

Adaptive Role Switching in Socially Interactive Agents for Children’s Language Learning

Huili Chen

B.S. Computer Science, University of Notre Dame (2016)

B.A. Psychology (minor: International Peace Studies), University of Notre Dame (2016)

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning
in partial fulfillment of the requirements for the degree of
Master of Science in Media Arts and Sciences
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2018

(c) 2018 Massachusetts Institute of Technology. All rights reserved.

Author **Huili Chen**
Program in Media Arts and Sciences,
School of Architecture and Planning
May 4, 2018

Certified by **Cynthia Breazeal**
Associate Professor of Media Arts and Sciences
Thesis Supervisor

Author **Tod Machover**
Muriel R. Cooper Professor of Music and Media
Academic Head, Program in Media Arts and Sciences

Adaptive Role Switching in Interactive Agents for Children’s Language Learning

by
Huili Chen

B.S. Computer Science, University of Notre Dame (2016)

B.A. Psychology (minor: International Peace Studies), University of Notre Dame (2016)

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning
on May 4, 2018 in partial fulfillment of the
requirements for the degree of
Master of Science in Media Arts and Sciences

Abstract

Learning language and literacy at a young age is important, as children’s early language ability can impact their later educational success [1][2]. However, one of the major barriers to early language and literacy learning for many children around the globe is a lack of resources in homes and schools. A variety of technological interventions, such as TV series and educational apps, were designed to help overcome such barriers and support children’s learning. However, not all of them necessarily provide children with conversational experiences, which have been found to significantly impact the children’s language-related neural development [3]. Among a variety of educational media, embodied interactive agents (e.g., social robots) seem to be an effective yet resource-efficient tool that can enable children to learn through conversational turn taking. Specifically, embodied interactive agents can serve as learning companions for young children and provide more interactive and immersive learning experience.

I explored how social robots could help promote children’s language and literacy learning. More specifically, I designed and computationally created a collaborative, engaging learning interaction between a robot and a child who play as peers. First, I designed a tablet-based literacy learning game called WordQuest using the design principles for educational games. Second, I developed a reinforcement learning model that enabled the robot to adaptively switch its collaborative roles (e.g., expert and novice roles) in a way that promoted children’s best learning. Third, I conducted an experiment with three conditions, which were fixed expert robot, fixed novice robot, and adaptive role switching robot, and tested on 60 children recruited from a local primary school in Boston. Last, I evaluated how the robot’s collaborative roles differentially affected children’s learning performance, engagement, and perception of the learning experiences. I found out that children across the three conditions all learned new words and had a very positive experience of playing WordQuest with the robot. In addition, children interacting with the adaptive robot consistently outperformed children from the other two conditions in terms of vocabulary acquisition and retention.

Thesis Supervisor: Cynthia Breazeal

Title: Associate Professor of Media Arts and Sciences

Adaptive Role Switching in Interactive Agents for Children's Language Learning

by

Huili Chen

The following people served as readers for this thesis:

Thesis Advisor:

Dr. Cynthia Breazeal
Associate Professor of Media Arts and Sciences
Personal Robots Group
MIT Media Lab

Adaptive Role Switching in Interactive Agents for Children's Language Learning

by

Huili Chen

The following people served as readers for this thesis:

Thesis Reader:

Dr. Rosalind Picard
Professor of Media Arts and Sciences
Affective Computing Group
MIT Media Lab

Adaptive Role Switching in Interactive Agents for Children's Language Learning

by

Huili Chen

The following people served as readers for this thesis:

Thesis Reader:

Dr. Maryanne Wolf
Professor of Citizenship and Public Service
Eliot-Pearson Department of Child Study and Human Development
Tufts University

Acknowledgement

Many people supported my research and contributed to the creation of this thesis. I would love to acknowledge all the support, help and efforts that each of them has generously provided.

First of all, I would love to thank my advisor, Dr. Cynthia Breazeal, for her guidance, ideas, support and vision for AI-driven social innovation. This thesis could not be possible done without her effort. Dr. Hae Won Park has been a great mentor in my life, and I really appreciate all the time she spent with me on brainstorming ideas, conceptualizing the projects and providing feedback on the project's interaction design, experiment design, and evaluation. Dr. Rosalind Picard provided invaluable insights on how to measure and capture children's engagement during their game play with a social robot. Dr. Maryanne Wolf shared her feedback on how to design the WordQuest game and child-robot interaction in a child-centered way to better foster children's language learning.

The Personal Robots Group is a really supportive lab group where I want to thank Sam Spaulding for his knowledge of social computing, and Safinah Ali for her knowledge in game design.

This research is funded by NSF IIP-1717362 and Huawei.

I sincerely want to thank my family and friends for their kind support, encouragement and empowerment. Last, I feel very grateful to have the noble souls (Dostoevsky, Camus, Tolstoy, Hesse) who enlightened me to see the beauty of life as my companions on this journey.

*“It was pride that changed angels into devils; it is humility that makes men as angels.”
---- St. Augustine*

Table of Contents

Abstract	2
Acknowledgement	6
Chapter 1: Overview	9
Introduction	9
Educational Media & Technology for Children’s Learning	10
Chapter 2: Related Work	15
Social Robots that Help Children Learn	15
Personalized Interaction for Children’s Learning	17
Chapter 3: Thesis Contributions	19
Social Robot’s Adaptive Role Switching	20
Child-Robot Game Play for Learning	22
Personalization & Adaptation	22
Chapter 4: Interactive Learning Framework & Architecture	24
System Design	24
WordQuest Game Platform	25
Child-Robot Interaction for Game Play	29
Robot’s Roles and Behaviors	30
Chapter 5: Pilot Study	33
Procedure	33
Observation & Discussion	34
Chapter 6: Computational Adaptive Role-Switching Model	36
Agent Model: Robot’s Role-Switching Policy	36
Reinforcement Learning Model’s Action Space, State Space, and Reward Function	37
Reinforcement Learning Model’s Value Function	39
Role-Switching Policy Training and Results	41
Chapter 7: Comparative Study	43
Overview	43
Conditions	43
Hypotheses	44
Participants	44
Robot Testbed	45
Methods	46
Procedure	46
Target Vocabulary words	47
Data Collection	48
Learning performance measures	49
Engagement and affect measures	49
Post-game enjoyment and perception measures	50
Chapter 8: Analysis & Results	52
Learning outcomes	52
Children’s Vocabulary Scores by Gender, Age and Prior Vocabulary Level	56
Acquisition of Individual Vocabulary Words	58
Children’s Perception of the WordQuest Interaction	60

Chapter 9: Discussion & Conclusion	63
Contributions	66
Future Work	67
Child-Robot Interaction Design.....	67
Long-term Child-Robot Personalization	68
References	69

Chapter 1: Overview

Introduction

Learning language at a young age is very crucial for children's later educational success (e.g., [1],[2]). It is also crucial for second language learners to start learning a language before the age of 10 in order to achieve proficiency similar to that of a native speaker [4]. However, a major problem faced by many children in the world today in learning a language is a lack of resources and developmentally important stimulation in their homes and schools. This lack of relevant resources and experiences for many young children leads to a detrimental effect on their language and literacy development. For example, preschool-age children raised in families with lower socio-economic status (SES) had significantly smaller vocabularies than children with a higher SES background, and these differences even magnified over time [2]. With technological advances, children's language environments can now be captured and automatically analyzed in a home setting via the Language Environment Analysis (LENA) system [5]. Using the LENA system, recent studies found out that the substantial disparities in SES not only exist in children's vocabularies but also in vocalizations, adult-child interactions and exposure to daily adult words [5]. Another recent study in 2018 found through a story-listening functional MRI task that children's conversational experience is associated with the activation of their left inferior frontal (Broca's area), which is crucial for language-related brain function[3]. The results from the study suggested that conversational turn taking – over and above SES – impacts young children's neural language processing [3]. From these studies and reports, we know that a lack of resources and inadequate early childhood education can negatively impact children's brain development and verbal aptitude.

Indeed, this challenge in children's language learning and development is a global issue. Normile (2017) recently wrote in *Science* that a huge urban-rural gap in educational achievement exists in China. In China, millions of rural children are left behind by their parents who migrate to the booming cities for work. These children grow up with their grandparents who had limited education themselves [6]. Normile (2017) states that these left-behind children have “little cognitive stimulation in the crucial first years of life” and had limited exposure to language and

literacy learning opportunities, due to lack of educational resources in rural areas and their grandparents' incapability of reading to them. As a result, China's rural children are both physically and mentally at an increasing disadvantage, according to Normile (2017). Other regions in the world, such as sub-Saharan Africa, are also facing similar challenges. Around 80.8 million children of ages 3 and 4 in low- and middle-income countries experienced low cognitive and/or socioemotional development in 2010 [7]. It was estimated in 2007 that 61% of children in Sub-Saharan Africa fail to meet their developmental potential [8]. In Mozambique, only 4 percent of children enroll in preschool, the majority of whom are from higher socio-economic status populations [9]. Thus, early childhood learning and development is a pressing global issue.

Educational Media & Technology for Children's Learning

Finding alternative ways to supplement language education for unprivileged populations is particularly crucial. To overcome this global challenge, a variety of media technologies and platforms are used to facilitate children's pre-school learning. These can fall into the categories of passive media (e.g., TV, radio, etc.) or interactive media (e.g., mobile apps, digital learning materials, etc.), or socially interactive media (e.g., remote tutors via video chat platforms, virtual agents, etc.).

Passive Media. Emerging in the early 1990s, a variety of passively viewed "educational" video collections (e.g., *Baby Genius*, *Sesame Street*) and television series (e.g., *Blue's Clues*, *Dora the Explorer*) became prevalent and easily accessible in the market [10][11][12][13]. Utilizing these media broadcast-based technologies had great potential, such as wide accessibility. However, prior research showed mixed results for passive education media's effectiveness in helping children learn. Some studies showed a positive effect of viewing broadcast programs on young children's vocabulary learning [14][15][16]. In one study, 3- and 5-year-old preschoolers who individually viewed animated programs on screen that introduced new words in a story context were able to comprehend some of the target words after two viewings [15]. The findings were replicated in another study with 3-year-olds [16]. However, some other studies showed that young children were not learning effectively from these popular "educational" screen media [17][18].

Although the effect of children viewing the videos or TV series alone was mixed, the presence of a competent and active co-viewer watching the videos or series with children significantly improve children's learning [19][20][21][22]. In one study, when watching *Sesame Street* with an adult who asked the children questions and provided them with feedback, 3- and 5-year-old children were able to better name and identify letters and numbers shown on the programs after the viewing [20]. These results showed that the passive educational media can be helpful in promoting children's learning with the presence of active parental co-viewing, but the educational effectiveness of the media alone is more questionable.

Interactive Digital Media. Interactive educational media emerged in the 2000s with the rising prevalence of computers and later mobile devices. Unlike videos or TV series, educational computer programs and mobile education apps support interactivity. Computers support interactive activities such as clicking and typing, and mobile devices have a variety of interactive features, including touch screen, GPS, QR codes/mobile tagging, and gestural input. For example, touch screen allows children to manipulate the educational contents on the screen by tapping, dragging, rotating and swiping the device. With these built-in features on mobile devices, contingent interactivity is very common in nowadays educational apps. One study done by the Joan Ganz Cooney Center in 2015 analyzed a sample of 183 top apps targeting childhood education in popular app stores [23]. They found out that 92% of the sample apps contained some form of animation, and 45% of the apps had interactive "hotspots," which make noise or animate when touched [23].

The contingent interactivity of educational apps that broadcast media lacks enables children to take a more active role in learning. For example, children can shape and change the progress of a digital game or a program based on how they respond to the game or program. They can also see the immediate feedback or effect of their responses, actions or answers. Prior research suggested that the interactive media has greater potential in promoting children's learning compared to passively watching videos [24][25]. In addition, this responsiveness of interactive media can also provide children a feeling of accomplishment, which they don't have when passively watching videos [26]. Furthermore, some well-designed educational apps that align with the science of how young children learn were found effective in children's language learning [27][28]. For

example, two vocabulary-focused apps designed for children ages 3 to 7 (*Martha Speaks* and *Super Why*) were found to increase young low-income children's vocabulary by up to 27% in a 2-week period [28]. Thus, the interactive media seems to hold promise for children's learning.

As mobile devices become more ubiquitous over time, interactive digital media's potential impact on the global literacy challenge is further magnified, even in low SES families and in developing countries. It is now possible to deploy inexpensive android-based devices (e.g., smartphones) at scale to different regions and areas in the world. One Laptop per Child [29] pioneered the deployment of inexpensive, robust hardware devices to developing countries. However, the devices lacked a breadth and depth of educational content. Later initiatives, such as the Global Literacy Project, leveraged the growing ubiquity of mobile devices and the ability to distribute a large repertoire of digital education media (e.g., educational apps, digitized books, educational videos) in these underserved areas to focus on the content and understanding how children can learn collectively [30][31][32][33][34]. For example, the Global Literacy Project highlighted the importance of child-driven learning process where children collaboratively play the educational apps as a group [32]. These core insights from child-driven learning were first articulated in the Hole-In-The-Wall experiment [35]. However, whereas the Hole-In-The-Wall studies focused on emergent computer literacy among children, the Global Literacy Initiative focused on early literacy skills. In principle, Internet connectivity to cloud-backend content education service can support the rapid deployment of a very large amount of educational materials far more cost effectively than the distribution of physical books, as well as supporting frequent updates to that content and tracking of student engagement to personalize learning.

Despite the substantial promise of interactive media technology in making a large-scale impact, it is still important to know how children use the technology. Without properly being used, it may still fail to delivery effective learning to children. A 10-year-long observational study by Neuman and Celano (2012) found out that children from a low socio-economic and an affluent neighborhoods used digital educational resources (computers and literacy software) in very different ways, even though the two groups in the study had adequate amount of the resources and invested similar effort in using them for learning [36][37]. According to the study, parents from the higher socio-economic neighborhood used the computer games as a tool to facilitate

lessons on literacy learning, whereas parents with a lower socioeconomic status let their children use the computers without interacting with them or giving them any feedback even when their children were struggling. From this study, Neuman and Celano (2012) indicated that the digital divide also existed in the intangible resources – the active participation of parents or adults in children’s literacy learning process [37]. Though the access gap in material resources may narrow enabled by the ubiquitous use of mobile devices, this “participation gap” affects how these educational media are used and how well children can learn using the media. Thus, to effectively foster children’s learning, social interaction seems to be a crucial component besides merely educational resource.

Socially Interactive Media. In past two decades, a variety of socially interactive media emerged to facilitate children’s learning. One type of socially interactive media is video-chat platforms such as Skype and FaceTime [38][39]. Video chat technology has both features of screen media and features of in-person communication. One feature that this communication tool and live interaction both have is social interaction. Children can directly engage with other children or adults on video chat. This platform enables children to communicate with others who don’t unnecessarily share physical space with them. It has been shown effective in helping young children learn what the on-screen person teaches in prior studies [40][41]. For example, Roseberry, Hirsh-Pasek and Golinkoff (2014) showed that the video chat technology supports children’s vocabulary learning due to the presence of its social contingency with an on-screen person [40]. According to the study, children only learned words in socially contingent interactions (in-person communication and video chat) but not in non-contingent prerecorded video training condition [40]. Given this great potential, some applications have been developed based on this in-person live video [42][43]. For instance, the app *Kindoma* allows shared reading between two people via video chat [42]. Another video chat platform, *VIP KID*, that allows American teachers to teach children in other countries English [43]. Both *Kindoma* and *VIP KID* provide video chat platforms where two partners (often a child and an adult) can see learning materials, each other’s face and each other’s touch interactions on the screen [42] [43]. Despite great efficacy, language learning through live video may be costly and inaccessible to children from disadvantaged neighborhoods, as it requires the presence of another person on the platform during learning.

Unlike video chat platforms, virtual agents and physical agents are two types of low-cost quasi-socially interactive media for learning. In prior research, virtual agents participated in a variety of learning activities with children together as a peer [44][45]. Another type of quasi-socially interactive media is physical agents such as social robots. Unlike virtual agents, social robots have a physical embodiment. They have been extensively used in early childhood education in the past decade in different contexts (e.g., storytelling, second language learning, literacy learning), and were shown as an effective tool for children's learning [46][47][48]. More details on how social robots help children learn is presented in Chapter 2.

Chapter 2: Related Work

For language learning, many educational technologies do not provide social context and interaction that is crucial for young children to effectively learn a language. According to Kuhl (2007, 2011), lacking social interaction with a partner may even impair young children's language learning [49][50]. Thus, the key questions are how to design educational technologies by drawing insights from the empirical research on how young children learn and grow, how to incorporate social components into educational technologies, and what media technologies have potential to best promote children's language learning.

Among all the platforms, socially interactive technology seems most promising, as it is able to integrate social components into young children's learning. Social robotics is a great representative of socially interactive technology that is both effective and resource-efficient. Prior research showed that social robotics is a more effective platform in improving people's learning, sustaining their engagement or leading to greater behavior change, when compared with virtual agents or non-socially-interactive devices (e.g., tablets) [51][52][53][54][55]. Although social robotics is still in its infancy, it is developing rapidly in the past decade and will soon become more easily accessible than ever before. For example, the first commercialized social robot for the home, Jibo, was just released in 2017 [56], and makes long-term deployment of this socially interactive technology for language and literacy learning in under-resourced areas possible in near future.

