

No, Your Other Left!

Language Children Use To Direct Robots

Deanna Kocher
Cornell University
Ithaca, USA
drk225@cornell.edu

Lea Sarmiento
Cornell University
Ithaca, USA
les238@cornell.edu

Samantha Heller
Cornell University
Ithaca, USA
sph66@cornell.edu

Yupei Yang
Cornell University
Ithaca, USA
yy884@cornell.edu

Tamar Kushnir
Cornell University
Ithaca, USA
tk397@cornell.edu

Keith Evan Green, *IEEE Senior Member*
Cornell University
Ithaca, USA
keg95@cornell.edu

Abstract—We present an analysis of how children between 4- and 9-years-old give directions to a robot. Thirty-eight children in this age range participated in a direction giving game with a virtual robot and with their caregiver. We considered two different viewpoints (aerial and in-person) and three different affordances (non-humanoid robot, caregiver with eyes closed, and caregiver with eyes open). We report on the frequency of commands that children used, the complexity of the commands, and the navigation styles children used at different ages. We found that pointing and gesturing decreased with age, while “left-right” directions and the use of distances increased with age. From this, we make several recommendations for robot design that would enable a robot to successfully follow directions from children of different ages, and help advance children’s direction giving.

Keywords—*human-human and human-robot interaction and communication, language and semantic reasoning; use of robots in applied settings such as autism therapy*

I. INTRODUCTION

In child-robot collaborations, the robot may need to respond to verbal cues and directions issued from the child. Young children readily engage with robots in a number of different ways, including verbally [1], [2]. Verbal direction giving has been explored in adult human-robot interaction, and largely focuses on interpreting natural language [3], [4]. Children, however, who are still developing their vocabularies and navigation skills, give significantly different directions [5], [6]. In a study comparing direction-giving across 6-12 year-old children [5], 6-8 year-olds struggled to give correct directions, and utilized landmarks and vague directions more often than 10-12 year-olds. Only a few of the 8 year-olds connected landmarks and directions in ways that resembled natural language from adults (e.g., *Go past the school and turn left.*) [5]. These same patterns are reflected by 6-8 year-olds in virtual environments as well as physical maps [6]. Thus, the ‘natural

language’ of a young child may require that a robot has different capabilities than a robot that navigates by interpreting the natural language spoken by adults.

In addition to differences in language patterns, there are a number of challenges with interpreting and recognizing children’s speech. There are several ways to address these challenges, including vocal tract shortening [7], training the robot with children’s voices [8], and emotion recognition and classification [9]. In addition, children may not direct a robot the same way they would direct another person, as adults often change their speech patterns when addressing robots [10], [11].

Overcoming these challenges, however, may have significant benefits for children’s spatial skill development. Verbal direction giving to a robot provides an opportunity for children to practice using spatial language. Such experiences may be beneficial, as children who have more experience with spatial language (both hearing it from others and using it themselves) improve their mental rotation skills over children who are exposed to less spatial language [12]. Children with more spatial language experience also show greater spatial problem-solving years later in their development [13]. Furthermore, direction giving requires perspective taking—an additional way for a child to practice mental rotations with left-right orientation [14]. In this way, giving directions to a robot may not only be necessary to help a robot navigate an environment, but may help the child develop their spatial skills.

However, few studies provide the detail necessary to understand how to design a robot that can follow children’s directions. Studies of children’s language processing [9] and spatial cognition [5], [6], do not provide an understanding of the language and strategies children would naturally use to direct an interactive robot. Such interactivity is not present in map navigation, and may significantly change how children give directions because they can adjust and correct their

directions as the robot moves. To address this, we focus on how children give directions in real-time interactions to understand the implications of children’s navigation on robot design.

II. SCOPE OF THIS RESEARCH

In this research, we assess how children give directions with an interactive “robot” agent to inform robot design. We focus on children 4-9 years old, to expand upon the current literature that typically starts with 6-year-olds (as we expect younger children may be more successful in interactive settings than in previous map studies) and stop short of 10-year-olds whose language begins to reflect that of adults. We consider the frequency with which children use certain commands, the complexity of the commands, and the navigation styles children utilize. We further consider how these commands and navigation styles vary with age, to understand the affordances and language processing abilities that robots may need to interact with children of different ages.

