A Social Robot Reading Partner for Explorative Guidance

Citation: Zhang, Xiajie, Breazeal, Cynthia and Park, Hae Won. 2023. "A Social Robot Reading Partner for Explorative Guidance."

As Published: https://doi.org/10.1145/3568162.3576968

Publisher: ACM|Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction

Persistent URL: https://hdl.handle.net/1721.1/150348

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

A Social Robot Reading Partner for Explorative Guidance

Xiajie Zhang
xiajie@media.mit.edu
Massachusetts Institute of Technology
Cambridge, MA, USA

Cynthia Breazeal
cynthiab@media.mit.edu
Massachusetts Institute of Technology
Cambridge, MA, USA

Hae Won Park
haewon@media.mit.edu
Massachusetts Institute of Technology
Cambridge, MA, USA

ABSTRACT

Pedagogical agent research has yielded fruitful results in both academic skill learning and meta-cognitive skill acquisition, often studied in instructional or peer-to-peer paradigms. In the past decades, child-centric pedagogical research, which emphasizes the learner's active participation in learning with self-motivation, curiosity, and exploration, has attracted scholarly attention. Studies show that combining child-driven pedagogy with appropriate adult guidance leads to efficient learning and a strengthened feeling of self-efficacy. However, research on using social robots for guidance in child-driven learning still remains open and under-explored. In our study, we focus on children's exploration as the vehicle in literacy learning and develop a social robot companion that provides guidance to encourage and motivate children to explore during a storybook reading interaction. To investigate the effect of the robot's explorative guidance, we compare it against a control condition in which children have full autonomy to explore and read the storybooks. We conduct a between-subjects study with 31 children aged 4 to 6, and the result shows that children who receive explorative guidance from the social robot exhibit a growing trend of self-exploration. Further, children's self-exploration in the explorative guidance condition is found correlated to their learning outcome. We conclude the study with recommendations for designing social agents to guide children's exploration and future research directions in child-centric AI-assisted pedagogy.

CCS CONCEPTS

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

KEYWORDS

child-robot interaction, child-centered pedagogy, educational technology, social robot, AI-guided education

ACM Reference Format:
Xiajie Zhang, Cynthia Breazeal, and Hae Won Park. 2023. A Social Robot Reading Partner for Explorative Guidance. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction (HRI '23), March 13–16, 2023, Stockholm, Sweden.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3568162.3576968

HRI '23, March 13–16, 2023, Stockholm, Sweden

ACM ISBN 978-1-4503-9964-7/23/03.
https://doi.org/10.1145/3568162.3576968

1 INTRODUCTION

Exploration has been critical in human development ever since infancy. Many psychologists believe that early childhood development is driven by children themselves through play and exploration in their environment [28, 29, 32]. For instance, in "effective pedagogy" [28], the author argues that an educator is to assist the learner in their learning process with engaging conversations for guidance. Moreover, Weisberg et al. [31] proposed "guided play" to explore the effect of limited guidance compared to fully child-driven play-based learning. They found that guided approaches have advantages over learning efficiency and positive social-emotional impact on the learner [9, 13]. However, challenges persist in theorizing a framework for the guided approach due to the complexity of the teacher-learner dual model and its personalization characteristics[35].

Social robots have great promises to provide personalized education [6]. Its affective, personalized, and playful interactions help the learner achieve an improved learning outcome with increased engagement [5, 15, 19]. In addition, social robots can provide effective social emulation for meta-cognitive adaptation, such as growing growth mindset, creativity, and curiosity [2, 16, 26]. Despite the fruitful results, less research in the area embraces child-centered pedagogy. Moreover, many pedagogical agent research studies meta-cognitive adaptation as a separate task from academic learning. We argue, from a learner-centric perspective, that the learner's motivation and meta-cognitive skills are important to be considered and supported within traditional learning contexts. Social robots have great potential to provide personalized guidance in child-centered learning. Our work explores using pedagogical agents for guided exploration, where the agent's behaviors are designed and personalized for children's self-exploration. Specially, we design robot behaviors to invoke and encourage children's exploration in a storybook reading context, and formulate a reinforcement learning model to learn a personalized policy for explorative guidance. Then, a between-subjects study is conducted to compare children's free-exploration and their exploration with the robot's explorative guidance. Our results show children have an increased amount of self-exploration with the explorative guidance from the social agent. Further, we find that children's exploration is related to their learning growth in the session. It shows social robots' potential to provide effective and personalized guidance supporting child-centric learning.

2 BACKGROUND

2.1 Free Exploration vs Guided Exploration

Exploration and play serve an essential role in children's learning process. Children acquire information about different objects' dynamic characteristics through active exploration of the environment. The inception engages their perceptual, motor-sensory, and cognitive abilities [4]. Many learning theories in psychology have established the importance of the learner's intrinsic motivation, exploration, and curiosity in learning [14, 23, 24]. For instance, in the constructivism theory of learning, instead of viewing the construction of knowledge as a result of instruction, it proposes that the learner's self-organization, motivated by reaching a cognitive equilibrium state, facilitates the learning process [14]. In addition, research in neuroscience suggests that human brains are intrinsically rewarded by seeking novel information [3, 17, 24]. Baranes et al. [3] find that trivia questions trigger participants' eye movements toward the area where they can find the answers. The advancement in learning theories and neuroscience has derived many learner-centric, exploration-and-play-oriented pedagogical approaches and technological interventions.

