PopBots

Leveraging Social Robots to Aid Preschool Children's Artificial Intelligence Education

by

Randi Williams

B.S. Computer Engineering
University of Maryland, Baltimore County, 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

June 2018

© MIT 2018

Author **Signature redacted**
_____
Program in the Media Arts and Sciences
May 18, 2018

Certified by **Signature redacted**
_____
Cynthia Breazeal
Associate Professor of Media Arts & Sciences
Thesis Supervisor

Accepted by **Signature redacted**
_____
Tod Machover
Professor of Media Arts and Sciences
Academic Head, Program in Media Arts and Sciences

PopBots:
Leveraging Social Robots to Aid Preschool Children's Artificial Intelligence
Education

by
Randi Williams

Abstract

Today's children are growing up with artificial intelligence (AI) devices such as voice personal assistants, home robots, and internet connected "smart" toys. In previous research, we have seen that children lack understanding of how modern AI devices work, making it difficult for them to engage in reflective and constructive interactions with the AI-enabled technology (Druga, Williams, Breazeal, & Resnick, 2017). This thesis explores how young children explore and create with AI, and how such activities influence children's perceptions of AI and their attitudes about themselves as engineers.

First, I discuss the design of *PopBots* -- the first hands-on toolkit developed for children ages 4-6 to explore and learn about AI. The social robot serves as both a programmable artifact as well as a window into understanding the machine learning algorithms. Accompanying this toolkit, I also developed a novel, developmentally-appropriate Preschool-Oriented Programming (POP) curriculum. The PopBots curriculum expands existing computational thinking curriculums by using creative learning activities to teach children three core AI concepts: rule-based systems, generative AI, and supervised machine learning.

Next, I evaluated the PopBots toolkit and curriculum with 80 pre-K and Kindergarten aged children from local schools. I found that young children can understand most of the AI concepts presented in the toolkit, but sometimes developmental factors like grade and Theory of Mind skills made a difference. After completing the PopBots curriculum, children developed an understanding of robots as "learning" machines. They also gained confidence in their ability to build their own robots. Overall, this work provided a highly engaging opportunity for children to explore robotics, AI and programming -- and ultimately see AI-based technology as something they can play a role in not just using but also creating.

PopBots

Leveraging Social Robots to Aid Preschool Children's Artificial Intelligence Education

by

Randi Williams

The following people served as readers for this thesis:

Signature redacted

Advisor: _____   _____
Cynthia Breazeal
Associate Professor of Media Arts & Sciences
MIT Media Lab

Signature redacted

Reader: _____   _____
Marina Bers
Professor of Child Study and Human Development
Tufts University

Signature redacted

Reader: _____   _____
Paul Harris
Professor of Education
Harvard Graduate School of Education

4

Acknowledgements

> We should always have three friends in our lives. One who walks ahead who we look up to and follow; one who walks beside us, who is with us every step of our journey; and then, one who we reach back for and bring along after we've cleared the way.
>
> – Michelle Obama

To Dr. Cynthia Breazeal

I cannot believe my luck in having you as an advisor! Everything about the past two years has been beyond my wildest dreams. When you interviewed me, I never imagined I would work in a robotics lab. When we were first discussing my project, I never imagined that young children could learn AI and I never imagined that I would build something so that they could. I never imagined that being a grad student could be so fulfilling for me as a person. Thank you for being an amazing advisor, I look forward to the PhD journey and developing a much more creative imagination.

To Dr. Marina Bers and Dr. Paul Harris

Thank you for choosing to be on my team; I am so fortunate to have had you both. I chose you as my readers because I am a huge fan of your minds as well as your work. I believed that you would help stretch and shape and stretch this thesis into something meaningful. You more than delivered. Thanks yo you, this work and my skills as researcher have evolved on multiple levels. Thank you!

To my family

Parents, thank you for life: both the one you gave me at the very beginning and the one you have nurtured every day of my life. My family and friends, thank you for listening, cheering me on, and feeding me on those days when I was struggling. I appreciate you all more than I could ever express.

To the Media Lab

I have the best research group and the best department admins in the world! Thank you to the Media Lab for welcoming, supporting me, pushing me, and teaching me to see the world in infinite dimensions.

To Prince George's County , Baltimore City, and D.C.

This thesis is dedicated to the communities that shaped me. Thank you for teaching your children to be strong, to dream bigger dreams, and to "reach back" for others. You inspire me.

Table of Contents

1 Introduction

Artificial intelligence (AI) is revolutionizing many industries, from healthcare and government to education and entertainment. As this technology becomes increasingly ubiquitous, it promises to change how we live, work, and play. The market for consumer AI technology (e.g., voice personal assistants, educational devices, and Internet connected smart toys) is rapidly growing. An estimated 31.2 million units will be sold in 2020, up from 6.5 million in 2015 (Tractica, 2015). Despite the explosive growth of AI in the public sector, very few consumers understand it. Qualtrics conducted a survey and found that only 10% of Internet users consider themselves experts of AI (eMarketer, 2017). In comparison, NPR and Edison Research estimated that 16% of Americans have a smart speaker in their home (Perez, 2018) and Verto Analytics found that 24% of the U.S. population use an AI-based agent in their smartphone at least 10 times per month (Hwong, 2017).

Organizations like OpenAI and Diversity.AI have already identified concerns with leaving AI knowledge in the hands of the few. For example, researchers have found that the lack of diversity in datasets has led to systems that are discriminative towards individuals based on their race, sex, income, age, and a number of other factors (Buolamwini & Gebru, 2018; Bolukbasi, Chang, Zou, Saligrama & Kalai, 2016; Neil, 2017). In order to fix this problem, we have to think about the AI pipeline: who has the skills to create technology in this field? Currently, the numbers are bleak. In 2017, a company named Tencent estimated that there are about 300,000 people in the *world* who have the skills necessary to work in the AI field (Tencent, 2017). If diversity in AI mirrors today's tech scene in Silicon Valley, then this group includes 75,000 women, 42,000 Asian Americans, and 45,000 African and Hispanic Americans (Ashcraft, Mclain, & Eger, 2016; Gee & Peck, 2016).

> There is a bias to what kinds of problems we think are important, what kinds of research we think are important, and where we think AI should go. If we don't have diversity in our set of researchers, we are not going to address problems that are faced by the majority of people in the world.
>
> - Timnit Gebru
> ""We're in a diversity crisis": cofounder of Black in AI on what's poisoning algorithms in our lives", 2018

Thousands of people cannot build technology that equitably addresses the concerns of billions. If AI is to avoid the pitfalls of bias, then it's important that we democratize AI now. When anyone can learn about and use AI in creative ways, then AI can become a tool for positive change.

How do we start changing the face of AI? One way is by empowering the youngest members of society, the next generation of technologists. First, it's important that we recognize that children growing up in the era of artificial intelligence will have a fundamentally different relationship with technology than us. They are growing up interacting with intelligent technologies at a much younger age and at the same time that they are learning to interact with other people. Reading and writing used to be the gateways to the Internet, but now new modalities like gesture, touch, and speech allow much younger children to access digital content and services.

Children tend to see intelligent machines as something in-between alive and not alive, smart and not smart. In 1980 through to the 2000, researchers observed children's interactions with interactive toys like Merlin, a tic tac toe AI, and Speak and Spell, a talking toy that quizzed children on spelling. These technologies were fairly straightforward; they had a limited number of preprorammed utterances and games. However, even this level of intelligence was sufficient for children to engage in debates about the perceived aliveness, intelligence, and emotions of the machines.