III. METHODS

A. Experimental Design

To assess how children give directions, we conducted a within-subjects experiment where children responded to open-ended prompts to direct a virtual robot or a person (e.g. a caregiver) in a treasure hunt game. As participants have been prohibited from in-lab studies due to pandemic “stay-at-home” orders, children participated in our study over video chat. Following completion of consent forms (and photo release for academic settings) by their guardians, participating children practiced giving directions in three different trials: directing a non-humanoid robot around obstacles from an aerial viewpoint (“robot” trial); directing a familiar caregiver with their eyes closed (“person-closed” trial), and directing a familiar caregiver with their eyes open (“person-open” trial). The parent trials were designed to supplement the data from the robot trial, as we expected children’s language to change in an in-person interaction, where they can turn themselves around and use familiar landmarks. The familiarity of working with a caregiver also identifies differences that may emerge in long term scenarios, when a child is more comfortable with a robot.

B. Procedure

All studies were conducted remotely using Zoom video chat. Children were asked to play three different versions of a game where they would practice giving directions to an agent (either a robot or a caregiver). Each child played the caregiver trials in the same order; order of the robot trial or caregiver trials were then counterbalanced to control for order effects. In each version of the game, the goal was to help the agent find the “treasure”: a shape (an obstacle) within the robot game in the “robot” trial, or an item the child hid in the “person-closed or open” trials.

a) Robot trial. For the robot trial, the experimenter utilized the screenshare function on Zoom to show the child the robot game (Figure 1). This allowed the experimenter to control the robot from the keyboard without the child’s knowledge. The interface was an overhead view of the robot with a number of different obstacles. All of the obstacles in the game were

stationary, and only the robot was teleoperated. The goal of teleoperation was to allow a child to try any command he/she wanted with equal responsiveness from the robot; in this way, the robot was able to follow any command from a child that the experimenter could understand. This was done with several guidelines: a) if the child gives a direction without specifying a distance, the robot moves in that direction until it hits an obstacle or is told to stop, b) all relative directions (e.g. “left”, “right”, “forward”) are interpreted relative to the robot’s orientation, and c) vague directions like “a little bit” or “more” moved the robot approximately one inch on a laptop screen. Specific directions “five spaces” were arbitrarily defined upon first use, and then scaled if the child used distance commands a second time (e.g. “20 spaces” is four times the distance the experimenter arbitrarily travelled initially for “five spaces”).

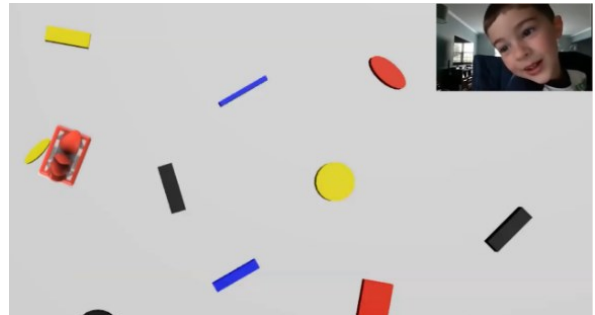


Fig. 1. Robot trial showing video chat setup. The child selects a colored shape (i.e. an obstacle) as the “treasure.”

Once the game was set up and the child confirmed that he/she could see the screen and the robot, the child was asked to pick a shape (i.e. an obstacle) on the screen to be the treasure. Once the child had chosen an obstacle, the child was told that the robot could follow directions and that the robot could hear him/her through the experimenter’s computer. The child was then asked to give the robot directions to get to the treasure. A child who struggled to give directions was prompted by the experimenter, who would ask the child what he/she thought the robot needed to do next. Children who tried to point were reminded that the robot can not see but can only hear him/her. Each child played this game two or three times (depending on the length of the trials), picking a new treasure each time with the reminder that the robot is “really smart, and can follow all kinds of directions.”