Though diverse in implementation based on the educational contexts, exploration or play-based pedagogy all share a similar main principle: the learner's control of the learning process [1, 13, 14, 31]. This type of pedagogy is vastly studied in both developmental learning and academic learning, although our work focuses on its application in skill-oriented academic environments. Evidence shows the benefits of exploratory learning include knowledge conceptualization, resilience, creativity, etc. [9, 10]. Child-centered pedagogical approaches are primarily divided into two categories: free- and guided-approaches. The free-approach, such as free-exploration, give children control of their learning processes. The guided-approach, on the other hand, refers to activities that include limited adult involvement to facilitate children's own learning [7]. For instance, Weisberg et al. [31] proposed a pedagogy named "guided play", combining children's autonomy with adult guidance to achieve a learning objective. In their proposal, the key components for an effective guided play practice are children's autonomy and guidance-focused scaffolding from the adult. The comparison between free-approach and guided-approach is a popular research topic. In 2011, a review done by Alfieri et al. [1] analyzes 164 pedagogical studies with "free-discovery" and "guided discovery" approaches and finds that the guided approach leads to improved learning. Further, when applied to geometric knowledge learning, researchers found stronger learning outcomes with the guidance [13]. In a study done by [9], the authors reported that "guided exploration" improved the learning outcome and increased the learner's feeling of self-efficacy Debowski et al. [9], suggesting that a trade-off exists between exploration and the complexity of the task according to the learner's ability. Guided-exploration provides a viable solution to the learner's effective problem-solving and social-emotional support.

However, challenges and opportunities are presented along with promising results in guided exploration. Theorizing and computationally modeling guided pedagogy suffers from the challenge imposed by the dynamic and interactive nature of the dual-person paradigm and its personalization characteristics based on each individual learner [35]. For instance, the effective scaffolding strategy in guided exploration differs depending on the learner's state and is under-studied. This motivates the development of personalized educational technology that searches for effective guidance strategies while giving learners space and autonomy in their learning.

2.2 Technological Design for Exploratory Learning

Social robots are studied as interactive educational technology due to their expressivity and potential for personalized interactions. Belpaeme et al. [6] presented a comprehensive review of the studies where a social robot was used to deliver an educational interaction with a learning-oriented objective. The author reviewed the efficacy of robot tutors based on the claimed cognitive outcome and affective outcome and found that in most of the reports, the robot's learning-oriented behavioral design yielded a positive effect. In addition, the pedagogical agent's personalization power has been proven effective at improving the learner's performance compared to corresponding non-personalized baselines [5, 15, 19]. Aside from academic learning, recent research shows that social robots have the potential to support children's meta-cognitive growth. Park et al. [26] built a cognitive architecture for a growth mindset social robot in problem-solving and found that children who interacted with the growth mindset robot exhibited stronger growth mindset behavior in problem-solving. Similarly, social robots have been used to help children become more creative and curious [2, 16]. Those approaches have advantages over potentially more persistent growth beyond the interaction and capitalize on children's self-improvement for other tasks.

Despite the fruitful results in academic learning and meta-cognitive skill adaptation with pedagogical agents, less research focuses on supporting children's meta-cognitive skills in traditional academic skill learning. For instance, most storybook companion robot studies focused on children's cognitive metrics as the main personalization target or the evaluation metric, such as the vocabulary level, story comprehension, or the perception of the robot [25, 37]. Elgarf et al. [12] studied creativity through the lens of storytelling; however, their study used storytelling as the creativity measurement, and the robot's meta-cognitive priming happened in an individualized activity.

Our work first makes a contribution by empirically studying the efficacy of using social robots to support child-centric learning with personalized guidance. Secondly, we investigate a learner's meta-cognitive skill development inside an academic reading context through the lens of self-exploration and literacy learning. Through this work, we hope to attract more research attention in social robotics to integrate different meta-cognitive skill attainment into academic learning settings and explore the robot's role in supporting the learner's self-efficacy in learning.

3 ROBOTICS SYSTEM OVERVIEW

The platform comprises three components–a social robot (Jibo), an android tablet with an interactive storybook app, and an ubuntu machine. The interactive storybook platform supplies the learning materials and interactable features that are designed to facilitate literacy learning. Robot Operating System (ROS) is used on the ubuntu machine to manage message passing between the robot, the tablet, and other software and algorithm results. Figure 1 shows an overview of the station.

(a) Robotics Station with Interactive Storybook

(b) A participant is learning with the robot at home

Figure 1: Robotics Station Overview

3.1 Interactive Storybook Reading

The interface and available functions for the interactive storybook are shown in Figure 2. The features in the storybooks are developed to support children's vocabulary and literacy learning phonologically and semantically. Interactive features in the storybook app are shown in Figure 2 (a). A "robot" button is placed on the bottom of the storybook scene for children to trigger robot behaviors, described in Section 3.2. Further, a keyword learning panel is provided for each keyword in the storybook with two learning materials: phonological decoding learning and word meaning explanation, shown in Figure 2 (c). The storybook selection is curated from a U.S. public charter school's storybook curricula for a kindergarten class. The audio-visual features and learning-support functionalities in digital books serve as an important medium for children's exploration [38]. In our storybook design, we identify mainly the following five behaviors as the means by which children can explore: (1) Scene Object Tapping (2) Tinker Text Tapping (3) Play Sentence Audio (4) Word Decoding and Pronunciation (5) Word Meaning Explanation Button.

3.2 Robot Platform and Robot Behaviors

A social robot, Jibo[1], is used in this study. Jibo is an embodied agent with auditory and visual features to support expressivity and interactivity. The robot listens, speaks, and has rich body language with animations to show a diverse range of affective capabilities.