> Marginal objects, objects with no clear place, play important roles. On the lines between categories, they draw attention to how we have drawn the lines...Computers, as marginal objects on the boundary between the physical and the psychological, force thinking about matter, life, and mind. Children use them to build theories about the animate and the inanimate and to develop their ideas about thought itself.
>
> – Sherry Turkle
> *The Second Self*, 1984

In the mid 2000s through early 2010, studies continued with new toys like Tamagotchi and Furby. These toys intentionally elicited nurturing behaviors. Teaching and caring for digital creatures led to children creating emotional bonds. They not only treated the devices as alive, but also said the devices loved and cared for them (Turkle, 2006). Children's attachment to more intelligent robots was also seen in research with humanoid robots such as Kismet, Cog, and Robovie. In these studies, children saw the robots as more emotionally and intellectually capable than they actually were. For example, children would find meaning in Kismet's nonverbal babbling and regarded Robovie as a moral being worthy of having rights.

Modern AI-enabled devices like voice personal assistants and smart toys blur the line between living and nonliving objects even more. These Internet-connected devices are a far cry from interactive toys of a decade ago. They can assess a vast array of information, content and services. They have state of the art speech recognition and natural language understanding and learn from experience

to continue to improve their accuracy and robustness. They can listen and speak to you from across the room, recognize who you are, answer a of questions through Internet search, perform tasks connected to digital services like access your calendar, tell you the weather, shop, tell jokes and stories, as well as convey their own opinions. They are connected to developer ecosystems so that the number of capabilities (skills) these devices can perform is increasing at a rapid pace. The voice assistant, Alexa, boasts around 10,000 skills to date.

Digital voice-enabled assistants are not only AI-enabled but also personified. These devices are intentionally designed with personalities to be perceived as ever more charming, lifelike, and intelligent. It is extremely likely that children will form very strong social bonds with these personified intelligent technologies that behave more like cognizant entities than machines. Therefore, it is imperative that we understand how AI may affect children -- now and in the future -- by understanding what children think about machines that think.

Public discussion about AI in the media and movies. How parents talk about AI to their children has dramatically changed as these technologies have become increasingly commonplace. For example, Aristotle, a product developed by Mattel, raised a lot of concerns from parents about smart toys and how they affect children's privacy and development (Campaign for a Commercial Free Childhood, 2017). Aristotle's goal was to be a parent's assistant and a child's companion with apps for lullabies, educational content, and child monitoring. However, parents were afraid of placing a surveillance device in their child's bedroom that would potentially allow a company to constantly gather data on a child. Parents were also concerned about the future implications of a child bonding with a "digital nanny.

In studies with children and voice personal assistants like Siri and Alexa and smart toys like Hello Barbie and Cognitoys Dino, researchers observed that children's interactions with the devices raised concerns in parents. For example, children would treat Siri and Alexa, not intended as toys for children, as both an information source and a companion (Lovato & Piper, 2015; McReynolds, Hubbard, Lau, Saraf, Cakmak, & Roesner, 2017). However, children would also abuse the devices, telling them mean things or rudely giving orders. Also, children were often unaware that their conversations were recorded and could be accessed by their parents and others later (McReynolds et al., 2017).

In our own studies with children and AI, we saw that children's perception of intelligent machines is strongly influenced by their preconceptions and individual interactions with devices (Druga, Williams, Breazeal, & Resnick, 2017). With respect to their preconceptions, we found that children's understanding of AI is influenced by media and conversations with their parents (Druga, Williams,

Park, & Breazeal, 2018). However, they lack technical knowledge, and their inability to understand the "minds" of AI devices leads to quick and faulty assumptions about the device's nature. In addition, the devices do little to make their inner workings transparent to users.

> "Alexa, what do sloths eat?" asked Jane[1]. The voice personal assistant replied, "I'm sorry. I don't know how to help you with that." Disappointment flashed across Jane's face, then she perked up again. "That's okay," she exclaimed, picking up a second Amazon Echo Dot, "I'll see if the other Alexa knows."
>
> – Paraphrased from Druga et al.
> "Hey Google, Is It OK If I Eat You?", 2017

Given this cultural context, we revisit the question of how children think about thinking machines, and how this can change not only through interaction -- but even by programming and playing with their creation. Children growing up with smart devices need to understand how they work and how to use them safely. However, there exists few tools and curricula for non-experts to learn about AI.

For the most part, AI curricula only exist for people at the high school and college level, or beyond. On the Internet, a number of online education platforms like Udacity, Coursera, and Udemy host courses for artificial intelligence topics. In addition, universities like MIT have free recordings of lectures and open problem sets for interested students. These courses generally cover reasoning, planning, search, constraints, supervised learning, unsupervised learning, representation, inference, and neural nets. Tools such as Keras, Amazon Web Services, IBM Watson, and Google Developers make it possible for anyone to leverage powerful computers to develop their own AI-related projects. Companies like Google and NVIDIA also have courses focused on producing more engineers that can work on the problems that the companies face. These courses focus on deep learning and high performance computing.

On the university level, 96 colleges or universities offer AI classes. These classes cover traditional AI topics applied to a different mediums such as simple LEGO robots, software simulation, and commercial (Kumar, 2004; McNally & Klassner, 2007; Koski, Kurhila, & Pasanen, 2008; Talaga & Oh, 2009; Burgsteiner, Kandlhofer, & Steinbauer, 2016). Research in this area has explored how these different mediums contribute to student's learning process. For example, programming within the bounds of simulation allows students to focus on the implementation of software, but programming on a robot has led to higher engagement and deeper understanding of the necessity of AI.

---

[1] Names of all children have been changed.

However, none of these options are appropriate for children or non-experiences. However, since Seymour Papert's LOGO Turtle in 1980, researchers have recognized the importance of making computational thinking skills something for everyone (Papert, 1980). This has led to a number of platforms that make computational thinking accessible to non-experts, from AppInventor, for developing mobile applications with a block-based language, to Scratch and ScratchJr., block-based programming platforms for children, to LEGO Mindstorms, a robotic platform for children ages 12 and up (Wolber, Abelson, Spertus, & Looney, 2014; Resnick, Martin, Berg, Borovoy, Colella, Kramer, & Silverman, 1998; Flannery, Silverman, Kazakoff, Bers, Bonta, & Resnick, 2013). Some of these platforms have been adapted to teach AI. For example, higher education classrooms use LEGO Mindstorms in their curriculum to teach path planning and simultaneous localization and mapping (SLAM) algorithms (McNally & Klassner, 2007; Talaga et al., 2009). In the Personal Robots Group, the Cognimates project allows children and parents to program a number of intelligent devices such as Jibo and Alexa (Druga et al., 2018). Finally, a software engineer and father in the UK designed a set of activities using Scratch and IBM Watson to help children learn about AI (Lane, 2016). Still, none of these platforms were designed for preschool-aged children.

1.1 Contributions of the Thesis

An important contribution of this thesis is the design iterations, development, and evaluation of a novel platform comprised of a developmentally appropriate toolkit, AI artifacts, hands-on activities, and assessments -- collectively called the *Preschool-Oriented Programming (POP) Platform and Curriculum*. This unique early AI education suite enables parents and educators to help young children, ages 4-6, explore and learn about how machines can make decisions, learn, and improvise.

