did not have a stable affect on the children's beliefs about whether the robot was alive. Descriptively, Table 3 shows that the younger children were less likely to attribute aliveness to the highly contingent dog. One possible explanation is that the contingent dog was threatening, and therefore, they were less inclined to want to believe it was alive.

The next question is whether children's belief that a robot dog was alive influenced their biological attributions. We conducted a second univariate analysis on the percentages of biological attributions. The factors were the children's age and whether the child said the robot was alive or not. As before, there was a strong effect of age; $F(2, 49) = 7.41$, $\text{MSE} = 1.5$, $p < 0.01$, such that younger children made more biological attributions. This analysis also indicates that if children thought the robot was alive, they were much more likely to attribute biological properties; $F(1, 49) = 30.5$, $p < 0.001$. There was also a strong interaction between age and alive; $F(2, 49) = 7.64$, $p < 0.01$. Figure 7 shows the effect. Older children who thought the robot was alive were much more likely to attribute biological properties compared to older children who did not think the robot was alive. In contrast, the young children attributed biological properties at the same rate, regardless of whether they said the robot was alive or not. Evidently, younger children do not associate "alive" with biological properties, and it seems probable that their category of "alive" is not well defined.[15]

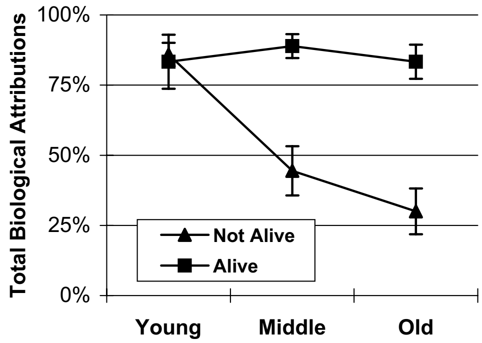

Fig. 7. Level of biological attribution by age and their belief that the robot was alive.

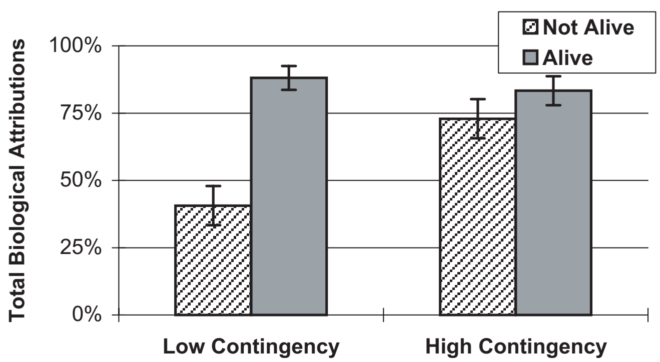

Fig. 8. Biological attributions broken out by contingency level and attribution of living.

There was also a modest interaction of contingency by alive; $F(1, 49) = 5.39$, $p < 0.05$. As shown in Fig. 8, children in the low contingency condition who attributed aliveness made the highest percentage of biology attributions, but children in the low contingency condition who did not attribute aliveness gave the lowest percentage of biology attributions. We do not have a strong single explanation of this effect.

The next set of analyses consider how children responded to each biology question and whether this varied by age and attribution of aliveness. Age, contingency level, and attribution of aliveness were crossed, and each of the six questions served as a separate dependent measure in a multivariate analysis.[e] The purpose of this analysis is to determine whether children show the same patterns of attribution for the different biological properties. That is, do the younger children attribute biological properties uniformly compared to older children, and does this vary by contingency level or attribution of aliveness?

[e]As before, each child only receives a "1" or "0" for each question, so an analysis of variance is not strictly appropriate. However, results are comparable to more arcane analyses.

To simplify matters, we only report significant multivariate effects of question type, and use Fig. 9 to show the patterns for the questions. There was a significant effect of question type; $F(5, 45) = 4.66$, Roy's Root $= 0.52$, $p < 0.01$. Children did not equally accept all biological properties for the robots. There was also a significant effect of whether children thought the robot was alive on their pattern of responses; $F(5, 45) = 2.59$, Roy's Root $= 0.29$, $p < 0.05$. Finally, there was a marginal interaction between age and alive on how children varied across the six biology questions; $F(5, 46) = 2.9$, Roy's Root $= 0.25$, $p = 0.06$.