Two robot behaviors are designed for the pedagogical agent within the storybook interaction context. The first is the **Prompting Behavior**. The design of prompting behavior is based on empirical results in adult-guided play studies [31] and practical principles in dialogic reading literature [33, 36]. It is a practical skill often used by adults to guide storybook reading activities. The behavior composes of four stages of a dialogic interaction: (a) the robot prompts a question related to the storybook content or a keyword; (b) the robot waits for the child's response; (c) the robot gives feedback or its version of the answer after the child's response; (d) contingent on children's response, the robot gives hints and encourages the child to answer again. We use google Automatic Speech Recognition (ASR) service to interpret children's speech with a focus on specific keywords for giving a hint (such as, "I don't know", "No", etc.) due to the limitation of current state-of-the-art child ASR models. The other type of robot behavior is the **Exploratory Demonstration Behavior**. Research shows that children can emulate social

[1]Jibo Robot, NTT Disruption. http://jibo.com

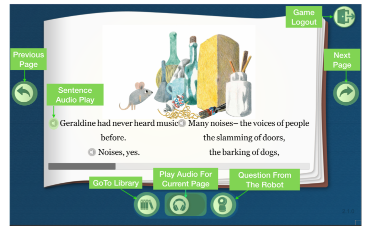

(a) Functions in Storybook Scene

(b) Storybook Interface and Text-Image Interactables

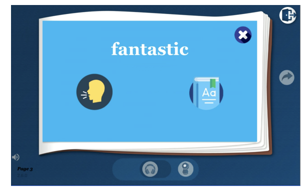

(c) Keyword Learning Panel

Figure 2: Interactable features in the Storybook Platform. (a) highlights all available features on the storybook scene, which includes basic utilities (page flipping) and interactable features (audio replay, robot interaction button); (b) shows the triggers between image and text for keyword learning; (c) presents the keyword learning panel

robots' behavior through interactions [2, 18]. For instance, a robot's creativity demonstration can be emulated to promote children's figural creativity. Taking insights from psychological theories for exploratory behavior and social emulation [4], the demonstrative exploration by an agent has the following key components:

(1) A display of curiosity and motivation to explore a certain stimulus (a keyword)
(2) Point out / show the means of exploration
(3) Carry out the exploration action
(4) Confirmation of the learning from the exploration behavior

In storybook reading, the robot's explorative behaviors are around the keywords as they serve a natural learning objective in literacy learning. Thus, for each keyword, an explorative demonstration consists of (a) displaying motivation to learn about word pronunciation or meaning; (b) finding and selecting the correct resource to attain the information; (c) demonstrating learning satisfaction. The prompts and responses in the open-ended dialogic interaction and word explanations are created by a team of educators and thoroughly reviewed for child-friendliness and engagement. In addition, a list of utterances to deliver the robot's intention (e.g., motivation, demonstration, etc.) is summarized in Appendix C.

3.3 Free Exploration and Explorative Guidance

In this study, we investigate the means and efficacy of using a social robot platform to provide explorative guidance for children. The robot's behaviors include open-ended questions and exploration demonstration, detailed in Section 3.2. As mentioned in Section 2.1, guided-exploration involves the personalization of guidance strategy depending on each learner. We now describe our approach to explorative guidance with a reinforcement learning framework.

3.3.1 Explorative Guidance Condition. Studies with guided exploration or any child-centered pedagogy with guidance are usually performed with an experienced adult. The theoretical framework for the guided approach is an ongoing challenge in pedagogy research due to the complexity of the dual modeling of the learner's and teacher's states. The adult must consider the learner's cognitive and emotional factors to decide on practical guidance for the specific learner. Taking insights from psychology literature on guided play [31], we designed the robot's explorative guidance by personalizing the guidance strategy and timed proactivity. The personalization is formulated with a reinforcement learning model.

For an autonomous agent to learn the best behavior policy for guided exploration, the information about the learner and their exploration environment is critical. Thus, we construct the state space of the Markov Decision Process (MDP) with three variables that span different aspects of the learner and their exploration environment: (1) The child's engagement level, categorized in three levels ($S_{engagement} \in [0, 1, 2]$). (2) Whether the child had explored on the last page ($S_{explored} \in [0, 1]$). (3) Whether there is an unknown keyword present on the page ($S_{unknown} \in [0, 1]$). In total, the state space has $3x2x2 = 12$ different states. This design includes children's affective, behavioral, and knowledge (cognitive) information for the agent's decision-making. The agent's actions consist of the two aforementioned robot behaviors–exploration demonstration and prompting. For the child's engagement, we measured it with a camera video feed and affect detection software–Affectiva [21]. The software returns a score for engagement in the $[0, 100]$ range. Then, the child's average engagement of a window is categorized into three levels based on their own lower and higher bounds. We use children's affect range in the pre-study session as their initial affect bounds.

The reward function is crucial for specifying the learning objective. In guided exploration, the agent's guidance must be rewarded based on children's self-exploration. Thus, we formulate the child's display of exploration with the amount and diversity of different explorative behaviors the child exhibits. Specifically, it is defined as Eq. 1.

$$
\begin{aligned}
R^t_{total} &= r^t_{part} + r^t_{expl} \\
r^t_{part} &= \begin{cases} 0.2 & \text{if responded after action } t-1 \\ 0 & \text{otherwise} \end{cases} \\
r^t_{expl} &= \frac{1}{Z} \sum_{b \in B}^{B} \log_2(C_b + 1), C_b \leq 3
\end{aligned}
\tag{1}
$$

$r_{part} \in [0, 0.2]$ stands for the reward for active participation. This reward is a one-time reward after each robot's interaction. When the child answers the robot's question or exhibits exploration after the robot demonstrates exploration, there is a positive reward for the agent's behavior. If the child is inactive after the agent's action at time $t-1$, no reward is assigned to the action. $r_{expl} \in [0, 2]$ stands for the reward for children's exploratory behaviors after the agent's action. r_{expl} is designed to account for both the amount of children's exploration and the diversity of their exploration. The set of exploratory behaviors captured by the app are summarized in Section 3.1 and are referenced as the in-exhaustive list of possible explorations a child might exhibit in this study. In the equation, B refers to the set of exploratory behaviors, and C_b means the count of certain exploratory behavior. The reward for exploratory behaviors r_{expl} is an averaged sum across all types of exploratory behaviors defined in Section 3.1. Intuitively, the reward increases when the child exhibits diverse exploratory behaviors or attempts multiple explorations. The count for each type of explorative behavior is limited to 3 to avoid over-counting an excessive behavior. The inbalanced rewards for participation and exploration are intentional for heavy-weighting on children's self-exploration.