Social Robots that Help Children Learn

Social robots can help young children develop and learn in different educational contexts (e.g., storytelling, coding, math) [46][47][48][57][58][59]. Due to its increasing accessibility and recent technological improvements, they have also been used for educational purposes more extensively than ever before. When compared with passive media technologies (e.g., on-screen media), socially interactive robots have demonstrated their better effectiveness in teaching students language skills in prior work [60][61]. Kanero *et. al* (2018) provided a thorough review of past research on using social robots for young children's language learning, and suggested that robots can more effectively teach children than other digital devices [62]. In addition, social

robotic technology provides opportunities for children to learn through peer modeling. Prior studies have shown that children mimic a peer-like social robot's behaviors and mindsets (e.g., robot's growth mindset and curiosity) [63][64]. Peer modeling, indeed, can be beneficial to the development of social interaction skills in young children, especially those with special needs. For example, peer-mediated interventions increase the social interaction skills in children with autism [65], and children with disabilities improve their social behavior skills through imitating typical peers [66]. Thus, social robotic technology provides social components essential for learning to happen and is a very promising tool for young children to learn and develop language and social skills.

Robot-as-Tutor. Social robots have also been designed to interact with children as either tutors or instructors [67][68][59]. Prior research on early childhood language learning showed that children learn more effectively when a competent person uses scaffolding strategies to teach a child as well as provide clear guidance, explanations and feedback to the child [69]. Thus, situating a social robot as a tutor holds promise in effectively teaching children's language.

Robot-as-Tutee. Children can also learn from teaching a robot how to learn. In different educational contexts, robots were situated as a supportive peer [46][70][71] or younger students whom the children need to help or teach [72][73]. This learning-through-teaching paradigm also has several benefits. First, a robot can learn in reverse about a child's ability level from children's demonstrations to the robot, and the information obtained from children's teaching can be used to build computational models to personalize learning contents matched to the child's knowledge level. Second, situating a robot as a younger learner also creates more fun and makes more acceptable for robot's technical limitations and mistakes [74].

Peer-based Interaction. Social, robots have been designed to engage children as a supportive peer [46][70][71], or even as a less knowledgeable peer whom a child can help teach [72][73]. This provides opportunities for children to learn through peer-peer inspired interaction and peer modeling. Prior studies have shown that children mimic a peer-like social robot's behaviors and mindsets (e.g., robot's growth mindset and curiosity) [63][64]. Peer modeling, indeed, can be beneficial to the development of social interaction skills in young children. For example, peer-

mediated interventions increase the social interaction skills in children with autism [65], and children with disabilities improve their social behavior skills through imitating typical peers [66]. Thus, social robotic technology provides social components essential for learning to happen and is a very promising tool for young children to learn and develop language and social skills.

Personalized Interaction for Children's Learning

Young children learn a language at different paces in different ways, which motivates the case for personalized learning experiences that support the child as needed. For example, prior studies showed that children as young as 20 months differed in terms of the style with which children approach the language learning task [75]. Another study showed that differences in learning styles among children still exist as late as 14 years old [76]. Thus, educational technologies need to personalize their interaction with individual children based on the children's learning styles and knowledge states with the attempt of promoting the best learning.

Compared with educational systems lacking personalization, some interactive systems that were able to adapt to individual students resulted in students' greater engagement and/or greater learning outcomes [77][78][79][80]. Similarly, personalizing child-robot interaction for language learning has great potential for improving the child's overall learning performance and engagement. For example, a robot tutor that provides personalized affective feedback to a child increases children's valence for second language learning over a duration of two months, whereas a non-personalizing robot does not bring such change to children [71]. Another study shows that personalizing a robot's actions to an individual student's strengths and weakness accelerates students' solving speed of grid-based logic puzzles and yields other significant benefits in educational human-technology interactions [81]. Other researchers investigated how a robot tutor could deliver personalized learning tasks and improve students' learning effectiveness by computationally modeling the child's knowledge states and acquisition level [82][83].

To create a personalized learning experience for students, a variety of computational models have been used. Some researchers model student's knowledge states using Bayesian Knowledge Tracing (BKT) [83][84], while others design action-based policy to personalize sequences of learning activities and teaching strategies through Partial Observable Markov Decision Processes

(POMDPs) [85], Multi-Armed Bandit [86] and Reinforcement Learning [87][88][89]. Empirical data showed that these above personalization models can enable the tutoring and teaching systems to help students achieve significantly better learning performance.

Chapter 3: Thesis Contributions

This thesis explores how to design socially interactive technology that can create social, supportive, playful and engaging contexts for young children's growth. While this thesis focuses primarily on literacy learning, this kind of child-technology interaction framework may also support other aspects of children's learning (e.g., attitude learning, math). In the thesis, a social robot is introduced as a learning peer and used to play a tablet-based collaborative game called WordQuest with a child. The thesis will talk about how to design this educational game play between a robot and a child to foster children's learning and engagement. Then, the thesis will introduce one innovative way to create peer-to-peer learning between the robot and the child, which is to computationally enable the robot to adaptively switch its role between a tutor role and a tutee role when playing the WordQuest game with the child.

I propose a novel educational technology system that seamlessly integrates the three fields: *social robotics*, *educational game design*, and *personalized learning*. More specifically, my thesis investigates:

the design and evaluation of a fully autonomous social robot that can adaptively switch its collaborative role and personalize its teaching strategies when playing a tablet-based vocabulary learning game with a child to improve vocabulary learning and engagement. I will compare this novel system against a robot that takes expert role only and one that takes a novice role only.

In this context, I will explore the following research questions:

First, I will explore **what collaborative roles a social robot can exhibit during its language game play with a child, and analyze the value of each for supporting children's learning.** I hypothesize that the two social roles (*expert* and *novice*) have differential values in children's learning. For example, a novice robot who asks a child to help it solve a task or correct its mistake may evoke the child's greater engagement and curiosity in learning, and give children more opportunities to practice. Conversely, an expert robot who demonstrates its knowledge to a child and provides hints to the child may accelerate the child's learning speed, deepen the child's

understanding of the knowledge, and reduce children's frustration evoked by learning difficulties during game play.

Second, I will explore **how to computationally implement a robot's role-switching policy** that intelligently utilize data on child's current learning states to figure out the most appropriate role to exhibit for next child-robot turn. To evaluate the effectiveness of this policy, I will compare how the performance of children who interact with this *adaptive role switching* robot with the performance of children who interact with only an *expert role* robot, and those who interact with only a *novice role* robot. This comparison in performance across the three roles (novice, expert, adaptive) will shed lights on how to design an effective personalized role-switching policy that can adapt to individual children in a future long-term interaction with the robot.

Social Robot's Adaptive Role Switching

Prior studies indicate differential advantages and benefits of robot's roles (*tutor*, *tutee*) for young children's learning (*chapter 2*). In work done so far, the robot's role was always fixed (either a *tutor* role or a *tutee* role) in this context, and social robots didn't switch roles when engaging in educational game play with a child. Given the differential advantages of both roles, Vogt *et al.* (2017) suggested it a desirable way to frame the robot as a peer yet act as more like a tutor that can scaffold children's learning based on "pedagogically well-established strategies" [90]. However, Vogt *et al.* (2017) did not address the possibility of a robot switching roles within learning interactions. Thus, the adaptive role-switching of social robots is still under-explored yet has distinct potential for promoting children's learning from the perspective of *the Science of Learning*. Children learn and reach learning goals more effectively in a flexible context that scaffolds their exploration, supports their questioning, and guides their discovery [91]. Since the robot's *tutor* role is able to provide scaffolding and guidance and the robot's *tutee* role provides children with opportunities to explore and question, we hypothesize that the adaptive role-switching would be superior to either of the fixed *tutor* and fixed *tutee* roles alone with respect to facilitating children's learning.

In addition, this adaptive role-switching may also create bidirectional, peer-to-peer learning opportunities, which would be fruitful for children's learning from the lens of *cooperative peer*

learning. In a cooperative learning setting, peers work together to accomplish shared goals. More specifically, Tharpe and Gallimore (1988) state that peer models facilitate children's assisted performance [92]. During interactions between the peers, the children are motivated to actively participate and guide each other in the direction of a shared endeavor, and collaborate in bi-directional nature [92]. Thus, enabling a social robot to switch its roles and interact with children as a collaborative learning peer may effectively promote children's learning.

Despite the great potential of social robot's adaptive role-switching, no studies examined this role-switching framework in the context of early childhood education, and only one study so far explored the use of robot's role-switching in the field of child-robot interaction [93].

Specifically, this study explored how a robot adaptively switched its role between leader and follower in order to intrinsically motivate children to learn creative dance [93]. In this study, the child first took a follower role and later switched to a leader role. The results showed that this role-switching system is an effective tool to motivate students to stay engaged in the activity [93].

Given that robot's role-switching for children's learning holds distinct promise yet is still under-explored, my thesis work is to enable a social robot to adaptively switch its roles (*expert tutor*, *novice tutee*) and adapt to the child's needs when playing a tablet-based vocabulary learning game with a child. For example, the robot could be an *expert tutor* when the child is struggling with understanding targeted learning contents. The *expert tutor* role would enable children to learn from the robot's demonstrations and rationale. In the same interaction, the robot could also switch to a *novice tutee* role when such a role would better facilitate the child's learning given the child's current knowledge and affective states. Being a learner, the robot could ask the child to help it solve a challenge, or intentionally make mistakes for the child to correct them. This *novice tutee* role allows children to learn by teaching a robot how to learn. Inspired by both the *learning by teaching* approach and the *learning from demonstration* approach, my research explores how to design the social robot's role switching policy and child-robot interaction framework in the context of language learning and how different ways of learnings (e.g., *adaptive role switching*, *learning by teaching*, *learning from demonstration*) affect children's learning and engagement.

Child-Robot Game Play for Learning

In prior studies, it is a commonly used interaction paradigm that a child and a robot play a game on a tablet [83][64][71][94]. However, they didn't explicitly show what principles are employed to guide the design of the tablet-based games the robot and child play with. The focus is more on the interaction design between the robot and the child rather than evaluating the educational value of the game itself. However, the design of tablet-based educational games can significantly affect children's learning and engagement as mentioned [27]. Thus, my work will explicitly incorporate principles for educational apps and for game design into the tablet-game the child and robot will play with, and evaluate the educational value of the game using existing frameworks in educational apps and game design. A rigorous design of the game activity would help promote the best learning of children, along with a robot learning companion.

Personalization & Adaptation

In work done so far, no studies have constructed a computational role-switching model in the context of early childhood learning. My work would be the first one to explore how to build a computational model enabling a robot to switch its role between *expert* and *novice* based on the child's learning states. Getting the inspiration from both affective computing and personalized teaching strategies for the intelligent tutoring system, I build a reinforcement learning model that dictates a robot's role-switching strategy. Reinforcement Learning is chosen mainly for its support for personalization. More specifically, it is trained across multiple students to obtain a baseline model and becomes adaptive and personalized to an individual student given the rewards from the student.

Since the comparative study in this thesis consists of only two 25-min learning sessions, the personalization effect of the role-switching policy may not be significant across individual children in the study. However, the data collected from it will be analyzed to understand on how to build more effective and engagement-promoting personalization/adaptation mechanisms for future long-term studies (e.g., 6-month studies with multiple repeated encounters), and what attributes or features about children may play critical roles in such mechanisms. Thus, this thesis

hopes to provide insights and suggestion on how to design long-term personalized role-switching policy for children's learning.

In sum, my contributions are as follows:

- New educational game app based on science of learning.
- New computational framework where to enable a socially interactive robot and a child to foster children's learning and engagement.
- New peer-to-peer learning paradigm between the robot and the child, which is to computationally enable the robot to adaptively switch its role between an *expert* role and a *novice* role when playing the game with the child.
- New comparative study showing differential impact of the robot's three collaborative roles (*expert, novice, adaptive role switching*) on children's learning and engagement.

Chapter 4: Interactive Learning Framework & Architecture

System Design

In this section I give an overview of my WordQuest integrated system for collecting multimodal data, analyzing children's learning state and modeling robot's role-switching behaviors during interactive game play. The child-robot interaction paradigm is displayed in Figure 1, and the full system architecture is outlined in Figure 2. The system consists of 5 major modules: Interaction Controller, WordQuest Game Platform, Agent Model, Student State and Data Collection. The Interaction Controller functions as a hub for the entire system, connecting Student State, Data Collection, and Agent Model modules together. When a child is playing the game with a robot, the child's learning performance and engagement level will be recorded via Student State Module and sent to the hub module Interaction Controller. Based on the data about student state and game state received from Interaction Controller, the Agent Model will generate the most appropriate role and corresponding behaviors for the next child-robot turn. Furthermore, the WordQuest Game Platform, controlled by the Interaction Controller, is the WordQuest app that both the child and robot directly interact with. Each of the modules will be explained in later subsections.

Last, this integrated system is very flexible and can be easily adapted to other child-robot interaction paradigms. For example, the Student State, Agent and WordQuest Game Platform modules are easily reconfigurable to support a variety of learning tasks and educational games. Sensors inside Data Collection modules can be easily added or removed.



Figure 1. Child-robot interaction paradigm.

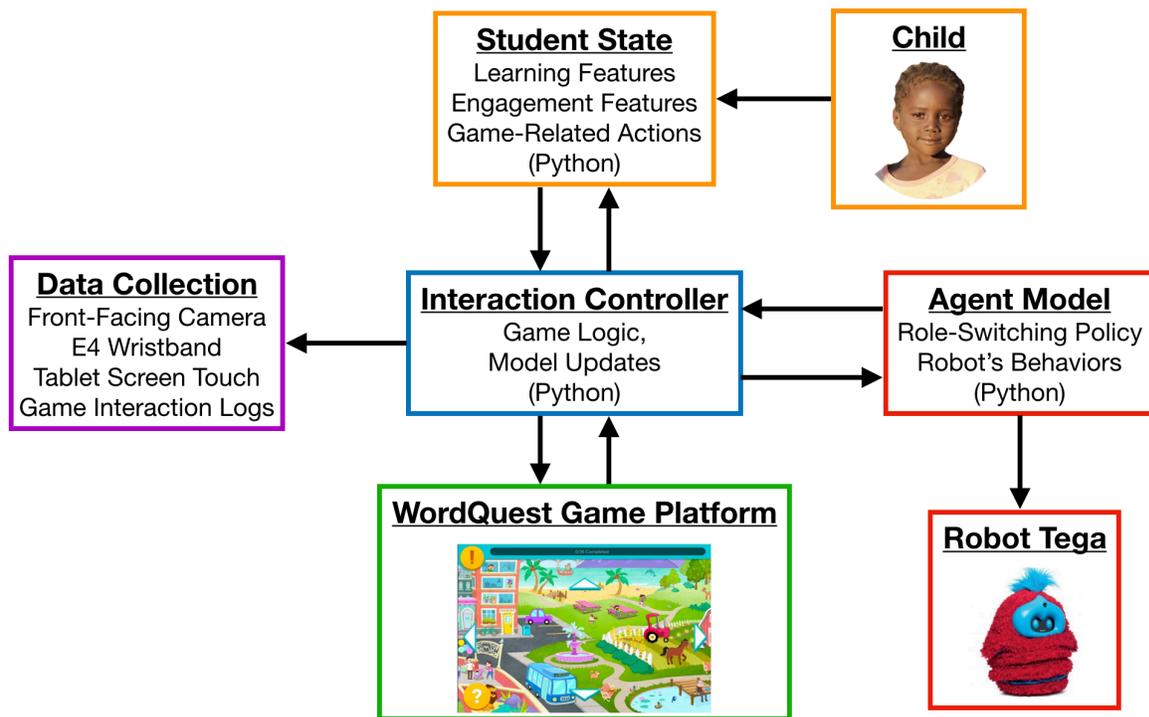


Figure 2. Diagram of interactive system for child-robot peer-learning.

WordQuest Game Platform

The WordQuest platform is an integrated system that enables a fully autonomous social robot to behave as a learning companion and provide social contexts for a child to learn English

vocabulary words. The child and robot work together to play the WordQuest game on a tablet. During the game play, the robot adaptively and decides what social roles (e.g., *expert*, *novice*) and actions it should take for next turn, based on the child’s learning state, so as to promote the child’s vocabulary learning and engagement.

The game is implemented in the Unity 3D game engine, which supports flexible updates and changes of learning curriculums. The game has a big background scene (e.g., outdoor scene), which contains a variety of 2D animated, clickable objects. Currently, the game has two scenes (outdoor and indoor), which are shown in Figure 3. A user can pan around the scene, zoom in and out, and click interactive objects. In total the game has around 50 clickable objects, which are mapped to the 50 basic English words preschoolers need to know or learn. In addition, the game has two game modes: Explore Mode and Mission Mode, as shown in Figure 3. In the Explore Mode, the user can click an object in the scene, and then hear its pronunciation as well as see its spelling in the game. In the Mission Mode, the user will receive a variety of game missions. Each game mission contains one advanced vocabulary word that the child needs to learn. For example, “gigantic” is the key vocabulary word in the mission “can you spy objects that are gigantic”. The child learns the meaning of the keyword in a game mission by searching objects and finding out whether those objects are correct ones for the mission. To find an object, the child needs to pronounce the English word representing that object. Their speech can be recorded and sent to SpeechAce for phoneme-level accuracy analysis [95]. Since a child and a robot play this tablet-based game together, the tablet app provides a shared context between the child and robot [46][70].



