What Are You Talking to?: Understanding Children's Perceptions of Conversational Agents

Ying Xu
University of California, Irvine
Irvine, U.S.
ying.xu@uci.edu

Mark Warschauer
University of California, Irvine
Irvine, U.S.
markw@uci.edu

ABSTRACT
Conversational agents (CAs) available in smart phones or smart speakers play an increasingly important role in young children's technological landscapes and life worlds. While a handful of studies have documented children's natural interactions with CAs, little is known about children's perceptions of CAs. To fill this gap, we examined three- to six-year-olds' perceptions of CAs' animate/artifact domain membership and properties, as well as their justifications for these perceptions. We found that children sometimes take a more nuanced position and spontaneously attribute both artifact and animate properties to CAs or view them as neither artifacts nor animate objects. This study extends current research on children's perceptions of intelligent artifacts by adding CAs as a new genre of study and provides some underlying knowledge that may guide the development of CAs to support young children's cognitive and social development.

Author Keywords
Conversational agents; perceptions; animacy; children; child-agent interactions

CCS Concepts
•**Human-centered computing** → **Empirical studies in HCI;**
•**Social and professional topics** → **Children;**

INTRODUCTION
As conversational agents (CAs) become increasingly prevalent in home life, whether through smart phones, tablets, or smart speakers, both scholars and the general public have noted young children's propensity to interact with them [6,10,31,47].

CAs are designed to take on many of the properties previously thought to be unique to humans. Specifically, CAs support natural spoken conversation, thus displaying a high level of intelligence. Moreover, some CAs have been designed as social companions for children [24,27] and are capable of provoking social reactions, such as empathy and trust [40]. However, CAs available in phones or smart speakers are neither anthropomorphic nor self-locomotive, making them physically different from a human dialogue partner. CAs' human-like capabilities without corresponding physical features create an intriguing research scenario for examining child-CA interaction. A handful of studies have found that children interacting with CAs utilize communication strategies similar to those normally used when interacting with a human interlocuter [46,50,57]. However, we know very little about children's perceptions during such interactions, particularly whether children attribute human properties to non-human CAs.

The question of how children perceive CAs is of interest to the fields of developmental psychology and human-computer interaction (HCI). First, this question is relevant to the long-standing focus within developmental psychology on the animate-inanimate (A-I) distinction in early childhood [41]. Given that CAs are highly interactive and intelligent, they may blur children's categorical distinctions between technological artifacts on the one hand and biological beings on the other [17, 48]. CAs' blurring of these boundaries may result in children categorizing CAs as neither artifacts nor living beings [42] or perceiving CAs as occupying some middling position along an animate-inanimate continuum [20,26]. Second, children's perceptions of CAs are of importance to HCI given that this field is keenly interested in developing CAs that simulate human-to-human communication [32,55]. Within the field of HCI research, children's behavioral interactions with CAs are typically used to evaluate whether CAs have gotten closer to that gold standard [16,50]. However, these studies fail to consider children's perceptions, which is another integral facet of children's experiences with CAs as perceptions shape behavioral interactions [28]. If children ascribe life characteristics to CAs, they may then engage in more natural communication patterns with those CAs. In contrast, if children view CAs as simply machines or tools, they may approach their interactions in a less natural way. Therefore, understanding children's perceptions may help make sense of previous research on child-CA interactions and provide a more complete picture of children's relationships with CAs. Moreover, children's perceptions of CAs' properties and human/non-human status may reveal children's expectations for consumer-level CAs, which could be helpful for development of future CAs.

The present study is grounded in and extends the developmental psychology and HCI lines of research above. We build on existing research of early childhood A-I distinction and extend its application to how children understand intelligent

CHI '20, April 25–30, 2020, Honolulu, HI, USA.

ACM ISBN 978-1-4503-6708-0/20/04 ...$15.00.
http://dx.doi.org/10.1145/3313831.3376416

artifacts, in particular, speaker-based CAs. We also explore children's perceptions, in particular whether children view CAs as human-like dialogic partners. Specifically, this study seeks to answer three questions. First, which domain do children perceive CAs as belonging to (e.g., artifact, living object, or something else)? Second, do children view CAs as possessing human-like cognitive, psychological, and behavioral properties? Third, how do children reason about whether CAs possess certain properties? To answer these three questions, 28 children aged 3 to 6 were invited to individually interact with a CA, after which we elicited their perceptions through a semi-structured interview and a drawing task. This study is intended to reveal children's perceptions behind their active engagement with a CA and offer theoretical and design implications. It is interesting to focus on children aged 3 to 6 primarily because children in this age group have developed a naïve framework of beliefs about living things, which emerged in the absence of formal instruction [33], and do not yet have a sophisticated conceptualization of computational objects [44].

RELATED WORK

Understanding of Animacy in Early Childhood

Children's understanding of the A-I distinction – the distinction between living and non-living things – is probably one of the most enduring questions in developmental psychology [41]. Indeed, the ability to recognize objects as animate or inanimate is thought to be a fundamental cognitive process since it provides the foundation upon which children categorize objects in the world [8, 41]. Children's primitive understanding of the A-I distinction begins in infancy, develops rapidly during early childhood, and matures in adolescence [38]. Melson et al., synthesizing research on the topic, suggested that a young child distinguishes between animate and inanimate things based on the child's perception of that thing's cognitive (thoughts), psychological (feelings or emotions), and behavioral (actions or speech) properties [35]. Inevitably, if an object displays either all or none of these properties, children find it less challenging to categorize the object as either animate or inanimate. However, objects that display only some of the properties are more likely to raise boundary questions for children. In other words, if A-I distinction is perceived as a continuum, some objects may be more clearly perceived to be on either end of the continuum, while some are perceived to fall in between [37].

Distinct properties play different roles in children's evaluation of whether an object is animate or inanimate. In studies on this topic, children are typically shown pictures of everyday objects and tasked with sorting them based on their membership in either category before then being asked to describe which of the object's properties informed that categorization. Two trends emerge from these studies. First, children tend to firmly, yet incorrectly, associate the ability to move physically with animacy itself. Second, children astutely understand that only animate things have the ability to think and feel [41]. Although these two trends have been observed in studies that did not involve artificial intelligence, it would not be surprising if children used these same principles to determine whether an intelligent artifact is animate or inanimate. As suggested in Edwards et al., when young children attempt to understand complex and novel technologies, they tend to apply a familiar schema they have developed from their daily lives [12].

Children's Perceptions of Intelligent Artifacts

A growing body of research has focused on children's perceptions of intelligent artifacts, especially computers, CAs, and robots [17, 25, 28, 36, 37, 45, 48]. These three artifacts represent objects that may elicit different levels of perceptions of animacy given their differing properties. If children conceive of the A-I distinction as a continuum rather than a dichotomy, computers would lie close to the inanimate end, robots close to the animate end, and CAs somewhere in between.