Model-free approaches in reinforcement learning are popular for Markov Decision Process (MDP) problems where the transition probability distribution is absent or difficult to model. In intelligent tutoring systems and pedagogical agent studies, researchers have employed such an approach for situations where the transition model of the learning task is hard to acquire [8, 20, 25]. In our work, we employ a model-free Q-learning [30] agent for online learning due to the difficulty in modeling children's behavior with guided exploration. In addition, the Upper Confidence Bound algorithm is applied for the action selection to balance the exploration-exploitation trade-off. We expand the discussion of the RL model's performance in Appendix A.

3.3.2 Free Exploration Condition. To evaluate the effect of guided exploration, we use a free-exploration condition for comparison. In the free-exploration condition, children's explorations are unfettered. They have complete autonomy and agency to read and explore the storybook app. Further, to account for the effect of the embodiment of the robot and its behaviors, children in the free-exploration condition have exposure to the same type of robot behaviors as in the guided-exploration condition. As a result, we employ an alternating approach to preserve children's full agency while giving them exposure to robot interactions. With the robot button on the tablet screen, children can request an interaction from the robot in the free-exploration condition. The interaction generated by the robot alternates mechanically between the prompting behavior and exploratory demonstration behavior. With this design, children in both conditions receive the robot's open-ended questions and exploration demonstration. In comparison, the explorative guidance receives proactive guidance customized based on their affective and behavioral states.

4 EXPERIMENT DESIGN

The experiment is designed to evaluate the impact of exploration demonstration and the efficacy of guided exploration when it is provided by a peer-like social robot. Our main hypotheses are:

- **H1a**: Children who interacted with the robot either in the free exploration or guided exploration group show an increasing trend in exploratory behavior.
- **H1b**: Children in the guided exploration condition show a more significant effect of interaction than children in the free exploration condition.
- **H2**: Children's learning outcome has a significant correlation to their exploratory behavior, and the trend is more significant in the guided exploration condition than in the free exploration condition.

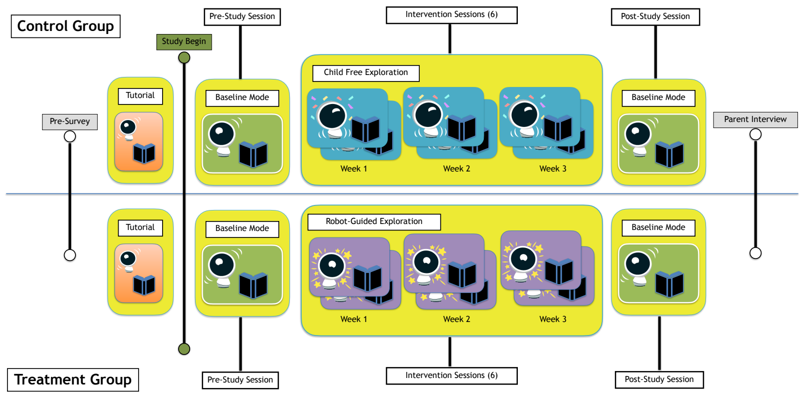

Figure 3: Overview of the study design. In the two baseline mode (pre-study & post-study sessions), the child leads and explores the storybook reading without much robot exploration (only the dialogic interaction). The intervention sessions for the control group is free exploration with full exposure to the robot's behaviors. For the intervention sessions in the experiment group, the robot provides explorative guidance.

4.1 Study Design

We designed a between-subjects A-B-A study to evaluate the effect of guided-exploration provided by the social robot. Participants were divided into two groups. We designed our comparative study motivated by the studies in pedagogical research where a free-approach, i.e., children have full control of their learning activity, is compared with a guided-approach [1, 9, 13]. In the control group, children's storybook interactions were free-exploration – they had freedom in the storybook reading, including when to interact with the robot. Specifically, children retained the agency to choose when the robot would present a question or a prompt for exploration, as described in Section 3.3.2. The robot in the treatment group proactively provided explorative guidance. The personalization algorithm that drove the robot's guidance policy is detailed in Section 3.3.1.

Within each condition, the two A-components were the pre- and post-study evaluation sessions. A baseline interaction mode was used in the evaluation sessions to understand children's natural interaction with the pedagogical agent without explorative demonstration or explorative guidance. The robot's behavior is limited to its prompting behavior and is activated by the child by pressing a button on the screen to give children the full agency to explore. In between the pre-study session and the post-study session are the intervention sessions.

The study design is shown in Figure 3. The entire station is shipped to participants' homes, and each participant goes through a comprehensive station setup procedure and onboarding process where the child receives a tutorial on how to use the storybook App with the robot from the experimenter before the study begins. All participants complete a pre-survey and relevant pre-screeners before officially starting the study. The study includes eight storybooks with different themes and genres, detailed in Table 1. The child is expected to complete the study in three to five weeks, depending on their reading pace. However, parents and the child have the autonomy on when to read with the robot in their homes. The study is fully autonomous, and the experimenters contact the parents weekly with the study progress updates through emails.

Session	Storybook	Illustration	Genre
Pre-Session	Farm Animals	PHO	INF
1st Intervention	The Legend of the Bluebonnet	DRW	NAR
2nd Intervention	The Little House (Part 1)	DRW	NAR
3rd Intervention	The Little House (Part 2)	DRW	NAR
4th Intervention	Homes Around the World	PHO	MIX
5th Intervention	Helpers in My Community (Part 1)	PHO	INF
6th Intervention	Helpers in My Community (Part 2)	PHO	INF
Post-Session	From Sheep to Sweater	PHO	MIX

Table 1: Storybook Summary. In Genre, INF, NAR, and MIX stand for informational, narrative, and mixed (description of each category is from [27]). In Illustration, PHO and DRW stand for photographic and hand drawn illustrations.