b) Person Trials. In the person trials, each child was asked to find a baseball-sized object in the room they were in and hide it somewhere in the same room (Figure 2). Participating children were asked to stay in front of their video chat device to prevent them from moving the person (a caregiver) around and keep them near the microphone so the experimenter could hear them. They were told that their caregiver was going to follow their directions, just like a robot, and they would have to help them find their hidden “treasure.” Each child played this game twice, once with the caregiver’s eyes closed (person-closed), and once with the caregiver’s eyes open (person-open). Children hid the object in a different place for each trial. When the caregiver had their eyes open, the child

was reminded that “Your [caregiver] can see things, so you can give [him/her] directions by what [he/she] can see.”



Fig. 2. Child showing the object they will hide as the “treasure” for the parent trials

C. Participants

Thirty-five children, 4-9 years old ($M = 7.33$, $SD = 1.76$; 19 Male, 16 Female) participated in the study. The total number of children who completed all tasks was 32 and 3 additional children completed at least one task. Seventeen children played the robot game first, and eighteen children played the person games first. Due to experimenter error, two children played with the person’s eyes closed twice, and three children played with the person’s eyes open first. All other children played with the person’s eyes closed first. Three additional children were allowed to participate after watching a sibling play the game but were excluded from analysis.

IV. CODING & ANALYSIS

All of the interactions were recorded and transcribed. From the transcriptions we developed three tiers of analyses: words, commands, and navigational strategies (Figure 3).




Fig. 3. The structure of the navigational style “Pointing”, showing the commands related to pointing, and example breakdown of words in the commands.

In order to be counted as a command, phrases issued by a child had to indicate a meaningful directive for a 3 degree of freedom mobile robot. From the transcriptions, we identified 52 meaningful commands from the 35 children included in our analysis. Each command was determined by the primary function of the phrase (e.g. “go left,” “turn left” and “go to the left” were all coded as “left”). Such commands ranged from vague (“go,” “turn”) to specific (“take X steps,” “turn X degrees”). Other commands considered were directional commands (*forward, straight, left*), gesture-related commands (*go there, this way, follow me*) and feedback-based commands (*yes, no, other way*). Each command was analyzed with a frequency (1), to control for the fact that children chose different treasures (which may have required more time or commands to navigate to), and that different children speak different amounts.

$$\frac{\text{instances of specific command in trial}}{\text{total commands in trial}} \quad (1)$$

Commands were then analyzed by word count, as greater word counts may be more difficult for a robot to understand than lower word count commands. In addition, we created a separate count of “distance words” that children used over the course of the interaction. This included a range of terms, from “a little bit” or “one more” to specific units like “go forward three [steps/spaces/rolls].” We also included specific angle references in this category, as a measure of how far the robot should turn.

From there, we identified five main navigation strategies that children used, and the prompts that related to this navigation strategy. The strategies we considered were a) turn-stop, b) left-right, c) pointing or guiding d) landmarks and e) absolute. The ‘turn-stop’ strategy indicates that children did not indicate directionality when telling the agent to turn, but waited until they reached the correct orientation and then had the agent stop and go forward. The left-right strategy indicates that children did provide directionality. The frequency of use for each strategy was determined using “key commands.” These commands were those which were the most frequently used for one strategy, and were clear indicators of the strategy being used (Table 1). The frequency of each command was summed to provide a ratio of how often children used the specific strategy relative to the total commands they used (as in Equation 1). Not all of children’s commands related to a clear navigation strategy (e.g. “go,” “the other one,” “yes,” “no”); these are reported on generally but are not considered in the analysis of navigational strategies.