For AI to be understandable to young children, opaque and abstract algorithms have to become hands-on and concrete. The POP Platform leverages a social robot, *PopBot*, as an AI artifact that children can code, train, and interact with to aid in their sense-making about AI concepts. Through direct interaction with the PopBot, children can see how robots sense, think, act and now, also learn. Social robots are a powerful platform for education for young children. They have been developed for a wide range of educational activities from early literacy, second language learning and STEM. More recent work has developed social robots to engage children as collaborative learning companions that can engage children using social and cognitive cues to create personalized learning experiences (Gordon, Spaulding, Westlund, Lee, Plummer, Martinez, Das, & Breazeal, 2016; Kennedy, Baxter, Senft, & Belpaeme, 2016; Keren & Fridin, 2014; Short, Swift-Spong, Greczek, Ramachandran, Litoiu, Grigore, & Scassellati, 2014; Breazeal, Harris, Deseteno, Westlund, Dickens, & Jeong, 2016; Gordon, Ackermann, & Breazeal, 2015; Gordon, Breazeal, & Engel, 2015; Kory & Breazeal, 2014). In this

curriculum, the PopBot helps children digest concepts as a friendly coach, leveraging dialogue to guide children and explain the algorithms that are running in the robot's "mind" (Jung, Martelaro, Hoster, & Nass, 2014). The social robot's behavior, mind, and actions are used to give children feedback for algorithms that are otherwise too abstract or opaque for them to grasp.

PopBots are programmed with a picture-based programming interface (for pre-readers) that encourages learning through exploration and coding. It is built on top of the Blockly Programming Language (Fraser, 2013). This constructionist approach allows children to build their understanding by reflecting on their own cognition (Bers, 2018; Resnick, 2017; Papert, 1980; Gordon, Ackermann, & Breazeal, 2015; Flannery et al., 2013). The use of creative learning activities such as storytelling, music, and games makes the platform flexible enough to appeal to children with a wide range of interests (Bers, 2008; Resnick, 2017, Flannery et al., 2013).

--

In the process of iteratively designing, developing, field testing and evaluating the effectiveness of the POP Platform and Curriculum for children 4-6 years old, I will answer the following questions:

1. How do young children come to perceive and understand AI concepts? What factors influence this process? (Chapter 2 & 7)
2. Is a social robot platform an effective way to deliver developmentally appropriate curriculum for young children to learn about AI concepts. What features support children's exploration, perception and understanding? (Chapter 3 & 7)
3. What can children learn about key AI concepts such as rule-based systems, supervised machine learning, and generative AI by coding and interacting with a social robot platform? How does this inform a developmentally appropriate curriculum for AI and its attributes? (Chapter 4 & 6)
4. What teaching methods are most effective for engaging and supporting young children with an early AI curriculum in classrooms? (Chapter 5 & 7)
5. How does learning about AI change children's perception of artificially intelligent devices and themselves as engineers? (Chapter 7)

I hope that this work and the POP Platform and Curriculum will help others make strides toward helping children learn about and have constructive interactions with the AI-enabled technologies around them.

2 Background and Prior Work

The questions posed by this work have roots in the fields of human-computer interaction, developmental psychology, computer science education, and gender equality in STEM. Central questions of this work are how children growing up with smart machines develop their understandings of them and, perhaps more importantly, how those understandings will develop over time. To find answers, we look at the methods and discoveries of studies in child-computer interaction over the past 40 years. These studies suggest that children relate to intelligent machines as living objects, items with a dual nature of inanimacy and consciousness. However, today's intelligent machines are quite opaque which raises questions about whether children may be negatively affected by interacting with them.

My approach to addressing this problem is to design a toolkit to help children develop an understanding of the technology around them, starting with young children who are developing their understanding of people at the same time as these intelligent machines. To create a developmentally appropriate curriculum for children ages 4 to 6, I take several cues from past computational toolkits developed for young children and from AI courses developed for older students. Given that this work focuses on young children and their psychological understandings of machines, I also look at how cognitive perspective-taking develops in children between the ages of 4 to 6-years-old and how this development may affect their understanding of AI in the curriculum. Finally, after using the curriculum, I expect that children will have a new perspective on themselves as engineers. This section explores prior work in all of these themes and how they shaped this thesis.

2.1 What Do Children Think about Thinking Machines

There is no simple answer to how children perceive artificially intelligent devices. Children's perceptions are dependent on their individual interactions with the devices, such as whether or not it says their name correctly, as well as social and cultural factors, like how those around them talk about the technology (Druga et al., 2017, Druga et al, 2018). It is clear that as children gain exposure and understanding of technology, their reasoning about these devices becomes more thoughtful and nuanced (Bernstein & Crowley, 2008). In this section, I examine the relationship between direct experience, knowledge, and natural development on how young children understand these machines. In particular, I look at how children's perceptions of robots may change after completing

the PopBots curriculum. As children begin to grow up with intelligent machines, our understanding and approach to deciphering how children think about them will have to continually develop.

Since the first studies analyzing children's attributions of animacy to smart toys in 1984, digital machines have evolved many times, and so have children and the way they interact with these machines. Home computers entered the consumer market in 1977, so children in the 1980s were interacting with them at a time when they were just starting to become widespread. Today, we have robots in our homes and schools, children within their first year of life are able to interact with tablets, small computers with magnitudes more computing power than the first computer. Despite these changes, some themes in child-computer remain consistent. Computers and intelligent machines are still objects of ambiguity with the power to engage children in debate about the nuances of being alive, children still anthropomorphize and personify these machines, children still display a range of probing behaviors towards these machines from nurturing to abuse, and children's perception of robots develops with as the child gets older and gains more experience with the machine.

In 1984, Sherry Turkle studied children's interactions with smart toys like Simon, Speak-and-Spell, and Merlin. She observed that smart toys turned children into "child philosophers" (Turkle, 1984). Children tend to see intelligent artifacts as somewhere between alive, placing them amongst animals and humans in terms of intelligence but amongst other inanimate objects in terms of aliveness (Bernstein & Crowley, 2008).

> "There is a constellation of attributes that children ascribe to robots which do not appear to mirror reasoning about canonical living entities."
> - Peter Kahn, "The New Ontological Category Hypothesis in Human-Robot Interaction", 2011

The toys Turkle studied were only intelligent enough to play simple games, however, children would still perceive mental and emotional states in the devices and therefore question whether or not the device was alive. Similarly, Jipson and Gelman explored the attributions that 3-5-year-old children compared to adults would make to items such as as dog, a robot dog, a starfish, a stuffed animal, a sensor box, and a car. They saw that having a face made a difference in whether children would refer to the item with a pronoun, and indicator that they regarded the item as a being (Jipson & Gelman, 2007). Three-year-old children, but not their older peers, would then use the presence of a face, intentional motion, and intelligence as factors to judge that the robot dog then must be alive. Intelligent machines break the rules and make it possible for abstract concepts like "aliveness" to be reasoned about in a concrete way (Ackermann, 2005; Beran, Ramirez-Serrano, Kuzyk, Fior, &

Nugent, 2011). The lifelike nature of machines makes them a useful metaphor for children to reflect on not only on how machines think, but on how they themselves think (Minsky, 2009). However, it's important to understand that children tend to form relationships with these machines and cannot necessarily analyze intelligent machines with impartiality. Regardless of attributions of aliveness, children see intelligent artifacts as social entities and become attached (Ackermann, 2000; Turkle et al., 2003; Turkle, 2006; Kahn, 2012).