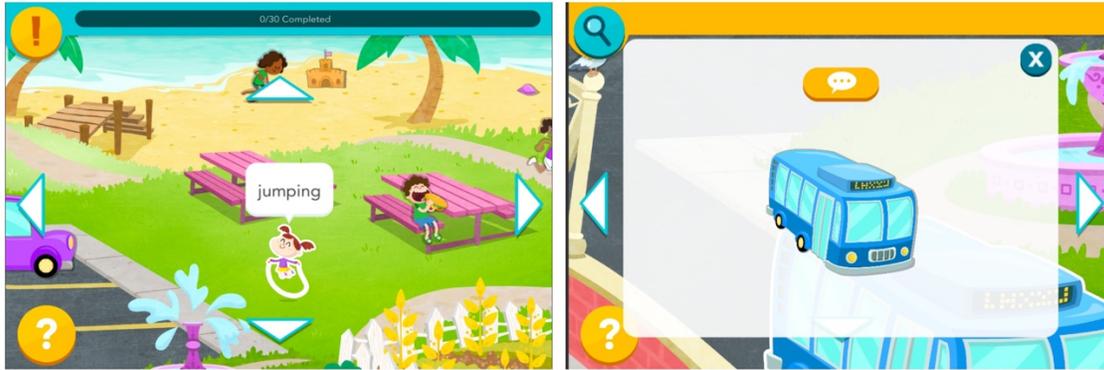


Figure 3. WordQuest game's interface and its game scenes

Today, research has provided a variety of guidelines and principles for designers who develop educational games targeting young child [96][27]. To design educational apps that promote the best learning, Hirsh-Pasek *et. al* (2015) suggested to use the *Science of Learning* as a guide for educational principles and proposed four evidence-based design principles for educational apps, which are “active,” “engaged,” “meaningful,” and “socially interactive” in the service of a learning goal [27]. The WordQuest game is designed under these four design principles for educational apps from the field of *Learning Science* [27].

- Active:** The first principle in Hirsh-Pasek *et. al* (2015) is “active learning”, which enables children to actively build their own knowledge rather than simply observing others and waiting for others to teach them. The WordQuest game has a variety of features that enable children to play an active role in learning vocabulary words. Instead of mechanically memorizing the meaning of a vocabulary word, the children can explore the screen space, make hypotheses about the word’s meaning and test them out, and then improve their hypotheses. For example, when learning the word “gigantic” in the mission “can you spy gigantic objects,” children can tap and pronounce different objects, have immediate feedback from the game on each attempt’s correctness, and discover that “gigantic” means “really big” by trial and error. This “trial and error” way of learning in the WordQuest game enables children to be minds-on and keep wondering about the meaning of a vocabulary word in a word learning task.

- **Engaged:** The second principle in Hirsh-Pasek *et.al* (2015) is “*engagement in the learning process*”, which can be achieved through *contingent interactions, extrinsic motivation and feedback, and intrinsic motivation*. The WordQuest game provides the above three engagement enhancers. Specifically, the game is responsive to every touch or swipe of the child, and the robot also performs contingent actions in response to the children’s learning-related actions. When the child spies an object, the game will display different animations to indicate whether the child’s attempt is correct or incorrect on the screen. The robot will also send contingent motivational messages (e.g., “Good job!” and “That is okay”) as a way to provide extrinsic motivation and feedback to children. When one mission is completed, the child and robot will also get one point as a reward. This scoring system is instrumental in tracking and visualizing the child’s learning progress.
- **Meaningful:** According to Hirsh-Pasek *et. al* (2015), the third principle is “*meaningful learning*,” which happens when children are “learning with a purpose, learning new material that is personally relevant, [or] linking new learning to preexisting knowledge.” The WordQuest provides game scenarios that children are very familiar with and can easily connect with their daily life. Specifically, the two scenes (outdoor and indoor scenes) contain all the objects children can see in their local communities or households (e.g., painting, lamp). Thus, when they learn vocabulary words embedded in game missions, they can be more easily reminded of the words. In addition, the vocabulary words they are learning are not conceptually foreign to children ages 4 to 7. For example, one vocabulary word is “*delighted*”, which is an advanced synonym of “*happy*.” In this sense, children are able to learn words that are relevant to their daily life and connect with the new words with their preexisting knowledge.
- **Socially interactive:** The last principle Hirsh-Pasek *et. al* (2015) proposed is “*social interaction*.” While our tablet-based WordQuest game does not alone provide social interactions like most of educational games, the presence of a social robot in the learning interaction make it possible for children to learn in a social context. Specifically, the robot serves as a learning companion for children, who can give both effective and affective feedback and personalize its responsivity to each child. For example, the robot

gives contingent praise and encouragement to a child when the child spies an object, a type of responsivity that has shown to facilitate language development in young children [97][40]. Thus, the child-robot interaction complements the limitation that the tablet-based WordQuest app has and promote children’s learning. The presence of a social robot is not to replace any interaction between a child and their caregivers but to make a standalone tablet-based educational app more socially interactive.

Child-Robot Interaction for Game Play

The child and robot collaborate together on solving game missions in WordQuest. They sit facing each other with a tablet in the middle for both of them to interact with, as shown in Figure 4. When the game starts, both the child and robot receive a game mission from the tablet (e.g., “*Can you spy vehicles in the scene?*”). Then the robot observes the child’s learning states and performance, and decides what roles to play and how to react. Each mission contains one keyword, of which meaning the child needs to learn. For example, the keyword is “*vehicle*” in the mission “*Can you spy vehicles in the scene?*” Then, the child needs to learn that “*vehicle*” refers to “*something that a person can ride in or steer*” after seeing objects “*bus*” and “*car*” are correct answers but object “*dog*” is not. Each mission may take 5-8 minutes, as the child and robot need to correctly find four target objects to complete the mission. When one mission is completed, another mission will pop up until all game missions are completed. One learning session consists of around 5 to 6 missions containing 5 to 6 keywords for the child to learn. During the game play, the robot provides help for the curricular content, affective support and feedback, and make the child’s learning experience fun and engaging.



Figure 4. Child-robot interaction paradigm



Figure 5. An example that shows how the child and the robot take turns finding objects that belong to the target vocabulary concept.

When working on a game mission, the child and robot take turns spying and retrieving objects. Each player has one chance of “spying” objects. If one player spies and pronounces an object, that player’s turn is finished, and it becomes the other player’s turn. An example of how the child and robot take turns spying objects and complete a game mission is displayed in Figure 5 above. During the child’s turn, the robot could choose to intervene and help if the child is struggling with finding a correct target object. During the robot’s turn, the robot may also ask the child for help and give the child a chance to spy an object for the robot. Last, the robot will always show happiness when a child finds a target object correctly, and show sympathy and express a growth mindset when the child fails, regardless of the social role the robot is playing.

Robot’s Roles and Behaviors

The robot can perform two different roles (*expert*, *novice*) during game play with the goal of maximizing the child’s learning and engagement. For example, the robot can become an *expert* and proactively demonstrate how to find a correct word and explain to the child the meaning of the target vocabulary word. On the other side, the robot may also play a *novice* role by showing intellectual curiosity, asking the child for help or spying incorrect objects. The robot’s specific behaviors for each role are displayed in Table 1 and Table 2. As shown from the two tables, the

expert-conditioned behaviors are more like a robot teacher giving instructions and informative feedback to the child during learning, and the novice-conditioned behaviors situate the robot more as a learning peer who is intellectually curious and persistent.

It is important to notice that all behaviors that the robot will exhibit when playing the WordQuest game with the child are learning-related and contingent on the child’s learning and engagement. The robot does not perform any behaviors that are not related to children’s learning, as prior research showed that extraneous information unrelated to learning may distract children from the learning task and disrupt their learning performance [98][99][100].

Robot’s Role	Behavior	Definition
Expert	Vocabulary	Explain to the child why the object the robot chooses is correct
	Explanation Speech	
	Keyword Definition Speech	Explain the meaning of the vocabulary word to the child
Novice	Asking for Help Speech	Asks whether the child can help the robot spy an object during the robot’s turn
	Asking for Explanation Speech	After the robot spies an incorrect object, the robot asks the child why the chosen object is not correct

Table 1. Robot’s roles and corresponding behaviors during robot’s turn

Robot’s Role	Behavior	Definition
Expert	Offering Help Speech	Offer to help the child spy a correct object
	Keyword Definition Speech	Explain the meaning of the vocabulary word to the child
	Hint Speech	Gives hints on the meaning of the vocabulary word to the child
Novice	Asking for Child’s Thoughts Process Speech	After the child spies a correct object, the robot asks why the child has chosen this object

Curiosity-driven speech	The robot shows curiosity in what the child is going to find
----------------------------	--

Table 2. Robot's roles and corresponding behaviors during child's turn.

Chapter 5: Pilot Study

I ran a pilot study with children ages 4-7 to find out the following:

- Is the game easy for children to learn to play?
- Are children able to discriminate clickable objects against non-clickable background images in the game?
- Are robot's behaviors ambiguous or confusing to children?
- How many vocabulary words children were able to learn before reaching the limit of their cognitive load?
- How long should each game session last?
- In what circumstances is children's learning most likely to occur?

In addition, the data collected from the pilot study were also used to guide the design of the robot's adaptive role-switching policy and train the policy model. The child-robot interaction environment for the pilot study is displayed in Figure 6 below.



Figure 6. A child is playing the WordQuest game with Tega

Procedure

A total of 25 children participated in the pilot study. Five children were not eligible for the study due to either their low age or low English proficiency. Sixteen children played the game with the Tega robot in the Personal Robots Group's study room, and 4 children played the game in a local

elementary school (JFK Elementary School). The sixteen children played the game for about 45 minutes, and the 4 children played the game for about 25 minutes. During the pilot study, all children interacted with a fully autonomous robot, and the robot has a random role-switching policy. In other words, the robot has a probability of 0.5 to switch its role between *expert* and *novice* every child-robot turn. After the pilot study, I collected a training dataset containing 143 episodes, in which a child and a robot finished one mission together. This dataset was then used to train the computational model for robot's adaptive role-switching policy.

Observation & Discussion

We found out that a few 3-year-olds who signed up for the study had difficulty understanding the game mechanics, but participants from ages 4 to 7 knew how to play the game after a 5-min practice round in which the experimenter demonstrated how to play. In addition, the first two children had some trouble discriminating clickable objects against non-clickable background images. In response to this confusion, I made the background images more faded out in the game. The rest of the pilot used this new version of the game. Regarding the robot's behaviors, I found out children were not confused about what the robot said and did in the game. A few children tried to ask the robot questions but the robot did not have a feature of listening and answering children's questions. In future studies, I would consider adding this feature to the robot. The initial pilot study was just one-shot encounter in which each child needed to learn 11 words. However, I found out that most children became disengaged or tired after 30 minutes of game play. Thus, I shortened the learning session to 25 minutes, and each session contained only 5-6 words.

Through observation, a few insights on how children could better learn a vocabulary word are generated. First, children need to wonder about a vocabulary word's meaning and try to find different objects before being informed of its meaning in order to deepen their understanding of the word. Second, they also need to find a few correct objects to reinforce their learning after they understand a vocabulary word's meaning. Third, it is not necessary for students to get more practice opportunities if they already know a vocabulary word's meaning prior to the mission, as increasing a learning session's overall time may likely decrease the child's attention and

engagement. These observation-based insights were later used to construct the robot's role-switching policy in *chapter 6*.

Chapter 6: Computational Adaptive Role-Switching Model

Agent Model: Robot's Role-Switching Policy

Reinforcement Learning (RL) originated from psychological studies of animal behavior. Animals learn complex behavior by learning to get rewards and avoid punishment. Thus, from the perspectives of computing, the goal of RL is to optimize an artificial system's behavior to rewards and punishments. A learning agent is situated in an environment and learns about the environment by executing actions and observing how these actions differentially change the state of the environment.

An RL model consists of three elements: (1) a state space that describes the states of the environment the agent is in, (2) an action space that contains a set of actions that the agent can perform to change the environment, and (3) a reward function that assigns numerical rewards and punishments. RL is traditionally defined as part of a *Markov decision processes* (MDP) [101]. An MDP is a tuple (S, A, P, R) such as S is a state space, A is an action space, and P and R are the distribution of probabilities and rewards respectively. Value-function RL methods can be in general categorized into two types: model-free methods (e.g., Q-Learning) and model-based methods (e.g., R-MAX) [102][103]. Model-based methods have advantages of sample efficiency, but learning a model of the domain, which is P and R functions, is required before training the RL model. Thus, the key to build an effective model-based RL methods and find an optimal policy is to construct an accurate model of the domain. Conversely, model-free methods learn the optimal policy without requiring the dynamics of the environment but need more experiences at the same time.

In our case, we are using RL methods to dictate how a social robot should behave when playing the WordQuest game with a child with the goal of maximizing children's learning performance. Since the robot interacts with children in the real-world environment, it is nearly infeasible to construct an accurate model of the environment without bringing model bias and compromising the model accuracy. Thus, model-free methods are favored over model-based methods despite its

requirement on a larger set of training data. To construct an effective model-free RL model, the effectiveness of state presentation S and reward function R is crucial. Thus, the insights derived from the pilot study were used to construct an effective S and R .

Reinforcement Learning Model's Action Space, State Space, and Reward Function

Given the above human expert knowledge, the RL model for robot's role-switching policy is constructed in the following way as illustrated in Figure 7. First, the action space consists of just two actions: $\{expert\ role, novice\ role\}$. Then, the state representation S consists of 3 dimensions: $\{current\ number\ of\ child-robot\ turns, current\ number\ of\ expert\ role\ occurrences, current\ number\ of\ correct\ objects\ spied\ by\ child\}$. The three state features are all discrete ranging from 0 to positive infinity. Since the goal of the model is to improve children's vocabulary acquisition, the listed three state features indicate the learning phase a child is currently in during the game play. For example, a state $\{5,1,0\}$ indicates that the child has probably already explored the game scene but still hasn't learned the vocabulary word yet, so an appropriate behavior for the robot to display at next turn is to an expert role that demonstrates to the child the meaning of the word.

To fully model a child's learning process, I use the observation-based insights above to derive the following reward function displayed in Table 3. The reward function in the robot's role-switching model. For the expert action, the total reward score consists of 2 components. First, r_1 is designed to ensure that the expert action becomes increasingly rewarding if the child makes a correct attempt this time but incorrect attempts in previous consecutive turns. Second, r_2 will be larger if the child already knows the meaning of the vocabulary word very well and makes few mistakes, which is reflected in the rate of child's correct attempts. If the rate is higher, the child is less likely to need extra practice and the robot's expert action will be more likely to perform to shorten the interaction time for this vocabulary word already acquired by the child. For the novice action, the total reward score consists of 3 components. First, if the child chooses to help the robot out, the child is very likely to be engaged and likes robot's novice behaviors during the turn. Thus, the novice action should be rewarded with r_3 . The reward r_4 is constructed to reward the novice action when the child just learns the vocabulary word's meaning and thus needs more practice and spies more objects. Third, a novice action will be more rewarded in the first few

game turns because the child is encouraged to explore the word’s meaning with the robot before the robot demonstrates to the child what correct objects are.

The total reward score is calculated:

$$r_{total} = r_1 + r_2 \text{ if RL action is EXPERT}$$

$$r_{total} = r_3 + r_4 + r_5 \text{ if RL action is NOVICE}$$

RL Action	Reward Condition	Reward Values
Expert Tutor	If child’s correct attempt	$r_1 = \text{the number of child's consecutive incorrect attempts}$
	If child’s correct attempt and the rate of child’s correct attempts is larger than 50%	$r_2 = 10 * (\text{the rate of child's correct attempts} - 0.5)$
Novice Tutee	Child’s help	$r_3 = 1$
	Child’s correct attempt	$r_4 = 4 - \text{the number of child's correct attempts}$
	Child’s incorrect attempt	$r_5 = 4 - \text{the number of child's total attempts}$

Table 3. The reward function in the robot’s role-switching model.

The RL model is applied to dictate the robot’s role-switching behavior as shown in Figure 7. As shown in Figure 7, the robot receives a role from the RL role-switching policy in the beginning of a child-robot turn. Then, the reward for this generated robot’s role (*RL action*) is calculated during the child’s turn in this child-robot turn. At the end of the child-robot turn, child’s state S is updated and used by the role-switching policy to generate a new role for the robot.