Figure 9 indicates the source of the effects. Young children make the most biological attributions across all questions and these attributions do not vary by whether

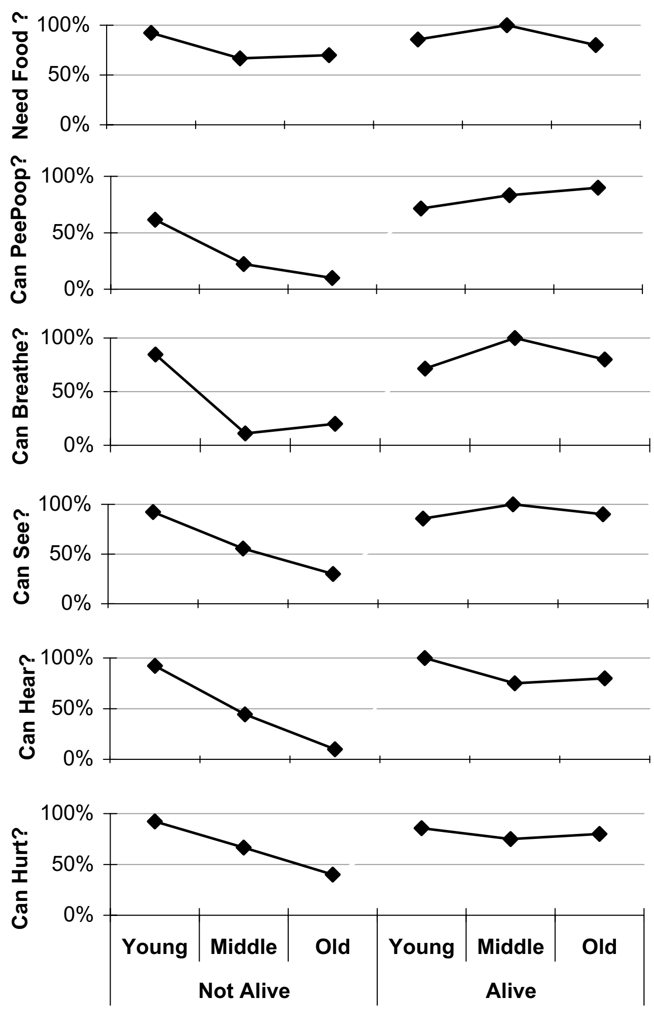

Fig. 9. Percentage of children who attributed a biological property broken about by age and belief in aliveness.

they believe the robot is alive. Middle and older children, on the other hand, have high attributions only when they say alive. However, there are exceptions. Children attribute the need for food, regardless of their age and belief in aliveness. And, whether the children believed the robot was alive or not had minimal influence on the rate children at each age level thought the robot could feel pain (hurt).

5.3. *Discussion*

In the second study, children either interacted with a robot that responded directly to their actions or just made various movements and sounds. The contingency of the robot had relatively modest effects on the children's biological attributions or their belief that the robot was alive. As the children increased in age they showed some sensitivity to the contingency of the robot, such that they made somewhat fewer biological attributions for the low-contingency robot. And, for the high-contingency robot, 10–15% more of the four- and five-year-olds thought it was alive. Perhaps these effects would reach significance in a subsequent study with modest changes including a larger sample size and a longer period of exposure. A subsequent study should also counter balance the type of robot by contingency level.