Children ascribe emotional states, mental states, intentionality, and morality to intelligent artifacts (Turkle, 1984; Scaife & van Duuren, 1995; Turkle et al., 2003; Turkle, 2006; Bernstein & Crowley, 2009; Melson, Kahn, Beck, Friedman, Roberts, Garrett, & Gill, 2009; Severson & Carlson, 2010; Oh & Kim, 2010; Fior, Nugent, Beran, Ramirez-Serrano, & Kuzyk, 2010; Kahn, Kanda, Freier, Severson, Gill, Ruckert, & Shen, 2012; Borenstein & Pearson, 2013). In Turkle's 1984 studies children said that the tic tac toe AI, Merlin, was "cheating" when it won too often or when a bug prevented them for making their move. The children believed that the toy both had emotions, motivation to win, and a will to purposely break the rules in pursuit of winning (Turkle, 1984). In 2003, Sherry Turkle, Cynthia Breazeal, Olivia Dasté, and Brian Scassellati observed how children interacted with two robots, Kismet and Cog. Kismet is a social robot designed with an emotionally expressive face and voice. It responds to human voices like an infant would.. Cog is an enormous humanoid robot with dextrous limbs and several degrees of freedom in its head and neck. Sixty children between the ages of 8 and 13 interacted with these two robots for about 15 minutes and quickly developed relationships with them. Researchers observed that children would relate to the robot in human terms, gendering the robots and labelling even nonsense sounds as higher order behavior like cooing in approval. Children projected mental and emotional states onto the robots. For Kismet especially, children would talk about how they could tell the robot understood them, cared about them, and, most interestingly, loved them. Even children's understanding of the robots as mechanical beings did not preclude their emotional connection. Researchers made a point to show the children how the robot was built and programmed and occasionally, the robot would break down in front of the child and a researcher would have to repair it. Still, children would relate to the robot like they would another person. Children's persistence in relating to robots as social beings despite being aware of their mechanical nature was also observed in children's interactions with popular smart toys in the 90s like My Real Baby and Furby and with today's voice personal assistants and smart toys (Turkle, 2006; Kahn et al., 2012; McReynolds et al., 2017; Druga et al., 2017). Although adults anthropomorphize inanimate objects too, for children the distinction between animate and inanimate is more ambiguous. Children nine and under attribute intelligent machines with mental and social characteristics that

older children and adults only give to people and animals (Severson & Carlson. 2010). This makes it unclear whether children truly believe that these objects are alive or not.

In favor of children just "pretending" these objects are alive, children may just treat robots as people in an attempt to understand them. We witnessed children's probing behaviors when interacting with intelligent machines (Druga et al, 2017). Children asked the robots questions like "Do you want an apple?" and "What are you?" just to see the device's response. Children use everything from asking lots of questions, to treating the robots like children, to even abuse to see how the robot will react (Turkle, 1984; Turkle et al, 2003; Turkle, 2006; Lovato & Piper, 2015; Nomura, Kanda, Kidokoro, Suehiro, & Yamada, 2017; Druga et al., 2017). As children gain experience interacting with this technology they develop their understanding of them (Scaife & van Duuren, 1995; van Duuren & Scaife, 1995; Severson & Carlson., 2010; Kahn et al., 2012; Druga et al., 2017). Mioduser and Levy (2007) conducted a study where they asked Kindergarten children to describe the actions of robot as it did increasingly complex tasks. When the robot was doing simple tasks children would use more mechanical descriptions of the robot "it's driving forward, then turning left." However, as the tasks became more complex children used more psychological descriptions "it wants to go to the light" of the robot's behavior. In a later study, Levy & Mioduser (2010) let children program the robots. They saw that even as behaviors became more complex children maintained giving technical descriptions of the robot's behavior. However, when describing robots they didn't program they gave more psychological descriptions. Teaching young children about programming changes the way they understand intelligent machines even if it does not necessarily change the way they talk about them.

Children continue to relate to intelligent machines as social beings today. Therefore, tech literacy is increasingly important. We have seen that children do not understand how smart toys and other intelligent technologies work. A study on children's interactions with smart toys found that children would tell robots personal information not realizing that their conversations with the toy were recorded by the toy (McReynolds et al., 2017). Parents in the study were disturbed by their children's trust in the toys and expressed interest in regulations that made privacy and permissions more clear. In our studies we saw that children were extremely trusting of smart toys (Druga et al., 2017). Our hope is that parents and educators can help their children gain a healthy awareness of how technology works. We have seen that children intuitive understandings of technology like robots, but these understandings can be shaped by their parents (Druga et al, 2018). Parents and educators need to help children become informed about this technology so that they can guide them towards having healthy relationships with it.

In this study, I am looking at how learning about AI changes children's perceptions. I expect that, at first, the young children will relate to robots as social and intellectual beings. After learning more

about how robot's minds work, children will see that robots are intellectual and social but in ways that are different from how people. I expect that children's attributions of intelligence and human-likeness to robots will change after the POP curriculum.

2.2 Young Children's Computational Thinking Platforms

There has been an increased emphasis on science, technology, engineering, arts, and math (STEAM) education in K-12 classrooms. The benefits of learning computational thinking at a young age range from increased metacognitive abilities (Clements & Gullo, 1984) to increased confidence in becoming an engineer (Sullivan, 2016).

> "Everyone should know how to program a computer because it teaches you how to think."
>
> > \- Steve Jobs

However, the benefits of teaching computational thinking to children were not always recognized. The general sentiment in the early 1970s was that computer programming was too advanced for children. Still, in 1974 Radia Perlman persisted and developed a computational thinking platform, TORTIS, to teach children about numbers, decomposition, debugging and "most important of all...that learning is fun" (Perlman, 1974). Since then, computational thinking platforms have had much, much success as commercial and research platforms. A number of platforms exist for children ages 4+ that include robots (e.g., *BeeBot, Cubetto, Think and Learn Code-A-Pillar, Sphero, Ozobot, Coji, Dash and Dot, LOGO turtle, RoboGan, Root,* and *KIBO*), educational apps (*ScratchJr., Lightbot, Kodable, the Foos,* and *LOGO*), websites (e.g., *Code.org, Scratch,* and *Tynker*), board games (e.g., *Robot Turtles*), and hardware (e.g., *Cubelets, Electronic Blocks,* and *Topobo*). Through these different mediums, children learn computational thinking skills such as sequencing, conditionals, decomposition, loops, sensing, debugging, assembling hardware, and functions. To make this accessible to young children, the platforms employ a number of engaging activities like playing games, solving puzzles, telling stories, making music, or creating artifacts.

One of the earliest computational thinking platforms for children was Seymour Papert's *Logo* programming language and *turtle robot* (Papert, 1980). The turtle robot's superpower was that it was the same size as the children, allowing them to put computational thinking in the context of their own physicality. When programming the turtle to draw shapes, children could first develop an algorithm for themselves to walk in the shape, then translate that to instructions for the robot. This perspective-taking and spatial reasoning allowed children to gain a foothold in programming while circumventing the need for instruction in more advanced mathematics.