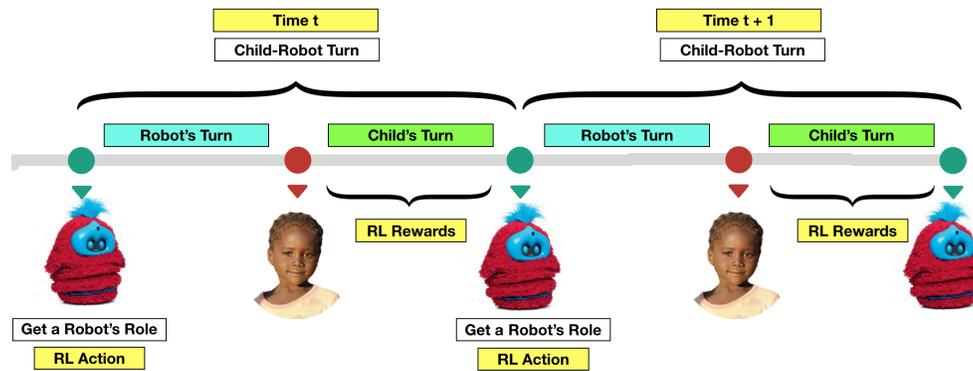


Figure 7. The reinforcement learning model is applied to dictate the robot's behaviors

Reinforcement Learning Model's Value Function

The value functions in reinforcement learning (RL) methods are commonly categorized into tabular solution methods and approximate solution methods [102]. For tabular solution methods, values of state-action pairs are represented as arrays or tables, and thus the methods often find the exact solutions. Conversely, approximate solution methods do not store individual state-action pairs; instead, they approximate the value for each state-action pair and are commonly applied to problems with large state spaces with limited computational resources.

In our case, the state space consists of 3 discrete state features, which can have up to 10, 4, 4 different values, respectively; the action space is 2. As a result, the model can have up to 160 discrete states and 320 state-action pairs. This number of state-action pairs make it implausible to use tabular solution methods to store each of 320 pairs and update the value function given that the amount of real-world child-robot game play data we have is limited and the collected data on state-action pairs in the value table would be sparse. Furthermore, some states may be barely encountered in the real interaction but still have a non-zero probability of being encountered. For example, most children are able to finish the game in 5 child-robot turns, though a few children may still take 10 turns to finish. To address the concerns and limitations, value function methods that are able to generalize from previous encounters with different states are favored. Because the approximate solution method is able to usefully generalize a limited subset of the state space to yield a good approximation for a larger subset of state space, it is used in the RL model as the robot's role switching policy.

In my case, a *stochastic gradient descent* (SGD) learning method for function approximation in value prediction is used, as SGD value approximation methods are one of the most widely used in various RL problems. The approximate value function $\hat{v}(s, \mathbf{w})$ is used to represent $v_\pi(s)$ as the RL model's value function where $\mathbf{w} = (w_1, w_2, \dots, w_d)^T$ and $s \in \text{state features}$. \mathbf{w} is updated at each timestamp when rewards r is received by a small amount to reduce the error between the real value and the estimated value:

$$w_{t+1} = w_t - \frac{1}{2} \alpha \nabla [v_\pi(S_t) - \hat{v}(S_t, \mathbf{w}_t)]^2$$

Regarding the specific architecture of $\hat{v}(s, \mathbf{w})$, a linear function of the weight vector, \mathbf{w} , is used, as linear methods are efficient in terms of both data usage and computation.

$$\hat{v}(s, \mathbf{w}) = \mathbf{w}^T \mathbf{x}(s) = \sum_{i=1}^d w_i x_i(s) \text{ where } \mathbf{x}(s) \text{ is a feature vector representing state } s.$$

As mentioned in *section 3.2.5.1*, the role-switching RL policy has 3 state features: the number of current turns, the number of expert role occurrences, and the number of child's correct attempts. Based on the three state features, polynomial features are constructed for the linear value approximation function using $x_i(s) = \prod_{j=1}^3 s_j^{c_{i,j}}$ where the natural state space s has 3 dimensions and $c_{i,j} \in \{0, 1, 2\}$. A polynomial base of 2 is chosen given its ability to take more complex interactions into account when compared with 1-dimensional feature vectors. However, the final feature vector is only a subset of all 3^3 possible quadratic features that can be generated from the 3-dimensional state space. The selection is done based on prior human expert beliefs about the nature of the approximation function to avoid potential overfitting issues. The final set of the selected features for the linear approximate value function is $\{s_1, s_2, s_3, s_1s_2, s_1s_3, s_2s_3, s_1^2s_2, s_1^2s_3, s_2^2s_3, s_2^2s_1, s_3^2s_1, s_3^2s_2\}$. To make the model's implementation easier, I built two SDG linear regression models, each representing one robot's role (expert, novice). The action (robot's role) that has a regression model with a higher value would be the best action predicted by the RL's value function.

Role-Switching Policy Training and Results

Q-learning is used as the control algorithm in the RL policy model [104] to learn the optimal action-value function, defined by

$$Q(S_t, A_t) \leftarrow Q(S_t, A_t) + \alpha [R_{t+1} + \gamma \max_a Q(S_{t+1}, a) - Q(S_t, A_t)]$$

The model's action selection method is the ϵ -greedy algorithm, shown below.

$$a_t = \begin{cases} a_t^* & \text{with probability } 1 - \epsilon \\ \text{random action} & \text{with probability } \epsilon \end{cases}$$

The role-switching policy model was trained on 143 episodes each containing data from a child and a robot starting a game mission to their completing the mission. Since the data had been collected prior to the model training, it was not possible to do a full on-policy training and a real action selection based on the ϵ -greedy algorithm. To address the above constraint, the model assumed that the ϵ -greedy algorithm always had a probability $1 - \epsilon$ to choose “exploitation” (the best action a_t^*) but chose “exploration” (random action) instead. After the training, each action's SDG linear regression model is displayed below in terms of its weights. The intercept values for the expert and novice regression models are 1.319 and 1.156 respectively.

These two trained models were used as base value function approximation models for adaptive role switching in the experiment. Then, the role-switching policy will further adapt to each child. During the experiment, the model uses the ϵ -greedy algorithm with a ϵ value of 0.25 and a discount factor of 0.5 with the goal of adapting the robot's role to each child. The pipeline for developing, training and testing the robot's role-switching model is outlined in Figure 8 below.

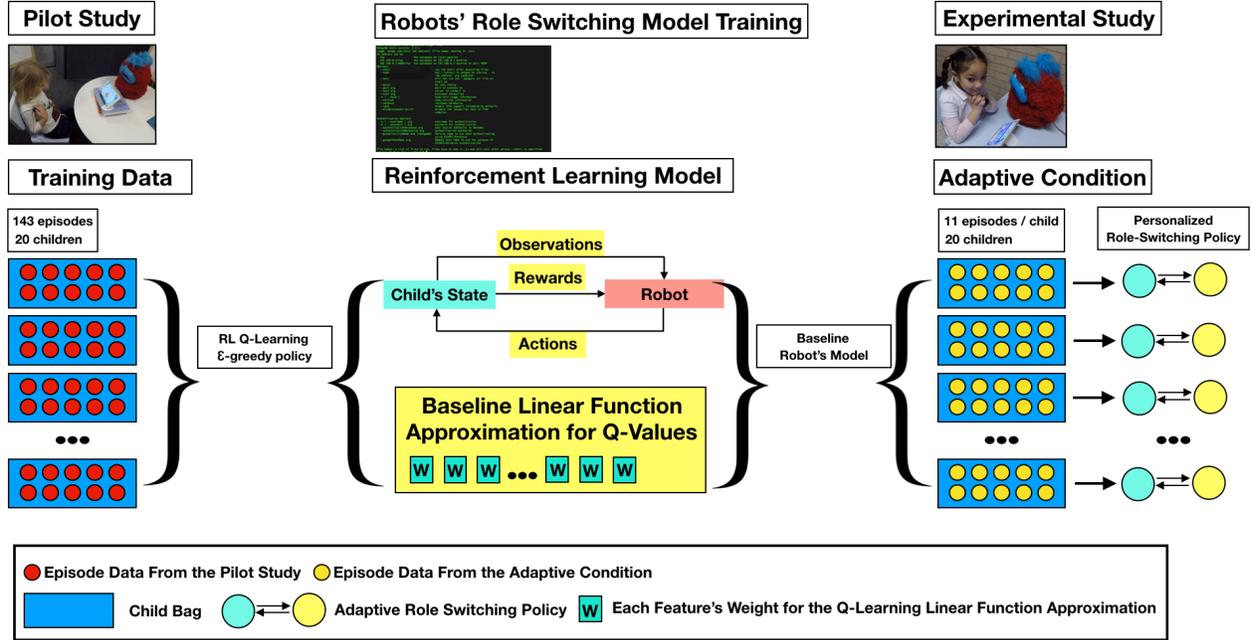


Figure 8. The pipeline for developing, training and testing the robot's role-switching model.

Regression	s_1	s_2	s_3	s_1s_2	s_1s_3	s_2s_3	$s_1s_2s_3$	$s_1^2s_2$	$s_1^2s_3$	$s_2^2s_3$	$s_2^2s_1$	$s_3^2s_1$	$s_3^2s_2$
Expert	-0.54	-0.86	-0.54	0.9	-0.58	-0.34	0.29	-0.35	-0.25	-0.10	0.03	0.30	0.58
Novice	-0.81	0.09	-0.11	0.26	0.19	-0.02	0.05	-0.16	0.07	-0.16	0.08	0.08	-0.02

Table 4. Weights for features in linear value approximation function.

Chapter 7: Comparative Study

Overview

The goal of this study was to investigate how an interactive agent (e.g., social robot) can interact with children in a bidirectional, peer-to-peer way to foster children's learning and engagement. A robot plays a collaborative tablet-based word learning game called WordQuest with a child, and the child learns a set of new words every time when they play the game with the robot. During the game play, the behaviors the robot exhibit may differ in terms of its demonstration of knowledge. Sometimes the robot behaves more like a novice peer who spies objects, makes mistakes, and explore the game scene without explicitly giving instruction and informative feedback to the child. In other situations, the robot behaves more like an expert tutor who gives instructions, demonstrates its knowledge, explains concepts, and offers help to the child. Thus, how do a robot's *novice* and *expert* behaviors affect a child's learning performance, learning attitudes, engagement, and perception of both the learning game and the robot? Furthermore, does playing with a robot that is able to adaptively switch its role between *novice* and *expert* lead to greater learning outcomes and greater engagement? Finally, how should a robot adaptively switch its role during its game play with a child to optimize the child's learning experience?

Conditions

Three conditions were designed and tested in the study. The first condition was a *Fixed Novice* condition (*Novice*, N). The robot always played a novice role during the game play with a child. It had a 50% chance of spying a correct object and 50% chance of spying an incorrect object, but it neither demonstrated to the child why an object was right/wrong nor explained the meanings of the key vocabulary words to the child. In this condition, the robot also asked the child questions and sometimes asked them for help.

The second condition was a *Fixed Expert* condition (*Expert*, E). In this condition, the robot behaved like a tutor and instructor with a greater language ability than the child. The robot verbalized the meanings of all vocabulary words, gave hints on what to spy, and spied correct objects every turn.

Last, the third condition was an *Adaptive Role Switching* condition (*Adaptive, A*), where the robot adaptively switched its role between novice and expert within each game mission based on how the child exhibited their knowledge of the vocabulary word.

Hypotheses

My hypotheses are the following:

- H1.** Children from all three conditions learn vocabulary words when playing the WordQuest game with a robot.
- H2.** Children in the *Adaptive (A)* condition recall most words in both the immediate post-test and delayed post-test, followed by children in the *Expert (E)* condition.
- H3.** Children in the *Adaptive* condition display highest overall engagement and highest overall positive affect during the course of the learning session among all children.
- H4.** Children in the *Adaptive* condition show the least amount of decline in engagement during the course of the learning session and highest engagement at the end of the session among all children.

Participants

Participants were 64 children ages 5-7 from a Boston-area elementary school. Thirty-four participants were from the K2 classrooms; 30 participants were from the 1st grade. Thirty-six were female; 28 were male. The average age for the 64 children is 5.91 (SD: 0.65), and 17, 37, 11 were 5-year-olds, 6-year-olds and 7-year-olds respectively. There were 33 Native Speaker and 31 ESL students. In total, 5 children were excluded from the data analysis. Two children had difficulty learning how to play the WordQuest game during the practice round and did not start the learning sessions. Three children did not finish the two learning lessons and post-assessments due to not-study-related reasons (e.g., early dismissal from school). One-way ANOVA and pairwise t-test were conducted across 3 conditions, and showed that both the age ($F=0.23$; $p=0.79$) and pre-word assessment score ($F=0.16$; $p=0.84$) were not significantly different across the 3 conditions. Table 5 indicates how many children were in each condition, when divided by their age, English proficiency, pre-test score, and sex. In addition, 12 from the 60 participants

who finished the study were reported by their teachers to have signs of either attention deficits or learning difficulties.

Condition	English Proficiency	Sex	Number of Children	Average Age (SD)	Pre-Assessment Score (SD)	Attention Deficits/Learning Difficulties
Fixed Expert	Native: 10 ESL: 11	F: 13 M: 8	21	5.85 (0.65)	2.44 (1.43)	3
Fixed Novice	Native: 10 ESL: 9	F: 11 M: 8	19	6.00 (0.74)	2.68 (1.37)	3
Adaptive Role Switching	Native: 11 ESL: 8	F: 9 M: 10	19	5.95 (0.60)	2.58 (1.43)	6

Table 5. This table indicates how many children were in each condition, when divided by their English proficiency, sex, age and pre-assessment score.

Robot Testbed

The Tega is an expressive, child-friendly and fluffy robot designed as a learning companion for young children in various educational settings such as schools and homes[105]. It has been used to teach children storytelling, learning attitudes and a second language in prior studies[64][71][48]. It is about 11 inches tall and has a plush exterior with different colors. In my study, I used a red Tega. Unlike most of other educational devices and robots, Tega has a child-like high pitched voice, slow speed of speech, exaggerated body and facial expressions, as well as child-like growth mindset. These peer attributes make Tega a perfect interactive agent for children to learn together with.

The robot used in the study is fully autonomous, including its gaze, facial expression, speech and actions. The robot asks children questions in the game play, so Google Automatic Speech Recognition (ASR) is used to detect and recognize children’s simple speech. A complex dialogue system between the robot and the child is not feasible giving the difficulties of children’s speech recognition and natural language processing. To avoid such difficulties, I simplified the questions Tega asked and the way to recognize children’s answers. When Tega asked “yes/no” question, Google ASR detected only individual keywords (e.g., “yes”, “yeah”, “no”, “nope”,

“sure”) or short phrases (e.g., “of course”) rather than entire sentences. When Tega asked open-ended questions, Google ASR detected only keywords or short phrases (e.g., “because”, “it is”) that indicate that the children give meaningful answers. Furthermore, the speech of the robot used in this project was scripted and recorded by a female voice. Then, the speech is pitch shifted to make Tega sound more like a peer.



Figure 9. Tega as the study's social robotics' system.

Methods

Procedure

The study used a between-subjects design, which has three conditions (*Expert* × *Novice* × *Adaptive*). To reduce variability within conditions and potential confounding, I counter-balanced children's age, sex, pre-word assessment score, and English proficiency across the three experimental conditions. Each child participated in two learning sessions spread over one week, one pre-test session and a 3-week delayed post-test session.

During the pre-test session, the children were given a pre-word assessment, which consists of the 11 vocabulary words the children would learn with the robot together in the rest of the two sessions. During the following two learning sessions, children played the WordQuest game with the robot for 20-30 minutes. The children did a practice round for 5 minutes with the help of an experimenter to learn how to play the game before they started the first learning session. After each learning session, the experimenter conducted a post-survey and assessment on the children to gauge their perception of the game and robot as well as to measure their learning outcomes. Then, each of the children were assessed on their retention of the learned 11 words 3 weeks after

they finished their second session. The overall procedure and the sequence of the experiment is shown in Figure 10 below.

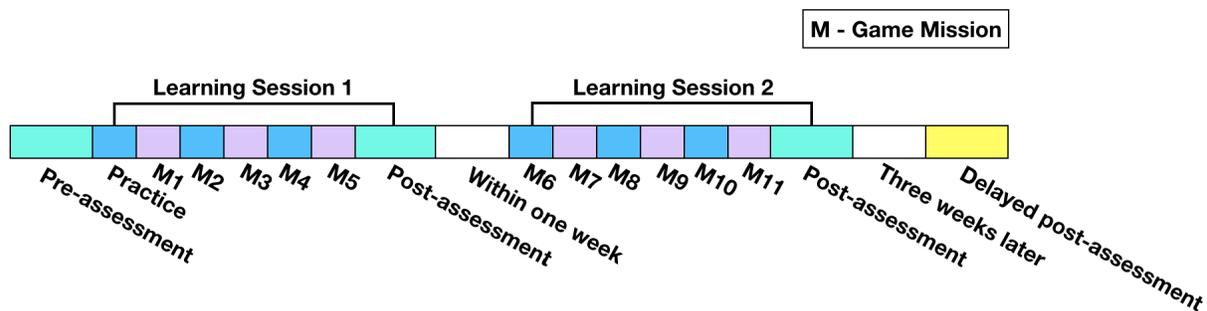


Figure 10. The procedure and sequence of the experiment.



Figure 11. Children play the WordQuest game with Tega.