TABLE I. NAVIGATIONAL STRATEGIES & COMMANDS

Strategy	Related Commands
Turn-Stop	Turn, Stop
Left-Right	Left, Right, Turn Degrees
Pointing	This way/That way, Go There/Right Here
Landmarks	Go to ..., (Object Reference) ^a
Absolute	Up, Down

^a Command for a general object reference with or without a navigational direction (i.e. “next to the fridge”) developed to count landmark references that may not have had an associated robot command

Further, we noted all commands that were prompted, either by the experimenter or by family members. We considered prompts for specific language (“Do you want it to go right, or left?”) as well as general prompts (“Which way should it go now?”). We removed children’s prompted commands from the final frequency analysis, to report only on the language children used of their own accord. We then considered the number of “unique commands”, which are the number of different commands the child used to navigate the robot or the parent on their own (non-prompted). Each participant interaction was coded by two reviewers independently, and disagreements were resolved through discussion. All correlations presented in this analysis are done with Pearson’s R.

V. RESULTS

A. General Results

Across all trials, children used an average of 6.16 unique commands ($SD = 3.16$). This was slightly higher for the robot trials ($M=7.57$, $SD=3.08$) and slightly lower for the person-

closed and person-open trials ($M=6.51$, $SD=2.86$; $M=4.38$, $SD=2.81$). The average number of words per command was 2.5 ($SD = 1.7$). This did not vary significantly between trials or with age, but was significantly correlated with landmark navigation ($r = 0.551$, $p < .001$). Children also used a combination of different navigation strategies in each trial. The average number of navigation styles children used (in any non-zero amount) was 2.14 across all trials ($SD = 1.02$).

Additionally, different children addressed the robot differently. Some directly spoke to the robot (i.e. “*Robot, go...*”) while others used one-word commands only, even for navigating to colors (i.e. “*BLUE!*”, “*forward!*”). Some even included encouragement or considered the robot’s preferences (i.e. “*be a good robot,*” “*go left, if you’d like!*”).

B. Commands and Strategies Across Trials

Children used a variety of commands, ranging from highly specific (“*go forward ten steps,*” “*go to the up-left diagonal!*”) to vague (“*turn,*” “*go!*”). The top ten commands from each trial are shown below, as well as the average across all trials (Figure 4).

Command	Robot	Po	Pc	Average
Forward	0.0667	0.1383	0.1478	0.1176
Right	0.1066	0.0571	0.0883	0.0840
Left	0.0916	0.0593	0.0744	0.0751
Down	0.0543	0.0172	0.0293	0.0336
Stop	0.1524	0.0375	0.0629	0.0842
Go to...	0.0363	0.0325	0.0140	0.0276
Up	0.0503	0.0424	0.0082	0.0337
Straight	0.0513	0.0392	0.0445	0.0450
Turn	0.0425	0.0421	0.0788	0.0545
No	0.0193	0.0275	0.0584	0.0351
Shape/Color Nav.	0.0901	0.1661	0.0987	0.1183
Yes	0.0183	0.0327	0.0250	0.0253
This/That Way	0.0261	0.0909	0.0592	0.0588

Fig. 4. Frequencies of the 10 most common commands used in each trial.

While these commands were the most common, there were also a number of vague and feedback commands that we identified. Vague commands were those that could have been interpreted in multiple ways, while feedback commands were in response to the agent. The top five vague and feedback commands are shown in Figure 5.

Feedback Commands				
Command	Robot	Po	Pc	Average
Wait	0.0125	0.0033	0.0030	0.0062
Keep Going	0.0060	0.0203	0.0134	0.0132
No	0.0193	0.0275	0.0584	0.0351
The other way/one	0.0056	0.0269	0.0272	0.0199
Yes	0.0183	0.0327	0.0250	0.0253
Vague Commands				
Turn	0.0425	0.0421	0.0788	0.0545
Move	0.0087	0.0237	0.0048	0.0124
Go	0.0179	0.0286	0.0203	0.0223
More	0.0187	0.0052	0.0245	0.0161
Sideways	0.0048	0.0063	0.0066	0.0059

Fig. 5. Frequencies of the 5 most common vague and feedback commands

Lastly, children used strategies differently in different trials. Children utilized landmarks the most often in the person open trial, while children used “turn-stop” and “left-right” more often in the robot and person closed trials (Figure 6). In general, children used “left-right,” “turn-stop,” and landmarks the most, while pointing and absolute directions were used less.