In addition, the storybook is selected and curated to a reasonable length to keep children engaged. Each story takes approximately 10 to 15 minutes to complete. To measure children's learning in each session, the child completes a PPVT-Style questionnaire about the keywords in the story before and after each storybook reading session. The PPVT-Style questionnaires are curated by an educational professional around keywords that appeared in the listed storybook, following the same style as the official PPVT test [11].

4.2 Participants

In total, out of 42 children (from 40 families) who were recruited, 33 children from 32 families completed the study. Families dropped out due to (a) family emergency (1); (b) time conflicts (5); (c) network or technical issues (2). In addition, 2 children were excluded due to consistent signs of disengagement (sessions started but not finished, etc.). A description of the demographic information was included in Appendix B. Before the study, participants' parents signed parental consent forms and filled out a survey with their demographic information and their children's reading habits and frequency. The

Group	n	Gender	Age	Grade	SES	Self-Read	Literacy Screener (out of 20)	
							Median	Mean ± SD
Control	16	F=8 M=8	M=5.312 SD=0.768	PreK = 7 K & Above = 9	High=11 Low=5	High=5 Mid=5 Low=6	20.0	19.25 ± 1.09
Treatment	15	F=8 M=7	M=5.133 SD=0.718	PreK = 8 K & Above = 7	High=10 Low=4 Unreported=1	High=9 Mid=2 Low=4	19.0	18.33 ± 2.52
Total	**31**	**F=16** **M=15**	**M=5.226** **SD=0.750**	**PreK = 15** **K & Above = 16**	High=21 Low=9 Unreported=1	High=14 Mid=7 Low=10	20.0	18.81 ± 1.97
Stat.			$p = .48$	$p = .56$	$p = .76$	$p = .20$	$p = .45$	

Table 2: Student Demographics Summary. SES stands for Socioeconomic status, measured by the family's annual income. Pre-study literacy scores statistics is from the Get-Ready-To-Read-Screener. No significant differences were found between the control and treatment groups with non-parametric Mann-Whitney U Test.

participants were randomly assigned into two study groups. 15 were in the treatment group, and 16 were in the control group. Mann-Whitney U test was applied to test statistical differences in the demographic information between the two groups. No significant differences between the two groups in terms of children's age, gender, grade, family socioeconomic status (annual family income), and children's self-reading efficacy (measured by parent-reported weekly reading time) were found. Children's pre-study reading level was assessed with the Get Ready To Read Screener [34]. The medians of the final score in the control and treatment groups were 20 and 19, respectively. The Mann-Whitney U test was applied to test the statistical difference in the pre-screener score. No significant difference between the two study groups was found. A summary of study participants' demographic and literacy screener data and statistical tests between the two groups are summarized in Table 2).

5 DATA ANALYSIS

The study persisted from 2 weeks to 4 weeks, depending on the participant's usage frequency. We first analyze the station usage between the two groups; then, we investigate the study's effect on children's self-exploration and learning outcomes.

5.1 Station Usage

Overall, the number of days for families to finish the eight sessions was 24.84 ± 12.53. The difference between the number of days for completion between the control and treatment group was not significant (Control: 23.44 ± 12.58; median 21.50, Treatment: 26.33 ± 12.74; median 25.00; Unit: Days). The **Frequency** was calculated by the average days between two sessions. The **Time Per Page** indicates the average amount of time spent on each page. On average, children read a book with the robot approximately every three days (Control: 2.93 ± 1.57 median 2.69; Treatment: 3.29 ± 1.59 median 3.13; Unit: Days). In an average intervention session, children in the control group spent around 1 minute reading a page while children in the treatment group spent 1.24 minutes (Control: 1.04 ± 0.32 median 0.99; Treatment: 1.24 ± 0.40 median 1.26; Unit: Minutes).

5.2 Interaction Effect and Exploration

We quantify children's exploration in a storybook session as the sum of their explorative behaviors (the list of behaviors is detailed in Section 3.1). Then, the sum is divided by the total amount of explorable features in the given storybook, including audio, keyword, and word-image triggers. Each session is separately calculated for the following analysis. Figure 4 shows the comparison of children's exploration during the pre-study session, the intervention sessions, and the post-study session. Mann-Whitney U test (M-W) revealed no significant difference between the two groups in the pre-study session ($U = 95.5, p = .746$). No statistical significance was observed either in the post-study session ($U = 100.5, p = .928$).

The Mann-Whitney U is applied to test the statistical differences between the medians of explorations in pre-study sessions and the intervention sessions in two study groups. The control group's median for pre-study sessions and intervention sessions are 0.075 and 0.048, respectively. M-W test results are $U = 796.50, p = .25$; In the treatment group, the median for the pre-study session and intervention sessions are 0.082 and 0.133, and the M-W results show a significant difference ($U = 311.00, p < .05$).

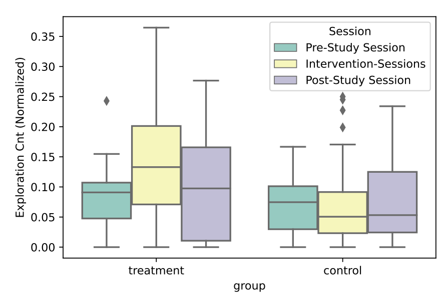

Figure 4: Exploration behavior comparison within and between study groups

To observe children's exploration progression on the time scale, Figure 5 and 6 show linear regression between children's exploration amount over the session timeline. We excluded the last session in this analysis because children in the treatment group were not used to using the robot button to activate robot prompts, resulting in an unfair comparison. Further discussion and reflection are provided in Section 6. The sessions on the x-axis include pre-study sessions and intervention sessions from 1 to 6. The linear regression line shows an evident ascend in the treatment group, suggesting a rise in children's exploration as the study progresses. In the control group, children's exploration shows smaller inconsistent variances with no sign of ascending or descending. Simple linear regression was used to test if the session number significantly predicted children's exploration. The treatment group's regression showed a statistically significant increasing trend ($R^2 = 0.051$, $F(1, 97) = 5.162$, $p = .0253*$). It was found that session number significantly predicted children's exploration ($\beta = 0.010$, $p = .025*$).