Target Vocabulary words

Eleven vocabulary words were selected for the WordQuest game play. These words include 3 color words (*lavender, crimson, azure*), 5 adjective words (*gigantic, minuscule, aquatic, delighted, recreational*), 2 object category words (*garment, vehicle*) and 1 verb (*soar*). Children learned 5 words from learning session 1, and 6 words from learning session 2. Table 6 shows the full word list. Those words are selected for the following reasons. First, they are advanced words that have simple synonym words children should be familiar with (e.g., blue, cloth, red, happy), so it would be easy for children to conceptually understand the meanings of the target vocabulary words. Second, all the children in the study understood at least basic level of English. Third, all these vocabulary words have been tested in the pilot study, and the children in the pilot study didn't have trouble comprehending their meanings. To accurately assess the impact of a social robot's different roles on children's learning performance, learning advanced words that children were not familiar with before the game play became necessary.

Vocabulary word list (session 1)		
gigantic	garment	azure
minuscule	lavender	

Vocabulary word list (session 2)		
vehicle	crimson	soar
recreational activity	aquatic animal	delighted

Table 6. List of target vocabulary words. Children learned 5 words in session 1 and 6 words in session 2

Data Collection

Multimodal data from multiple channels were collected in the study. Figure 12 below displays how data were collected. Front-face video of children’s interactions with the robot and the video from the robot’s view was recorded with two USB cameras. All touch actions on the tablet screen in the WordQuest game (e.g., dragging, panning, tapping), children’s in-game performance (e.g., successful attempts, incorrectly spied objects), and robot’s verbal and nonverbal behaviors (e.g., question asking, demonstration) were recorded. Children’s physiological states during game play were also recorded using E4 sensors. Tega’s built-in microphone was used to detect children’s answers when Tega asked questions.

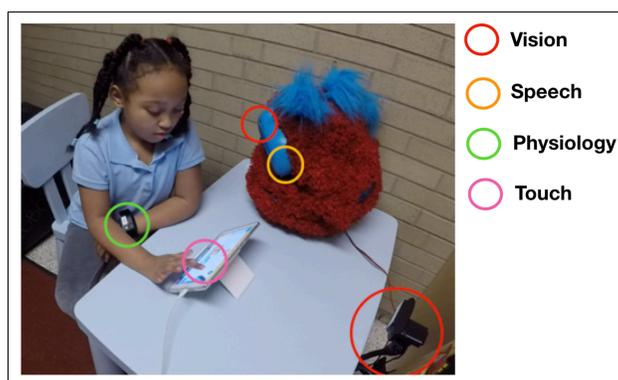


Figure 12. Multimodal data are collected during the two learning sessions.

Learning performance measures

Children were assessed on the target vocabulary words during the pre-test word assessment before playing the WordQuest game. For each target word, four images were shown to the child. The child was asked to point to the image matching the target word. Then, the experimenter asked the child about the meaning of each target vocabulary word to avoid any false positive and false negative answers. After each learning session, each child was assessed on the target vocabulary words they just learned in that particular session. This measure was used to determine what vocabulary words each child had already learned before the study and what words they acquired after the game play. Figure 13 below shows the interface for this vocabulary assessment. Furthermore, a child's learning performance is also measured in their interaction duration for each vocabulary word's mission, which is the duration between the time they start a mission and the time they end this mission.

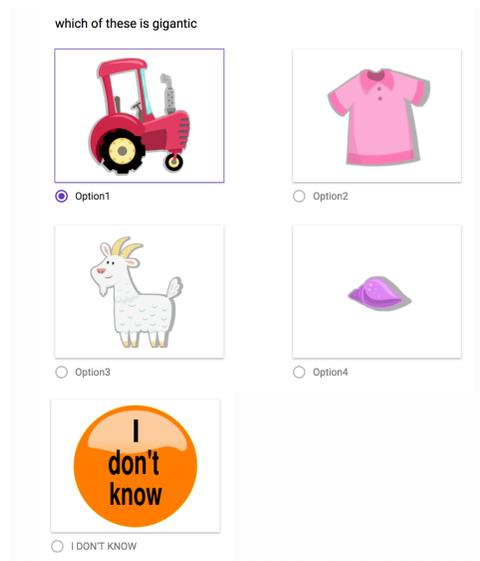


Figure 13. The interface for the target vocabulary word assessment.

Engagement and affect measures

Children's engagement and affective-states were measured in the following way. Children's engagement was captured by Affdex SDK and their touch interactions on the tablet [106]. Affdex SDK is a commercially available automated affect recognition software. It takes in video data containing a person's front face, and extracts 15 physical expressions from facial features, and predicts the likelihood of an exhibited emotion. In my study, it is used to estimate child's

engagement level, as it can return a score for engagement level ranging from 0 to 100 with a high value indicating high engagement. In addition to Affdex SDK, child's engagement level is also measured by their touch interactions with the tablet. The touch interactions include how often the child pans the screen, how often they zoom in and out on, and how often they click an object on the screen. Last, the child's positive affect and valence/arousal are measured through child's facial expression captured Affdex SDK, which is able to estimate the degree of valence (positive and negative emotion) and expressiveness (intensity of an expression) based on the child's facial features.

Children's engagement and affective states were measured at specific time incidents. Table 7 below describes the time incidents where children's engagement was measured as well as the corresponding time interval.

Incident	Time Interval (sec)
Start of child's turn	2
Start of robot's turn	2
Robot taps one object	3
Result return of child's attempt	2
Result return of robot's attempt	2

Table 7. Incident during which child's engagement and affective states were measures.

Post-game enjoyment and perception measures

Children were asked questions regarding their enjoyment of the game play sessions and perception of the game and the robot. The questions were inspired by prior research on measuring children's long-term relationships with social robots [107]. The questions were then modified to tailor to my study. The finalized questions in the study are:

1. How did you feel before the game play?
2. How did you feel when you played with Tega?
3. How did you feel when you played the WordQuest game?
4. How easy/hard do you think the game is?
5. Did you learn new words today? How many?
6. How much do you like/dislike Tega?

7. How smart do you think Tega is?
8. Are you smarter than Tega, or Tega is smarter than you?
9. How much did Tega help you learn new words?
10. How much did you help Tega learn new words?
11. It is more fun playing the WordQuest game with Tega or without Tega?

The children were asked the questions above after each learning session. The questions about children's feelings were presented with emojis, and children were asked to choose one of the emojis that best matched their own feeling, as shown in Figure 14. Other questions were presented on a 7-item scale, as shown in Figure 14.

How did you feel when you played with Tega? (AFTER GAME PLAY) *


 Very Happy


 Happy


 Just okay


 Unhappy


 Very unhappy

I am smarter than Tega / Tega is smarter than me *

I am smarter

1 2 3 4 5 6 7

Tega is smarter

I am smarter

Tega is smarter

Figure 14. Survey on child's feeling and perception after the game play.

Chapter 8: Analysis & Results

A generalized linear model was applied to predict various measurements of learning outcomes and engagement. The predictor variable was contrast-coded as ordered values [-1,0,1], *Novice*, *Expert* and *Adaptive* respectively for all the measurements[108].

Learning outcomes

The scores of children's vocabulary assessment by experimental condition were reported in Table 8 and visualized in Figure 15. The ANOVA analysis and the trend analysis were conducted on each vocabulary score category. The results from the trend analysis for each score category were displayed in Figure 16.

As shown in the table, the mean difference in vocabulary scores between the immediate post-test and pre-test were 4.79 (SD=2.51), 3.57 (SD=2.33), 1.95 (SD=2.27) for *Adaptive*, *Expert* and *Novice*, respectively. A generalized linear model was calculated to predict the trend of children's immediate vocabulary score change across the three conditions. A contrast-coded generalized linear regression model showed a statistically significant trend of increase found in the immediate vocabulary change in the order (*Adaptive* > *Expert* > *Novice*) with $F(2, 56) = 13.85$, $p < 0.001$. The Levene Test's score is 0.5 with $p = 0.6$, and the Shapiro Test's score is 0.97 with $p = 0.24$. The one-way ANOVA showed a statistical difference in the vocabulary score change across the three conditions, $F(2, 56) = 6.023$, $p < 0.005$.

Then, the mean difference in vocabulary scores between the delayed post-test and pre-test were 3.47 (SD=2.14), 2.67 (SD=1.94), 2.05 (SD=1.36) for *Adaptive*, *Expert* and *Novice*, respectively. A contrast-coded generalized linear regression model showed a statistically significant trend of increase in the delayed score change across the three conditions in the order (*Adaptive* > *Expert* > *Novice*), $F(2, 56) = 5.45$, $p = 0.023$. The Levene Test's score is 1.95 with $p = 0.14$, and the Shapiro Test's score is 0.97 with $p = 0.09$. The one-way ANOVA showed a statistical difference in the delayed score change across the three conditions, $F(2, 56) = 2.70$, $p = 0.076$.

Condition	Pre-Test Score (SD)	Post-Test Score (SD)	Delayed Post-Test Score (SD)	Diff (Post-Test, Pre-Test) (SD)	Diff (Delayed Post-Test, Pre-Test) (SD)
Novice	2.68 (1.38)	4.63 (1.74)	4.74 (1.65)	1.95 (2.27)	2.05 (1.36)
Expert	2.43 (1.43)	6.00 (2.07)	5.10 (1.95)	3.57 (2.33)	2.67 (1.94)
Adaptive	2.58 (1.43)	7.37 (2.41)	6.05 (2.14)	4.79 (2.51)	3.47 (2.14)

Table 8. The average scores of children's vocabulary assessment by experimental condition.

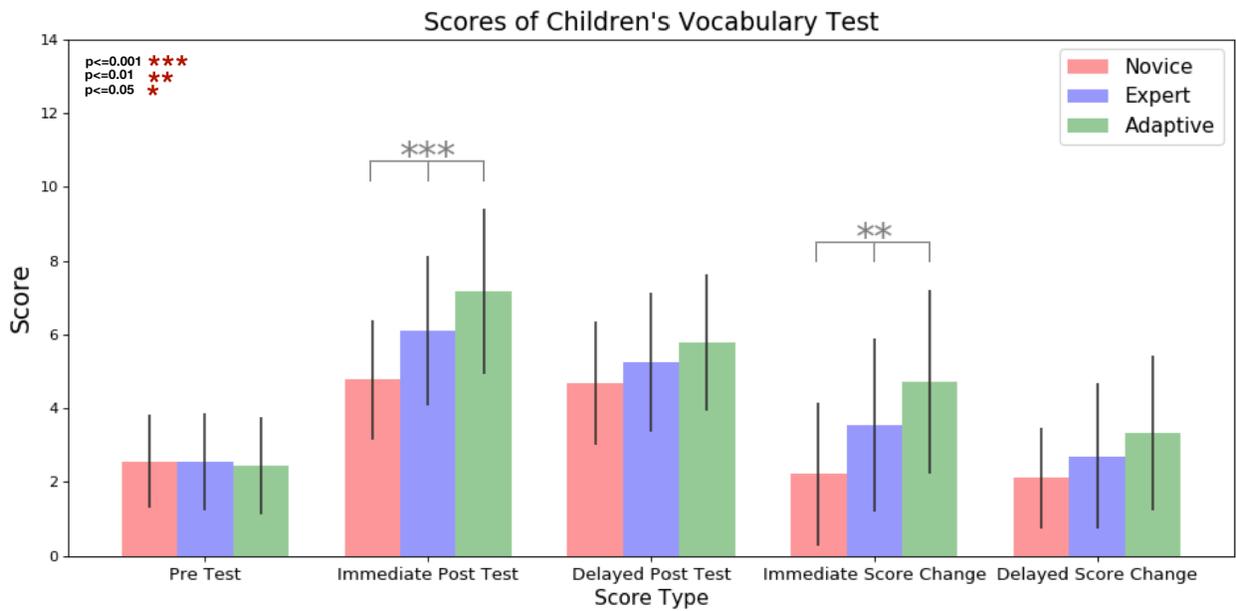


Figure 15. Average scores of children's vocabulary assessments across the 3 experimental conditions. The ANOVA analysis showed significantly results in the immediate posttest and immediate score change.

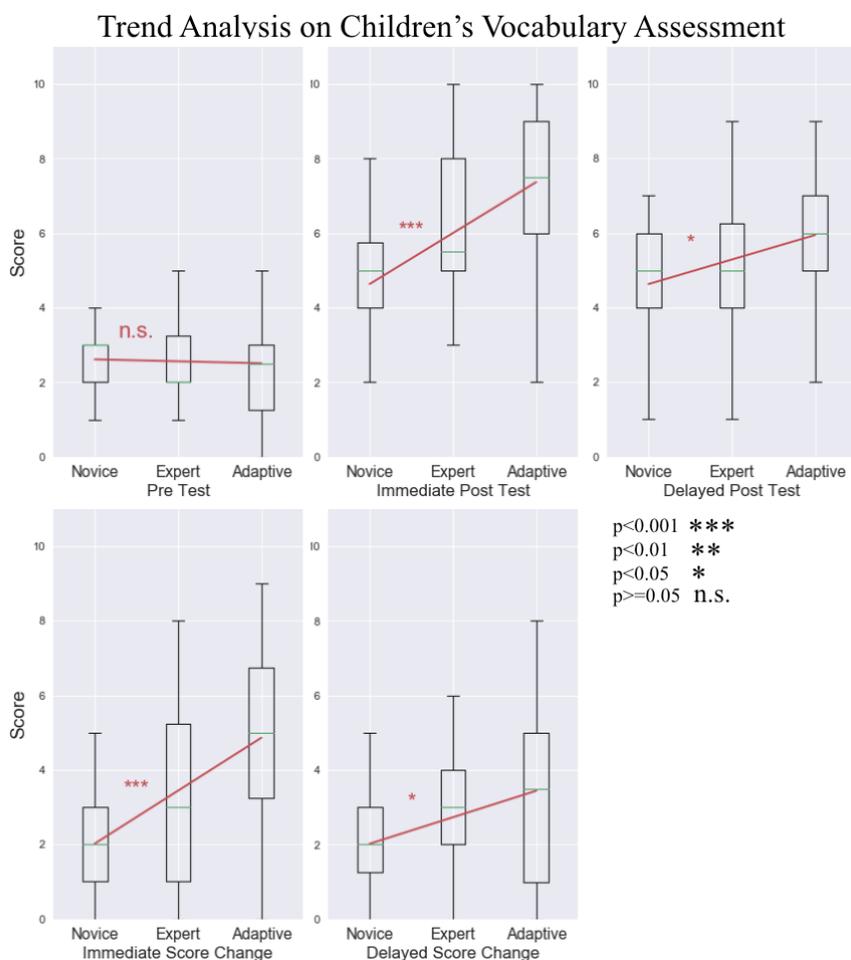


Figure 16. Trend Analysis on the scores of children's vocabulary assessment. The increase of vocabulary scores is shown in the order (Novice < Expert < Adaptive) for all score categories except the pre-test vocabulary assessment.

The above results showed that the children in the *novice* condition learned least amount of words among the three conditions. This result is not surprising, as previous research on children's learning showed that children may not get enough support for effective learning if they attempt to learn through unguided or un-scaffolded exploration and discovery [109].

To understand children's learning on an individual level, Figure 17 below displays the scores of individual student's vocabulary assessment over time by condition. For the *adaptive* and *expert* conditions, students tended to learn more new words in the immediate post-test but also tended to forget more words three weeks later. However, this pattern was not that clear in the *novice* condition. Children in the *novice* condition seemed to better retain the very few words they learned, and some of them even scored higher in the delayed post-test than they did in the

immediate post-test. One possible explanation for this observation is that children were in a more engaged and positive affective state when doing the delayed post-test than when doing the immediate post-test. It is very likely that children in the *novice* condition might have experienced much frustration and anxiety due to the lack of help and instruction during their learning sessions. Thus, they might have been cognitively overloaded and tired after each learning session and failed to correctly perform well in the immediate post-test. Another potential reason is that children might have discussed the meanings of the target vocabulary words they failed to comprehend with their friends or family after the study, and this reinforcement of their vocabulary learning then showed its effect in the 3-week delayed post-test.

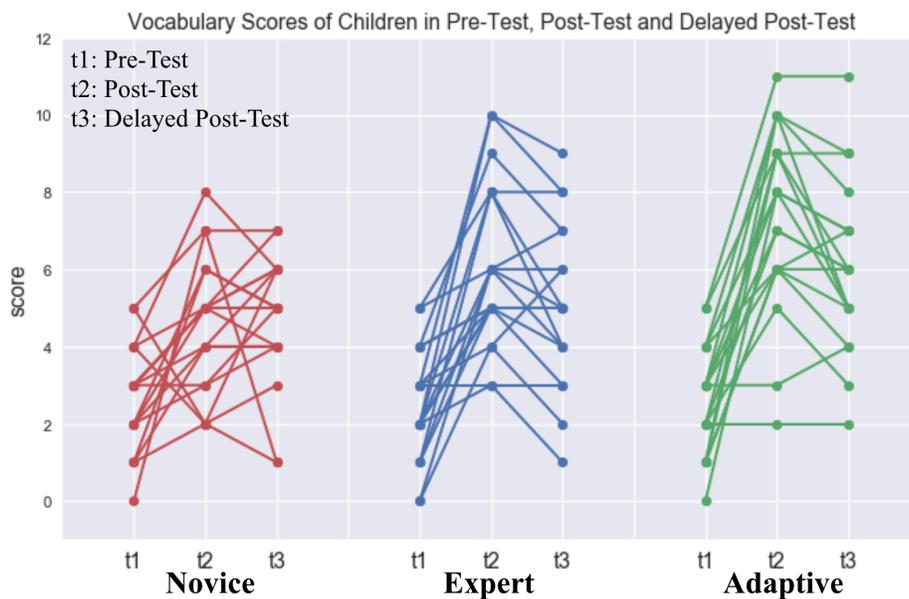


Figure 17. Vocabulary scores of individual students change over the three vocabulary assessment sessions by condition.