Fig. 6. Navigation style in each trial – robot, person-closed, person-open – and the average

C. Navigation Differences by Age

Several variables varied significantly with children’s age. First, the frequency of prompts (total prompts/total commands) decreased with age ($r = -.586$, $p < .001$). Many of these prompts came from the caregivers during the person trials, where young children would give a vague command (such as “*turn*” or “*move!*”) and parents would ask leading questions to understand what to do. Prompts given during the robot game were focused on reminding the child of the task, and asking them for further commands (“*what should it do next?*”).




Fig. 7. Child using her hands to determine the robot’s left and right

Additionally, we found age differences in how children gave directions. Distance words increased significantly with age ($r = 0.499$, $p < .002$), with variations of “*more*” and “*a little bit*” being more common than specific units of distance, like “*two feet*” or “*five steps*”. Navigation styles also changed with age, as pointing decreased ($r = -0.581$, $p < .001$) and left-right navigation increased ($r = 0.389$, $p < .021$) with age. Several younger children attempted to use the left-right strategy, despite admitting that they did not know their left and right. In contrast, when the robot responded to vague commands like “*turn,*” children who knew their left and right would occasionally revert to a turn-stop approach. Interestingly, children who used the left-right strategy also occasionally told the agents to go to their “*other right*” or “*other left.*”



Fig. 8. A child pointing to direct their parent to the hidden treasure

VI. DISCUSSION

This study provides insights on the language and navigation strategies children use to give directions to a robot. In the study, children are able to interact with a robot that can (through teleoperation) respond to commands given by children, as well as a caregiver that followed commands “like a robot.” While some differences emerged between trials, these results should be considered altogether, as the trials were meant to complement one another and account for gross navigation styles, commands, and interactions children use to navigate. Our findings suggest that robots should have different capabilities and responses to directions from children of different ages, to both allow for successful interactions and encourage children to advance their direction-giving skills.

A. “Natural Language” Processing for Children’s Directions

Unlike adults, whose “natural language” for direction giving often involves landmarks, distances, and directions, we found that children give much simpler commands. The majority of children in our study successfully navigated an agent through an environment using less than 10 unique commands, with an average of 2-3 words per command. The exception to this was when children navigated by landmarks, which may have required significantly more words per command. This low word count, was despite the fact that children were told the robot was “really smart” and that they could tell it whatever they wanted. Rather, the simplicity of commands may in part be due to the interactivity of the robot, where children used single step commands to get immediate feedback. Such language is much simpler than that of adults and this simplification is ideal for the already challenging task of child language processing.

B. The Benefit of Directing a Robot

While interpreting children’s voices poses significant challenges, our study presents several benefits of having children practice directions with a robot. First, children used the most unique commands and navigation strategies in the robot trial. This may have been because children were unsure of what the robot could and could not understand, and they tried a number of different commands to try to figure this out. Such potential for “discovery” may encourage children to try and challenge the robot, while simultaneously challenging their own direction giving abilities.

Moreover, studies with 6-8 year-olds suggest that children often give incorrect directions [5]; in our study, children were

able to assess the robot’s actions as they directed it, and could correct themselves at any point. Further, the response of the agents to vague and feedback commands allowed children as young as four years old to navigate the agent through space. Such interactive capabilities allow children to understand their own direction giving skills in a way they could not when they do not receive feedback.

C. Robot Design for Children of Different Skill Levels

From these findings, we make several recommendations for robots that can enable them to successfully follow children’s directions. We present conclusions for competency levels, rather than age groups, as children’s navigational abilities may not develop in explicit stages with age. These competency levels are “beginner,” “intermediate,” and “advanced.” At each of these levels, the robot should be able to offer some prompting to the child (explicit verbal prompts or otherwise) to motivate them to use more advanced navigation skills. In addition, the robot should be trained on feedback commands for all skill levels, as all children make occasional mistakes that may require quick feedback.