Further statistical analysis of the difference between the treatment and control group by each session is performed with Mann-Whitney U for a non-parametric test of two populations [22]. The results are summarized in Table 3. All of the intervention sessions between the two groups, except for the second session, were significantly different in the between-group comparison.

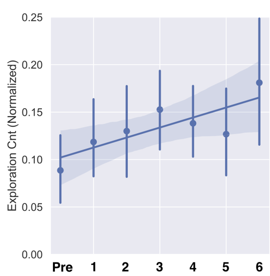

Figure 5: Linear Regression on Session in Treatment

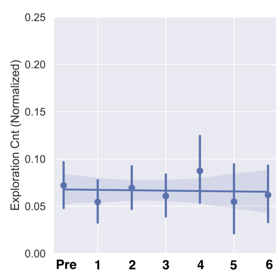

Figure 6: Linear Regression on Session in Control Group

5.3 Exploration and Learning

Aside from exploration, children's learning outcome is another critical criterion for the evaluation. At the beginning and end of each storybook session, children completed a pre- and post-session vocabulary assessment of the target vocabulary in the storybook. The exploration resources provided in the storybook were developed to help children learn the target vocabulary. We extracted the accuracy rate of children's pre- and post-session assessment results.

To investigate whether exploration was related to the learning gain, we calculated the learning growth by checking the difference between the pre- and post-session assessment results. If the child answered a question correctly in the pre-assessment but got it wrong in the post-assessment, we recognized it as a guessed answer and marked it as incorrect. During all the intervention sessions, children in the treatment group had an average of 0.096 ± 0.16 in learning growth, and the control group had an average of 0.073 ± 0.09. The treatment group learning gain ($Median = 0.067$) was higher than the control group ($Median = 0.048$), but the Mann-Whitney test revealed that this difference was not statistically significant ($U = 2774.5$, $p = .84$).

Next, Spearman's rank correlation was computed to assess the relationship between children's exploration and learning gain. In the treatment group, there was a significant positive correlation between exploration and learning ($r = 0.22$, $p = .046*$), but no correlation was found in the control group ($r = -0.01$, $p = .92$). This correlation suggests that children's exploration in the group with explorative guidance positively correlates with their learning outcome.

6 DISCUSSION

Explorative behaviors are essential in early childhood learning as a self-motivated information-seeking process. In pedagogical research, guided exploration has been proven to effectively nudge children's exploration and lead to efficient learning compared to free exploration, along with social and emotional benefits, such as an improved feeling of self-efficacy [9, 13, 31]. Our work motivates to build social pedagogical agents that support children's exploration with personalized guidance.

Social Robot's Guided Exploration is Effective at Increasing Children's Exploratory Behavior.
Although children's exploration in the pre-study and post-study baseline interactions did not show significant differences, the comparison between the pre-study session and the intervention sessions shows that children in the guided-exploration group exhibited more exploratory behaviors than the control group. Further, the linear regression result between children's exploration and the session number shows that only guided exploration led to behavioral change in children, even though the social agent demonstrated exploration in both conditions; proactively in the treatment group and reactively in the control group. Thus, hypothesis **H1a** was not supported, and social emulation alone could not explain children's exploration growth. However, it was shown that personalized guidance provided by the robot in the guided exploration group played a significant role in increasing children's exploratory behavior over time, which supported hypothesis **H1b**.

Despite the growing trend in the treatment intervention sessions, children's exploration in the post-study session did not retain that growth. The reason could have been that the study length, i.e., the length and amount of exposure to guided exploration wasn't sufficient, but upon a deeper look into children's explorative behavior patterns, we noticed that they seldom used the robot button that activated the robot's exploration guidance behavior. This is likely due to the treatment group children's unfamiliarity with the use of the button compared to children in the control group. Another trend to note is the 4th and 5th sessions. Children's explorative growth in the treatment intervention was consistent except for a sudden drop in these two sessions. This drop might be due to the change in the storybook genre. The storybook genre transitioned into more informational content with real-life photo illustrations, compared to narrative-style story content with hand-drawn illustrations in prior sessions. Education research shows that preschool children's exposure to different storybook genres is heavily skewed towards narrative stories [27]. The change in the storybook to a less familiar genre might have introduced a learning period in which children focus their effort on comprehending the new style, and as

Session	Treatment (avg±sd)	Control (avg±sd)	M-W test
Pre-Study Session	0.083 ± 0.061	0.072 ± 0.053	$U = 95.5, p = .746$
1st Intervention Session	0.113 ± 0.086	0.054 ± 0.041	$U = 45.0, p = .044^{*}$
2nd Intervention Session	0.129 ± 0.105	0.074 ± 0.048	$U = 83.5, p = .163$
3rd Intervention Session	0.140 ± 0.079	0.058 ± 0.047	$U = 42.5, p = .002^{**}$
4th Intervention Session	0.129 ± 0.086	0.064 ± 0.046	$U = 51.0, p = .031^{*}$
5th Intervention Session	0.120 ± 0.097	0.039 ± 0.047	$U = 44.5, p = .028^{*}$
6th Intervention Session	0.176 ± 0.123	0.056 ± 0.054	$U = 45.0, p = .014^{*}$
Post-Study Session	0.113 ± 0.102	0.070 ± 0.072	$U = 100.5, p = .928$

Table 3: Between-group comparison of children's exploratory behavior in each session using Mann-Whitney U Test.

they become more familiar with the new genre, they are motivated to further explore the book as the trend shows in session 6.