On one end of the spectrum, robots appear to possess the cognitive, psychological, and behavioral capacities that elicit perceptions of animacy. They may move, learn, communicate, self-organize, and respond to emotions in humans. As such, it is not surprising that children who regularly interact with robots often view them as animate objects [17, 28, 48]. Correspondingly, children tend to perceive robots as possessing all of the properties associated with animacy [23, 35]. For example, Melson and colleagues found that the majority of children affirmed that AIBO, the robotic dog, had mental states, social awareness, and moral standing [35]. Similarly, Beran and colleagues suggested that a significant proportion of children in their study ascribed cognitive, behavioral, and psychological characteristics to robots [1]. This implies that children impose their own understanding of human nature onto such technological devices and see them as possessing similar capabilities. However, Beran's study also noted that children's assigning animacy to robots is driven more by robots' physical movements rather than by their intelligence [1]. When children were asked why they considered the robot to be a living being, most children pointed out the robot's humanoid appearance and its seeming ability to move spontaneously. This is consistent with children's firm association of animacy with the ability to move [51]. Given that smart speakers lack mobility and anthropomorphic embodiment, it is unclear whether the findings from robot studies will hold true for speech-only CAs.

On the other end of the spectrum, computers, while demonstrating some level of cognitive capability, typically lack the psychological and behavioral properties that children emphasize when evaluating an object's animacy. One study looked at the properties young children ascribed to computers [52]. Interestingly, while a considerable proportion of children believed that computers were capable of performing tasks that required intelligence, almost all children viewed computers as lacking psychological and behavioral capabilities. A second study also found that although children believed computers possessed moderate to high intelligence capability, they did not view computers as living objects [45]. Through analyzing children's drawings of how they think computers might look inside, Mertala suggested that children tended to view computers as machines, as most children depicted computers as technological objects such as monitors, wires, or keyboards [36]. Much like computers, CAs lack both the mobility and anthropomorphic embodiment that children readily associate with animacy. However, CAs' ability to engage in natural spoken conversation may lead children to ascribe CAs with cognitive and psychological properties similar to those in humans.

Children's Perceptions of Speech-only CAs

Only two studies, to our knowledge, exist which speak to children's perceptions of smart speaker-based CAs. In the first study, Druga and colleagues examined how children perceive CAs' psychological properties [10]. The authors asked children to interact with different smart speakers during both free and structured play and found that most children viewed the CAs as friendly and genuine. This finding supports the idea that, although CAs lack mobility and anthropomorphic embodiment, children still view them as possessing psychological properties similar to those of robots. In the second study, Lee, Kim, and Lee asked participants with previous experience interacting with Amazon Alexa or Google Home devices to draw what they thought a CA looked like [29]. The participants produced drawings that fell within four general categories: human, speaker, system, and space object. This finding confirms that CAs' unique combination of features may induce some users to view CAs as living beings rather than inanimate objects. However, while the study included participants ranging in age from 4 to 51, it did not distinguish between drawings produced by young children and those produced by older children or adults. In addition, the study focused exclusively on domain membership perceptions and did not examine perceptions of CA properties.

The current study expands on each of these two projects by examining children's perceptions of CAs' domain membership and their cognitive, psychological, and behavioral properties, while also exploring children's explanations for these perceptions.

METHOD

Participants

Our participants consisted of 28 typically developing children between the ages of 3 and 6 recruited from preschools and afterschool programs in a university community. These participants were drawn from a larger study, with the total number of participants designated so as to have sufficient statistical power. From that larger pool, 28 were in this "conversational agent" condition. We believe that 28 is a suitable sample size for the methodology used in the study.

Parents or guardians completed a brief survey on demographic characteristics. According to parent reports, the mean age of the participants was 4.7 years, and 54% were girls. Nine children (32%) were identified as White. Nineteen children (68%) spoke only English at home, and the rest of them were bilingual or spoke English as a second language, but all children possessed sufficient oral English proficiency for daily conversation. Twenty-one percent of the participants had never interacted with a CA, 36% had done so less than once a month, 18% between weekly and monthly, and 25% had daily interaction with a CA.

Description of the Interaction Tasks

The interaction tasks provided children with direct and in-the-moment experience with a CA. We noted that this approach is among the three common methods utilized by prior research. One approach relies on children's past experiences with technology (e.g., [29]); however, very young children are less able to accurately recall past experiences [43]. The second approach involves showing children videos of how a technology works (e.g., [26]); however, young children mostly learn from authentic, direct experiences rather than events they indirectly witness [15]. We believe that allowing children to interact with the CA during the study session will provide them proximal and in-person experiences, and thus we will be better able to elicit their perceptions.

Each child's interaction entailed three sessions with a Google Home Mini device and lasted approximately 40 minutes in total. The child first had a structured personal conversation with the CA, then played a structured narrative game, and finally had an unstructured dialogue. These three sessions mirrored the common interaction experiences a child typically would have with CAs in their everyday lives [11, 46]. The three sessions were carried out in one sitting, and all children were able to complete the sessions.

Personal conversation

In the first session, the CA asked children their age, favorite color, and a simple animal question (i.e., which animal has a really long neck?). The CA was programmed to repeat children's responses. For example, when a child tells the CA that he or she is five years old, the CA responds, "Wow, you're five years old. You are such a big kid!" When a child tells the CA that his or her favorite color is blue, the CA responds, "Great choice! My favorite color is also blue." In cases where the CA failed to understand a child's voice input, whether due to fuzzy pronunciation, an irrelevant response, or some other issue, the CA adopted a fallback mechanism to move the conversation forward. Such mechanisms included more general, neutral responses that did not directly repeat the child's answers (e.g., "Great choice! That is my favorite color too.").

Narrative game

In the second session, the CA read a ten-minute fantasy story and asked the children 10 story-related open-ended questions throughout. The CA gave responsive feedback based on a child's answer, either praising the child for a correct answer or encouraging the child to try again after an incorrect answer. In the latter case, the CA provided hints or rephrased the original question into a multiple-choice format, with the goal of simulating how an engaging adult would scaffold children's learning and conversation during shared reading. A fallback mechanism was triggered if the CA failed to capture a child's response twice in a row. The CA gave a general, neutral response without repeating the child's response and moved the story forward.

Unstructured dialogue

In the final session, children were encouraged to freely talk to the CA and ask the CA any questions they would like. Common topics included math (e.g., "What is one hundred plus one hundred?"), culture facts, personal questions (e.g., "How old are you?"), and the child's sharing of personal information (e.g., "My favorite princess is Elsa."). These topics corroborated findings from [31].

Procedure
Each child met individually with a trained experimenter in a designated quiet area at the child's school. At the beginning of each session, the experimenter introduced the interaction task as a game and described the Google Home Mini device as "Google." During the interaction sessions, the experimenter sat beside the child, interfering only if technical issues interrupted the child's interactions with the CA (e.g., Internet or battery issues). In the case that a child asked the experimenter questions or initiated comments, the experimenter simply answered the question or replied "okay," but avoided elaborating or extending the conversation. After the child completed all sessions, the experimenter administered a semi-structured interview and a drawing task to elicit children's perceptions, as discussed below.

Measures
Semi-structured interview
The interview protocol was derived from Beran et al. [1] and modified based on our review of the literature yielding key prompts to assess animitic attributes. Each child responded to a series of questions assessing the child's perceptions of the CA's domain membership and cognitive, psychological, and behavioral properties. Regarding domain membership, we asked children to generally describe what they thought Google was by posing the question, "What were you talking to?" After the child responded, the experimenter followed up and asked the child to further elaborate his/her answer, by posing the question, "What is XX (repeating the child's preceding response)?" We used open-ended questions to elicit children's perceptions of the CA's domain membership because we wanted to allow children to freely express their ideas without a predefined framework imposed on them [48].