a) *Beginners*. To work successfully with the earliest navigators, a robot should be multimodal, able to interpret gestures or identify landmarks in a space. This includes (but may not be limited to) pointing, following a child, and identifying objects and navigating to them. Children in this category are less likely to identify specific directions, and may rely on more vague commands and feedback to navigate. Thus, a robot that works with these children should be well trained on commands such as “forward,” “turn,” “go,” “stop,” “yes,” and “no.” Training on distances (even vague distances like “a little” and “more”) may not be necessary, as these navigators primarily offered directions of travel; thus, the robot should move in a given direction until the child offers feedback. Prompting should be common, either to prompt feedback commands from the child or remind them to give commands.

b) *Intermediate*. In this category, minimal gesture response should be used to encourage children to use verbal commands. The robot should recognize the various vague commands indicated in the beginner level, as well as more specific directions like “left,” “right,” and “diagonal.” Strategies of navigation should account for children using advance commands (e.g. “diagonal”) but not fully specifying these commands (e.g. “the forward left diagonal”). Voice recognition and motion patterns should be calibrated for quick feedback, as these children may still be learning their left and right, or need to give quick feedback on commands they did not fully specify (e.g. “not that diagonal!,” or “right - I mean left”). The robot should be trained for vague distances, as children may use them occasionally. Prompting should be used to encourage children to give directions at this level, especially to encourage using distances and left-right orientation.

c) *Advanced*. A robot for an advanced direction-giver should respond primarily to specific verbal commands. If the robot has a good language and vision processor, it may respond to landmarks, but it should not respond to vague gestures - these navigators are capable of being far more specific. In addition, the robot should be calibrated to vague and specific distances,

such as “five feet” or “80 degrees”. Prompting for this skill level should focus on moving children from using vague commands like the turn-stop strategy and unspecified distances, to specific directions with distances.

Rather than programming a robot to be fully multi-modal, we recommend these transitions as children in our study were both eager to try the directions they thought they should use (trying left and right, even without knowing which is which) as well as slipping back to more basic navigation strategies when the robot responded (using “turn-stop” after having told the robot to “turn left forty five degrees and go forward”). Each of these modalities and stages presents something meaningful for children to learn, as gesturing has been shown to improve spatial skills on different tasks [15], as does practice with left-right orientation [14] and spatial language exposure [12].

D. Future Work

Future work should explore how a mobile robot can transition through the development process, and how a robot could identify the level that a child is at from the directions they give. This would allow a robot to work with a number of young navigators, and adapt to the modality most appropriate for them. Further, child speech recognition of simpler commands should be reported on, as young children’s voices may still be difficult to interpret, even with two- or three- word commands.

E. Limitations

One of the limitations of our study was the inconsistency of how caregivers responded to and prompted commands in the person trials. Because we could not predict the language children would use during the task, caregivers were told to follow the directions provided by the children to the best of their ability. In some cases, this meant children could navigate by landmark when their caregivers had their eyes closed, because they knew where things were in the room. In contrast, other caregivers were very meticulous, requiring children to specify right or left every time they were told just to turn. Such differences could have altered the frequencies of children’s commands. Additionally, because we could not control for the size of the room and the location where the child hid the treasure, some interactions were too short to provide good data, even with frequencies. This could have skewed our results in situations where the child only gave a few commands, as the frequency of language use would be higher when only these few commands were used. This was most prominent in the “person-open” trial, as some children were able to say “go to [where the treasure is]” in just one command. Lastly, connectivity differences may have impacted the quality of children’s interactions with the robot. We cannot explicitly report on this as we could not see how well the robot movements tracked on their screens.

VII. CONCLUSION

In this paper, we presented findings on how children give directions to agents in interactive situations, and how their direction-giving can inform robot design. Unlike non-interactive studies that suggest that young children navigate primarily by landmarks, our study finds that young children use a combination of strategies to navigate in interactive settings,

including a turn-stop strategy, left-right directions, pointing, absolute orientations, and landmarks. Further, we explored the characteristics of children’s “natural language” when giving directions, to understand the language processing requirements of a robot. Using these trends, we present suggestions for robot design and interactivity for children at different levels of navigational ability. In this way, a robot can ensure that children of different skill levels can direct a robot in a way that comes naturally to them, while the robot can simultaneously understand their navigation skills and help them advance.