Children's Exploratory Behavior Leads to Learning Gains with Robot's Guided Exploration

Children's exploration in the robot-guided condition correlates with their learning outcome, supporting hypothesis **H2**. This correlation suggests that the robot's guided exploration helps children to initiate learning-oriented exploration. This result shows that guided exploration by a social agent is a malleable approach to promoting learning with child-centered pedagogical principles.

Design Recommendations for Robot Guided Exploration Interaction for Children

Research in developmental psychology and empirical pedagogy has not yet concluded a unified framework of the effective components in adult-guided child-centered learning. Through this study, we draw design recommendations for future pedagogical agent research for guided exploration. First, **diversify robot's behavioral design** to promote children's exploration. In our study, the robot's open-ended question around the storybook contents and keywords and demonstration of curiosity and exploration both provided means for children to interact with storybooks in multiple dimensions and socially emulate being curious and motivated to explore. Secondly, **the timing and amount of guidance are critical**. In the guided-exploration group, the robot's behavior is driven by a personalized policy that chooses the most supportive guidance behavior for the child. The personalization approach is essential in guided exploration because pedagogical research shows that guidance effective for one child might not work as well for others due to interpersonal differences [13]. Moreover, the pedagogical agent's **proactive guidance** is more effective than "reactive" guidance that only responds to the child's requests. Lastly, researchers need to be **mindful of external factors that may affect children's engagement and motivation to explore**, such as the storybook's genre, illustration style, and difficulty levels. Our study used a pre-curated storybook collection and fixated the order to avoid environmental variance for the same session. However, we recommend future research studying children's exploration in storybook interactions to examine the storybook-related factors to ensure the balance between explorability and uncertainty.

7 FUTURE WORK AND CONCLUSION

This work investigates a social robot's supporting role in an unprecedented long-term home setting with child-driven pedagogy in story reading. A reinforcement learning algorithm was proposed for training the robot's personalized explorative guidance policy. The algorithm was rewarded by a diverse set of exploratory behaviors of the child in response to the robot's guidance. Compared to child-led learning, our result showed that children who interacted with a proactive explorative guidance robot showed increased self-explorations. Further, their exploration was correlated with learning gain. We hope this work raises researchers' attention to social robots' supporting role in children's meta-cognitive skill development. With that, we suggest the following future work.

The current study focused on comparing the effect of free exploration and guided exploration to learn if the results in pedagogy literature between adults and children are replicated in a child-robot study. Due to the study design, this paper could not discuss the personalization effect in isolation. Future work will include a factorial study design to disentangle the effect of each element in the active explorative guidance–proactiveness and personalization–provided by the robot. We acknowledge that the assessment for learning in this study could be strengthened and validated with a more sophisticated assessment design, for instance, a delayed post-study vocabulary assessment. Similarly, the quantitative metrics for children's exploration in our analysis could further be improved by combining qualitative analysis of exploration. Lastly, this work focused on the behavioral interventions provided by the robot for explorative guidance through social emulation and prompting. However, the motivation to explore has both cognitive and emotive factors. Understanding children's cognitive and emotive aspects and providing social-emotional support with well-designed dialog would be critical for the robot's guidance in child-robot interaction.

8 ACKNOWLEDGEMENT

This work was supported by a National Science Foundation grant, DRL-1734443 and by Inclusive AI Literacy and Learning gift grant from DP World. We sincerely thank Prof. Robin Morris and Prof. Gary Bingham for their advisement on the early literacy curriculum and study design. We also thank Jon Ferguson, Kejia Patterson, and Meng Xi for system development, study content management, and logistics support that enabled this long-term home deployed study.

REFERENCES

[1] Louis Alfieri, Patricia J Brooks, Naomi J Aldrich, and Harriet R Tenenbaum. 2011. Does discovery-based instruction enhance learning? *Journal of educational psychology* 103, 1 (2011), 1.

[2] Safinah Ali, Hae Won Park, and Cynthia Breazeal. 2021. A social robot's influence on children's figural creativity during gameplay. *International Journal of Child-Computer Interaction* 28 (2021), 100234.

[3] Adrien Baranes, Pierre-Yves Oudeyer, and Jacqueline Gottlieb. 2015. Eye movements reveal epistemic curiosity in human observers. *Vision research* 117 (2015), 81–90.

[4] Abigail Baxter and Harvey Switzky. 2008. *Exploration and Curiosity*. 460–470. https://doi.org/10.1016/B978-012370877-9.00058-X

[5] Paul Baxter, Emily Ashurst, Robin Read, James Kennedy, and Tony Belpaeme. 2017. Robot education peers in a situated primary school study: Personalisation promotes child learning. *PloS one* 12, 5 (2017), e0178126.

[6] Tony Belpaeme, James Kennedy, Aditi Ramachandran, Brian Scassellati, and Fumihide Tanaka. 2018. Social robots for education: A review. *Science robotics* 3, 21 (2018), eaat5954.

[7] Erica Danniels and Angela Pyle. 2018. Defining play-based learning. *Encyclopedia on early childhood development* (2018), 1–5.

[8] Lucie Daubigney, Matthieu Geist, and Olivier Pietquin. 2013. Model-free pomdp optimisation of tutoring systems with echo-state networks. In *Proceedings of the SIGDIAL 2013 Conference*. 102–106.

[9] Shelda Debowski, Robert E Wood, and Albert Bandura. 2001. Impact of guided exploration and enactive exploration on self-regulatory mechanisms and information acquisition through electronic search. *Journal of Applied Psychology* 86, 6 (2001), 1129.

[10] Daniel A DeCaro, Marci S DeCaro, and Bethany Rittle-Johnson. 2015. Achievement motivation and knowledge development during exploratory learning. *Learning and Individual Differences* 37 (2015), 13–26.