Regarding the cognitive, psychological, and behavioral properties of the CA, a sample question on the cognitive property was, "Do you think Google is smart?" A sample psychological question was, "Can Google like others as a friend?" A sample behavioral question was, "Can Google see?" For each of the property items, we first asked an affirmation/negation question to elicit children to respond "yes" or "no." Then we followed up with an open-ended question ("How do you know that?"; "Why do you think/not think so?"). The purpose of the follow-up question was to prompt the child for justification of the preceding response.

The interviews were video-taped and transcribed verbatim.

Drawing task
After the semi-structured interview, children were given a drawing task, in which they were asked to draw what they thought was inside the Google Home Mini device. The experimenter, who provided a standard unlined sheet of paper and markers, sat beside the child and posed questions to encourage children to describe what they were drawing as they drew. This kind of drawing task is widely used to understand children's states of mind and internal perceptions. It is based on the premise that children will draw what they are thinking about and that children can better express their thoughts through drawings than through words [9]. The drawing task was video-taped and transcribed verbatim.

Analytic Strategies
To analyze video and interview transcription data, we used a hybrid approach to thematic analysis [49]: we incorporated both a data-driven inductive process and a deductive process where we referenced the relevant frameworks outlined in previous studies to inform our coding. For the domain membership, we referenced the framework in Khan et al. [21] and Kim et al. [26]; for the justification of property attribution, we referenced the framework in Duuren and Scaife [52] and Melson et al. [35]. The inductive process produced a set of a priori codes that came from children's responses to interviews, and the deductive process allowed us to re-formulate our codes based on existing theories.

To establish inter-rater reliability, two coders were involved in the coding process: Coder One coded data from all participants, and Coder Two coded data from 30% to establish reliability. The two coders met weekly for one month to calibrate their coding. Specifically, of the child participants Coder One analyzed each week, Coder Two randomly selected 30% to perform the coding. Discrepancies in coding were used to iteratively refine the coding protocol until an inter-rater reliability of 85% was achieved.

Interview data
For the open-ended question on what children thought they were talking to (i.e., the domain membership question), we categorized children's responses into three groups: artifacts, living objects, and a residual category for all other descriptions based on the framework in Kahn et al. [21] and Kim et al. [26].

For the affirmative questions on property attributions, each question was coded as an affirmation (e.g., "yes," "I think so," or nodding) or a negation (e.g., "no", "I don't think so," or shaking head). In some instances, children had difficulty deciding on a response, so we created a separate category ("I don't know") to capture this type of response.

To code children's verbal response to the open-ended follow-up questions regarding their justification for their property attribution, we developed a scheme with 9 codable categories, derived from protocols used in Duuren and Scaife [52] and Melson et al. [35]. Children's justifications for thinking CAs possessed a certain property were classified as 1) **domain references** if the child relied on the domain they perceived the CA to belong to (e.g., "It is just a machine."); 2) **analogical reasoning** if the child compared the CA to other familiar objects and pointed to either similarities (e.g., "It is like phones so it can talk.") or differences (e.g., "It is not like a human so it can't remember well."); 3) **biological references** if the child indicated the CA possessed or lacked body parts or internal organs (e.g., eyes, heart); 4) **physical feature references** if the child mentioned the material the CA was made of (e.g., "It is plastic.") or the appearance of the CA (e.g., "It is orange, and a person can't be orange."); 5) **mental state references** if the child pointed to the CA's mental state, such as knowing, perceiving, and emotion (e.g., "It learns a lot of things"; "It was trying to be kind and nice to me. That's her personality."); 6) **behavioral references** if the child mentioned what the CA did (e.g., "It just read stories to me.") or how the CA behaved (e.g., "It just listened to me nicely."); 7) **reciprocity** if the child believed the CA's properties were results of others' actions, in particular, the child's own actions; 8) **mechanical references** if the child believed the CA's properties were the result of human programming (e.g., "It is made to be smart."); and 9) **fantasy reasoning** if the child attributed the CA's properties to magic or a supernatural power (e.g., "Google uses magic to listen"; "It is a witch."). One response could be coded for multiple justifications, and off-topic responses and "I don't know" were coded as invalid.

Drawing data

Data generated from the drawing task were intended to supplement the findings from interview data. Hence, we combined the two data sources when presenting the findings on children's perceptions of the CA's domain membership and properties. Given that most of children's drawings are hard to interpret without referring to their explanations, we annotated each drawing sample based on the child's verbal explanation. The drawing samples as well as the accompanying verbal accounts were coded in relation to how the CA's domain membership and cognitive, psychological, and behavioral features were exhibited in them.

RESULTS

Domain Membership

Our first research question focuses on which domain children perceived the CA as belonging to. Children's interview answers can be grouped into three domains: 1) artifacts, 2) living objects, and 3) a residual category, which was assigned for any description that is neither artificial nor living (see Table 1).

Table 1. Children's Domain Membership Categorization of the CA in Interview and Drawing

Interview	Drawing	Counts
artifacts	artifacts	13
	residual	3
residual category	residual	6
	artifacts	1
living beings	living beings	5

Over half of the children (n = 16, 57%) conceived the CA as an artifact. In interviews, children provided differing levels of specificity when describing CAs: some children broadly described the CA as a "device," "machine," or "tool," while others described the CA as a specific object, such as a "phone," "speaker," "CD-player," "robot," or "app." A small proportion of children (n = 5, 18%) viewed the CA as a living object. All of these children described it as "human," and one child specifically said that the CA was a "girl." A considerable proportion (n = 7, 25%) of children indicated that the CA was neither an artifact nor a living object. A variety of responses were grouped together into this residual domain, such as "Google is some sort of girl," "something very special that can talk like us but not a person," "a sound we can't see," "magic things to talk," and "Google is Alexa, Alexa is Google. They are not other things."

We also analyzed children's drawings using the same three coding categories as interview data. Half of the children's drawings presented the CA as a technological artifact (n = 14, 50%). These drawings suggested the children did not understand CAs as simple objects but instead as complex machines that contained multiple components [29]. These drawings typically depicted the actual shape of the smart speaker the children had interacted with (i.e., a circle). But within that outer shape, these children drew clusters of multiple objects, including wires, microphones, speakers, electricity, batteries, light bulbs, radios, or nails. The prevalence of these components may be because they were either visible (e.g., light and nails) or familiar to children from other devices they had experiences with (e.g., batteries and wires). Figure 1a contains several rounds of pink wires and a red line that connects the wires and makes them "work together," and Figure 1b contains a microphone, wires, and electricity within a circle. Each of these drawings presented the CA as a connected system with all of its parts functioning synergistically. One child noted as he pointed to the "wires," "microphones," "holes," "plugs," and "connectors" he drew, the CA has "a lot of things. All of these help it speak."

Almost one fifth of the children illustrated the CA as a human face or human-like figure (n = 5, 18%). However, none of these drawings depicted a complete human figure, but all contained the most vital elements of a human from a child's perspective. For example, Figure 2a only illustrates the CA as a face with two eyes, a nose, and a mouth yet without a body, and Figure 2 displays a girl who does not have arms or legs. Such incompleteness in human figures may reveal that although these children were inclined to identify CAs as living things, the children were also aware of some typical human features that the CA lacked.