The results also indicated that children's belief in whether their robot was alive had a large influence on their biological attributions. Though it made no difference for the young children, it had a very large effect for the older children. If the older children thought the robot was alive, their biological attributions were indistinguishable from the young children. They attributed all the biological properties at high levels. One interpretation of this finding might be that pre-school children do have a theory of being alive, because once they think a robot is alive, they confer to it all the biological properties. However, we think this is a mistaken interpretation. If the children had a strong sense of alive, then otherwise comparable children who thought the robot was not alive should have rejected all the biological properties. Instead, the older children who did not believe the robot was alive rejected some biological properties, but not others. The children were quite willing to accept that robots needed food, and nearly 50% of the five-year-olds thought the robot could feel pain. These piecemeal changes in their biological attributions to a robot they thought was not alive indicates these children neither have a theory of alive nor a coherent concept of which biological properties come together (e.g. the children said the robot needed food, but did not defecate).

An interesting question is why the older children, who thought the robot was not alive, still maintained that the robot needed food and to some degree thought it felt pain. One possibility involves the prevalence of food and pain in stories around stuffed animals and dolls. The narrative for play often involves activities such as feeding the baby and dining, as well as treating and healing.[25] Sarbin[26] proposes that the degree of involvement with a narrative corresponds to the reality of that narrative (cf. Ref. 27). So, although the children may believe the robot is not alive, the strength of their familiar narrative has a stronger sense of reality. This strong sense of reality drives that particular piece of their network of beliefs regardless of their other beliefs. If children do have a coherent theory of animate properties, it is not evident in these studies.

6. Conclusion

Two studies explored three- to five-year-old children's "animistic intuitions." The studies examined whether children attributed intelligence, biology, and agency to entertainment robots. These three classes of animate behaviors are typically associated with real, living beings. A practical question was whether the realism of the robots influenced the children's attributions. In one experiment, children watched the behaviors of robots with no direct interaction, while in the second experiment, children had direct physical interaction with the robots. The studies indicated that there was a mild influence of the realistic behavior of the robots. A moving robot received more animate attributions than one that did not move, and for older children, the interactive responsiveness of the robot had a modest, but non-significant effect on their attributions. In general, the results suggest that improving the realism of the robots does not have a tremendous effect on children's conceptual beliefs.

This result, however, does not mean that the realism of a robot does not have other effects. Recall that our primary argument is that children do not have a very well-developed concept of what it means to be alive or animate. By this argument, the place to look for the effects of robot realism is not going to be in children's ideas or beliefs. It may be better to look at other variables, for example, are the children more attracted to realistic robots or do children engage in longer interactions with realistic robots?

We suspect, however, that realism *per se* is probably less important for pre-school children. Rather, it is more important to include features that enable the children to bring to bear familiar schemas so they can sustain a productive interaction[28] with the robot. Unlike a doll or stuffed animal, the robot responds to children. This means the child's ability to pretend is constrained by what the dog will do in response. Until such time that robots have the intelligence to flexibly respond to children's interactive bids, children will have to follow a well-known script (e.g. a tea-party script) so the children can be sure to stay within the repertoire available to the dog. If children cannot bring to bear a strong schema, the children will try a number of interactive bids to see what emerges, but the robot will not be able to respond flexibly and the children will get frustrated.

The studies were also intended to develop a portrayal of child development around the notions of "animate" and "alive." The studies differ from other developmental studies because they looked at how children learned to *let go* of animacy attributions. This has the merit of seeing how children handle evidence or beliefs that partially conflict with their prior beliefs. If children have coherent knowledge, then "falsifying" evidence or beliefs should make them let go of other purportedly related beliefs about animacy. This was not the case in the current studies. Children merrily attributed some properties and not others.

The studies provided a clear portrayal of how children of different ages treat entertainment robots. The youngest children broadly attributed animate properties to all the robots. This attribution is unlikely to be the result of any coherent theory. For example, whether or not they said the robot was alive, they offered the same high level of biological attribution. Rather, the results seem to be due to a general tendency to over-extend all the facts that they know about dogs. Notably, the exceptions to their over-extension of what they know about dogs involved questions that had a moral component. For example, the youngest children were the least likely of the age groups to say the robot could jump on the couch when nobody was there. It seems likely that they were not reasoning about "could" for this question, but rather, they were reasoning about "should." A good dog would not jump on the couch. Perhaps a good design for an entertainment robot for young children would be a "naughty" dog that learns to be good. This is a very familiar schema children could bring to bear.