To further understand the change of children's vocabulary scores across the 3 conditions, I visualized the distribution of the score change immediately after the study and the change 3 weeks after the study in Figure 18. The distribution of the top performing students who learned the most words or retained the most words was displayed in Table 9. As shown in the figure and table, the top performing children in the *Adaptive* condition consistently outnumbered the children from the other two conditions in four measures. This result indicates that interacting with the *adaptive* robot is more likely to promote children's best learning than with either the *expert* robot or the *novice* robot.

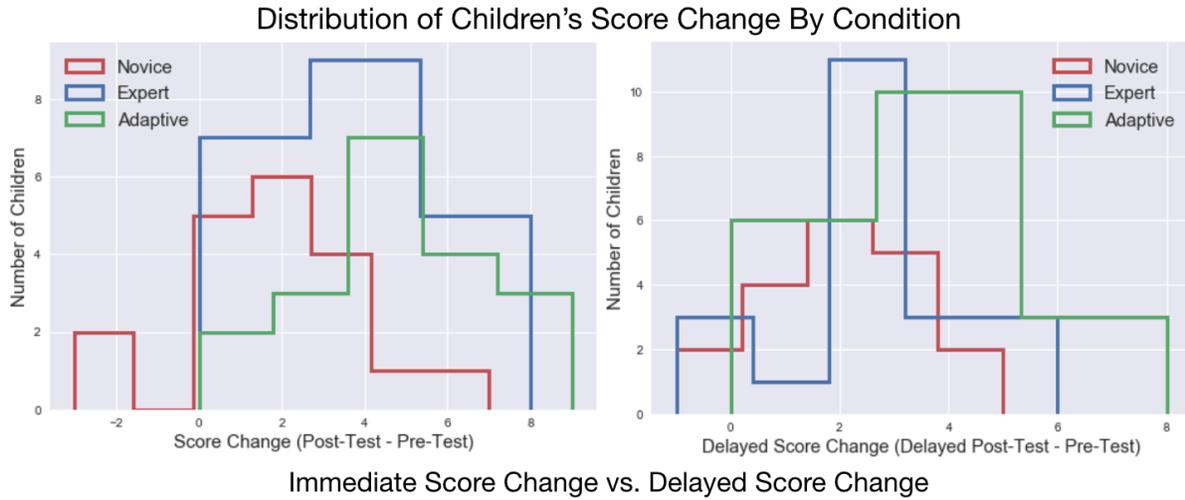


Figure 18. Distribution of children's vocabulary score change by condition.

Condition	Immediate Score Change		Delayed Score Change	
	Top 9 students	Top 20 students	Top 11 students	Top 18 students
Novice	1	2	1	2
Expert	3	7	3	6
Adaptive	5	11	7	10

Table 9. Distribution of top performing children by condition.

Children's Vocabulary Scores by Gender, Age and Prior Vocabulary Level

The change of children's vocabulary score immediately after the study and 3 weeks after the study were compared based on their gender, age, and prior vocabulary level. The results are displayed in Table 10. The unpaired t-test showed that no significant result in both children's immediate score change and delayed score change between the male students and the female students. The one-way ANOVA analysis showed no significant difference across the 3 age conditions, $F(2, 56) = 1.99, p = 0.15$.

Category	Condition	Number of Participants	Immediate Score Change (SD)	Delayed Score Change (SD)
Gender	Female	33	3.70 (2.55)	3.09 (1.78)
	Male	26	3.12 (2.58)	2.30 (2.01)

Age	5	15	2.73 (2.76)	2.67 (2.29)
	6	33	4.03 (2.38)	3.03 (1.55)
	7	11	2.64 (2.42)	2.45 (2.23)
Prior Vocabulary Level	Top (pre-test score: 3-5)	24	2.50 (2.27)	2.17 (1.72)
	Bottom (pre-test score: 0-3)	29	4.48 (2.34)	3.24 (1.99)
Prior Vocabulary Level	High (pre-test score: 3-6)	24	2.50 (2.27)	2.17 (1.72)
	Middle (pre-test score: 2-3)	15	3.53 (2.16)	2.53 (1.71)
	Low (pre-test score: 0-2)	14	5.50 (2.10)	4.00 (2.00)

Table 10. Children's vocabulary score change by gender, age, and prior vocabulary level.

To understand how children's prior vocabulary level affects their vocabulary learning during the game play, I categorized children based on their vocabulary score during the pre-test. The children were first categorized into two groups: *top 50%* of which students have pre-test scores ranging from 3 to 5, and *bottom 50%* of which students have pre-test scores ranging from 0 to 3. The learning performance of two groups were displayed in the left plot in Figure 19 below. The children from the *bottom 50%* group significantly learned more words than the children from the *top 50%* group immediately after the study and 3 weeks after the study, given that the t-test results between the two groups were 3.05 ($p=0.004$) and 2.04 ($p=0.047$) respectively. To further gauge how much the vocabulary improvement is biased toward children with low prior vocabulary level, I categorized the children into 3 groups: *high* (pre-test scores ranging from 3 to 5), *middle* (pre-test scores ranging from 2 to 3), *low* (pre-test scores ranging from 0 to 2). The results showed that the *low* group has the highest immediate score change and delayed score change, as shown in the right plot in Figure 19. The Shapiro test score is 0.97 ($p=0.30$), and Levene test score is $F=0.049$, $p=0.95$. The one-way ANOVA on the immediate score change shows a significant difference across the *low*, *middle* and *high* groups with $F(2,56)=7.80$, $p=0.001$. The unpaired t-test results between the *low* group and *middle* group and between the *low* group and *high* group were both significant, which were $t=2.40$ ($p=0.024$) and $t=3.93$ ($p<0.001$) respectively. Furthermore, the same significant results were found in the delayed score change. Specifically, the one-way ANOVA result across the three conditions was $F(2,56)=4.49$, $p=0.016$. In addition, the *low* group was also significantly higher in the delayed score change from both the middle group and high group with $t=2.05$ ($p=0.049$) and $t=2.90$ ($p=0.006$), respectively. The results indicate that children with a wide range of vocabulary knowledge are all

able to play the game and learn new words. Even more, the WordQuest interaction system is most beneficial to children with a low prior vocabulary level.

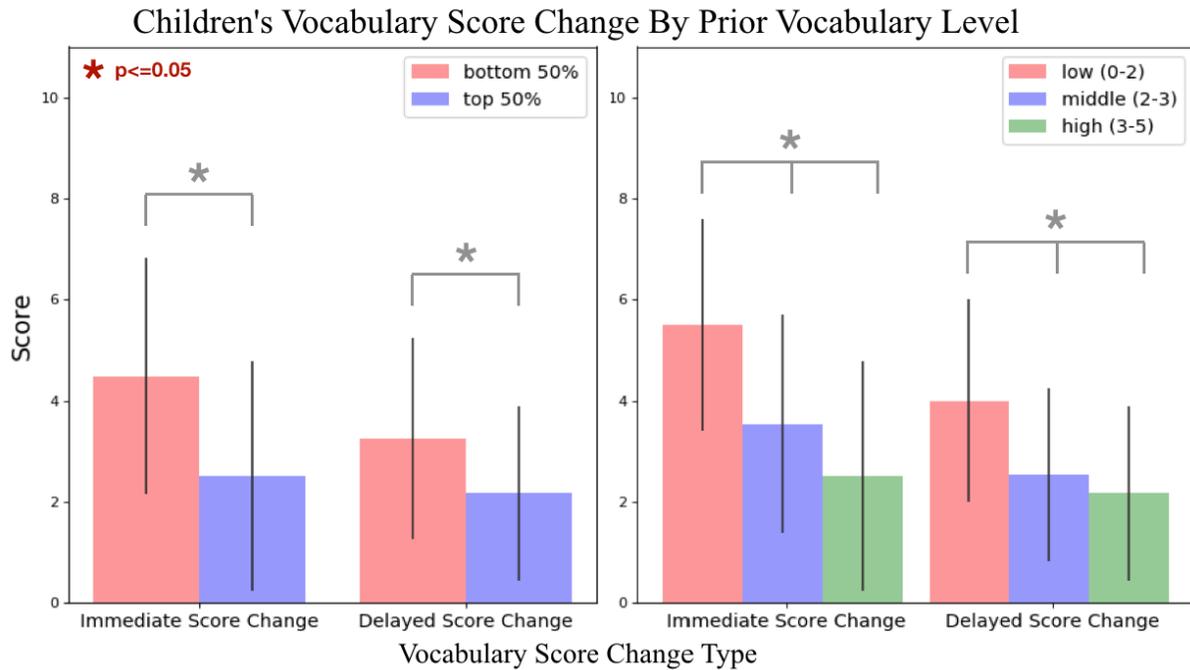


Figure 19. Children's vocabulary score change by children's prior vocabulary level.

Acquisition of Individual Vocabulary Words

Table 11 below displays the number of children who knew the meanings for the vocabulary words before the study and after the study. Among the 11 vocabulary words, the top 3 words that the participants had known before the learning sessions were “gigantic”, “vehicle”, and “recreational activity.” The top 3 words, of which meanings the participants knew after the two learning sessions were “gigantic”, “vehicle” and “lavender,” regardless of whether they had known the words prior to the study. Furthermore, the top 3 words that the children learned after the learning sessions were “delighted”, “lavender” and “crimson,” and the top 3 learned words three weeks after were still “delighted”, “crimson” and “lavender.” It is interesting to notice that only 9 students knew the meaning of “garment” immediately after the study and only 3 still remembered it 3 weeks after, indicating that this word is a relatively hard word to learn and remember for all children.

The number of students who knew each word declined three weeks after the study for most of the words except “gigantic” and “azure.” The number of students who knew “gigantic” increased from 48 to 51 three weeks after, but the change was small. However, the number of students who knew “azure” increased from 11 to 22 three weeks after the study. One potential explanation is that the children may have exposure to “gigantic” and “azure” in class at school or in the home during this 3-week time interval, and the meanings of the two words were reinforced. Figure 20 below displays the change in the number of students who knew the vocabulary words across pre-test, post-test and delayed post-test.

Vocabulary word	Num. children who knew the word before the study	Num. children who knew the word immediately after the study	Num. children who knew the word 3 weeks after the study	Num. children who learned the word after the study	Num. children who learned the word 3 weeks after the study
azure	1	11	22	10	21
gigantic	39	48	51	9	12
minuscule	7	26	16	19	9
garment	1	9	3	8	2
lavender	14	40	36	26	22
vehicle	47	57	55	10	8
delighted	5	43	31	38	26
crimson	2	28	25	26	23
soar	15	28	25	13	10
aquatic	6	29	19	23	13
recreational activity	25	38	29	13	4

Table 11. Children’s acquisition of individual vocabulary words.

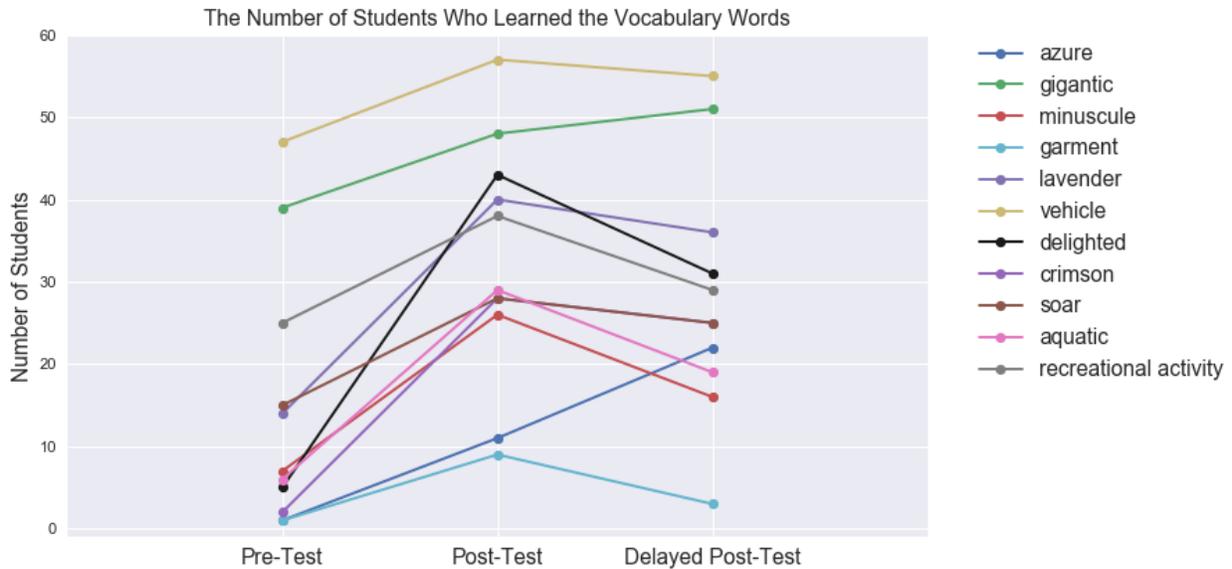


Figure 20. The number of students who knew, learned and retained the vocabulary words over time.

Children’s Perception of the WordQuest Interaction

Children’s perception of the WordQuest Interaction was measured in the post-assessments following the learning sessions. Figure 21 displays the results of their perception across two learning sessions by experiment conditions. The first three survey questions on children’s feeling were measured on a 5-item scale ranging from 1 to 5 with 5 indicating “super happy” and 1 indicating “super unhappy.” The figure shows that all children felt happy after each game play, regardless of the experiment conditions they were in. The rest of the questions (Q4-Q13) were measured on a 7-item scale with 1 indicating a complete agreement on the first statement and 7 indicating a complete agreement on the second statement in the question. Regarding the easiness of the game, all children perceived the game as *not hard* or *somewhat easy*. Specifically, the children from the *Novice* and *Adaptive* conditions felt the game was harder in the first session but became easier in the second session, while the children in the *Expert* condition gave the most consistent, lowest scores to this question on the game’s hardness. This result indicates that the WordQuest interaction enabled the children to feel sense of mastery even when the majority of the words they were learning were unknown, challenging words.

The top 3 questions to which answers the three conditions differed most were Q8 (“*Who is smarter? Tega or me?*”), Q7 (“*How smart is Tega?*”), and Q10 (“*Did I help Tega learn new words?*”). Regarding Q8, the children in the *Expert* condition thought Tega was somewhat

smarter than themselves, and this perception was reinforced further in the second condition. Conversely, the children from both the *Novice* and *Adaptive* condition thought Tega was a little bit smarter in the first session, but later thought they were somewhat smarter or at least equally smart in the second session. This result indicates that Tega had some level of smartness in children's eyes, but Tega's role-related behaviors could also affect children's perception of how smart Tega was in relative to themselves. This result also sheds light on how to personalize a robot's behaviors to individual children in future work. For example, the robot can exhibit more novice behaviors (e.g., making mistakes or asking the child for help) when interacting with children having low self-esteem, as the robot's novice behaviors tend to lead children to think they are smarter than the robot.

Regarding Q7, most of the children thought Tega was very smart, especially the children in the *Expert* condition. The children in the *Adaptive* condition thought Tega became smarter in the second learning session, whereas the children in the *Novice* condition thought Tega became less smart in the second session. Regarding Q10, all children thought that they helped the robot out but at varying degrees. Specifically, the children in the *Expert* condition rated their helpfulness to the robot lowest, while the children in the *Novice* condition rated their helpfulness highest. This result suggests that the children assessed how helpful they were to the robot based on what role-related behaviors the robot exhibited. Thus, when designing the WordQuest interaction for children with low self-esteem or on the autism spectrum, a robot's novice behaviors may be very helpful in building up children's confidence and engagement.

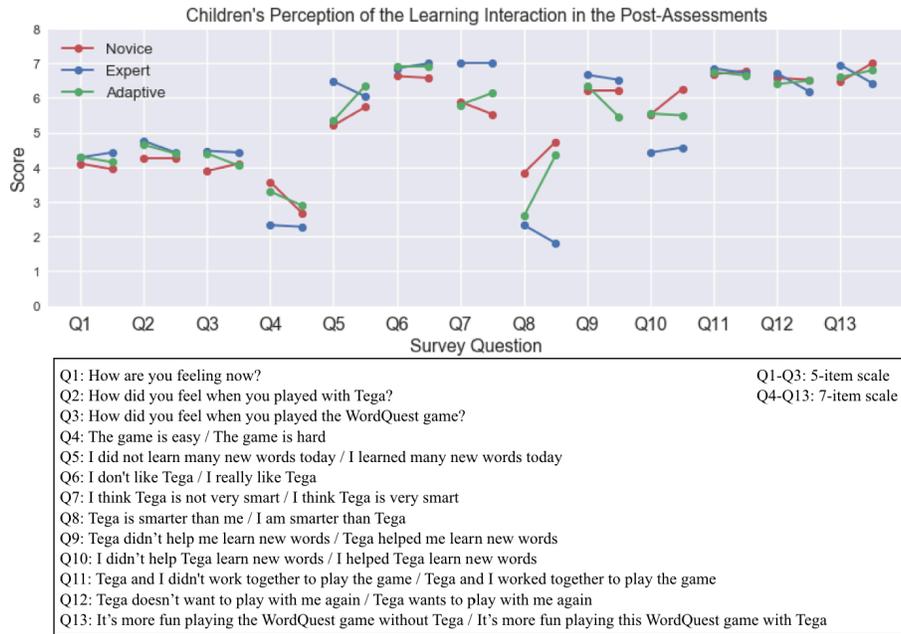


Figure 21. Children's perception of the WordQuest interaction measured in the post-assessments.