The remainder of the drawings (n = 9, 32%) contained a mixture of representations of human and artifact elements or representations that could not be clearly categorized into either domain. For the drawings that contained both an artifact component and a living object component, children typically included a round outer shape similar to those drawings that represent the CA as a technological object, but included human figures or human body parts inside. For example, in Figure 3a, the rectangle and wavy lines represent a speaker and wires that "bring different parts together, so it won't fall apart," while inside that speaker is a human figure and the foods he/she requires. For the drawings that could not be categorized into either technological objects or living beings, children represented the CA as a wide range of varying things. For example, a child drew the CA as an apple with juice, flesh, and seed because the Google Home device was orange and the flashing lights looked like seeds. Another child drew the CA as lightning because there was a storm sound during the narrative game. A third child combined a variety of shapes and colors to represent the CA as sound (Figure 3b). This child indicated the CA was "rainbow sound" which is "happy and

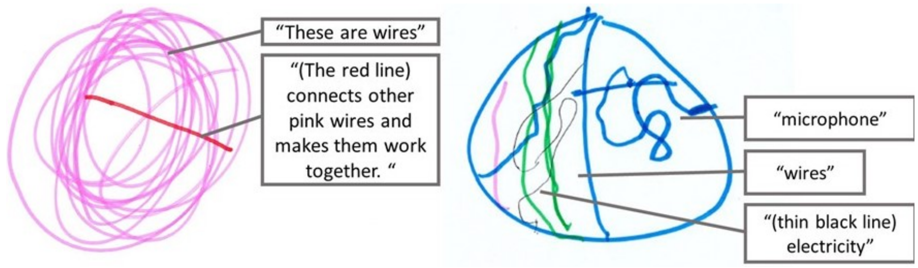

Figure 1. Figures 1a (left) and 1b (right): Drawings that illustrate Google Home Mini as artifacts

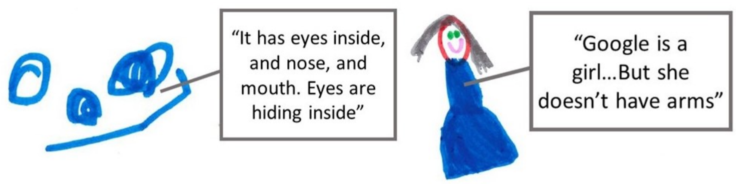

Figure 2. Figures 2a (left) and 2b (right): Drawings that illustrate Google Home Mini as living objects

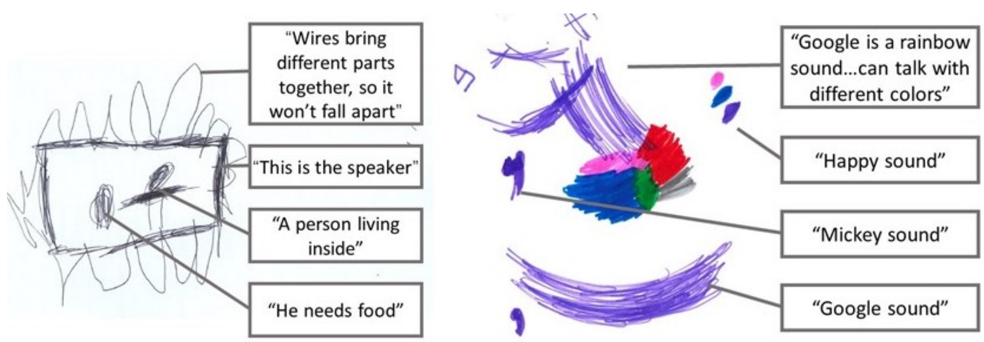

Figure 3. Figures 3a (left) and 3b (right): Drawings that illustrate Google Home Mini as a combination of artifacts and living objects or as neither artifacts nor living objects

smart." Overall, this group of drawings was centered narrowly on certain micro-level features of the CA that stood out to each child (e.g., the CA's color, its lights, or a sound it made).

For the majority of children, their drawings corroborated their interview answers (n = 24, 86%). Every child who categorized the CA as human in the interview also drew the CA with clearly anthropomorphic features. Instances in which drawings differed from interview responses only occurred among children who did not identify the CA as a human.

Property Attribution

Our second research question explored children's attribution of cognitive, psychological, and behavioral properties to CAs (Table 2). The vast majority of children believed that the CA possessed cognitive ability; these children stated that the CA was smart (n = 26, 93%) or that it could remember their previous conversation well (n = 24, 86%). A slightly smaller majority of all children attributed psychological properties to the CA, indicating that the CA could like others as a friend (n = 18, 64%) and feel emotion (n = 19, 68%). Lastly, in terms of behavioral properties, the majority of children indicated that the CA possessed speech-related capabilities (listening, n = 25, 89%; talking, n = 26, 93%), but only a quarter of these children believed that the CA could see (n = 7, 25%).

Table 2. Children's Attribution of Cognitive, Psychological, and Behavioral Properties to the CA

	Yes	No	I don't know
Cognitive			
Smart	26 (92.8%)	1 (3.6%)	1 (3.6%)
Remember	24 (85.7%)	1 (3.6%)	3 (10.7%)
Psychological			
Like	18 (64.2%)	6 (21.4%)	4 (14.3%)
Emote	19 (67.8%)	6 (21.4%)	3 (10.7%)
Behavioral			
See	7 (25.0%)	21 (75.0%)	0%
Listen	25 (89.3%)	3 (10.7%)	0%
Talk	26 (92.8%)	1 (3.6%)	1 (3.6%)

We found that children's drawings also contained these cognitive, psychological, and behavioral elements. However, given the inherent difficulty in visually representing these three elements, we focus here exclusively on those drawings where children provided relevant clarifications.

In Figure 4a, for example, a child indicated the CA's cognitive properties by drawing letters within the CA (i.e., the letters "c" and "w") to signify that the CA is "smart and knows a lot of things," and in Figure 4b, another child wrote her age (4) using her favorite color (pink) explaining that the CA put this information in its memory. Figure 5a and Figure 5b shows two children's drawings indicating psychological properties. In Figure 5a, a child drew a heart and a smiley face, commenting that the CA "knows if I am happy" and is "sometimes happy but sometimes not." In Figure 5b, a child drew a rainbow, a smiley face, and rain drops, commenting that the CA "has a rainbow inside that makes it laugh and happy" and "rains inside if Google is sad." Representations of behavioral properties were rarer. In Figure 6, a child drew a large mouth and said, "This is why Google can talk so loud."