Starting at four years and increasing into five years, children begin to develop a meaningful distinction between alive and not alive. If children believe a robot is alive, their attributions still look very much like young children. But, they differ from young children, because if they think a robot is not living, they are much less likely to attribute biological properties. Their notion of alive is becoming connected to the attributes of animacy. At the same time, they do not have a fully developed concept of animate or alive. The relevant evidence comes from their willingness to attribute some qualities of animate things but not others. For example, the children believed that the robot needs food but does not defecate. Rather than having a strong model of what it means to be animate, the children appeared to be slowly shifting which facts should apply to the strange category of entertainment robot. Nearly all of the five-year-olds, for example, believed that the motionless robot needed a remote control to get "animated," but over 50% were still willing to say that it had a heart. The fact that the children were making piecemeal adjustments to their understanding does not mean the children were not searching for coherence. Perhaps children are theory builders,[29] and the concepts of alive and animate are simply difficult.

Boundary objects, like the monsters Dracula and Frankenstein, have long revealed the weaknesses in people's belief systems. Though people know they are only fictions and it is only a movie, the combination of categories such as alive and dead can trigger responses that defy people's beliefs. Perhaps children are not so different from adults. Regardless of the quality of one's belief system, the place to look for the affects of entertainment robots may not be in the world of concepts and words. It may be in affective responses that can operate regardless of beliefs.

Acknowledgments

The authors greatly appreciate all the help and support from our collaborators at Sony, AIST, Omron, Keio University, and Stanford University. Special thanks to all the children at Bing Nursery School. This research was supported by the National Science Foundation under Grants Nos. REC-0231946 & SLC-0354453. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

1. C. Bartneck and J. Forlizzi, Shaping human-robot interaction: Understanding the social aspects of intelligent robot products, in Extended Abstracts *Human Factors in Computer Systems (CHI)* (ACM Press, 2004), pp. 1731–1732.
2. B. Robins, K. Dautenhahn, R. te Boekhorst and A. Billard, Effects of repeated exposure to a humanoid robot on children with autism, in *Proc. Cambridge Workshop on Universal Access and Assistive Technology* (Springer-Verlag, 2004), pp. 225–236.
3. N. Suzuki, K. Kakehi, Y. Takeuchi and M. Okada, Social effects of the speed of hummed sounds on human–computer interaction, *Int. J. Human-Comput. Studies* **60** (2004) 455–468.
4. D. Wada, T. Shibata, T. Saito and K. Tanie, Analysis of factors that bring mental effects to elderly people in robot assisted activity, in *Proc. Int. Conf. Intelligent Robots and Systems, IROS* (IEEE Press, 2002), pp. 1152–1157.
5. Y. Bar-Cohen and C. Breazeal, *Biologically Inspired Intelligent Robots* (SPIE Press, Washington, 2003).
6. B. Reeves and C. Naas, *The Media Equation* (CSL1 Publications, Stanford, CA, 1996).
7. B. Robins, K. Dautenhahn, R. te Boekhorst and A. Billard, Robots as assistive technology—Does appearance matter? in *Proc. IEEE Int. Workshop Robot and Human Interactive Communication (RO-MAN)* (IEEE Press, 2004).
8. F. Michaud and S. Caron, Roball, the rolling robot, *Autonom. Robots* **2**(12) (2002) 211–222.
9. S. Y. Okita, D. L. Schwartz, T. Shibata and H. Tokuda, Exploring young children's attributions through entertainment robots, in *Proc. IEEE Int. Workshop Robot and Human Interactive Communication (RO-MAN)* (IEEE Press, 2005).
10. P. H. Kahn, B. Friedman, D. R. Perez-Granados and N. G. Freir, Robotic pets in the lives of preschool children, in *Proc. Conf. Human Factors in Computer Systems (CHI)* (ACM Press, 2004).
11. T. Kanda, T. Hirano, D. Eaton and H. Ishiguro, Interactive robots as social partners and peer tutors for children: A field trial, *Human-Comput. Interaction* **19**(2) (2004) 61–84.
12. R. Fox and C. McDaniel, The perception of biological motion by human infants, *Science* **218** (1982) 486–487.
13. B. Friedman and P. H. Kahn, Human agency and responsible computing: Implications for computer system design, *J. Syst. Soft.* **17** (1992) 7–14.
14. S. C. Johnson, A. Booth and K. O'Hearn, Inferring the goals of a nonhuman agent, *Cogn. Develop.* **16**(1) (2001) 637–656.
15. S. Carey, *Conceptual Change in Childhood* (MIT Press, Cambridge, MA, 1985).
16. K. Inagaki and G. Hatano, *Young Children's Naive Thinking about the Biological World* (Psychology Press, New York, 2002).
17. A. Gopnik and A. N. Meltzoff, *Words, Thoughts, and Theories* (Bradford MIT Press, Cambridge, MA, 1997).
18. H. M. Wellman, *The Child's Theory of Mind* (MIT Press, Cambridge, MA, 1990).

