



What Do Children and Parents Want and Perceive in Conversational Agents? Towards Transparent, Trustworthy, Democratized Agents

Jessica Van Brummelen
jess@csail.mit.edu
Massachusetts Institute of Technology
Cambridge, Massachusetts, USA

Mingyan Claire Tian
mt1@wellesley.edu
Wellesley College
Wellesley, Massachusetts, USA

Maura Kelleher
maurakel@mit.edu
Massachusetts Institute of Technology
Cambridge, Massachusetts, USA

Nghi Hoang Nguyen
nghin@mit.edu
Massachusetts Institute of Technology
Cambridge, Massachusetts, USA



Figure 1: A portion of child participant responses during an ideation design session about their ideal conversational agents.

ABSTRACT

Historically, researchers have focused on analyzing Western, Educated, Industrialized Rich and Democratic (WEIRD), adult perspectives on technology. This means we may not have technology developed appropriately for children and those from non-WEIRD countries. In this paper, we analyze children and parents from various countries' perspectives on an emerging technology: conversational agents. We aim to better understand participants' trust of agents, partner models, and their ideas of "ideal future agents" such that researchers can better design for these users—for instance, by ensuring children do not overtrust agents. Additionally, we empower children and parents to program their own agents through educational workshops, and present changes in perceptions as participants create and learn about agents. Results from the study (n=49) included how children felt agents were significantly more human-like, warm, and dependable than parents did, how overall participants trusted agents more than parents or friends for correct

information, how children described their ideal agents as being more artificial than human-like than parents did, and how children tended to focus more on fun features, approachable/friendly features and addressing concerns through agent design than parents did, among other results. We also discuss potential agent design implications of the results, including how designers may be able to best foster appropriate levels of trust towards agents by focusing on designing agents' competence and predictability indicators, as well as increasing transparency in terms of agents' information sources.

CCS CONCEPTS

- **Human-centered computing** → **Natural language interfaces**; *User models*; *User interface programming*;
- **Social and professional topics** → *Children*; *Age*; *Cultural characteristics*; *K-12 education*;
- **Computing methodologies** → *Intelligent agents*.

KEYWORDS

conversational agents, chatbots, virtual assistants, conversational AI, non-WEIRD and WEIRD, parents, trust, partner models, agent personification, computational action, technology democratization

ACM Reference Format:

Jessica Van Brummelen, Maura Kelleher, Mingyan Claire Tian, and Nghi Hoang Nguyen. 2023. What Do Children and Parents Want and Perceive in Conversational Agents? Towards Transparent, Trustworthy,



This work is licensed under a Creative Commons Attribution International 4.0 License.

IDC '23, June 19–23, 2023, Chicago, IL, USA
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0131-3/23/06.
<https://doi.org/10.1145/3585088.3589353>

Democratized Agents. In *Interaction Design and Children (IDC '23)*, June 19–23, 2023, Chicago, IL, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3585088.3589353>

1 INTRODUCTION

Conversational artificial intelligence (AI)—or the ability of a computer program to understand human language and respond accordingly—is ripe with potential. Imagine a conversational agent engaging children in learning history with a virtual Rosa Parks, or an agent providing constant, accurate healthcare answers to those in need. With recent major advances in natural language processing and automatic speech recognition these ideas are not far-fetched [8, 10, 15, 21, 24, 29, 58].

Nonetheless, current agents, like Google Home, Apple’s Siri and Amazon Alexa, still misrecognize speech and misunderstand intent [5, 5, 49]. For instance, researchers found speech recognition systems by Amazon, Google, IBM and Microsoft did substantially worse when recognizing black speakers versus white [27]. Others have found significant gender biases in embeddings [7, 63]. Biases in AI systems are widespread, and if users are not aware of such flaws, there could be serious implications, including misinformation being spread, human bias being compounded, and users unwittingly acting on incorrect advice [45].

Ideally, agents would be developed to portray the reality of their abilities and limitations to their users through effective design. In a study with AI decision-aids, researchers describe how if users are too averse to technology’s advice and information, they cannot truly benefit from using the technology. However, if they are too appreciative, users may make ill-informed decisions when technology presents incorrect information [20]. By portraying conversational agents in an honest way through design, discrepancies between users’ expectations of agents—or their agent “partner models”—and the reality of agents can be reduced, which can also reduce user frustration [17].

In our study, we investigate users’ perceptions of agents, including their partner models (defined in this context as dialogue systems’ perceived competence and communicative ability [17]) and trust. The results revealed how for certain aspects of agents—including warmth, human-likeness and dependability—children perceived agents differently than parents. Participants’ general trust of agents’ correctness (compared to other people’s and systems’ correctness), however, was similar for both children and parents. In general, people trusted agents more than their friends and parents. Based on these results and others, we discuss agent design recommendations to foster appropriate levels of trust of agents.

Historically, human-computer interaction research has largely recruited participants from Western, Educated, Industrialized, Rich and Democratic (WEIRD) countries, who comprise less than 12% of the world’s population [22, 30]. This means many of the design recommendations developers use are likely biased towards this population. Furthermore, a large portion of software developers reside in WEIRD countries [18, 26, 59], meaning technology development is likely further biased towards the WEIRD population. In order to address this, and develop technology meaningful and relevant to more of the world, researchers have developed different strategies. One strategy involves including more participants from non-WEIRD countries and developing recommendations based on wider demographics [48]. We utilize this strategy through involving

participants from non-WEIRD and WEIRD countries, investigating their perceptions of agents, and asking them how they envision their “ideal conversational agents”. The results and recommendations aim to provide agent designers with perspectives from those from different countries and generations.

Another strategy to reduce the gap between non-WEIRD- and WEIRD-centric technology is to empower those from non-WEIRD countries to develop their own technology. There are a number of tools that help enable nearly anyone to develop technology, many of which utilize visual or block-based coding [23]. These tools have largely been born out of the constructionist movement in education, which encourages the use of low-floor, high-ceiling programming tools to empower a wide variety of people to learn to program, including other underrepresented groups in the technology sector, like children [43]. Scratch, for instance, allows children to program their own web-based animations using block-based coding [44]. Other low-floor platforms enable users to develop conversational agents, including the Flow Editor and Alexa Blueprints [2, 25]. The MIT App Inventor platform allows users to develop fully-fledged apps, which can be deployed to mobile devices’ app stores [60], as well as conversational agents, which can be deployed to Amazon Alexa devices, through “ConvoBlocks” [35, 51, 52, 57].

In this paper, we aim to democratize conversational agent technology to young learners from various countries and their parents through an educational intervention with the ConvoBlocks platform. This intervention empowers students to develop their own agents. We adopt ConvoBlocks in our study, as it is open-source and has a low barrier to creating deployable agents [51, 52]. Through constructionist workshops with this tool, we inform participants about how agents work and technology’s societal impact. Our contributions include a novel study of partner models and trust of agents as children and parents learn about agents; a study of how children and parents envision the future of agents; and a discussion of the potential implications of the results on how developers design conversational agents.

1.1 Research Questions

Through engaging children and parents from various countries in conversational agent and societal impact curriculum, including agent-development, learning, and design sessions, we aimed to answer the following research questions:

RQ1: How do children and parents perceive Alexa with respect to partner models [17] and trust before, during and after conversational agent development and societal impact activities?

RQ2: How do children and parents envision the future of conversational agents?

We discuss the results of these research questions with respect to conversational agent design. (Note that due to space constraints, we address additional research questions related to pedagogy from this study in another paper [57].)