Children's Justification of Property Attributions

We then analyzed children's explanation for their attribution of cognitive, psychological, or behavioral properties to the CA. The most frequently occurring justifications across all properties referred to the CA's presumed behaviors, biological features, mental states, or the reciprocal relationships between the child and the CA (see Table 3). Unique patterns of justifications also appeared when children were deciding whether

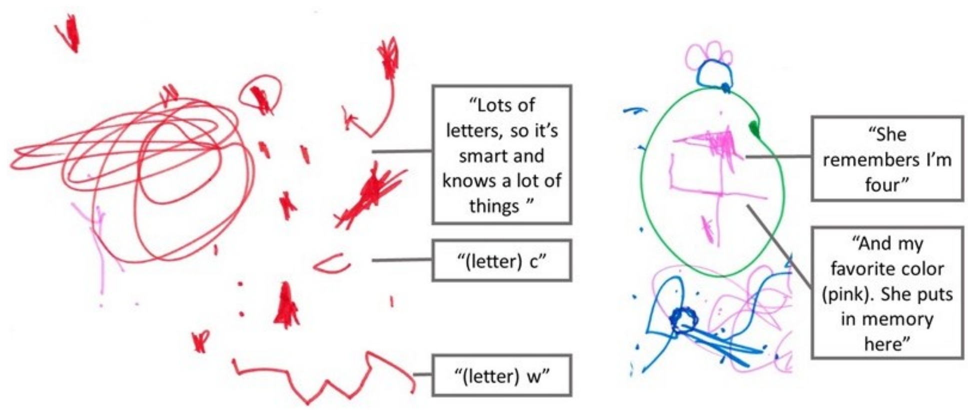

Figure 4. Figures 4a (left) and 4b (right): Drawing samples that contain cogntive elements

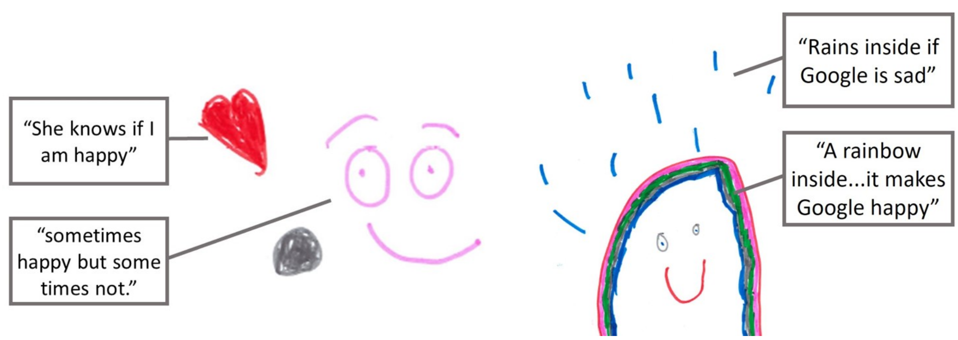

Figure 5. Figures 5a (left) and 5b (right): Drawing samples that contain psychological elements

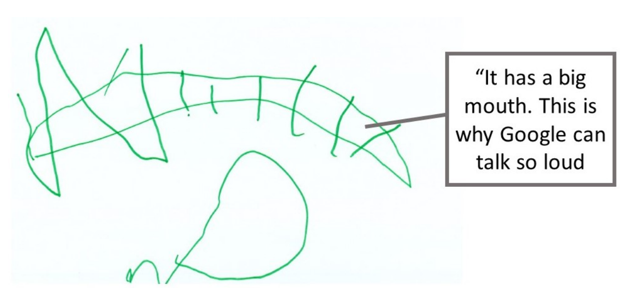

Figure 6. A drawing sample that contains behavioral elements

Table 3. Children's Justifications of the CA's Properties. Bolded Numbers Indicate Salient Justification Patterns within Each Property

	Total	Cognitive		Psychological		Behavioral		
		Smart	Rmb.	Like	Emote	See	Listen	Talk
Domain references	5	0	2	0	0	2	1	0
Analogical reasoning	13	1	3	4	3	0	1	1
Biological references	27	1	1	0	0	***13***	***6***	***6***
Physical references	10	0	0	0	0	4	2	4
Mental state references	21	4	4	6	6	1	0	0
Behavioral references	28	***12***	***7***	2	0	1	4	2
Reciprocity	24	0	5	***8***	***9***	0	2	0
Mechanical references	17	3	0	0	0	0	***6***	***8***
Fantasy reasoning	11	0	0	0	0	***3***	***5***	***3***

the CA had any of the three properties (see bolded, italicized numbers in Table 3):

- Cognitive properties were most frequently justified through behavioral references,
- psychological properties most frequently through references to reciprocity, and
- behavioral properties through biological references, mechanical causality, and fantasy reasoning.

In terms of cognitive property attribution, children commonly relied on their observation of the CA's behaviors, particularly its communication techniques. Two techniques we programmed were frequently mentioned by children as a sign of cognition, namely the repetition strategy that allows the CA to repeat what a child has said and the fallback strategy that ensures that the CA always responds to the child to prevent communication breakdowns. For example, a child stated that Google had a good memory because it "just repeated what I told her," and another child mentioned that "Google was smart because it always talks back to me."

When justifying attribution of psychological properties, children frequently referred to reciprocity; their own actions led to the CA's affective reactions. For example, one child commented that the reason why she thought Google liked her was because she was nice to Google, and another child said that Google may have felt sad when she was not listening to the story. These comments suggest that children believed the CA could reciprocate socially or emotionally on children's behaviors.

A more complex pattern was frequently observed in children's justification of behavioral properties. Children tended to first search for the biological features commonly associated with a particular behavior (i.e., eyes to see, ears to hear, and a mouth to talk). When children could not justify their attribution of a particular property through biological references (e.g., when the CA can listen or talk but doesn't have ears or a mouth), they tended to resolve the conflict through mechanical explanations or fantasy reasoning. For example, one child noted that "we installed a speaker so it can talk without a mouth," while another child said that the CA "talks with a magic mouth we can't see."

DISCUSSION

CAs as Humans, Artifacts, and What Else?

Our first research question examined children's categorization of the CA's domain membership. We found that some children associated the CA with artifacts or humans, while some children provided answers that did not fit either of these two categories. Children's categorizing CAs as artifacts or humans is consistent with the traditional A-I distinction proposed in the developmental psychology research [56]. The traditional distinction suggests that children develop their understanding of living or non-living things during early childhood and then use this schema to classify things they encounter in their daily lives [56]. In addition to these two domains, an ambiguous status of CAs among the A-I distinction was also demonstrated by children in this study, one that does not map onto artifacts or living beings. In the interview, a number of children suggested that CAs are something unique. This was further evidenced in children's drawings, with a considerable portion representing a combination of human and artifact elements. As Kahn et al. suggest, such findings may imply that children's interactions with intelligent artifacts have led to a "new ontological category" that is cutting across prototypic categories of animate and inanimate [21, 26]. However, as Kahn further pointed out, the English language may not be well equipped to characterize or talk about this new category [20], and children may thus turn to use familiar, yet less accurate, terms to describe their perceived domains of CAs. Moreover, if we conceive of these categories as existing on a continuum, our evidence suggests that this new ontological category may be closer to the technological artifact side of the continuum [52]. Every child who spoke of the CA as human also drew the CA accordingly, but approximately 20% of the other children exhibited some level of inconsistency in their depictions of the CA during the interview and drawing sessions. Taken all together, our evidence suggests that a strict distinction between animate and inanimate may fail to accurately capture children's conceptions of CAs that appear to be more nuanced and multifaceted.