More recent platforms such as *KIBO* (Sullivan, Elkin, & Bers, 2015; Elkin, Sullivan, & Bers, 2016; Bers, 2018) and *ScratchJr.* (Flannery et al., 2013; Strawhacker, Lee, Caine, & Bers; 2015; Strawhacker, Lee, & Bers, 2017; Portelance, Strawhacker, & Bers, 2015; Portelance & Bers, 2015; Leidl et al., 2017) have been designed to help young children develop computational thinking skills with creative, open-ended activities. *KIBO* is a robot that uses tangible programming blocks with icons . This hands-on, customizable robot gives immediate feedback through motion, lights, and sounds. The robot is configurable with a range of sensors and a design platform. These make it possible for children to customize the robot. Researchers designed a curriculum and assessments for the KIBO robot to assess how children in K-2nd grade learn computational skills such as sequencing, loops, and conditionals. They found that children could learn computational thinking with the *KIBO* platform, but that their understanding was not as complete as their older counterparts (Sullivan, 2016). Younger children understood concepts more slowly and made less complex programs. Children in Kindergarten were significantly less likely than children in 2nd grade to use control blocks for loops and message blocks for timing in their programs. Instead, children used more trigger blocks and motion blocks, blocks that were very exciting. When designing platforms for Kindergarten children, immediate, exciting feedback and a more limited number of blocks are important.

ScratchJr. is a free app for home use, developed by researchers at Tufts University, that allows children to explore computational thinking by programming stories and games on a computer (Flannery et al., 2013). It is recommended for children between the ages of 5-7 years. While the platform does include tutorials and instructions, most of the learning occurs while children engage in free play. The platform allows children to create animated scenes with onscreen characters, and to program them using an icon-based block-based programming language. Researchers analyzed what kinds of blocks children used, when they used the app, and how their programs became more complex over time (Leidl et al., 2017). They found that some of the most popular features of *ScratchJr.* were ones where children could add their own recordings and pictures. Children seemed to deeply care about being able to personalize their projects. They also found that about one-fifth of the time children triggered a block, it was the forward motion block, reinforcing what was seen with *KIBO*. The takeaways from this project are that children could go from tutorials to designing their own programs with block-based languages on a tablet.

The *Social Robot Toolkit* (SoRo) is a programming interface that sought to engage children ages 4-8 years old with computational thinking concepts through interaction with a social robot (Gordon et al., 2015). A set of vinyl stickers represented different triggers, actions and relationships. The programming sticker elements were mostly speech and emotion related to emphasize the nature of

the robot as a social other. Children composed their programs by placing these stickers on a smooth surface and showing them to the robot. This resulted in a highly interactive way for children to create rules for the social robot to "learn" and follow, more of a conversation than programming. The social robot served as a unique platform because it elicited a social response from children. Researchers saw that children would explore computational concepts by predicting the robot's mental state, "teaching" it social responses and emotive reactions.

Platform/Tech	Ages	Programming Interface	Programming Medium	Activities
Logo/Turtle (Papert, 1980)	5+	Text	Physical turtle robot	Problem-solving
ScratchJr. (Flannery et al., 2013)	5-7	Digital blocks, primarily iconic	On-screen stage	Open-ended creation, task completion
KIBO (Bers, 2018)	4-7	Tangible blocks, primarily iconic	Customizable KIBO robot	Open-ended creation, storytelling, art
SoRo (Gordon et al., 2015)	4-8	Tangible iconic stickers	Social robot (Dragonbot)	Open-ended creation
POP Curriculum (this work)	4-6	Digital blocks, completely iconic	Buildable, social robot	Open-ended creation, storytelling, music, games

Table 1: Summary of relevant children's computational thinking platforms

To my knowledge, *PopBots* is the first platform that uses a programmable, interactive robot that teaches children about artificial intelligence. One of the key questions in this work is how can we deliver a developmentally appropriate curriculum for young children to learn about AI and what features promote children's exploration and, as a result, understanding. The POP curriculum takes a number of design cues from the successful projects shown in Table 1. In particular, using a block-based, primarily icon programming interface, using a social robot that children can customize and build themselves, and using open-ended, creative learning tasks to teach the concepts. By leveraging these different components of previous platforms I hope to expand the list of computational concepts that are accessible to young children.

2.3 AI Curriculum and Platforms for K-12

It is important to distinguish platforms that develop computational thinking skills from those that help children understand AI. Using the syllabus of MIT's Artificial Intelligence course as a guide, the field of AI compromises of designing and implementing algorithms in which computers can plan, learn, reason, problem-solve, represent knowledge, perceive, move or manipulate objects with control, have social intelligence, and be creative.

> "Artificial intelligence is the science of making machines do things
> that would require intelligence if done by men."
> - Marvin Minsky, 1966

In his Artificial Intelligence course, Dr. Patrick Henry Winston admits that "The hard thing about AI is that once you've done it, it's no longer AI" (P.H. Winston, September 6, 2017). In other words, once an AI algorithm is developed that can do something "intelligent" that algorithm no longer seems intelligent. That said, the computational thinking skills described in the previous section are different from the AI concepts that students are taught with the tools I will describe here. AI education goes beyond teaching how robots (and computers) sense, think, and act -- it talks about how robots learn, make decisions, create, and understand their world.

For the most part, classes that teach AI are for students in high school and up. These classes use a variety of strategies such as robots, simulation, and non-programming examples to help students learn about AI. Aside from courses for older students, I also found two platforms that use Scratch to teach children machine learning, *Machine Learning for Kids* and *Cognimates* (Lane, 2016; Druga, 2018). However, Scratch is designed for children ages 7 and up. I also found two blog posts and a kickstarter campaign by instructors who describe AI to children ages 10 and up using videos and non-programming activities (Slavin, 2016; Chen, 2017; Milford, 2018). To my knowledge, there are no platforms or curricula in existence today that truly allow young children to program, train, build, and learn about AI platforms.

In the higher education courses on AI, students program and test AIs that cover the full range of AI topics described above (Imberman, 2003; Kumar, 2004; Kumar, Kumar, & Russell, 2004; Martin, 2007; McNally & Klassner, 2007; Koski et al., 2008; Talaga & Oh, 2009; Burgsteiner et al., 2016). Instructors have found that robots can be extremely effective in helping students concretize their understanding. Building robots may be more difficult, but students say that the projects are more engaging and that putting algorithms on a hardware platform helps students see the algorithms in action and notice how small changes to algorithms lead to different behaviors (Kumar, 2004; Koski et

al., 2008; Talaga & Oh, 2009). However there are other approaches to running algorithms like using non-programming and software-only approaches (Kumar et al., 2004). Even with these approaches it is important that students get a chance to work with and tinker with the algorithms from the ground up.

Approaches to teaching AI to children using videos and non-programming examples have also been successful in helping children gain an appreciation for the technology (Slavin, 2016; Chen, 2017; Milford, 2018). Slavin explains AI to children by giving examples of the kinds of things that computers are good at (computation) and the kinds of things that are harder (making decisions). Chen goes a bit more in-depth to show all of the different kinds of AI, ones that make decisions and ones that complete complex tasks or relate emotionally to humans. Milford's kickstarter campaign contains a book full of examples of AI in the real world and its success (raising over $2,700USD in less than 30 days) demonstrates parents' interest in helping children understand these concepts. However, while pictures and videos may be useful for showing children different kinds of AI, most ideas are most likely too abstract for young children to grasp and apply to the kinds of technology that they see around them. The projects by Lane and Druga are a step in the right direction (Lane, 2016; Druga, 2018). These use Scratch to allow children to use different kinds of AI in their own projects. For example, children can train their own object recognizers or create conversational AIs to make games and smart home controllers themselves.