2 BACKGROUND AND RELATED WORK

2.1 Trust and Relationships with Agents

Because conversation is one of the most intuitive, primary methods humans use to communicate with each other, conversational interfaces are uniquely positioned to inspire relational interactions with technology [40, 47]. For instance, an agent recently won a

Peabody Award for engaging in “emotional interactions, empathy, and connection” [13]. Furthermore, researchers have found correlations between human-agent relationship development and increased trust of agents [47]. Considering how trust is a key factor in misinformation spread [46, 61], we decided to specifically investigate people’s trust of agents’ correctness in this study. We also chose to emphasize children’s trust in this study, as the risks associated with misinformation spread could be particularly acute with children, especially since they do not have the same critical analysis skills as adults [28, 50].

Other studies have investigated people’s trust of conversational agents’ correctness. One example includes a study in which clinicians decide whether or not to utilize agents’ advice on diagnoses [20]; another includes a study in which customers decide whether or not to follow agents’ recommendations [33]. Nonetheless, few studies have investigated children’s or those from non-WEIRD countries’ trust of agents [19]. Even fewer have investigated how this trust may change through educational interventions. One example includes a study in which children engage in social robot curriculum, including modules on conversational AI, computer vision and societal impact, among others [16]. If participants engaged in the societal impact module, their trust of the robot generally decreased [16]. Another example includes a study with ConvoBlocks in which students engaged in curriculum entirely focused on conversational agents, including their societal impact. In this study, researchers did not find any significant differences in trust through the curriculum. They did, however, observe concerning correlations between children’s perceived friendliness and trust of agents [56]. In both of these studies, however, the researchers only investigated general trust.

Many researchers have developed methods to investigate specific aspects of trust, such that developers can better assess which aspects of their technology affect such trust [11]. In our study, we adopt McKnight and Chervany’s widely-used model, which has four main components: (1) **competence**, (2) **benevolence**, (3) **integrity** and (4) **predictability** [34]. In our study, we found children most often referred to competence and predictability when discussing trust. We discuss potential implications of this on agent design in later sections.

2.2 Partner Models and Perceptions of Agents

People’s partner models, or mental models of their conversational partners, can significantly affect how they interact with agents. For instance, researchers have found that people make different language choices depending on their initial expectations of partner models [14, 17]. Partner models can be described in terms of three main dimensions: (1) **competence and dependability**, (2) **human-likeness**, and (3) **cognitive flexibility** [14, 17]. Designing agents that produce partner models that align with the capabilities of the agent (e.g., producing a partner model of perceived limited flexibility, if the agent is truly limited in flexibility), could help minimize user frustrations and ease conversation [17]. However, a deep understanding of conversational agent users’ partner models—and especially children’s partner models—is not reflected in the literature [17, 19].

Certain studies have investigated children’s general perceptions of conversational agents. For instance, one study found that the majority of 5-6 year old children considered agents to be friendly, alive, trustworthy, safe, funny, and intelligent [32]. Another study investigated 3-10 year old children’s perceptions, and found that children had different perceptions of agents’ intelligence depending on the modality of interaction with conversational agents. Others found students perceived agents to be more intelligent and felt closer to them after learning to program them [56]. None of these studies specifically investigated children’s partner models of agents.

2.3 Diversity of Agent Design Perspectives

In the past few years, a large number of researchers have developed much-needed conversational agent design guidelines [4, 12, 38, 39, 62]. In developing such guidelines, researchers have gained insight from classical human-computer interaction research, like Nielsen and Norman [37], to pop-culture icons, like the Star Trek agent [4]. The number and breadth of recent agent design guidelines shows the importance of improving conversational agent user experience; however, the vast majority of human-computer interaction research these guidelines are based on are heavily biased towards WEIRD, adult perspectives [22, 30, 41, 42, 48]. To begin filling this gap, more research needs to investigate perspectives from children and those from non-WEIRD countries. In our study, we investigate perspectives on agents and the future of the technology from such underrepresented groups. Through this research, we aim to increase the diversity of perspectives in conversational agent design and provide a stepping stone for future agent design considerations.

3 PROCEDURE

3.1 Developing Agents with the ConvoBlocks Platform

ConvoBlocks is an open-source, block-based programming platform within the App Inventor environment, which allows nearly anyone to program conversational agents [35, 51, 60]. To do so, students first define their agent’s *invocation name* (e.g., “My Carbon Footprint Agent”), *intents* (e.g., groups of phrases like, “Calculate my carbon footprint”, “What’s my carbon footprint?”, etc.) and *entities* (e.g., information units like number of miles driven, kilowatts of energy used, etc.) the agent should be able to recognize. Through the process of agent development, students learn conversational agent terminology and concepts, which are described in-detail in the appendix [53]. Next, students define how the agent responds to the defined intents (e.g., “You have a carbon footprint of 11 tonnes/year”). They can do so using the web pages shown in Figure 2. After this, students can test their agent on ConvoBlocks, or deploy their agents to any Alexa-enabled devices, like the Alexa mobile app or an Echo Spot [52].

3.2 Workshops

As shown in Table 1, the workshops consisted of two 3-hour Zoom classes taught in English by three researchers, and two professionals working in the area of technology impact. Additionally, approximately four teaching assistants were available to answer questions and provide technical help in Zoom rooms at any given time. Each

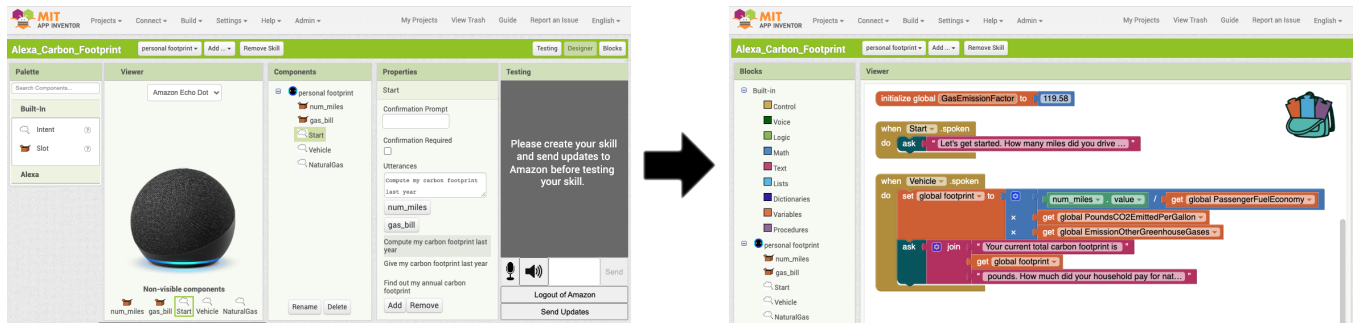


Figure 2: Two web pages from ConvoBlocks [35], allowing users to define invocation names, intents and entities, and then program agents’ responses to intents.

child-parent pair engaged in the workshops on their own Zoom account and a computer in their own environment (e.g., home). The first day of the curriculum taught participants to program agents that responded to questions about carbon footprints, as shown in Figure 4. Instructors led participants step-by-step through two conversational agent development tutorials. Participants received PDF versions of the tutorials, such that they could complete them at their own pace. They also received a third “challenge tutorial” PDF, which they could attempt if they finished early. The code for the third tutorial was explained at the end of the first day. The group also completed an ideation session on the first day. They responded to prompts about what their “ideal” agent would look like, sound like, do, and say (among other prompts) using a virtual whiteboard (with separate sections for children and parents). Sections of the whiteboard are shown in Figure 1 and Figure 3. The researchers provided approximately 20 minutes for the participants to add ideas to the whiteboard on their own. Afterwards, the researchers gave a brief summary to the participants about what they noticed on the whiteboard.