Highlighted Cognition and Speech Properties

Our second research question explored what properties children perceive the CA to possess. We found that children, overall, assigned many animistic abilities to the CA. Further, children understood CAs as possessing a unique constellation of properties: almost all of the children in this study ascribed cognitive and speech-related behavioral properties to the CA, while fewer children ascribed psychological and non-speech related behavioral properties.

Children overwhelmingly believed the CA to possess cognitive and speech abilities. This is not surprising as the ability to converse intelligently is a defining feature of CAs [32]. Nevertheless, we note that the ability to listen and talk can be considered from either a cognitive or behavioral perspective [34]. A cognitive perspective emphasizes the mental aspects involved in listening and talking (i.e., understanding, interpreting, responding), while a behavioral perspective highlights the CA's actions (i.e., listening quietly, making sounds). It would have been difficult, if not impossible, to truly elicit the underlying perspective of the young children in this study. However, our evidence suggests a strong association between cognitive and speech-related behavioral attributions: almost all children believed the CA possessed abilities in these two categories. Also, when children justified attributing cognitive abilities to the CA, they commonly referred to the CA's speech behaviors, in particular, the communication strategies we programmed. Taken together, this implies that children in this study may understand the CA's speech behaviors as indicators of cognitive abilities.

Children's responses were more heterogeneous regarding the CA's psychological abilities; slightly over half of the children believed the CA had the ability to like others and feel emotions. In the context of the broader literature, this proportion places CAs somewhere between robots and computers in terms of perceived psychological properties. While Melson et al. [35] and Weiss et al. [53] suggested most children believed robots could experience happiness and sadness, Scaife and van Durren found that only 20% of five-year-olds attributed those same abilities to computers [45]. Children's differing perceptions of psychological properties of computers, CAs, and robots may be largely due to these artifacts' varying expressive abilities [54]. As discussed in Johnstone and Scherer, spoken languages utilize acoustic qualities (i.e., tones, pauses, pitch) that convey affective and social signals which go beyond the content of the speech [19]. In this study, we speculate that the CA's ability to engage in natural spoken conversation may have led the majority of children to ascribe psychological properties to the CA despite its lack of embodiment.

Justification of CA Properties

Our third research question looked at the explanations children provided when justifying their attribution of properties to the CA. Overall, we identified nine distinct strategies children used to decide whether the CA possessed certain abilities.

Two strategies have already been described in prior research: children may regard the CA's capabilities as either a result of programming (i.e., mechanical references) or as a result of natural intelligence (i.e., mental state references) [30]. The former perspective ascribes no intentions to the artifact and considers its ability to arise from human design, while the latter ascribes intentions and awareness to the artifact itself. Relatedly, some children used justification that focused on reciprocity – the contingencies between the child's actions and those of the CA. While some justifications hinged on what appeared to be automatic responses (e.g., the CA listened only because the child spoke more loudly), many justifications hinted at a perceived social reciprocity (e.g., the CA liked a child because she was nice to the CA). These latter justifications, overlapping with mental state references, suggest that children might view CAs as psychological entities and form social relationships with CAs, a speculation which expands previous findings that show children form relationships with robots [22, 53].

Evidence from our study suggests another perspective involving fantasy reasoning. Some children relied on magical thinking or supernatural justification to explain how the CA could have speech abilities yet lack the human body parts necessary for those abilities.

Another form of justification strategy uses empirical observation, referring to children's focus on CAs' behaviors and physical and biological features. This justification strategy echoes Rucker and Pinkwart's assertion that children's actual interactions with intelligent artifacts impact the way children construct mental models of the aliveness status of such objects [44]. We expect that the increase in children's experiences with CAs may lead to more nuanced views [2].

Lastly, some children relied on the perceived domain of CAs and used such perception as a premise to reason their abilities (i.e., domain references and analogical reasoning). They explained that CAs possessed certain abilities because they belonged to a certain animate domain. However, such strategy did not occur frequently in our study. The infrequency of this strategy may be due to CAs' straddling the boundaries between animacy and inanimacy, thus creating difficulties for children to firmly associate them with either domain in the first place [18].

Tentative Age Trend

Although not one of the research questions, we noticed a possible association between age and children's identification of domain membership. The oldest children in our study (i.e., 6-year-olds) all described the CA as a technological object in the interview and in the drawings, while the answers from younger children (i.e., 3- to 5-year-olds) were mixed. This may be due to the older children's more advanced understanding of programmable machines. Younger children have less awareness of this concept and tend to make sense of computational objects by personifying them [7]. Druga et al. provided evidence for this hypothesis in their study of children who observed a robot solving a maze problem [11]. One third of the children between the ages of four to seven credited the robot's successfully solving the problem to its innate cognitive capability, while none of the children aged eight to ten did. The latter group was much more likely to think that the robot was programmed to perform such strategies. However, the children's age range within each group (i.e., 4-7 years and 8-10 years) was wide, which may have obscured developments occurring within each group of children, particular the younger group. Future studies, if carefully structured, may produce more nuanced findings regarding young children's development of perceptions about CAs.

Design Implications

A number of CAs are being developed to provide young children with learning opportunities or social companionship. Findings from our study may help improve the design of such CAs in two ways. First, children's recognition of the CA's cognitive capabilities is an encouraging sign, as children have been found to selectively seek information and learn from those they believe to be intelligent [13]. As demonstrated in Breazeal et al. [5], preschoolers are more willing to trust the information from smarter robots that can provide contingent responses. These children remembered more information from and talked more with the contingent robot than with a non-contingent robot. As such, CAs should be best designed so as to elicit children's attribution of cognitive abilities. In our study, we identified some reoccurring communication techniques that children recognized as a sign of being smart, including the repetition and fallback mechanism we programmed. Children commonly commented that the CA was able to remember and understand because it "repeats" what they said; they also said that the CA always responds to them (even in the case when the CA actually failed to understand). These two techniques both amplified the CA's role as an active interlocutor that is capable of engaging in contingent interactions. Developers may consider incorporating these two communication techniques.

Second, as compared to children's overwhelming recognition of CAs' cognitive properties, children's attitudes regarding whether CAs are psychological entities were mixed. This challenges researchers to develop CAs that children are more willing to engage with socially [3, 4]. Even though the disembodiment of CAs such as Alexa and Google Assistant may prevent them from leveraging the full range of psychological cues (i.e., facial expressions and body language) [39], researchers can compensate for this lack by improving on such CAs' conversational expressiveness [14]. CAs may be designed to talk explicitly about their emotions or leverage natural acoustic features (i.e., tone, prosody, speech speed), which may more consistently elicit children's affective reactions and thus may be more likely to lead children to treat CAs as psychological entities. However, while it is important to increase the human-likeness of CAs, we note that children's attribution of human-like qualities to CAs may potentially open children up to undue influence (e.g., misinformation). As such, designers should keep in mind to ensure the content appropriateness of CAs' conversations.

Limitations and Future Work

We noticed three potential limitations during the course of the study. First, our study provided children with opportunities to interact with a CA within a controlled environment. However, children's perceptions may be less about CAs in general and more about the particular CA they interacted with. We addressed this issue by designing the study to cover the typical interactions a child would have with a CA. Second, children's perceptions may be associated with their differing levels of prior experience with conversational technologies. While we did not formally test for this relationship, anecdotal evidence in our study suggests such a relation. For example, when answering a question about whether the CA could listen, one child replied yes and explained that from his previous experience at home, calling Alexa's name would always wake her up. Third, while we suspect a relation between children's overall development and their perceptions of CAs, the small sample size of this study limits our ability to statistically examine this relationship. Future studies should be carried out with larger sample sizes to tackle this issue.