AI concepts are easiest for students to learn when they have the right tools for hands-on learning. AI courses in higher education emphasize the importance of letting students build their own projects. Younger children definitely need a hands-on approach to grasp ideas as well. A social robot is a powerful platform that can relate to children as a social being. By interacting with the robot, children will be able to see how it solves problems in its "mind."

2.4 Theory of Mind Development: Thinking about Other's Thinking

Perspective-taking skills, or the ability to view the world from another's vantage point is rooted in one's Theory of Mind skills (Barnes-Holmes, McHugh & Barnes-Holmes, 2014). Theory of Mind (ToM) is defined as one's "ability of inferring mental states of self and others" (Wang, 2015; Premack & Woodruff, 1978). People use ToM skills to do everything from making sense of other's behavior to predicting what a person will do next. However, this ability is not innate; ToM develops progressively from infancy until children are age 6 or 7 (Apperley, 2012). Barnes-Holmes et. al (2014) outline children's development of ToM as five distinct levels:

"Level 1 of this model involves simple visual perspective-taking, and concerns the fact that different people can see different things.

Level 2 is referred to as complex visual perspective-taking, and is concerned with the ability to know that people can see the same things differently.

At Level 3, visual features are believed to play a less salient role in perspective-taking, and individuals at this level come to understand the principle that 'seeing leads to knowing.'

The fourth level of understanding informational states involves true beliefs and predicting actions on the basis of a person's knowledge.

The fifth and most complex level of knowledge of informational states from a ToM perspective involves the understanding of false belief and predicting actions on the basis of beliefs that are false rather than true." (p. 16-17)

Wellman and Liu compiled multiple studies to develop a series of tasks to assess the fifth level of children's ToM skills. The tasks use storytelling and targeted questions to measure children's understanding of diverse desires, diverse beliefs, knowledge access, contents false belief, explicit false belief, belief emotion, and real-apparent emotion (Wellman & Liu, 2004).

Key ToM abilities that children encounter in this curriculum are false-belief understanding -- that another person may believe something different, and knowledge access understanding -- that a person may not know something that you know, and explicit false belief understanding -- knowing how a character will behave given its knowledge state. These ToM skills are still developing in 3-5-year-old children. Wellman and Liu saw that 73% of children could correctly answer a knowledge access question and that 59% of children could correctly answer a false belief one, and 57% of children could correctly do explicit false belief.

People naturally anthropomorphize inanimate objects like robots because we are naturally wired for social relationships (Breazeal & Scassellati, 1999; Breazeal, 2004). Previous work in the Personal Robots group shows that children are sensitive to social cues like backchanneling, attentive signs, and voice expressivity when interacting with a social robot (Breazeal & Scassellati, 1999; Kory & Breazeal, 2014; Gordon, Breazeal, Engel, 2015; Breazeal, Harris, DeSteno, Kory, Dickens, Jeong, 2016; Gelsomini, Park, & Breazeal, 2017; Lee, Breazeal, DeSteno, 2017; Park, Gelsomini, Lee, Brealzeal, 2017; Westlund, Jeong, Park, Ronfard, Adhikari, Harris, DeSteno, Breazeal, 2017). When robots behave as social agents, children learn better from them.

In the POP curriculum, children relate to the robot as a social being to understand the algorithms that are running in its "mind". Children are asked to understand the robot's perspective making it

likely that their comprehension of AI in this curriculum will be dependent on their ToM skills. I hypothesize that ToM development will be an even better predictor of how children perform with the PopBots curriculum than children's age.

2.5 Developing Empowered STEM Identity

Gender stereotypes related to STEM begin as young as age 5 (Cvencek & Meltzoff, 2011; Sullivan, 2016). Sullivan studied how gender played a role in what careers and toys kindergarten through 2nd grade children thought were most appropriate for people. Children drew on aesthetics, gut instinct, role-modeling, and personal experience to draw conclusions about different objects. Boys perpetuated gender stereotypes more than girls, but both were sensitive to them (Sullivan, 2016). Fortunately, Sullivan and other researchers also found that doing an engineering curriculum had a positive effect on girl's and boy's confidence in themselves as engineers (Cunningham & Lachapelle, 2010; Sullivan, 2016). These researchers used an Engineering and Science Attitudes Assessment to evaluate what children thought before and directly after completing a developmentally appropriate engineering curriculum. Boys were significantly more likely to say that they would like to be an engineer when they grew up than girls. After doing the curriculum with KIBO, girls and boys increased in saying they would like to be engineers (Sullivan, 2016). Engineering is Elementary found similar results with third through fifth graders. At first, boys were more likely than girls, but after doing a STEM curriculum, both increased (Cunningham & Lachapelle, 2010). Long-term, researchers have also found that students who participate in STEM programs at a young age, they are more likely to be confident in their STEM abilities when they head to college (Aschbacher, Li, & Roth, 2010)

Curriculums that encourage children to pursue STEM first help children see how what they are learning connects with something in their life, and next empower them to use that learning in a meaningful way (Resnick, 2017). In 2015, the European commission studied families with high digital technology use and found that most children under the age of 8 are technology consumers rather than producers (van Kruistum & van Steensel, 2016). They posit that this may be because children do not have the tools they need to create with technology or the opportunity to see how creating technology can be personally meaningful.

The POP platform is a tool for young children to create their own intelligent robots and thereby see how they can become creators of this technology. I hypothesize that children learning about how robots learn will lead to them seeing the importance of the technology. Then, getting hands-on experience building their own robots will lead to children developing an identity in engineering.

3 The POP Platform: Design & Technical Development

My goals in designing the POP platform were to make a low-cost system that could relate the nuances of AI algorithms to young children in a developmentally appropriate manner. This charge, along with inspiration from prior computational thinking and social robot platforms led to the development of the system as it exists today. The platform consists of an Android phone with the PopBot app installed, a tablet with the PopBlocks app installed, a LEGO WeDo 2.0 set, and LEGO blocks. The PopBot app displays a robot face on the Android phone. This chapter presents the POP Platform design to help answer **Research Question 2**:

Is a social robot platform an effective way to deliver developmentally appropriate curriculum for young children to learn about AI concepts. What features support children's exploration, perception and understanding?