The second day included presentations and group discussions about societal impact of technology. Participants gathered in groups of 2-4 children with their parents for the discussions. The presentations encouraged participants to think about the positive and negative impact of technology; the discussions explored how technology could help address world problems, like sustainability, with an emphasis on conversational agents as part of the solution. In the final activity, small groups of participants presented their proposed solutions to the entire group. They had the opportunity to design conversational agents, which they could demonstrate in their presentations. Overall, the workshops aimed to teach participants conversational agent concepts described in the appendix [53], and focused specifically on eight of the concepts: *Training, Intents, Agent modularization, Entities, Events, Testing, Turn-taking, and Societal impact and ethics*. (For detailed content from the workshops, including the tutorials, refer to the thesis in [54].)

4 THE STUDY

4.1 Participants

Study participants came from various backgrounds (non-WEIRD and WEIRD countries), various generations (children and parents),

Table 1: The order of activities and workshop agenda. All activities were completed in English over Zoom.

Time	Activity
Day 1	
25 min	Pre-survey & Introduction
45 min	Tutorial 1: Build a Carbon Footprint Question & Answer Agent
5 min	Break
20 min	Envisioning Future Agents Ideation Session
50 min	Tutorial 2: Build a Single Turn Carbon Footprint Calculator Agent
20 min	Tutorial 3 Overview: Multi-Turn Carbon Footprint Calculator Agent
15 min	Mid-survey & Close
Day 2	
30 min	Session 1: Technology, Sustainability Societal Impact & Mindset Changes
30 min	Discussion & Final Project Development with Teams
10 min	Break
30 min	Session 2: How Should We Develop the Future of Technology & Agents?
30 min	Discussion & Final Project Development with Teams
30 min	Final Presentations
20 min	Post-survey & Close

and various prior experiences (e.g., programming, AI and conversational agent experience). Interest forms for the study were sent to educational email lists worldwide (e.g., the AI4K12 email list [1]). In the workshops, 49 participants ($n_{total}=49$) completed research consent forms, and completed at least 1 of the 3 surveys that were given before ($n_{pre}=46$), during ($n_{mid}=40$), and after ($n_{post}=35$) the study. According to the demographics survey, children comprised 58.7% of participants (age average=13.96, SD=1.829), parents comprised 41.3%, WEIRD comprised 50% (age average=26.45, SD=19.24), and non-WEIRD comprised 50% (age average=25.48, SD=15.18). Participants came from Indonesia, Iran, Japan, India, U.S, Singapore, Canada, and New Zealand. Twenty participants identified as female, 25 identified as male, and 1 identified as non-binary. Fourteen participants had no prior programming experience, 6 only had visual (or blocks-based) programming experience, and 26 had text-based

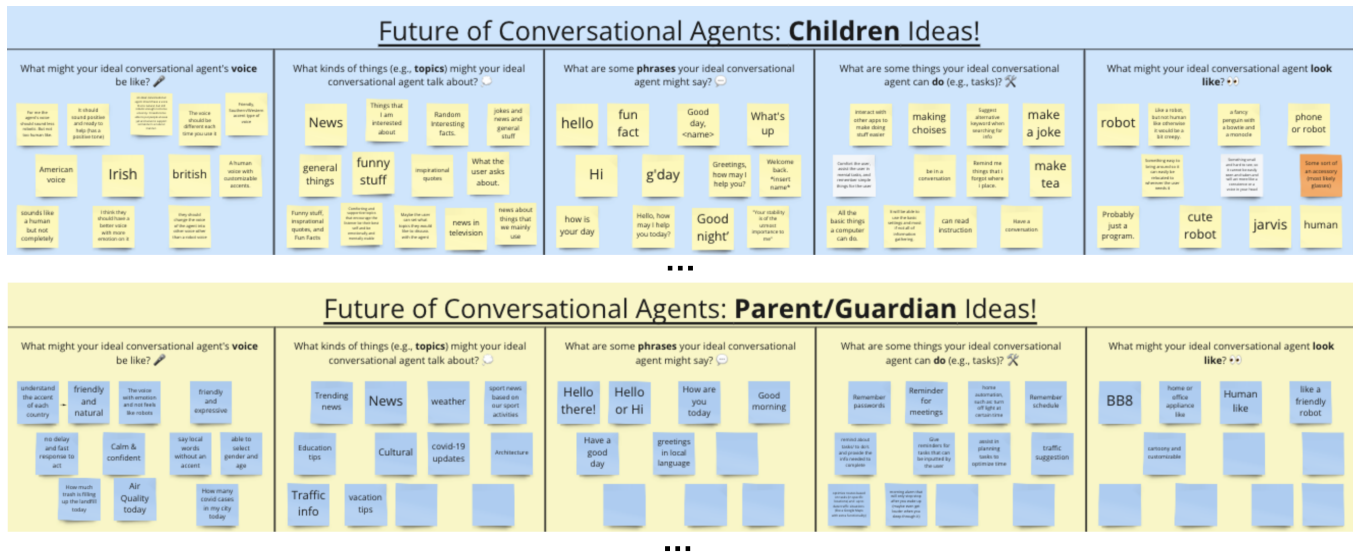


Figure 3: Approximately half of children’s and parent’s responses on the virtual whiteboard during the ideation session about their ideal conversational agents. The questions for each section are as follows: “What might your ideal conversational agent’s voice be like?”, “What kinds of things might your ideal conversational agent talk about?”, “What are some phrases your ideal conversational agent might say?”, “What are some things your ideal conversational agent can do?”, and “What might your ideal conversational agent look like?”. There was additionally space for “Other ideas”, not shown here.

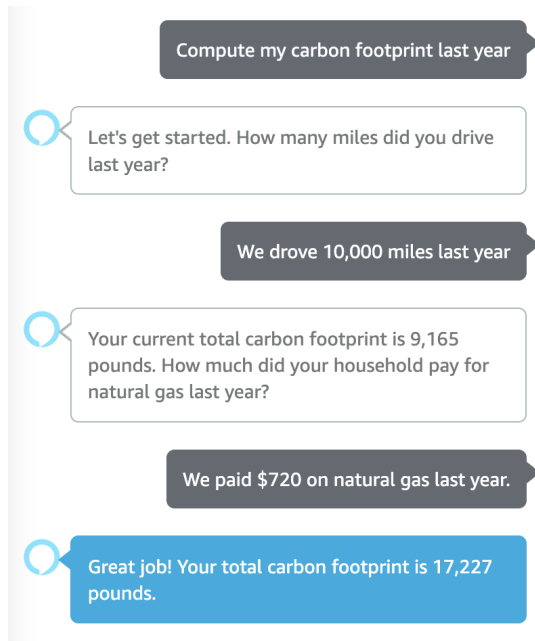


Figure 4: An example conversation with the agent developed in the workshop tutorials.

programming experience. Thirty-eight participants reported typically using conversational agents in their first language; 8 reported typically using them in another language. Demographics numbers broken down by survey can be found in Table 2. In this paper, unless

Table 2: Number of participants and subsets of participants who filled out the each of the surveys.

	Pre	Mid	Post
Total	46	40	35
Children	27	24	21
Parents	19	16	14
Non-WEIRD	23	18	17
WEIRD	23	22	18

otherwise specified, the term 'participants' refers to all participants in the study (i.e., children and parents).