ACKNOWLEDGMENTS

We would like to thank Sam Fickel for developing the robot game used in this study. We also thank all children, parents, and friends who participated in the study and helped recruit participants virtually.

REFERENCES

- [1] P. Lucking, K. Rohlfing, B. Wrede, and M. Schilling, “Preschoolers’ engagement in social interaction with an autonomous robotic system,” *2016 Jt. IEEE Int. Conf. Dev. Learn. Epigenetic Robot. ICDL-EpiRob 2016*, pp. 210–216, 2017.
- [2] D. Kocher, T. Kushnir, and K. E. Green, “Better Together: Young Children’s Tendencies to Help a Non-Humanoid Robot Collaborator,” in *Proceedings of the Interaction Design and Children Conference*, 2020, pp. 243–249.
- [3] C. Matuszek, D. Fox, and K. Koscher, “Following directions using statistical machine translation,” *5th ACM/IEEE Int. Conf. Human-Robot Interact. HRI 2010*, pp. 251–258, 2010.
- [4] Y. Wei, E. Brunskill, T. Kollar, and N. Roy, “Where to go: Interpreting natural directions using global inference,” *Proc. - IEEE Int. Conf. Robot. Autom.*, pp. 3761–3767, 2009.
- [5] M. Blades and Ise Medlicott, “Developmental differences in the ability to give route directions from a map,” *J. Environ. Psychol.*, vol. 12, no. 2, pp. 175–185, 1992.
- [6] P. Jansen-Osmann and G. Wiedenbauer, “The representation of landmarks and routes in children and adults: A study in a virtual environment,” *J. Environ. Psychol.*, vol. 24, no. 3, pp. 347–357, 2004.
- [7] I. Dologlou, “A novel feature transformation for vocal tract length normalization in automatic speech recognition,” *IEEE Trans. Speech Audio Process.*, vol. 6, no. 6, pp. 549–557, 1998.
- [8] E. Booth, J. Carns, C. Kennington, and N. Rafla, “Evaluating and Improving Child-Directed Automatic Speech Recognition,” *Proc. 12th Lang. Resour. Eval. Conf.*, no. May, pp. 6340–6345, 2020.
- [9] A. Batliner *et al.*, “‘You stupid tin box’ - Children interacting with the AIBO robot: A cross-linguistic emotional speech corpus,” *Proc. 4th Int. Conf. Lang. Resour. Eval. Lr. 2004*, pp. 171–174, 2004.
- [10] K. Fischer, K. Foth, K. Rohlfing, and B. Wrede, “Is talking to a simulated robot like talking to a child?,” *2011 IEEE Int. Conf. Dev. Learn. ICDL 2011*, 2011.
- [11] E. S. Kim, D. Leyzberg, K. M. Tsui, and B. Scassellati, “How people talk when teaching a robot,” *Proc. 4th ACM/IEEE Int. Conf. Human-Robot Interact. HRI ’09*, pp. 23–30, 2008.
- [12] M. Casasola, W. S. Wei, D. D. Suh, P. Donskoy, and A. Ransom, “Children’s Exposure to Spatial Language Promotes Their Spatial Thinking,” *J. Exp. Psychol. Gen.*, vol. 149, no. 6, pp. 1116–1136, 2020.
- [13] S. M. Pruden, S. C. Levine, and J. Huttenlocher, “Children’s spatial thinking: Does talk about the spatial world matter?,” *Dev. Sci.*, vol. 14, no. 6, pp. 1417–1430, 2011.
- [14] R. Rigal, “Right-left orientation, mental rotation, and perspective-taking: When can children imagine what people see from their own viewpoint?,” *Percept. Mot. Skills*, vol. 85, no. 3 PART 1, pp. 831–842, 1996.
- [15] S. B. Ehrlich, S. C. Levine, and S. Goldin-Meadow, “The importance of gesture in children’s spatial reasoning,” *Dev. Psychol.*, vol. 42, no. 6, pp. 1259–1268, 2006.