3.1 PopBot Construction Kit

Children can construct their own LEGO characters around the phone. The phone is held in a 3D printed case (Fig. 1) that attaches to regular and LEGO DUPLO blocks. The connection between the phone and the LEGO WeDo hub is a wireless bluetooth low energy (BLE) connection. This means that children can arrange robots into virtually any form using a combination of regular LEGO and LEGO DUPLO blocks. However, the LEGO WeDo hub only provides two ports. Hence, the number of motors was limited to two, but fewer if children chose to use LEGO WeDo sensors.
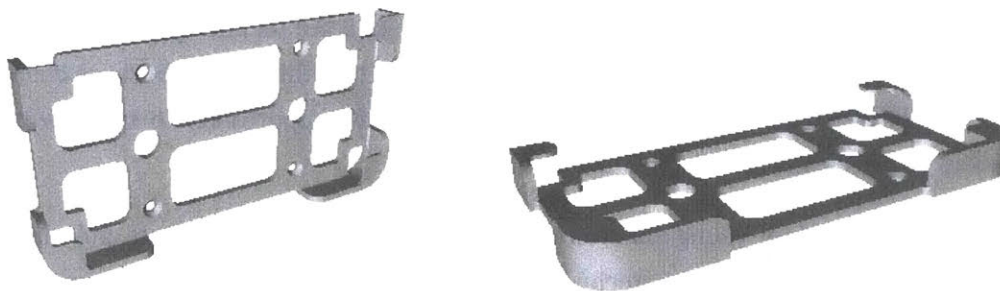
Figure 1: 3D printed phone holder for PopBot, attached to LEGO blocks

The robot is programmable, but also has autonomous functionality and plays an active role in the curriculum. The robot's "mind" becomes a metaphor through which children can examine the current state of the different algorithms. The robot gives feedback about itself either automatically when it wants to make a process transparent (e.g, if the child forgets to connect the WeDo hub and tries to use it), or if the robot makes a move in the AI activity. The second kind of feedback requires that algorithms in the AI activities be implemented such that the robot can explain its actions verbally.

Since the target ages of children for the POP platform is slightly lower than the recommended ages for using LEGO blocks, I provided a number of pre-built LEGO artifacts that could inspire children to build their own creations from there. Along with the pre-built artifacts I developed introduction programs to give children examples of social robot behavior. Five types of robot forms were developed.

1. The earliest artifact was a LEGO Wall-E robot, Fig. 2, where the phone attached in place of Wall-E's head. This is how the robot got its name Mall-E. The wheels of the Wall-E body were connected to LEGO WeDo motors. This body came with a "greeting" program[2] and a "shy" robot program. In the greeting program, a tilt sensor was attached to the robot's arm and a child could shake the hand of the robot to trigger the robot introducing itself. The "shy" robot program consisted of using the phone's built-in light sensor to trigger the robot's SCARED animation, and then the motors move the robot backwards, away from the bright light.

Figure 2: Wall-E body for PopBots

---

[2] PopBot Wall-E greeting program https://youtu.be/83U7dwQ_TvY

2. The next artifact was a spinning head. An actuated head allows children to develop gestures for emotions such as disagreement, sadness, and emotions. The example program[3] involved the robot spinning with a flag while shouting "Wee!"
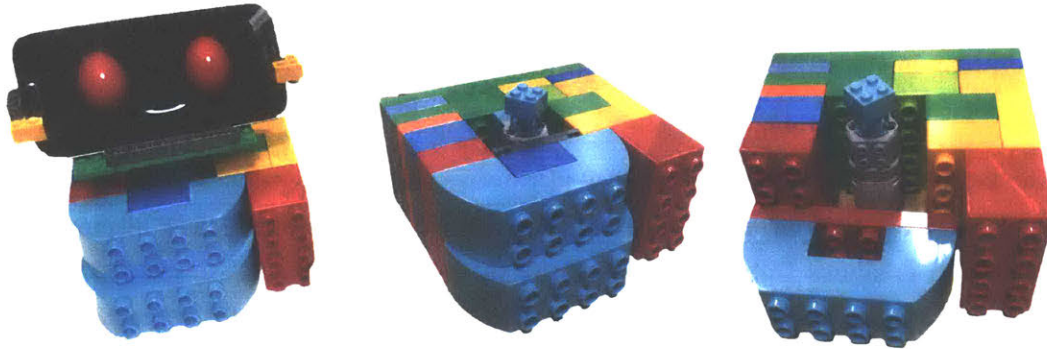
Figure 3: Spinning head body for PopBots

3. I also developed a rock, paper, scissors game playing robot, where a robot servo arm could point to its move, shown in Fig. 4 (right). Note that the LEGO WeDo motors are continuous rotation servos rather than standard servos, meaning that the arm tended to drift over time. The other arm configuration shown in Fig. 4 (left) is used for expressivity.

Figure 4: PopBots with arms used for pilot studies

4. Note that some of robot forms I developed were not actuated to support character play. For instance, I developed a robot animal form and a robot humanoid form that could be used as characters in a play with LEGO DUPLO characters.

---

[3] PopBot spinning head program https://youtu.be/c2vDSCjH3WQ

Figure 5: PopBot humanoid and animal forms used for pilot studies

5. For the final evaluation, I developed a car form that was easily reproducible. This design took advantage of the continuous LEGO WeDo motors. This robot form had a more machine-like appearance that the animal or humanoid forms.

Figure 6: PopBot car form used for final study

3.2 POP Platform Software Architecture

Fig. 7 shows the big picture of the software architecture. The main modules I developed included the robot controller (R1D1), AI modules, and rule processing system. I leveraged R1D1's animations library, R1D1's tools library, the LEGO WeDo SDK, and the TopCodes SDK to build core functionality.
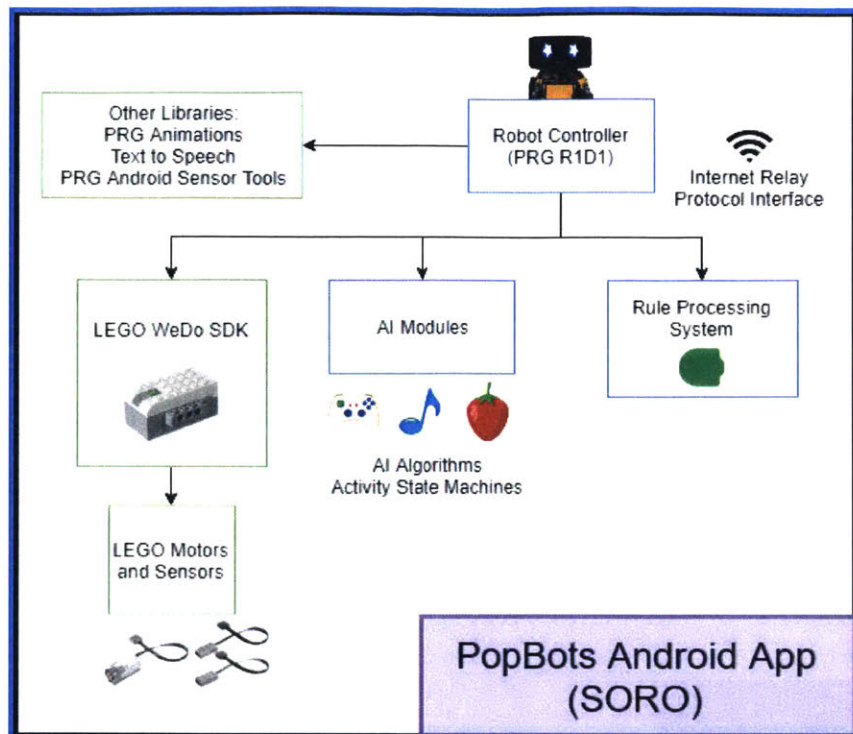
Figure 7: Overview of PopBot software architecture

The rule processing system is the interface between the *PopBot App* and the *PopBlocks App*. It contains the robot's memory of rules and makes calls to R1D1 to execute commands. In addition to a queue for sequential robot commands, the processing system does threading, contains a tutoring system, gives status information, and controls idle behavior. Each of the functions shown in Table 2 contributes to the role of the robot as an intelligent agent and a helpful companion.