4.2 Data collection

As shown in Table 1, there were three surveys. These surveys were administered through an anonymous online collection form. On each of the surveys, we asked participants about their trust and partner models of conversational agents, and self-identification as programmers through Likert scale and short answer questions. For example, we asked students to respond to the prompt “Conversational agents (e.g., Siri, Alexa, Google Home) say things that are...” using a 5-point scale from “Always Right” to “Always Wrong”. We also asked students to write sentence responses to questions like, “Please explain why you think conversational agents say things that are right/wrong”. We derived the survey questions from McKnight and Chervany’s work on trust [34] and Doyle et al.’s work on partner models [17]. On the mid- and post-survey, we additionally asked participants if their opinions had changed. On the pre-survey, we additionally asked them about their demographics. Children

and parents completed the surveys separately. We collected participants' "ideal agent" ideas from the virtual whiteboards, which we separated into child and parent sections. Figures 1 and 3 show portions of the virtual whiteboards.

4.3 Data analysis

To analyze the Likert scale data, we used Mann-Whitney U tests, Wilcoxon signed-rank tests, and independent and paired t-tests, depending on the sample and distribution of the data. We identify statistical significance in Figures using star symbols (i.e., "*" for $p \leq .05$, "***" for $p \leq .01$ and "****" for $p \leq .001$). The analysis was within-subjects for comparing across surveys (e.g., pre- vs. post-survey child trust results) and between-subjects for comparing results within one of the surveys (e.g., child vs. parent pre-survey trust results).

To analyze the responses to the short-answer questions and the prompts during the design session, we used a coding reliability approach to thematic analysis [9]. Three researchers tagged each section of the data and reconvened to agree on common sets of themes, including guidelines and definitions for each theme. The theme definitions are shown in the appendix [53]. The researchers completed three rounds of coding such that the Krippendorff's Alpha between all researchers was $\alpha \geq .800$ [3]. We aggregated the tagged data by union between researchers, and organized them with respect to the child and parent categories.

5 LIMITATIONS AND FUTURE WORK

In this paper, we focus on voice-based agents due to humans' long history of voice-based interactions and how this mode of interaction may cause agents to seem especially personified (and likely especially trustworthy [47, 56]). Nonetheless, future research may investigate people's perceptions of text-based agents, as they are also common and have great potential for societal impact. Since we specifically used the voice-based agent of Amazon Alexa (as this is the only current type of agent the ConvoBlocks platform supports [35]), its default persona could have biased people's perceptions of agents. Future research could investigate how developing agents with different voices and on different platforms affects perceptions.

Another limitation includes how we leave the definition of "accurate" partner models and "appropriate" levels of trust to future research, and only investigate how participants' perceptions of these change in our study. An additional limitation includes the context of the study. Since the participants engaged in the workshops in their home environment over Zoom, other factors in their environment could have affected the results. Future research could verify the results of this study in other environments.

This study aimed to diversify participants by including those from non-WEIRD countries; however, future research could further include those from non-WEIRD countries who do not speak English and are not as familiar with WEIRD-designed tools. It could also include perspectives from additional non-WEIRD countries, neurodiverse perspectives, and perspectives from people of different gender identities. With more diverse perspectives, researchers could adapt and extend current conversational agent design guides to better address the world's population.

6 RESULTS AND DISCUSSION

This section describes the results most relevant to agent design recommendations. We describe other results (e.g., most relevant to pedagogy recommendations) in [54, 57].

6.1 Partner Model

Sixty-two percent of overall participants indicated they felt their partner models changed through the programming activity in their long-answer responses, as shown in Table 3. Alongside the results that, on average, participants successfully learned to create 2-3 ($\bar{x}=2.26$, $\bar{x}_{child}=2.30$, $\bar{x}_{parent}=2.18$) agents during the workshops, this indicates that by developing a greater understanding of how agents work, people's feelings towards agents also change. For instance, after the workshops, participants thought of agents as more of friends than co-workers (pre/post: $\bar{x}=3.58, 3.24$; $t(32)=2.15$; $p=.039$). This may indicate developing agents with the ability to educate users about themselves may be valuable if one wants the agent to develop friendly relationships with users. Such education is also valuable in terms of increasing AI transparency [55, 56].

In terms of children and parents, before ($\bar{x}=2.74, 2.11$; $U(44)=167$; $p=.018$) and after ($\bar{x}=2.79, 2.13$; $U(38)=112$; $p=.0093$) the programming activity, children thought Alexa was more human-like than parents did. They also thought Alexa was warmer than their parents did before ($\bar{x}=2.70, 3.37$; $U(44)=170.5$; $p=.021$), during ($\bar{x}=2.96, 3.56$; $U(38)=129.5$; $p=.034$) and after ($\bar{x}=2.62, 3.50$; $U(33)=81.5$; $p=.011$) the workshops. After the programming activity, they thought Alexa was more dependable than their parents did ($\bar{x}=3.82, 3.14$; $U(16)=21$; $p=.039$). This may indicate children generally have a more positive view on agents, and may develop relationships [47] with agents more readily than parents would. This could be concerning, considering children's vulnerability, and the potential for agents to provide incorrect information [6]. Designers may want to consider designing agent personas to foster appropriate relationship building (e.g., whether that means shifting perceptions from co-worker to friend or vice-versa) and therefore trust, as described in [47].

In terms of gender, male participants felt Alexa was more like a friend (pre/post: $\bar{x}=3.74, 3.26$; $W(18)=8$; $p=.039$) after the workshops than they did before. There were no significant differences in female participants' opinions overall in terms of the partner model through the workshops. This may indicate that males' perceptions of agent friendliness may more readily change through interaction than females' perspectives; however, participants' perceptions could also have been affected by the default gender (female) of the Alexa agent's voice. Future research may investigate how agent relationship formation changes depending on agent and participant gender.

With respect to prior experience, before the workshops, participants who had text-based programming experience thought Alexa was less competent than those who had no programming experience did ($\bar{x}=2.73, 2.07$; $W(16)=0$; $p=.038$). This, in addition to how the majority of participants indicated they felt their partner models changed after learning to program agents (see Table 3), indicates programming knowledge contributes to perception changes about agents. Thus, when designing agents, it may be important to

consider the target users' programming knowledge (e.g., designers may want to ensure agents intended for programmers are especially competent).

With respect to language, at all times throughout the workshop, participants who used conversational agents in their first language thought Alexa was more human-like than those who used them in another language. Before the workshop activities, they also thought Alexa was more correct than those who used it in another language ($\bar{x}=4.03, 3.00$; $U(44)=52$; $p=5.50 \times 10^{-4}$). This may be due to agents misunderstanding accents, causing Alexa to seem more artificial and less correct. Design implications of this may include ensuring agents understand end-users first language(s) where possible, training agents to recognize diverse accents where possible, or designing agents to recognize user frustration (e.g., when a user repeats something louder) and engage using especially attentive personas in these cases.

Table 3: Percent of long-answer responses indicating a shift in participants' perceptions of agent partner models through the programming activity.

Subset	Changed	Did not change	Ambiguous
Overall participants	62%	35%	3%
Children	67%	33%	0%
Parents	54%	38%	8%

6.2 Trust

In the long-answer responses, we found overall participants' reasoning for their levels of trust towards agents leaned towards the aspect of competence on both the pre- (Table 4) and mid-survey (Table 5). The next two aspects participants most often mentioned were predictability and then integrity. We found no responses indicating participants considered the benevolence aspect of trust with respect to conversational agents. Thus, when considering how to design agents with accurate levels of trustworthiness, designers may want to focus on the aspects of agents' competence, then predictability and then integrity. Designers may also want to specifically focus on creating agents to be transparent in terms of the source of the agent's information, including human data, the internet and other sources, as these were the themes participants most often referenced for changes in their trust. This is shown in Figure 6.