Chapter 9: Discussion & Conclusion

The results presented above examined the hypotheses proposed in this work.

- H1: Playing the WordQuest game with a socially interactive robot helped children learn new vocabulary words in a context where the majority of target vocabulary words are challenging words to them.
- H2: Children in all conditions were engaged and had a positive learning experience.
- H3: Children in the *Adaptive* condition learned the most words after the learning sessions and recalled the most words three weeks after the learning sessions.
- H4: The number of vocabulary words learned by children immediately after the study and 3 weeks after the study was in the order *Adaptive* > *Expert* > *Novice*.

H1: The results support the hypothesis. The majority of children successfully finished the two learning sessions without an experimenter's intervention, learned new words, and correctly recalled some learned words 3 weeks after the learning sessions. This result is supported in particular by the finding that even children with attention issues or learning difficulties learned and successfully recalled some new words. Another finding that strongly supports the hypothesis is that children with a low prior vocabulary level also improved their vocabulary after the study. Indeed, our results showed that children with a lower vocabulary level before the study showed significantly greater vocabulary improvement in both immediate post-test and 3-week-delayed post-test than children with a higher vocabulary level.

The result indicates that children are able to learn and stay engaged in this child-robot game, even if the vocabulary words are all very challenging. This result is not surprising, as the interaction is designed in a way to promote children's curiosity-driven learning, cultivate their growth mindset, and scaffold their learning. For example, the robot always gave positive feedback and show excitement when a child successfully retrieved an object; it encouraged the child to figure out a word's meaning by trial and error when the child didn't know the meaning of the word. Thus, children with a low prior vocabulary level felt much less intimidated by the game with the presence of a socially supportive robot; thus, they felt safe, comfortable, confident even though learning task was indeed very challenging. This interaction design explains why

children with low pre-test scores could learn vocabulary effectively but doesn't explain why children with high pre-test scores did not learn as well as the children with low pre-test scores. One explanation behind the difference in their learning performance is that some vocabulary words may be inherently harder to learn than others. From our results on children's acquisition of individual words (*section 4.3.3*), some target words (e.g., "garment") were harder to learn than others for all the children because only very few students successfully learned these words. Since the *high-vocabulary-level* children already knew easier target words (e.g., "lavender", "delighted"), the remaining target words they needed to learn were just inherently sophisticated words (e.g., "garment", "azure"), it was harder for them to improve the number of new words they learned than the children who didn't even know the easier target words (e.g., "lavender").

H2: The results confirmed the hypothesis. Children across the 3 conditions all felt happy when playing the game with the robot, as reflected in the results on their self-reported feeling. In addition, the majority of them thought it was more fun playing the game with the robot than playing the game alone and wanted to play the game with Tega again next time. Last, they perceived the game as *moderately easy* despite the fact that they didn't know most of the target words in the game. This result showed that the WordQuest system was able to provide a highly engaging, playful and joyful learning experience for young children even when the learning task itself is highly challenging to them.

H3, H4: The results support the hypotheses. While children in all conditions learned new vocabulary words and enjoyed the game, children who played with an *adaptive* robot improved more than children who played with either an *expert* or *novice* robot. The children in the *Adaptive* condition consistently outperformed the other two conditions across multiple measures. The trend analysis on children's immediate vocabulary score change and delayed score change shows a significant increase of vocabulary score among *Adaptive*, *Expert* and *Novice* conditions. In addition, nearly half of the top performing children who learned and recalled the most words were from the *Adaptive* condition. Why was a socially adaptive robot able to significantly improve children's vocabulary even when the majority of vocabulary words in the game were unknown to them? According to Vygotsky 1978, a learner will learn best when in the zone of proximal development where the learner is assisted by a teacher or peer with a higher skill set

[110]. When a robot adaptively switched its role between *novice* and *expert* based on children's current learning state, it controlled how much assistance it provided to the child to better ensure that the child received enough but not too much challenge. This theory may explain why an adaptive robot outperformed either the novice or the expert robots in helping young children learn vocabulary.

General Discussion: From the results, it is found that children with a range of learning abilities and vocabulary levels were able to effectively improve their vocabulary after playing the game with a social robot for only two 25-min sessions. In addition, they really enjoyed the game and found it moderately easy, even when they didn't know the majority of the vocabulary words prior to the study. This effect on children's learning performance was even more significant when they were interacting with a socially adaptive robot that switched its role between *expert* and *novice* during the game play. From the theory of flow, balancing challenge and capability is essential for achieving the optimal state of learning [111]. Csikszentmihalyi (1991) argues that a learner will become frustrated if experiencing too much challenge; conversely, the learner will become bored when not being challenged enough. Without the presence of a robot learning companion, the WordQuest game alone was too challenging for most of the children to play. Thus, the presence of the robot was to reduce the objective and subjective challenge of the game and build up the children's action capabilities. The *expert* role of the robot offered assistance, guidance and constructive feedback to children, so it helped reduce the learning difficulty of the game as well as facilitated children's acquisition of the vocabulary words. On the other hand, the *novice* role of the robot reduced the perceived difficulty that the children felt by showing a growth mindset, cheering the children up, making mistakes and asking the children for help. With such a *novice* learning companion, the children might have felt more comfortable, confident and patient in face of the challenge in the game. Furthermore, an adaptive robot was able to switch between the *novice* role and the *expert* role every child-robot turn based on the children's learning state to keep a dynamic balance between the child's perceived challenge and capability during the game play. Thus, an adaptive robot fostered children's learning from the theory of flow. Furthermore, Vygotsky's scaffolding theory also suggested that a child's learning outcomes are maximized when a task is challenging enough but not too difficult to do for the child [112]. An ability to move toward a state of better flow and to adaptively adjust the learning

task's difficulty level may be the reason why children in the *adaptive* condition outperformed other children.

Contributions

Regarding the technical contribution, I developed a child-robot WordQuest system that allows children to learn English words when playing a tablet-based learning game with a robot. This WordQuest system I created integrated game design, social robotics, and ideas from the science of early childhood's learning together to deliver a playful, engaging and personalized language learning platform for young children. Furthermore, the WordQuest system can be easily extended to incorporate other learning games or new sets of vocabulary words, and the WordQuest app can be used with and without the presence of a physical robot. Thus, this WordQuest system provides a scalable platform that can be used to do long-term studies in the future. Following this, I designed a robot's role-switching policy model based on the reinforcement learning model. This policy model dictates what role and behaviors the robot should exhibit based on the child's learning states. This behavior policy model can be easily adapted and applied to other child-robot interaction contexts.

In addition to the technical contribution, my work also sheds light on how to design educational technology for young children and child-technology interaction to foster their learning. The design of the WordQuest learning interaction was inspired by the field of learning science, and its educational value was also evaluated using the design principles for educational apps [27]. My work is the first that directly compares the differential impact of robot's roles on children's learning, as well as the first to study the effect of a robot's adaptive role-switching on young children's learning. It thus sheds light on how to better design a robot's adaptive roles in order to promote children's best learning and engagement. The study's results showed that the WordQuest system was able to teach children English words effectively. Children, even with signs of attention deficits, were able to learn English words in a 25-min game play, and remember some of the newly learned words three weeks after the study. This effect in learning performance is most significant when the robot could adaptively switch its role. Thus, my work

proposed a new child-technology learning paradigm that enables young children to learn new vocabulary words effectively and joyfully.

Future Work

Child-Robot Interaction Design

My current work consisted of only two 25-minute learning sessions per child. Future work plans to extend the child-robot game play to continue longer term. Thus, successfully sustaining children's engagement becomes critical for children's long-term learning, especially after the novelty effect of social robots on children fades out. A variety of ways can be employed to improve the WordQuest game design and child-robot interaction paradigm to make children's learning experience more playful and engaging. For example, the WordQuest game can be further improved in terms of its game mechanics, dynamics and aesthetics using Hunicke, LeBlanc, Zubek's MDA framework for game design [113]. As listed in the MDA framework, the *fantasy* and *narrative*, which refer to game as make-believe and game as drama respectively, are two important components to make a game more aesthetically appealing. These two elements create a more immersive game experience and make the game a more coherent story. These two components are currently lacking in the current version of the WordQuest game and can be added in future version. One way to add *fantasy* and *narrative* is described here. The robot and the child can be situated in a fantastic space where they play as aliens from another planet who are visiting the earth and learning about different aspects of people's daily life. This storyline will connect all individual learning sessions. There will thus be a series of themes for them to learn about the earth (e.g., animals and weather) and people's life (e.g., activities). Before the start of each learning sessions, a narrative will be added to inform the children of what theme today's session is, why they receive this theme and what they will learn out of it. At the end of each learning session, another narrative will be added to recap what the child and robot have learned during the session and what they will learn next time. Adding a storyline to the long-term child-robot WordQuest interaction will help keep the game engaging and playful to children.

Long-term Child-Robot Personalization

The current study consists of only two sessions, but personalization over two interactions limits the amount and variety of data that can be collected. To bring a significant impact to children's learning, the WordQuest learning interaction needs to become a long-term interaction with the same individual over multiple sessions and tasks. A long-term interaction between a child and a robot would also enable the robot to learn more about each individual child and adapt its interaction with the child to best promote the child's learning. Thus, I am planning to conduct a long-term study with multiple learning sessions in order to further explore how the robot's role-switching policy can adapt to individual children and analyze the impact of interacting with a personalized robot on children's long-term learning and engagement.

In addition, the Personal Robots Group has a suite of language learning games that allow a child and a robot to play together. Each of the game has a different interaction paradigm and addresses a different learning goal of a common curriculum. The WordQuest learning system is one of them, which helps build up children's vocabulary and cultivate the growth of their inferential thinking. Another one is the Tap game, in which a robot and a child race against each other to read out loud words popped up on the tablet. Although the learning goal for the Tap game is to improve children's pronunciation, it shares the same curriculum that the WordQuest game is using. This curricular similarity between games make it possible to build transfer models of users across educational games and achieve learning personalization both within and between games. For example, one way to achieve this task transfer is to use data from an individual's previous interactions to inform the robot's behaviors and roles during the same individual's interactions with new games. With a suite of learning games, I will be able to explore how to personalize how the robot should behave across a variety of games to promote each child's best learning over long-term interactions.

References

- [1] M. M. Páez, P. O. Tabors, and L. M. López, “Dual language and literacy development of Spanish-speaking preschool children.,” *J. Appl. Dev. Psychol.*, vol. 28, no. 2, pp. 85–102, Mar. 2007.
- [2] B. Hart and T. R. Risley, *Meaningful differences in the everyday experience of young American children*. ERIC, 1995.
- [3] R. R. Romeo *et al.*, “Beyond the 30-Million-Word Gap: Children’s Conversational Exposure Is Associated With Language-Related Brain Function,” *Psychological Science*, 2018.
- [4] J. K. Hartshorne, J. B. Tenenbaum, and S. Pinker, “A critical period for second language acquisition: Evidence from 2/3 million English speakers,” *Cognition*, May 2018.
- [5] J. Gilkerson *et al.*, “Mapping the early language environment using all-day recordings and automated analysis,” *Am. J. Speech-Language Pathol.*, vol. 26, no. 2, pp. 248–265, 2017.
- [6] D. Normile, “China’s childhood experiment.,” *Science*, vol. 357, no. 6357, pp. 1226–1230, Sep. 2017.
- [7] D. C. McCoy *et al.*, “Early Childhood Developmental Status in Low- and Middle-Income Countries: National, Regional, and Global Prevalence Estimates Using Predictive Modeling,” *PLoS Med.*, vol. 13, no. 6, 2016.
- [8] S. Grantham-McGregor, Y. B. Cheung, S. Cueto, P. Glewwe, L. Richter, and B. Strupp, “Developmental potential in the first 5 years for children in developing countries,” *Lancet*, vol. 369, no. 9555, pp. 60–70, 2007.
- [9] W. Bank, “Mozambique - Education Sector Support Project (English),” Washington, DC, 2011.
- [10] “Baby Genius.” .
- [11] “Sesame Street.” [Online]. Available: <https://www.sesamestreet.org/>.
- [12] “Blue’s Clues.” [Online]. Available: <http://www.nickjr.com/blues-clues/>.
- [13] “Dora the Explorer.” [Online]. Available: <http://www.nickjr.com/dora-the-explorer/>.
- [14] M. L. Rice, A. C. Huston, R. Truglio, and J. Wright, “Words from ‘Sesame Street’: Learning Vocabulary While Viewing.,” *Dev. Psychol.*, vol. 26, no. 3, pp. 421–428, 1990.

- [15] M. L. Rice and L. Woodsmall, "Lessons from television: children's word learning when viewing.," *Child Dev.*, vol. 59, no. 2, pp. 420–429, 1988.
- [16] M. L. Rice, J. C. Buhr, and M. Nemeth, "Fast mapping word-learning abilities of language-delayed preschoolers.," *J. Speech Hear. Disord.*, vol. 55, no. 1, pp. 33–42, 1990.
- [17] M. B. Robb, R. A. Richert, and E. A. Wartella, "Just a talking book? Word learning from watching baby videos," *Br. J. Dev. Psychol.*, vol. 27, no. 1, pp. 27–45, 2009.
- [18] J. S. DeLoache *et al.*, "Do babies learn from baby media?," *Psychol. Sci.*, vol. 21, no. 11, pp. 1570–1574, 2010.
- [19] G. Salomon, "Effects of Encouraging Israeli Mothers to Co-observe ' Sesame Street ' with Their," *Child Dev.*, vol. 48, no. 1971, pp. 1146–1151, 2013.
- [20] R. A. Reiser, M. A. Tessmer, and P. C. Phelps, "Adult-child interaction in children's learning from 'Sesame Street,'" *Educ. Commun. Technol.*, vol. 32, no. 4, pp. 217–223, 1984.
- [21] S. Ball and G. Bogatz, *The first year of Sesame Street: An Evaluation*. Princeton, NJ: Educational Testing Service, 1970.
- [22] G. A. Strouse, K. O'Doherty, and G. L. Troseth, "Effective coviewing: Ppreschoolers' learning from video after a dialogic questioning intervention," *Dev. Psychol.*, vol. 49, no. 12, pp. 2368–2382, 2013.
- [23] S. Vaala, A. Ly, and M. H. Levine, "Getting a read on the app stores: A market scan and analysis of children's literacy apps," New York, New York, USA, 2015.
- [24] H. L. Kirkorian, K. Choi, and T. A. Pempek, "Toddlers' Word Learning From Contingent and Noncontingent Video on Touch Screens," *Child Dev.*, vol. 87, no. 2, pp. 405–413, 2016.
- [25] A. R. Lauricella, T. A. Pempek, R. Barr, and S. L. Calvert, "Contingent computer interactions for young children's object retrieval success," *J. Appl. Dev. Psychol.*, vol. 31, no. 5, pp. 362–369, 2010.
- [26] D. A. Christakis, "Interactive media use at younger than the age of 2 years," *JAMA Pediatrics*, vol. 168, no. 5, pp. 399–400, 2014.
- [27] K. Hirsh-Pasek, J. M. Zosh, R. M. Golinkoff, J. H. Gray, M. B. Robb, and J. Kaufman, "Putting Education in 'Educational' Apps: Lessons From the Science of Learning," *Psychol. Sci. Public Interes. Suppl.*, vol. 16, no. 1, pp. 3–34, 2015.