[11] Lloyd M Dunn and Leota M Dunn. 1965. Peabody picture vocabulary test–. (1965).

[12] Maha Elgarf, Natalia Calvo-Barajas, Patrícia Alves-Oliveira, Giulia Perugia, Ginevra Castellano, Christopher Peters, and Ana Paiva. 2022. " And then what happens?" Promoting Children's Verbal Creativity Using a Robot. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction*. 71–79.

[13] Kelly R Fisher, Kathy Hirsh-Pasek, Nora Newcombe, and Roberta M Golinkoff. 2013. Taking shape: Supporting preschoolers' acquisition of geometric knowledge through guided play. *Child development* 84, 6 (2013), 1872–1878.

[14] Catherine Twomey Fosnot and Randall Stewart Perry. 1996. Constructivism: A psychological theory of learning. *Constructivism: Theory, perspectives, and practice* 2, 1 (1996), 8–33.

[15] Goren Gordon and Cynthia Breazeal. 2015. Bayesian active learning-based robot tutor for children's word-reading skills. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 29.

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

[17] Celeste Kidd, Steven T Piantadosi, and Richard N Aslin. 2014. The Goldilocks effect in infant auditory attention. *Child development* 85, 5 (2014), 1795–1804.

[18] Jacqueline M Kory-Westlund and Cynthia Breazeal. 2019. A long-term study of young children's rapport, social emulation, and language learning with a peer-like robot playmate in preschool. *Frontiers in Robotics and AI* 6 (2019), 81.

[19] Daniel Leyzberg, Samuel Spaulding, and Brian Scassellati. 2014. Personalizing robot tutors to individuals' learning differences. In *2014 9th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 423–430.

[20] Hsuan-Ta Lin, Po-Ming Lee, and Tzu-Chien Hsiao. 2015. Online pedagogical tutorial tactics optimization using genetic-based reinforcement learning. *The Scientific World Journal* 2015 (2015).

[21] Daniel McDuff, Abdelrahman Mahmoud, Mohammad Mavadati, May Amr, Jay Turcot, and Rana el Kaliouby. 2016. AFFDEX SDK: a cross-platform real-time multi-face expression recognition toolkit. In *Proceedings of the 2016 CHI conference extended abstracts on human factors in computing systems*. 3723–3726.

[22] Patrick E McKnight and Julius Najab. 2010. Mann-Whitney U Test. *The Corsini encyclopedia of psychology* (2010), 1–1.

[23] Monica Nilsson, Beth Ferholt, and Robert Lecusay. 2018. 'The playing-exploring child': Reconceptualizing the relationship between play and learning in early childhood education. *Contemporary Issues in Early Childhood* 19, 3 (2018), 231–245.

[24] P-Y Oudeyer, Jacqueline Gottlieb, and Manuel Lopes. 2016. Intrinsic motivation, curiosity, and learning: Theory and applications in educational technologies. *Progress in brain research* 229 (2016), 257–284.

[25] Hae Won Park, Ishaan Grover, Samuel Spaulding, Louis Gomez, and Cynthia Breazeal. 2019. A model-free affective reinforcement learning approach to personalization of an autonomous social robot companion for early literacy education. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 687–694.

[26] Hae Won Park, Rinat Rosenberg-Kima, Maor Rosenberg, Goren Gordon, and Cynthia Breazeal. 2017. Growing growth mindset with a social robot peer. In *Proceedings of the 2017 ACM/IEEE international conference on human-robot interaction*. 137–145.

[27] Jill M Pentimonti, Tricia A Zucker, and Laura M Justice. 2011. What are preschool teachers reading in their classrooms? *Reading Psychology* 32, 3 (2011), 197–236.

[28] Roland G Tharp. 2006. Four hundred years of evidence: Culture, pedagogy, and Native America. *Journal of American Indian Education* (2006), 6–25.

[29] Lev Semenovich Vygotsky. 1978. Mind in Society: The Development of Higher Psychological Processes, ed. *Michael Cole et al.(Cambridge, MA: Harvard UniversityPress, 1978). More information about Brian's system, Synanim, can be found at http://www. synanim. com* (1978).

[30] Christopher JCH Watkins and Peter Dayan. 1992. Q-learning. *Machine learning* 8, 3 (1992), 279–292.

[31] Deena Skolnick Weisberg, Kathy Hirsh-Pasek, Roberta Michnick Golinkoff, Audrey K Kittredge, and David Klahr. 2016. Guided play: Principles and practices. *Current directions in psychological science* 25, 3 (2016), 177–182.

[32] Ann Weisler and Robert R McCall. 1976. Exploration and play: Resume and redirection. *American Psychologist* 31, 7 (1976), 492.

[33] Grover J Whitehurst. 1992. Dialogic reading: An effective way to read to preschoolers.

[34] Grover J. Whitehurst and Christopher J. Lonigan. 2002. "Get Ready to Read! Screening Tool." 2002.

[35] Yue Yu, Patrick Shafto, Elizabeth Bonawitz, Scott C-H Yang, Roberta M Golinkoff, Kathleen H Corriveau, Kathy Hirsh-Pasek, and Fei Xu. 2018. The theoretical and methodological opportunities afforded by guided play with young children. *Frontiers in psychology* 9 (2018), 1152.

[36] Andrea A Zevenbergen and Grover J Whitehurst. 2003. Dialogic reading: A shared picture book reading intervention for preschoolers. *On reading books to children: Parents and teachers* 177 (2003), 200.

[37] Zhao Zhao and Rhonda McEwen. 2022. " Let's read a book together" A Long-term Study on the Usage of Pre-school Children with Their Home Companion Robot. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction*. 24–32.

[38] Marcy Zipke. 2017. Preschoolers explore interactive storybook apps: The effect on word recognition and story comprehension. *Education and Information Technologies* 22, 4 (2017), 1695–1712.