Participants overall (and child and parent subsets) prior to, during and after the workshops, generally trusted Google, Alexa and newspapers significantly more than both parents and friends to report correct information. Figure 5 shows this trend. In other words, participants tended to trust technology more than people for correct information. This may indicate an overtrust of Alexa, depending on the actual correctness of the device (although we leave this as a question for future research). Since different agents show varying levels of correctness [31], different agents should be trusted differently. To foster such levels of trust, which match agents' actual trustworthiness, as mentioned previously, designers may want to focus on the competence aspect of their agents, as well as ensure transparency in terms of agents' sources of information.

Trust of people and technology to give correct information

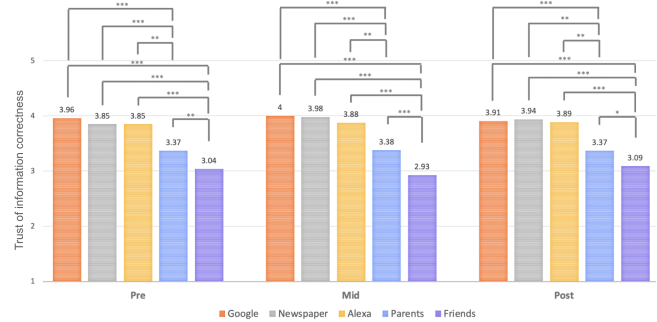


Figure 5: The mean responses for participants overall when rating Google, the newspaper, Alexa, parents and friends on a 5-point scale in terms of trust of information correctness. Reproduced from the thesis, [54].

Table 4: Percent of long-answer responses indicating different aspects of McKnight and Chervany's trust model when participants discussed their opinions on trust of conversational agents on the pre-survey.

Subset	Competence	Integrity	Predictability	Benevolence
Overall	39%	25%	36%	0%
Children	34%	30%	36%	0%
Parents	48%	17%	35%	0%

Table 5: Percent of long-answer responses indicating different aspects of McKnight and Chervany's trust model when participants discussed their opinions on trust of conversational agents on the mid-survey.

Subset	Competence	Integrity	Predictability	Benevolence
Overall	43%	23%	34%	0%
Children	37%	26%	37%	0%
Parents	52%	17%	30%	0%

OVERALL REASONING FOR TRUST LEVELS




Figure 6: Overall participants' responses to the question asking about their reasoning for their opinions on trust of agents in terms of counted tag frequency. (See the appendix [53] for descriptions.)

As shown in Figure 7, after the programming activity, children trusted Alexa to be more correct than parents did ($\bar{x}=4.04, 3.63$; $U(38)=127.5$; $p=.023$). Children also trusted agents to report correct information more after the societal impact activity than before

(mid/post: $\bar{x}=2.60, 2.35$; $t(19)=2.52$; $p=.021$). This indicates children may more readily find conversational agents more trustworthy through increased interaction. Thus, it may be especially important to consider the factors affecting children’s trust in human-agent interaction. As shown in Table 4, agent predictability was the most influential trust factor before the programming activity, and afterwards, predictability was tied with competence. Future research may investigate how to affect children’s perceptions of agent competence and predictability through agent design (e.g., through using particular agent diction, like ‘maybe’ or ‘perhaps’, when providing answers).




Figure 7: Children and parents’ responses when asked to rate their trust of Alexa’s correctness on a 5-point Likert scale after the programming activity. Reproduced from the thesis, [54].

6.3 Ideal Agents

In terms of thematic analysis of the ideation session (see Figures 1 and 3), participants described their ideal conversational agents with more task-oriented (75%) than non-task oriented (or socially-oriented; 25%) language, and used slightly more human-like (55%) than artificial (45%) descriptions, as shown in Figure 8. (See the appendix [53] for example task vs. non-task oriented, and human-like vs. artificial descriptions.) The subsets of children and parents also showed the same tendency towards human-like and task-oriented agents, albeit with slightly different proportions. Children commented relatively more on how conversational agents should be artificial (52%) than parents did (30%); Parents had relatively more task-orientation (82%) than children (71%).

Participants’ perspectives may have been influenced by how current agents tend to be task-oriented, rather than truly conversational or social [12]. That said, participants still included social (non-task) oriented agent attributes in their responses (e.g., having agents ask about how users feel)—despite this being rare in current commercial agents [12]. Thus, designers may want to include some social abilities in their task-based agents.

In terms of human-likeness, participants—especially children—mentioned how it is important for agents to be artificial (e.g., “Like a robot, but not human like otherwise it would be a bit creepy”), emphasizing the need for designers to consider the uncanny valley [36], or to balance the human-likeness of agents with artificiality. Other concerns emerged about information security (e.g., “[Agents] should only be able to access information on the internet (not take actions like creating an account)”), emergency preparedness (e.g., “[It should be able to] get help in emergencies”), ensuring agents

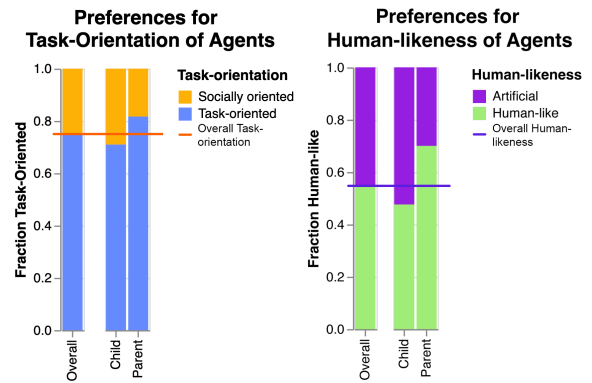


Figure 8: The number of phrases indicating a preference for either task-oriented or non-task oriented (i.e., socially-oriented) agents (left) and a preference for either human-like or artificial (e.g., robotic) agents (right) normalized and grouped by various subsets of the participants.

can provide emotional support (e.g., “[It should] encourage the listener be their best self and be emotionally and mentally stable”), and ensuring agents do not instill fear (e.g., “[It shouldn’t be] too intimidating and absolutely freak me out every time I see it”, “It needs to be able to put people at ease”), among other concerns. Interestingly, children responded with relatively more concerns about agents than parents did, as shown in Figure 9. Thus, designers should consider addressing user concerns when designing agents, including (and especially) agents intended for children.

Other themes that emerged from participants describing their ideal agents are shown in Figure 9, from most to least frequent. Three of the themes indicate participants want future conversational agents to be user-oriented (*Convenient, Personalized, and Proactive*); three indicate a desire for enjoyable interactions (*Approachable/friendly, Familiar or pop-culture related, and Fun*); and two indicate a desire for emotional intelligence (*Addresses concerns and Culturally intelligent*). The final theme, *Basic features*, indicates participants want future agents to include the typical features current agents have, like the ability to play music or get the weather. Detailed descriptions of each theme are in the appendix [53].

As shown in Figure 9, parents tended to focus more on *personalized features* and *pop-culture or familiar features* than children, whereas children tended to focus more on *fun features, approachable/friendly features, and addressing concerns* (as previously mentioned) than parents. Designers may want to take this into consideration when designing agents for children or parents.