- [28] C. Chiong and C. Shuler, “Learning: Is there an app for that? Investigations of young children’s usage and learning with mobile devices and apps.,” *Joan Ganz Cooney Cent.*, p. 34, 2010.
- [29] “One Laptop per Child.” [Online]. Available: <http://one.laptop.org/>.
- [30] T.-H. T. T.-H. Tan and T.-Y. L. T.-Y. Liu, “The mobile-based interactive learning environment (MOBILE) and a case study for assisting elementary school English learning,” *IEEE Int. Conf. Adv. Learn. Technol. 2004. Proceedings.*, pp. 4–8, 2004.
- [31] M. Kam, V. Rudraraju, A. Tewari, and J. Canny, “Mobile Gaming with Children in Rural India : Contextual Factors in the Use of Game Design Patterns,” in *Electrical Engineering*, 2007, pp. 292–301.
- [32] C. Breazeal, R. Morris, S. Gottwald, T. Galyean, and M. Wolf, “Mobile Devices for Early Literacy Intervention and Research with Global Reach,” in *Proceedings of the Third (2016) ACM Conference on Learning @ Scale - L@S '16*, 2016, pp. 11–20.
- [33] S. Gottwald, R. Morris, M. Wolf, and T. Galyean, “Bringing the Bottom Billion into Basic Literacy: How We Can and Why We Must,” *New Dir. Child Adolesc. Dev.*, vol. 2017, no. 158, pp. 93–104, 2017.
- [34] M. Wolf, S. Gottwald, C. Breazeal, T. Galyean, and R. Morris, “‘I Hold Your Foot:’ Lessons from the Reading Brain for Addressing the Challenge of Global Literacy Maryanne,” in *Children and Sustainable Development*, 2016, pp. 225–238.
- [35] S. Mitra, “Minimally Invasive Education for mass computer literacy,” *Conf. Res. Distance Adult Learn. Asia*, pp. 21–25, 2000.
- [36] S. B. Neuman and D. C. Celano, “Worlds apart: One city, two libraries, and ten years of watching inequality grow,” *Am. Educ.*, vol. Fall, pp. 13–23, 2012.
- [37] S. B. Neuman and D. C. Celano, *Giving our children a fighting chance: Poverty, literacy, and the development of information capital*. New York, New York, USA: Teachers College Press, 2012.
- [38] “Skype.” [Online]. Available: <https://www.skype.com/en/>. [Accessed: 09-May-2018].
- [39] “Facetime.” [Online]. Available: <https://itunes.apple.com/us/app/facetime/id414307850?mt=12>. [Accessed: 09-May-2018].
- [40] S. Roseberry, K. Hirsh-Pasek, and R. M. Golinkoff, “Skype Me! Socially Contingent Interactions Help Toddlers Learn Language,” *Child Dev.*, vol. 85, no. 3, pp. 956–970,

- 2014.
- [41] G. A. Strouse, G. L. Troseth, K. D. O’Doherty, and M. M. Saylor, “Co-viewing supports toddlers’ word learning from contingent and noncontingent video,” *J. Exp. Child Psychol.*, vol. 166, pp. 310–326, 2018.
 - [42] “Kindoma.” [Online]. Available: <http://kindoma.com/>. [Accessed: 09-May-2018].
 - [43] “VIP KID.” [Online]. Available: <https://www.vipkidteachers.com/>.
 - [44] K. Ryokai, C. Vaucelle, and J. Cassell, “Virtual peers as partners in storytelling and literacy learning,” *J. Comput. Assist. Learn.*, vol. 19, no. 2, pp. 195–208, Jun. 2003.
 - [45] K. Ryokai, C. Vaucelle, and J. Cassell, “Literacy learning by storytelling with a virtual peer,” *Proc. Conf. Comput. Support Collab. Learn. Found. a CSCL Community - CSCL ’02*, p. 352, 2002.
 - [46] J. Kory, “Storytelling with robots: Effects of robot language level on children’s language learning,” Massachusetts Institute of Technology, MA, 2014.
 - [47] M. Fridin, “Storytelling by a kindergarten social assistive robot: A tool for constructive learning in preschool education,” *Comput. Educ.*, vol. 70, pp. 53–64, Jan. 2014.
 - [48] H. W. Park, M. Gelsomini, J. J. Lee, and C. Breazeal, “Telling Stories to Robots,” *Proc. 2017 ACM/IEEE Int. Conf. Human-Robot Interact. - HRI ’17*, pp. 100–108, 2017.
 - [49] P. K. Kuhl, “Is speech learning ‘gated’ by the social brain?,” *Dev. Sci.*, vol. 10, no. 1, pp. 110–120, 2007.
 - [50] P. K. Kuhl, *Social Mechanisms in Early Language Acquisition: Understanding Integrated Brain Systems Supporting Language*. Oxford University Press, 2011.
 - [51] D. Leyzberg, S. Spaulding, M. Toneva, and B. Scassellati, “The Physical Presence of a Robot Tutor Increases Cognitive Learning Gains,” *34th Annu. Conf. Cogn. Sci. Soc.*, no. 1, pp. 1882–1887, 2012.
 - [52] S. Jeong, C. Breazeal, D. Logan, and P. Weinstock, “Huggable: Impact of embodiment on promoting verbal and physical engagement for young pediatric inpatients,” in *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2017, pp. 121–126.
 - [53] C. D. Kidd and C. Breazeal, “Robots at home: Understanding long-term human-robot interaction,” in *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS*, 2008, pp. 3230–3235.

- [54] M. Fridin and M. Belokopytov, “Embodied Robot versus Virtual Agent: Involvement of Preschool Children in Motor Task Performance,” *Int. J. Hum. Comput. Interact.*, vol. 30, no. 6, pp. 459–469, 2014.
- [55] J.-H. Han, M.-H. Jo, V. Jones, and J.-H. Jo, “Comparative Study on the Educational Use of Home Robots for Children,” *J. Inf. Process. Syst.*, vol. 4, no. 4, pp. 159–168, 2008.
- [56] “Jibo.” .
- [57] M. Gordon, E. Rivera, E. Ackermann, and C. Breazeal, “Designing a relational social robot toolkit for preschool children to explore computational concepts,” in *Proceedings of the 14th International Conference on Interaction Design and Children - IDC '15*, 2015, pp. 355–358.
- [58] L. Brown and A. M. Howard, “Engaging children in math education using a socially interactive humanoid robot,” in *IEEE-RAS International Conference on Humanoid Robots*, 2015, vol. 2015–Febru, no. February, pp. 183–188.
- [59] J. Kennedy, P. Baxter, and T. Belpaeme, “The Robot Who Tried Too Hard: Social Behaviour of a Robot Tutor Can Negatively Affect Child Learning,” in *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*, 2015, no. 801, pp. 67–74.
- [60] J. Han, M. Jo, S. Park, and S. Kim, “The Educational Use of Home Robots for Children,” in *2005 IEEE International Workshop on Robots and Human Interactive Communication*, 2005, pp. 378–383.
- [61] Eun-ja Hyun, So-yeon Kim, Siekyung Jang, and S. Park, “Comparative study of effects of language instruction program using intelligence robot and multimedia on linguistic ability of young children,” in *RO-MAN 2008 - The 17th IEEE International Symposium on Robot and Human Interactive Communication*, 2008, pp. 187–192.
- [62] J. Kanero, V. Geçkin, C. Oranç, E. Mamus, A. C. Küntay, and T. Göksun, “Social Robots for Early Language Learning: Current Evidence and Future Directions,” *Child Dev. Perspect.*, Jan. 2018.
- [63] G. Gordon, C. Breazeal, and S. Engel, “Can Children Catch Curiosity from a Social Robot?,” in *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction - HRI '15*, 2015, pp. 91–98.
- [64] H. W. Park, R. Rosenberg-Kima, M. Rosenberg, G. Gordon, and C. Breazeal, “Growing

- Growth Mindset with a Social Robot Peer,” *Proc. 2017 ACM/IEEE Int. Conf. Human-Robot Interact. - HRI '17*, pp. 137–145, 2017.
- [65] D. Kamps *et al.*, “Peer Training to Facilitate Social Interaction for Elementary Students with Autism and Their Peers,” *Except. Child.*, vol. 68, no. 2, pp. 173–187, Jan. 2002.
- [66] A. N. Garfinkle and I. S. Schwartz, “Peer Imitation,” *Topics Early Child. Spec. Educ.*, vol. 22, no. 1, pp. 26–38, Jan. 2002.
- [67] C. Chang, J. Lee, P. Chao, C. Wang, and G. Chen, “Chang, C.-W., Lee, J.-H., Chao, P.-Y., Wang, C.-Y., & Chen, G.-D. (2010). Exploring the Possibility of Using Humanoid Robots as Instructional Tools for Teaching a Second Language in Primary School.,” *Educ. Technol. Soc.*, vol. 13, no. 2, pp. 13–24, 2010.
- [68] J. Kennedy, P. Baxter, E. Senft, and T. Belpaeme, “Social robot tutoring for child second language learning,” in *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2016, pp. 231–238.
- [69] D. Matthews, E. Lieven, and M. Tomasello, “How toddlers and preschoolers learn to uniquely identify referents for others: A training study,” *Child Dev.*, vol. 78, no. 6, pp. 1744–1759, 2007.
- [70] N. Freed, “‘This is the fluffy robot that only speaks french’: Language use between preschoolers, their families, and a social robot while sharing virtual toys,” Massachusetts Institute of Technology, 2012.
- [71] G. Gordon *et al.*, “Affective Personalization of a Social Robot Tutor for Children’s Second Language Skills,” *Thirtieth AAAI Conf. Artif. Intell.*, Mar. 2016.
- [72] D. Hood, S. Lemaignan, and P. Dillenbourg, “When Children Teach a Robot to Write,” in *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction - HRI '15*, 2015, pp. 83–90.
- [73] F. Tanaka and S. Matsuzoe, “Children Teach a Care-Receiving Robot to Promote Their Learning: Field Experiments in a Classroom for Vocabulary Learning,” *J. Human-Robot Interact.*, vol. 1, no. 1, pp. 78–95, Aug. 2012.
- [74] T. Kanda, T. Hirano, D. Eaton, and H. Ishiguro, “Interactive robots as social partners and peer tutors for children: a field trial,” *Human-Computer Interact.*, vol. 19, no. 1, pp. 61–84, 2004.
- [75] E. Bates, “Individual Differences at 20 Months: Analytic and Holistic Strategies in

- Language Acquisition,” *J. Child Lang.*, vol. 10, no. 2, pp. 293–320, 1983.
- [76] D. Horgan, “How to answer questions when you’ve got nothing to say,” *J. Child Lang.*, vol. 5, no. 1, pp. 159–165, 1978.
- [77] G. CASTELLANO, I. LEITE, A. PEREIRA, C. MARTINHO, A. PAIVA, and P. W. MCOWAN, “Multimodal Affect Modeling and Recognition for Empathic Robot Companions,” *Int. J. Humanoid Robot.*, vol. 10, no. 1, p. 1350010, 2013.
- [78] S. D’mello and A. Graesser, “AutoTutor and affective autotutor,” *ACM Trans. Interact. Intell. Syst.*, vol. 2, no. 4, pp. 1–39, 2012.
- [79] J. B. Janssen, C. C. Van Der Wal, M. A. Neerinx, and R. Looije, “Motivating children to learn arithmetic with an adaptive robot game,” in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2011, vol. 7072 LNAI, pp. 153–162.
- [80] M. Ahmad, O. Mubin, and J. Orlando, “A Systematic Review of Adaptivity in Human-Robot Interaction,” *Multimodal Technol. Interact.*, vol. 1, no. 3, p. 14, Jul. 2017.
- [81] D. Leyzberg, S. Spaulding, and B. Scassellati, “Personalizing robot tutors to individuals’ learning differences,” in *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction - HRI ’14*, 2014, pp. 423–430.
- [82] T. Schodde, K. Bergmann, and S. Kopp, “Adaptive Robot Language Tutoring Based on Bayesian Knowledge Tracing and Predictive Decision-Making,” in *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction - HRI ’17*, 2017, pp. 128–136.
- [83] S. Spaulding, T. Aviv, and C. Breazeal, “Affect-Aware Student Models for Robot Tutors,” *Proc. 15th Int. Conf. Auton. Agents Multiagent Syst. (AAMAS 2016)*, pp. 864–872, 2016.
- [84] T. Schodde, K. Bergmann, and S. Kopp, “Adaptive Robot Language Tutoring Based on Bayesian Knowledge Tracing and Predictive Decision-Making,” in *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction - HRI ’17*, 2017, pp. 128–136.
- [85] A. N. Rafferty, E. Brunskill, T. L. Griffiths, and P. Shafto, “Faster Teaching via POMDP Planning,” *Cogn. Sci.*, pp. 1–43, 2015.
- [86] B. Clement, D. Roy, P.-Y. Oudeyer, and M. Lopes, “Multi-Armed Bandits for Intelligent Tutoring Systems,” Oct. 2013.

- [87] A. Bennane, “Adaptive Educational Software by Applying Reinforcement Learning,” *Informatics Educ. - An Int. J.*, vol. 12, no. 1, pp. 13–27, 2013.
- [88] M. Chi, K. Vanlehn, D. Litman, and P. Jordan, “Empirically evaluating the application of reinforcement learning to the induction of effective and adaptive pedagogical strategies,” *User Model. User-adapt. Interact.*, vol. 21, no. 1–2, pp. 137–180, 2011.
- [89] A. Iglesias, P. Martínez, R. Aler, and F. Fernández, “Learning teaching strategies in an Adaptive and Intelligent Educational System through Reinforcement Learning,” *Appl. Intell.*, vol. 31, no. 1, pp. 89–106, 2009.
- [90] P. Vogt, M. de Haas, C. de Jong, P. Baxter, and E. Krahmer, “Child-Robot Interactions for Second Language Tutoring to Preschool Children,” *Front. Hum. Neurosci.*, vol. 11, 2017.
- [91] L. Darling-Hammond *et al.*, *Powerful learning: What we know about teaching for understanding*. San Francisco: Jossey-Bass, 2008.
- [92] R. G. Tharpe and R. Gallimore, *Rousing Minds to Life: Teaching, Learning, and Schooling in a Social Context*. New York, New York, USA: Cambridge University Press, 1988.
- [93] R. Ros *et al.*, “A Motivational Approach to Support Healthy Habits in Long-term Child–Robot Interaction,” *Int. J. Soc. Robot.*, vol. 8, no. 5, pp. 599–617, 2016.
- [94] J. K. Westlund and C. Breazeal, “The Interplay of Robot Language Level with Children’s Language Learning during Storytelling,” in *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction Extended Abstracts - HRI’15 Extended Abstracts*, 2015, pp. 65–66.
- [95] “SpeechAce.” [Online]. Available: <https://www.speechace.com/>.
- [96] D. Mak and D. Nathan-Roberts, “Design considerations for educational mobile apps for young children,” in *Proceedings of the Human Factors and Ergonomics Society*, 2017, vol. 2017–Octob, pp. 1156–1160.
- [97] J. Gros-Louis, M. J. West, and A. P. King, “Maternal responsiveness and the development of directed vocalizing in social interactions,” *Infancy*, vol. 19, no. 4, pp. 385–408, 2014.
- [98] C. Chiong and J. S. DeLoache, “Learning the ABCs: What kinds of picture books facilitate young children’s learning?,” *J. Early Child. Lit.*, vol. 13, no. 2, pp. 225–241, 2013.
- [99] J. Parish-Morris, N. Mahajan, K. Hirsh-Pasek, R. M. Golinkoff, and M. F. Collins, “Once

- upon a time: Parent-child dialogue and storybook reading in the electronic era,” *Mind, Brain, Educ.*, vol. 7, no. 3, pp. 200–211, 2013.
- [100] R. Barr, L. Shuck, K. Salerno, E. Atkinson, and D. L. Linebarger, “Music interferes with learning from television during infancy,” *Infant Child Dev.*, vol. 19, no. 3, pp. 313–331, 2010.
- [101] R. Sutton and A. Barto, *Introduction to Reinforcement Learning*, 1st ed. Cambridge, MA, USA: MIT Press, 1998.
- [102] R. S. Sutton and A. G. Barto, *Reinforcement learning : an introduction*, vol. 9, no. 5. 2013.
- [103] R. I. Brafman and M. Tennenholtz, “R-MAX - A general polynomial time algorithm for near-optimal reinforcement learning,” in *IJCAI International Joint Conference on Artificial Intelligence*, 2001, pp. 953–958.
- [104] C. WATKINS and P. Dayan, “Q-LEARNING,” *Mach. Learn.*, vol. 8, no. 3–4, pp. 279–292, 1992.
- [105] J. K. Westlund *et al.*, “Tega: A social robot,” in *ACM/IEEE International Conference on Human-Robot Interaction*, 2016, vol. 2016–April, p. 561.
- [106] “Affdex SDK.” [Online]. Available: <https://developer.affectiva.com/metrics/>.
- [107] J. M. Kory Westlund, H. Park, R. Williams, and C. Breazeal, “Measuring Young Children’s Long-term Relationships with Social Robots,” in *In Proceedings of the 17th ACM Interaction Design and Children Conference (IDC)*, 2018.
- [108] M. J. Davis, “Contrast coding in multiple regression analysis: Strengths, weaknesses, and utility of popular coding structures,” *J. Data Sci.*, vol. 8, pp. 61–73, 2010.
- [109] R. E. Mayer, “Should There Be a Three-Strikes Rule against Pure Discovery Learning? The Case for Guided Methods of Instruction,” *American Psychologist*, vol. 59, no. 1. pp. 14–19, 2004.
- [110] L. S. Vygotsky, “Mind in society: The development of higher psychological processes,” *Mind Soc. Dev. High. Psychol. Process.*, vol. Mind in So, p. 159, 1978.
- [111] M. Csikszentmihalyi, “Flow: The psychology of optimal experience: Steps toward enhancing the quality of life,” *Des. Issues*, vol. 8, no. 1, p. 80, 1991.
- [112] L. Vygotsky, “Interaction between learning and development,” in *Readings on the Development of Children*, 1978, pp. 34–40.

- [113] R. Hunicke, M. LeBlanc, and R. Zubek, “MDA: A Formal Approach to Game Design and Game Research,” *Work. Challenges Game AI*, pp. 1–4, 2004.