CONCLUSION

CAs, such as Apple Siri, Google Assistant, and Amazon Alexa, play an increasingly important role in young children's technological landscapes and life worlds. While a handful of studies have documented children's natural interactions with CAs, little is known about children's perceptions of CAs. To fill this gap, we examined three- to six-year-olds' perceptions of CAs' domain membership and properties, as well as their justifications for these perceptions. Overall, these three research questions yielded converging evidence that children sometimes take a more nuanced position and spontaneously attribute both artifact and animate properties to CAs. At least some children appeared unwilling to describe the CA as either a living being or an artifact. These children described the CA as either being a combination of these two categories or fitting into some other third category. Additionally, children appeared to consistently conceive of CAs as possessing a unique constellation of animate properties while lacking others. Almost all of the children in this study ascribed cognitive and speech-related behavioral properties to the CA, while fewer children ascribed psychological and non-speech related behavioral properties. This also reflects children's dilemma in determining CAs' animacy domains. Examination of children's justifications for their perceptions further revealed nuanced reasoning. Taken together, these findings extend current research on children's perceptions of intelligent artifacts by adding CAs as a new genre of study and also provide some underlying knowledge that may guide the development of CAs to support young children's cognitive and social development.

ACKNOWLEDGEMENT

The authors would like to thank the dedicated research assistants in the Digital Learning Lab at the University of California, Irvine for assisting with data collection and data coding. We also thank the children, families, and childcare centers that participated in this study.

REFERENCES

[1] Tanya N Beran, Alejandro Ramirez-Serrano, Roman Kuzyk, Meghann Fior, and Sarah Nugent. 2011. Understanding how children understand robots: Perceived animism in child–robot interaction. *International Journal of Human-Computer Studies* 69, 7-8 (2011), 539–550.

[2] Debra Bernstein and Kevin Crowley. 2008. Searching for signs of intelligent life: An investigation of young children's beliefs about robot intelligence. *The Journal of the Learning Sciences* 17, 2 (2008), 225–247.

[3] Jason Borenstein and Yvette Pearson. 2013. Companion robots and the emotional development of children. *Law, Innovation and Technology* 5, 2 (2013), 172–189.

[4] Cynthia Breazeal. 2009. Role of expressive behaviour for robots that learn from people. *Philosophical Transactions of the Royal Society B: Biological Sciences* 364, 1535 (2009), 3527–3538.

[5] Cynthia Breazeal, Paul L Harris, David DeSteno, Jacqueline M Kory Westlund, Leah Dickens, and Sooyeon Jeong. 2016. Young children treat robots as informants. *Topics in cognitive science* 8, 2 (2016), 481–491.

[6] Kaitlin L Brunick, Marisa M Putnam, Lauren E McGarry, Melissa N Richards, and Sandra L Calvert. 2016. Children's future parasocial relationships with media characters: the age of intelligent characters. *Journal of Children and Media* 10, 2 (2016), 181–190.

[7] David Cameron, Samuel Fernando, Abigail Millings, Roger Moore, Amanda Sharkey, and Tony Prescott. 2015. Children's age influences their perceptions of a humanoid robot as being like a person or machine. In *Conference on Biomimetic and Biohybrid Systems*. Springer, 348–353.

[8] Alfonso Caramazza and Jennifer R Shelton. 1998. Domain-specific knowledge systems in the brain: The animate-inanimate distinction. *Journal of cognitive neuroscience* 10, 1 (1998), 1–34.

[9] Kara Chan. 2006. Exploring children's perceptions of material possessions: a drawing study. *Qualitative Market Research: An International Journal* 9, 4 (2006), 352–366.

[10] Stefania Druga, Randi Williams, Cynthia Breazeal, and Mitchel Resnick. 2017. Hey Google is it OK if I eat you?: Initial explorations in child-agent interaction. In *Proceedings of the 2017 Conference on Interaction Design and Children*. ACM, 595–600.

[11] Stefania Druga, Randi Williams, Hae Won Park, and Cynthia Breazeal. 2018. How smart are the smart toys?: children and parents' agent interaction and intelligence attribution. In *Proceedings of the 17th ACM Conference on Interaction Design and Children*. ACM, 231–240.

[12] Susan Edwards, Andrea Nolan, Michael Henderson, Ana Mantilla, Lydia Plowman, and Helen Skouteris. 2018. Young children's everyday concepts of the internet: A platform for cyber-safety education in the early years. *British journal of educational technology* 49, 1 (2018), 45–55.

[13] Paul L Harris and Kathleen H Corriveau. 2011. Young children's selective trust in informants. *Philosophical Transactions of the Royal Society B: Biological Sciences* 366, 1567 (2011), 1179–1187.

[14] Marcel Heerink, Ben Kröse, Vanessa Evers, and Bob Wielinga. 2010. Relating conversational expressiveness to social presence and acceptance of an assistive social robot. *Virtual reality* 14, 1 (2010), 77–84.

[15] Katharine E Heintz and Ellen A Wartella. 2012. Young children's learning from screen media. *Communication Research Trends* 31, 3 (2012), 22.

[16] Jennifer Hill, W Randolph Ford, and Ingrid G Farreras. 2015. Real conversations with artificial intelligence: A comparison between human–human online conversations and human–chatbot conversations. *Computers in Human Behavior* 49 (2015), 245–250.

[17] Jennifer L Jipson and Susan A Gelman. 2007. Robots and rodents: Children's inferences about living and nonliving kinds. *Child development* 78, 6 (2007), 1675–1688.

[18] Jennifer L Jipson, Selin Gülgöz, and Susan A Gelman. 2016. Parent–child conversations regarding the ontological status of a robotic dog. *Cognitive development* 39 (2016), 21–35.

[19] Tom Johnstone and Klaus R Scherer. 2000. Vocal communication of emotion. *Handbook of emotions* 2 (2000), 220–235.

[20] Peter H Kahn, Nathan G Freier, Batya Friedman, Rachel L Severson, and Erika N Feldman. 2004. Social and moral relationships with robotic others?. In *RO-MAN 2004. 13th IEEE International Workshop on Robot and Human Interactive Communication (IEEE Catalog No. 04TH8759)*. IEEE, 545–550.

[21] Peter H Kahn, Batya Friedman, Deanne R Perez-Granados, and Nathan G Freier. 2006. Robotic pets in the lives of preschool children. *Interaction Studies* 7, 3 (2006), 405–436.

[22] Peter H Kahn Jr, Heather E Gary, and Solace Shen. 2013. Children's social relationships with current and near-future robots. *Child Development Perspectives* 7, 1 (2013), 32–37.

[23] Peter H Kahn Jr, Takayuki Kanda, Hiroshi Ishiguro, Nathan G Freier, Rachel L Severson, Brian T Gill, Jolina H Ruckert, and Solace Shen. 2012. "Robovie, you'll have to go into the closet now": Children's social and moral relationships with a humanoid robot. *Developmental psychology* 48, 2 (2012), 303.