7 SUMMARY

Based on the results of how children and parents’ trust and partner models changed through learning about conversational agents, we recommend taking the following results into consideration when designing conversational agents:

- With respect to partner models:
 - Education about agents increased users’ feelings of friendship towards agents

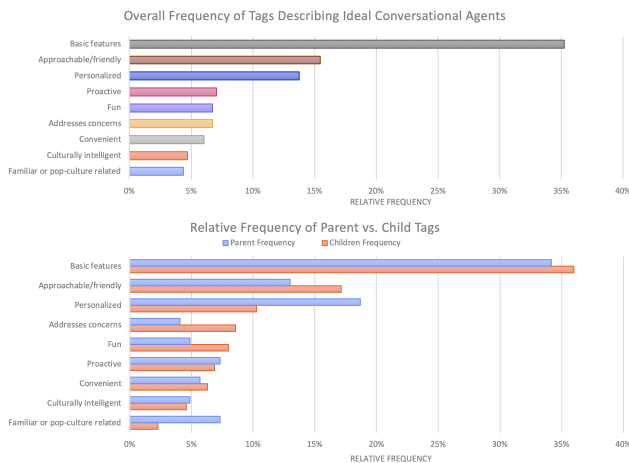


Figure 9: Bar charts showing the relative frequency of phrases tagged with particular themes for overall participants (top) and parents vs. children (bottom). Reproduced from the thesis, [54].

- Children felt agents were more human-like, warm, and dependable than parents did at various times during the workshops
- Male users' feelings of friendship towards agents seemed to change more readily than females' feelings
- Users with more programming experience felt agents were less competent
- Those using agents in their first language felt agents were more human-like and correct than those using agents in a language other than their first
- With respect to trust:
 - Users generally trusted technology more than people for correct information (which might indicate an overtrust in this technology)
 - Users reasoned about their trust towards agents most often with respect to agents' competence
 - Users frequently mentioned how learning about agents' information sources changed their trust of agents
 - Children's trust of agents increased through education
- With respect to what they want to see in their "ideal agents":
 - Users described their ideal agents with more task- than social-orientation
 - Parents had more task-orientated descriptions than children
 - Children commented relatively more on how conversational agents should be artificial than parents did
 - Users had concerns about the uncanny valley, information security, emergency preparedness, emotional support, and intimidation, among other concerns, with respect to agent design
 - Users wanted agents to have the basic current features that typical commercial agents have today, as well as be user-oriented, enjoyable, and emotionally intelligent

- Parents tended to focus more on *personalized features* and *pop-culture or familiar features* than children, whereas children tended to focus more on *fun features*, *approachable/friendly features*, and *addressing concerns* than parents

8 CONCLUSIONS

This study investigated how people of various backgrounds (WEIRD and non-WEIRD, as well as different generations) perceive agents in terms of partner models and trust, and how they envision their ideal agents. The results (summarized in Section 7) showed how partner models and trust can differ between children and parents, and change through learning about and how to program agents. These results led to discussion about how agent designers can be aware of children and parents' perceptions while designing. For instance, developing agents with the ability to educate users about agents' inner-workings could result in friendlier human-agent relations, as well as increase agent transparency. However, since relationship-building can increase trust of given information [47, 56], agents are not always correct [6], and people tended to trust agents' correctness more than humans', designers may want to provide users with indicators of agents' actual accuracy. This may include designing agents to be transparent in terms of the source of the agent's information, since participants most often referenced this when describing changes in their trust.

Other discussion included how designers may want to align their agent designs with children's and parents' ideas for "ideal agents". For instance, participants wanted agents to be user-oriented, enjoyable, and emotionally intelligent, as well as have the basic features already found in current commercial agents. When designing for children, designers may want to emphasize *fun features*, *approachable/friendly features*, and *addressing concerns*, as these were mentioned more frequently by children than by parents. We describe these themes in detail in the appendix [53].

Additional pedagogical-related results from this study are discussed in [57] and [54], including how educators may want to further emphasize the societal impact of conversational agents and encourage critical thinking when students engage with these agents. There are many opportunities to continue this research, as described in Section 5. We hope that through researchers' continued development of studies with diverse participants, and by developers' utilization of recommendations, we will increasingly design conversational agents "for all".

9 SELECTION AND PARTICIPATION OF CHILDREN

In this study, we recruited fifty-five children who took part in our educational workshops. Participants in the workshops were not required to participate in the research. For the children who did participate in the study (n=27; ages 11-17), each child completed a child assent form written in language appropriate for their age level. A parent or guardian of each child completed a parental consent form for the child, in addition to an adult consent form for themselves, if they participated in the study. The forms explained the study procedure, data collection methods, processes to keep their data confidential, and the research goals. We followed institutional

recommendations before, during and after the study, including anonymization and data security procedures.

Recruitment involved providing information about the study on a website in English, and sending this information and links to the website to educational email lists world-wide (e.g., the AI4K12 email list [1]). Due to the complexity of the coding activities and experience with students of various ages during prior pilot studies, we only included participants within the age range of 11 to 17. Participants were not paid to take part in the study, but could keep the agents they developed online on the ConvoBlocks website and use them later. Participants did not need prior programming experience or an Alexa-enabled device to participate. The only requirements were a computer with Zoom installed and access to the internet. The research study was approved by MIT's Institutional Review Board prior to the study.

ACKNOWLEDGMENTS

We would like to thank the participants in this study, the Brave Conversations organizing committee, the Amazon Future Engineer members, and the MIT App Inventor team, who made this study possible. A special thanks goes to Hal Abelson: Thank you for your dedication to empowering people of all ages to understand, develop, and take computational action with technology.