19. F. C. Keil, *Concepts, Kinds, and Cognitive Development* (MIT Press, Cambridge, MA, 1989).
20. A. A. diSessa, Knowledge in pieces, in *Constructivism in the Computer Age*, eds. G. Forman and P. Pufall (Lawrence Erbaum Associates, Hillsdale, NJ, 1988), pp. 49–70.
21. S. A. Gelman and G. M. Gottfried, Children's causal explanations of animate and inanimate motion, *Child Develop.* **67** (1996) 1970–1987.
22. C. M. Massey and R. Gelman, Preschooler's ability to decide whether a photographed unfamiliar object can move itself, *Develop. Psychol.* **24** (1988) 307–317.
23. M. Taylor, *Imaginary Companions and the Children Who Created Them* (Oxford University Press, New York, 1999).
24. R. W. Mitchell, Imaginative animals, pretending children, in *Pretending and Imagination in Animals and Children*, ed. R. W. Mitchell (Cambridge University Press, Cambridge UK, 2002), pp. 3–22.
25. D. Fernie, The nature of children's play, *ERIC Digest* **ED307967** (1988) 2–3.
26. T. R. Sarbin, Believed-in imaginings: A narrative approach, in *Believed-In Imaginings*, eds. J. de Rivera and T. R. Sarbin (Americal Psychological Association, Washington DC, 1998), pp. 15–30.
27. A. Lillard, Just through the looking glass: Children's understanding of pretense, in *Pretending and Imagination in Animals and Children*, ed. R. W. Mitchell (Cambridge University Press, Cambridge, UK, 2002), pp. 102–114.
28. D. L. Schwartz, K. B. Pilner, G. Biswas, K. Leelawong and J. Davis, Animations of thought: Interactivity in the teachable agents paradigm, to appear in *Learning with Animation: Research and Implications for Design*, eds. R. Lowe and W. Schnotz (Cambridge University Press, UK, in press).
29. S. Vosniadou and W. F. Brewer, Mental models of the earth: A study of conceptual change in childhood, *Cogn. Psychol.* **24** (1992) 535–585.

Sandra Y. Okita is a Ph.D. candidate at the Cognition and Technology Laboratory (Stanford Center for Innovations in Learning, SCIL) at Stanford University. She earned her M.A. degree in Psychology at Stanford. Her research focuses on learning, teaching, how technology can facilitate the cognitive development of children, the role of robots in the educational setting, and how robot interaction can influence the behavior of young children.

Daniel L. Schwartz is Associate Professor of Education at Stanford University, leading the Cognition and Technology Laboratory (Stanford Center for Innovations in Learning, SCIL). His research area is at the intersection of cognitive science, education and computer science. His research interests are in agent technology, imagery and action, and how the facility for spatial thinking can inform and influence processes of learning, instruction and assessment. The research laboratory web page is *aaalab.stanford.edu*.