[24] Junko Kanero, Vasfiye Geçkin, Cansu Oranç, Ezgi Mamus, Aylin C Küntay, and Tilbe Göksun. 2018. Social robots for early language learning: Current evidence and future directions. *Child Development Perspectives* 12, 3 (2018), 146–151.

[25] Nobuko Katayama, Jun'ichi Katayama, Michiteru Kitazaki, and Shoji Itakura. 2010. Young children's folk knowledge of robots. *Asian Culture and History* 2, 2 (2010), 111.

[26] Minkyung Kim, Soonhyung Yi, and Donghun Lee. 2019. Between living and nonliving: Young children's animacy judgments and reasoning about humanoid robots. *PloS one* 14, 6 (2019), e0216869.

[27] Yanghee Kim, Amy L Baylor, and E Shen. 2007. Pedagogical agents as learning companions: the impact of agent emotion and gender. *Journal of Computer Assisted Learning* 23, 3 (2007), 220–234.

[28] Jacqueline M Kory-Westlund and Cynthia Breazeal. 2019. Assessing children's perceptions and acceptance of a social robot. In *Proceedings of the 18th ACM International Conference on Interaction Design and Children*. ACM, 38–50.

[29] Sunok Lee, Sungbae Kim, and Sangsu Lee. 2019. What does your Agent look like?: A Drawing Study to Understand Users' Perceived Persona of Conversational Agent. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, LBW0143.

[30] Sharona T Levy and David Mioduser. 2008. Does it "want" or "was it programmed to..."? Kindergarten children's explanations of an autonomous robot's adaptive functioning. *International Journal of Technology and Design Education* 18, 4 (2008), 337–359.

[31] Silvia B Lovato, Anne Marie Piper, and Ellen A Wartella. 2019. Hey Google, Do Unicorns Exist?: Conversational Agents as a Path to Answers to Children's Questions. In *Proceedings of the 18th ACM International Conference on Interaction Design and Children*. ACM, 301–313.

[32] Ewa Luger and Abigail Sellen. 2016. Like having a really bad PA: the gulf between user expectation and experience of conversational agents. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 5286–5297.

[33] Peter J Marshall and Kimberly Brenneman. 2016. Young Children's Developing Understanding of the Biological World. (2016).

[34] Donald Meichenbaum. 2017. Teaching thinking: A cognitive behavioral perspective. In *The Evolution of Cognitive Behavior Therapy*. Routledge, 85–104.

[35] Gail F Melson, Peter H Kahn Jr, Alan Beck, Batya Friedman, Trace Roberts, Erik Garrett, and Brian T Gill. 2009. Children's behavior toward and understanding of robotic and living dogs. *Journal of Applied Developmental Psychology* 30, 2 (2009), 92–102.

[36] Pekka Mertala. 2019. Young children's perceptions of ubiquitous computing and the Internet of Things. *British Journal of Educational Technology* (2019).

[37] TASSOS A Mikropoulos, PLOUSSIA Misailidi, and FOTINI Bonoti. 2003. Attributing human properties to computer artifacts: developmental changes in children's understanding of the animate-inanimate distinction. (2003).

[38] John E Opfer and Susan A Gelman. 2011. Development of the animate-inanimate distinction. *The Wiley-Blackwell handbook of childhood cognitive development* 2 (2011), 213–238.

[39] Catherine Pelachaud, Rodolphe Gelin, Jean-Claude Martin, and Q Anh Le. 2010. Expressive gestures displayed by a humanoid robot during a storytelling application. *New Frontiers in Human-Robot Interaction (AISB), Leicester, GB* (2010).

[40] Amanda Purington, Jessie G Taft, Shruti Sannon, Natalya N Bazarova, and Samuel Hardman Taylor. 2017. Alexa is my new BFF: social roles, user satisfaction, and personification of the amazon echo. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 2853–2859.

[41] David H Rakison and Diane Poulin-Dubois. 2001. Developmental origin of the animate–inanimate distinction. *Psychological bulletin* 127, 2 (2001), 209.

[42] 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.

[43] David La Rooy, Margaret-Ellen Pipe, and Janice E Murray. 2007. Enhancing children's event recall after long delays. *Applied Cognitive Psychology* 21, 1 (2007), 1–17.

[44] Michael T Rücker and Niels Pinkwart. 2016. Review and discussion of children's conceptions of computers. *Journal of Science Education and Technology* 25, 2 (2016), 274–283.

[45] Michael Scaife and Mike van Duuren. 1995. Do computers have brains? What children believe about intelligent artifacts. *British Journal of Developmental Psychology* 13, 4 (1995), 367–377.

[46] Alex Sciuto, Arnita Saini, Jodi Forlizzi, and Jason I Hong. 2018. Hey Alexa, What's Up?: A mixed-methods studies of in-home conversational agent usage. In *Proceedings of the 2018 Designing Interactive Systems Conference*. ACM, 857–868.

[47] Sofia Serholt and Wolmet Barendregt. 2016. Robots tutoring children: Longitudinal evaluation of social engagement in child-robot interaction. In *Proceedings of the 9th nordic conference on human-computer interaction*. ACM, 64.

[48] Rachel L Severson and Stephanie M Carlson. 2010. Behaving as or behaving as if? Children's conceptions of personified robots and the emergence of a new ontological category. *Neural Networks* 23, 8-9 (2010), 1099–1103.

[49] Jon Swain. 2018. A hybrid approach to thematic analysis in qualitative research: Using a practical example. *Sage research methods* (2018).

[50] Anuj Tewari and John Canny. 2014. What did spot hide?: a question-answering game for preschool children. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1807–1816.

[51] Patrice D Tremoulet and Jacob Feldman. 2000. Perception of animacy from the motion of a single object. *Perception* 29, 8 (2000), 943–951.

[52] Mike van Duuren and Michael Scaife. 1996. "Because a robot's brain hasn't got a brain, it just controls itself"—Children's attributions of brain related behaviour to intelligent artefacts. *European Journal of Psychology of Education* 11, 4 (1996), 365.

[53] Astrid Weiss, Daniela Wurhofer, and Manfred Tscheligi. 2009. "I love this dog"—children's emotional attachment to the robotic dog AIBO. *International Journal of Social Robotics* 1, 3 (2009), 243–248.

[54] Travis J Wiltshire, Emilio JC Lobato, David R Garcia, Stephen M Fiore, Florian G Jentsch, Wesley H Huang, and Benjamin Axelrod. 2015. Effects of robotic social cues on interpersonal attributions and assessments of robot interaction behaviors. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 59. SAGE Publications Sage CA: Los Angeles, CA, 801–805.

[55] Julia Woodward, Zari McFadden, Nicole Shiver, Amir Ben-hayon, Jason C Yip, and Lisa Anthony. 2018. Using co-design to examine how children conceptualize intelligent interfaces. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 575.

[56] Kristyn Wright, Diane Poulin-Dubois, and Elizabeth Kelley. 2015. The animate–inanimate distinction in preschool children. *British Journal of Developmental Psychology* 33, 1 (2015), 73–91.

[57] Ying Xu and Mark Warschauer. 2019. Young Children's Reading and Learning with Conversational Agents. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, CS10.