REFERENCES

- [1] AI4K12. 2020. The Artificial Intelligence for K-12 initiative. <https://ai4k12.org/>. Accessed: 2023-01-19.
- [2] Amazon. 2021. Skill Blueprints. <https://blueprints.amazon.com/>. Accessed: 2021-09-14.
- [3] Ron Artstein and Massimo Poesio. 2008. Inter-Coder Agreement for Computational Linguistics. *Computational Linguistics* 34, 4 (12 2008), 555–596. <https://doi.org/10.1162/coli.07-034-R2> arXiv:<https://direct.mit.edu/coli/article-pdf/34/4/555/1808947/coli.07-034-r2.pdf>
- [4] Benett Axtell and Cosmin Munteanu. 2021. Tea, Earl Grey, Hot: Designing Speech Interactions from the Imagined Ideal of Star Trek. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 249, 14 pages.
- [5] John Barryman. 2014. 18 Siri FAILs That Will Ruin Your Day. <https://www.ranker.com/list/siri-fails-iphone-sucks/john-barryman>. Accessed: 2021-09-01.
- [6] Timothy W Bickmore, Ha Trinh, Stefan Olafsson, Teresa K O'Leary, Reza Asadi, Nathaniel M Rickles, and Ricardo Cruz. 2018. Patient and Consumer Safety Risks When Using Conversational Assistants for Medical Information: An Observational Study of Siri, Alexa, and Google Assistant. *J Med Internet Res* 20, 9 (04 Sep 2018), e11510. <https://doi.org/10.2196/11510>
- [7] Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in neural information processing systems* 29 (2016), 4349–4357.
- [8] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* 1 (2021), 214 pages.
- [9] Virginia Braun, Victoria Clarke, Nikki Hayfield, and Gareth Terry. 2019. Thematic Analysis. In *Handbook of research methods in health social sciences*, Pranee Liamputtong (Ed.). Springer, Singapore, Singapore.
- [10] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [11] Jin-Hee Cho, Kevin Chan, and Sibel Adali. 2015. A Survey on Trust Modeling. *ACM Comput. Surv.* 48, 2, Article 28 (oct 2015), 40 pages. <https://doi.org/10.1145/2815595>
- [12] Leigh Clark, Nadia Pantidi, Orla Cooney, Philip Doyle, Diego Garaialde, Justin Edwards, Brendan Spillane, Emer Gilmartin, Christine Murad, Cosmin Munteanu, Vincent Wade, and Benjamin R. Cowan. 2019. *What Makes a Good Conversation? Challenges in Designing Truly Conversational Agents*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300705>
- [13] Jessie Cohen, Nicole Kerr, and Trina Dong. 2022. The Peabody Awards Announce Winners for Digital and Interactive Storytelling. <https://peabodyawards.com/stories/the-peabody-awards-announce-winners-for-digital-and-interactive-storytelling/>. Accessed: 2022-06-20.
- [14] Benjamin R Cowan, Holly P Branigan, Habiba Begum, Lucy McKenna, and Eva Szekely. 2017. They Know as Much as We Do: Knowledge Estimation and Partner Modelling of Artificial Partners. In *The Annual Meeting of the Cognitive Science Society (COGSCI)*. COGSCI, London, UK, 6.
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL. NAACL*, Minneapolis, MN, 4171–4186.
- [16] Daniella DiPaola. 2021. *How does my robot know who I am?: Understanding the Impact of Education on Child-Robot Relationships*. Master's thesis. Massachusetts Institute of Technology.
- [17] Philip R Doyle, Leigh Clark, and Benjamin R. Cowan. 2021. What Do We See in Them? Identifying Dimensions of Partner Models for Speech Interfaces Using a Psychological Approach. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 244, 14 pages. <https://doi.org/10.1145/3411764.3445206>
- [18] Ben Frederickson. 2018. Where Do The World's Software Developers Live? <https://www.benfrederickson.com/github-developer-locations/>. Accessed: 2022-09-12.
- [19] Radhika Garg, Hua Cui, Spencer Seligson, Bo Zhang, Martin Porcheron, Leigh Clark, Benjamin R. Cowan, and Erin Beneteau. 2022. The Last Decade of HCI Research on Children and Voice-Based Conversational Agents. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 149, 19 pages. <https://doi.org/10.1145/3491102.3502016>
- [20] Susanne Gaube, Harini Suresh, Martina Raue, Alexander Merritt, Seth J. Berkowitz, Eva Lerner, Joseph F. Coughlin, John V. Guttag, Errol Colak, and Marzyeh Ghassemi. 2021. Do as AI say: susceptibility in deployment of clinical decision-aids. *npj Digital Medicine* 4, 1 (19 Feb 2021), 31. <https://doi.org/10.1038/s41746-021-00385-9>
- [21] Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, et al. 2020. Conformer: Convolution-augmented Transformer for Speech Recognition. In *21st Annual Conference of the International Speech Communication Association (INTERSPEECH 2020)*. INTERSPEECH, Shanghai International, 5.
- [22] Joseph Henrich, Steven J. Heine, and Ara Norenzayan. 2010. Most people are not WEIRD. *Nature* 466, 7302 (01 Jul 2010), 29–29. <https://doi.org/10.1038/466029a>
- [23] Robert Holwerda and Felienne Hermans. 2018. A Usability Analysis of Blocks-based Programming Editors using Cognitive Dimensions. In *2018 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, Portugal, 217–225. <https://doi.org/10.1109/VLHCC.2018.8506483>
- [24] Jeremy Howard and Sebastian Ruder. 2018. Universal Language Model Fine-tuning for Text Classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Melbourne, Australia, 328–339.
- [25] Matthew Huggins, Anastasia K Ostrowski, Andrew Rapo, Eric Woudenberg, Cynthia Breazeal, and Hae Won Park. 2021. The Interaction Flow Editor: A New Human-Robot Interaction Rapid Prototyping Interface. *arXiv preprint arXiv:2108.13838* 1, 1 (2021), 8.
- [26] IDC Corporate USA. 2022. Changing the way the world thinks about the impact of technology on business and society. <https://www.idc.com/>. Accessed: 2022-09-12.
- [27] Allison Koenecke, Andrew Nam, Emily Lake, Joe Nudell, Minnie Quartey, Zion Mengesha, Connor Toups, John R. Rickford, Dan Jurafsky, and Sharad Goel. 2020. Racial disparities in automated speech recognition. *Proceedings of the National Academy of Sciences* 117, 14 (2020), 7684–7689. <https://doi.org/10.1073/pnas.1915768117> arXiv:<https://www.pnas.org/content/117/14/7684.full.pdf>
- [28] Deanna Kuhn. 1999. A Developmental Model of Critical Thinking. *Educational Researcher* 28, 2 (1999), 16–46. <https://doi.org/10.3102/0013189X028002016>
- [29] Kushal Lakhotia, Evgeny Kharitonov, Wei-Ning Hsu, Yossi Adi, Adam Polyak, Benjamin Bolte, Tu-Anh Nguyen, Jade Copet, Alexei Baevski, Adelrahman Mohamed, and Emmanuel Dupoux. 2021. On Generative Spoken Language Modeling from Raw Audio. *Transactions of the Association for Computational Linguistics* 1, 1 (Feb. 2021), 20 pages. <https://hal.inria.fr/hal-03329219>
- [30] Sebastian Linxen, Christian Sturm, Florian Brühlmann, Vincent Cassau, Klaus Opwis, and Katharina Reinecke. 2021. How WEIRD is CHI?. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 143, 14 pages. <https://doi.org/10.1145/3411764.3445488>
- [31] Gustavo López, Luis Quesada, and Luis A. Guerrero. 2018. Alexa vs. Siri vs. Cortana vs. Google Assistant: A Comparison of Speech-Based Natural User Interfaces. In *Advances in Human Factors and Systems Interaction*, Isabel L. Nunes (Ed.). Springer International Publishing, Cham, 241–250.

- [32] 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* (Boise, ID, USA) (IDC '19). Association for Computing Machinery, New York, NY, USA, 301–313. <https://doi.org/10.1145/3311927.3323150>
- [33] Siu Man Lui and Wendy Hui. 2010. Effects of Smiling and Gender on Trust Toward a Recommendation Agent. In *2010 International Conference on Cyberworlds*. IEEE, Singapore, Singapore, 398–405. <https://doi.org/10.1109/CW.2010.26>
- [34] D. Harrison McKnight and Norman L. Chervany. 2001. What Trust Means in E-Commerce Customer Relationships: An Interdisciplinary Conceptual Typology. *International Journal of Electronic Commerce* 6, 2 (2001), 35–59. <https://doi.org/10.1080/10864415.2001.11044235> arXiv:<https://doi.org/10.1080/10864415.2001.11044235>
- [35] MIT App Inventor. 2022. ConvoBlocks. <https://alexa.appinventor.mit.edu/>. Accessed: 2022-09-14.
- [36] Masahiro Mori, Karl F. MacDorman, and Norri Kageki. 2012. The Uncanny Valley [From the Field]. *IEEE Robotics Automation Magazine* 19, 2 (2012), 98–100. <https://doi.org/10.1109/MRA.2012.2192811>
- [37] Christine Murad, Cosmin Munteanu, Leigh Clark, and Benjamin R. Cowan. 2018. Design Guidelines for Hands-Free Speech Interaction. In *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct* (Barcelona, Spain) (*MobileHCI '18*). Association for Computing Machinery, New York, NY, USA, 269–276. <https://doi.org/10.1145/3236112.3236149>
- [38] Christine Murad, Cosmin Munteanu, Benjamin R. Cowan, and Leigh Clark. 2019. Revolution or Evolution? Speech Interaction and HCI Design Guidelines. *IEEE Pervasive Computing* 18, 2 (2019), 33–45. <https://doi.org/10.1109/MPRV.2019.2906991>
- [39] Christine Murad, Cosmin Munteanu, Benjamin R. Cowan, and Leigh Clark. 2021. Finding a New Voice: Transitioning Designers from GUI to VUI Design. In *CHI 2021 - 3rd Conference on Conversational User Interfaces* (Bilbao (online), Spain) (*CHI '21*). Association for Computing Machinery, New York, NY, USA, Article 22, 12 pages. <https://doi.org/10.1145/3469595.3469617>
- [40] Clifford Ivar Nass and Scott Brave. 2005. *Wired for speech: How voice activates and advances the human-computer relationship*. MIT press Cambridge, Cambridge, MA.
- [41] Jakob Nielsen. 1994. Enhancing the Explanatory Power of Usability Heuristics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, Massachusetts, USA) (*CHI '94*). Association for Computing Machinery, New York, NY, USA, 152–158. <https://doi.org/10.1145/191666.191729>
- [42] Don Norman. 2013. *The design of everyday things: Revised and expanded edition*. Basic books, New York, NY.
- [43] Seymour A Papert. 2020. *Mindstorms: Children, computers, and powerful ideas*. Basic books, New York.
- [44] Mitchel Resnick, John Maloney, Andrés Monroy-Hernández, Natalie Rusk, Evelyn Eastmond, Karen Brennan, Amon Millner, Eric Rosenbaum, Jay Silver, Brian Silverman, et al. 2009. Scratch: programming for all. *Commun. ACM* 52, 11 (2009), 60–67.
- [45] Elayne Ruane, Abeba Birhane, and Anthony Ventresque. 2019. Conversational AI: Social and Ethical Considerations.. In *AICS*. AICS, Ireland, 104–115.
- [46] Haeseung Seo, Aiping Xiong, and Dongwon Lee. 2019. Trust It or Not: Effects of Machine-Learning Warnings in Helping Individuals Mitigate Misinformation. In *Proceedings of the 10th ACM Conference on Web Science* (Boston, Massachusetts, USA) (*WebSci '19*). Association for Computing Machinery, New York, NY, USA, 265–274. <https://doi.org/10.1145/3292522.3326012>
- [47] William Seymour and Max Van Kleek. 2021. Exploring Interactions Between Trust, Anthropomorphism, and Relationship Development in Voice Assistants. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 371 (oct 2021), 16 pages. <https://doi.org/10.1145/3479515>
- [48] Christian Sturm, Alice Oh, Sebastian Linxen, Jose Abdelnour Nocera, Susan Dray, and Katharina Reinecke. 2015. How WEIRD is HCI? Extending HCI Principles to Other Countries and Cultures. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems* (Seoul, Republic of Korea) (*CHI EA '15*). Association for Computing Machinery, New York, NY, USA, 2425–2428. <https://doi.org/10.1145/2702613.2702656>
- [49] Reddit Users. 2015. Alexa misunderstandings... https://www.reddit.com/r/amazonecho/comments/41159u/alexa_misunderstandings/. Accessed: 2021-09-01.
- [50] Charis V, Chaudron S, Di Gioia R, Vuorikari R, Escobar Planas M, Sanchez Martin JI, and Gomez Gutierrez E. 2022. *Artificial Intelligence and the Rights of the Child: Towards an Integrated Agenda for Research and Policy*. Scientific analysis or review KJ-NA-31048-EN-N (online). Joint Research Centre, Luxembourg (Luxembourg). [https://doi.org/10.2760/012329\(online\)](https://doi.org/10.2760/012329(online))
- [51] Jessica Van Brummelen. 2019. Conversational Artificial Intelligence Development Tools for K-12 Education. http://appinventor.mit.edu/papers/JessVBPublications/AAAI_Fall_Symposium_2019_SMCConvAI_Final.pdf. In *2019 AAAI Fall Symposium*. AAAI, Washington, DC, 8.
- [52] Jessica Van Brummelen. 2019. *Tools to Create and Democratize Conversational Artificial Intelligence*. Master's thesis. Massachusetts Institute of Technology, Cambridge, MA.
- [53] Jessica Van Brummelen. 2022. Appendix: What Do WEIRD and Non-WEIRD Conversational Agent Users Want and Perceive? Towards Transparent, Trustworthy, Democratized Agents. <https://gist.github.com/jessvb/fa1d4c75910106d730d194ffd4d725d3>. Accessed: 2023-01-01.
- [54] Jessica Van Brummelen. 2022. *Empowering K-12 Students to Understand and Design Conversational Agents: Concepts, Recommendations and Development Platforms*. Ph.D. Dissertation. Massachusetts Institute of Technology, Cambridge, MA.
- [55] Jessica Van Brummelen, Tommy Heng, and Viktoriya Tabunshchik. 2021. Teaching Tech to Talk: K-12 Conversational Artificial Intelligence Literacy Curriculum and Development Tools. In *2021 AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI)*. AAAI, Virtual, 9 pages.
- [56] Jessica Van Brummelen, Viktoriya Tabunshchik, and Tommy Heng. 2021. "Alexa, Can I Program You?": Student Perceptions of Conversational Artificial Intelligence Before and After Programming Alexa. In *Interaction Design and Children* (Athens, Greece) (*IDC '21*). Association for Computing Machinery, New York, NY, USA, 305–313. <https://doi.org/10.1145/3459990.3460730>
- [57] Jessica Van Brummelen, Mingyan Claire Tian, Maura Kelleher, and Nghi Hoang Nguyen. 2023. Learning Affects Trust: Design Recommendations and Concepts for Teaching Children—and Nearly Anyone—about Conversational Agents. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. AAAI, DC, USA, 9.
- [58] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. Advances in neural information processing systems, Long Beach, California, 5998–6008.
- [59] John K. Waters. 2014. Idc Study Counts the World's Developers: 11 Million Pros. <https://cacm.acm.org/news/171546-idx-study-counts-the-worlds-developers-11-million-pros/fulltext>. Accessed: 2022-09-12.
- [60] David Wolber, Harold Abelson, and Mark Friedman. 2015. Democratizing Computing with App Inventor. *GetMobile: Mobile Comp. and Comm.* 18, 4 (jan 2015), 53–58. <https://doi.org/10.1145/2721914.2721935>
- [61] Xizhu Xiao, Porismita Borah, and Yan Su. 2021. The dangers of blind trust: Examining the interplay among social media news use, misinformation identification, and news trust on conspiracy beliefs. *Public Understanding of Science* 30, 8 (2021), 977–992. <https://doi.org/10.1177/0963662521998025> arXiv:<https://doi.org/10.1177/0963662521998025> PMID: 33663279.
- [62] Xi Yang and Marco Aurisicchio. 2021. Designing Conversational Agents: A Self-Determination Theory Approach. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 256, 16 pages. <https://doi.org/10.1145/3411764.3445445>
- [63] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Ryan Cotterell, Vicente Ordonez, and Kai-Wei Chang. 2019. Gender Bias in Contextualized Word Embeddings. *NAACL-HLT* (1) 1 (2019), 629–634.