The programmable actuators on the robot include the robot's animations, speech and sound, eye color, LED, and detachable motors. The robot's sensors are tilt, proximity, light, and touch. Additionally, the programming blocks include controls like delays and for loops. The phone can also use its camera to read special barcodes called Topcodes.

Rule Processing System Function	Function Rationale
Sequential Command Execution	Sequential Command Execution allows the robot to execute a line of code of multiple blocks. The robot automatically detects when it is finished completing a previous line of code and can proceed to the next line. This makes it possible for children to program complex behavior
Threaded Command Execution	Sometimes, multiple triggers happen at once, or a child wants more than one thing to execute. Using threads allows the robot to do a motion like moving its motor at the same time as starting to speak and cycling through its eye colors. Also, sometimes multiple sensors are triggered at the same time. When this happens, threading allows the robot to seem like it is responding to both triggers at the same time.
Help Command	If a child is unfamiliar with a block, they can press the block and the robot will either act it out or describe the block to the child and how to use it.
Status Information	When a child uses a new block (such as a motor block), but has not yet connected any motors -- the robot can help the child get the motors configured first. This "self-awareness" helps the robot seem more sentient.
Idle Behavior	Children occasionally get distracted and stop playing with the robot. The robot can perform idle behaviors to get the child's attention and to seem more lifelike. While a normal toy sits and waits for the child, this robot eagerly asks the child what activity they should do next.

Table 2: Overview of main functions in rule processing system

3.3 The PopBot App

The PopBot App runs on a Samsung smartphone and serves as the expressive face of the PopBot. The code for the PopBot app was adapted from the code used for the Tega robot developed by the Personal Robots Group. The original Tega App already contained core functionality to trigger animations, interface with phone sensors, and play recorded audio files. For the PopBot app, I added the ability to interface with LEGO WeDo, programmatically trigger sensors and new actions (such as changing eye color), record new sounds, play synthesized text to speech (TTS), read barcodes, and process commands from the tablet.

The opening screen of the phone is a black page with the name of the main class "SoroController." First, the user should input the number of the robot into the top right corner of the screen. The number of the robot controls the initial color of the robot's eyes, the tablet the robot connects to,

and the LEGO WeDo hub that the robot connects to. The robot face takes approximately 20 seconds to initialize due to a number of animations loading in the background. When the robot awakes, it yawns and says, "Hi there. Ready to have some fun?" From that point on, the robot is ready to receive commands from the tablet.
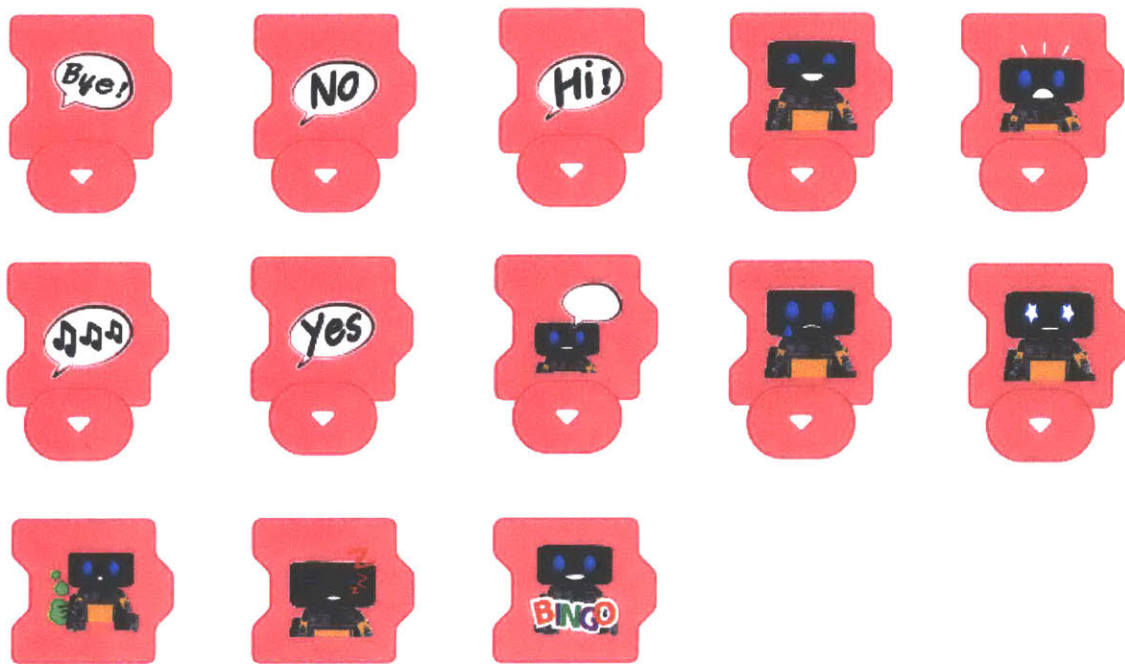
Figure 8: PopBlocks for robot animations

The PopBot face is the same face as the Tega robot[4]. It comes with a library of 50+ expressive animations and more than 1,000 human-recorded sound and speech files. From these, I selected seven expressive animations and five speech files, as shown in Fig. 8. I chose four different emotions (excited, sad, interested, and scared) and three different comedic animations (fart, sleepy, and dance to bingo). Rather than use the pre-recorded speech files, I opted for a text to speech engine with a high pitch and slow speed, so that it would sound genderless and ageless, yet still enunciate enough

---

[4] Personal Robots Group: Tega http://robotic.media.mit.edu/portfolio/tega/

to be understood. While the text-to-speech engine was produced intelligible speech, it was not particularly expressive and sounded artificial.

3.4 PopBlocks App & Data Logging

The PopBlocks App runs on the tablet. Children use it to program the PopBot and do the AI activities. The programming interface is based on Scratch blocks, and each AI activity has a custom interface (as described in Chapter 4). The programming blocks are all picture-based to accommodate children who cannot read yet.
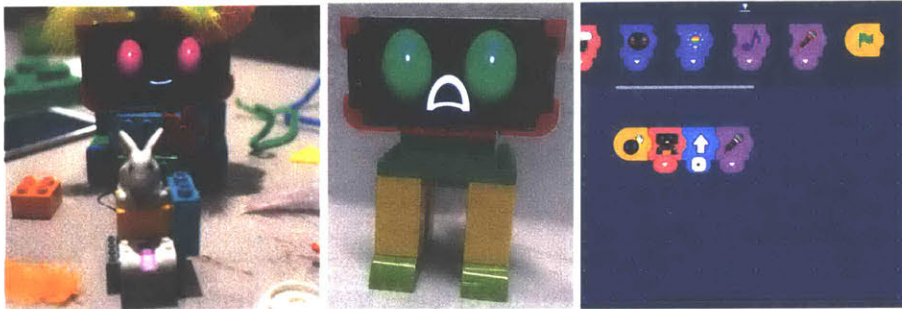
Figure 9: Decorated and animating PopBots with PopBlocks interface

A tablet can control one PopBot at a time. The tablet logs all information including the time, the current activity the child is using, and any buttons the child presses. These log files are used to capture children's interactions with the platform for later analysis. Furthermore, logging information is also sent real-time to a server (Fig. 10) to allow an instructor to see what activities are actively used on the tablet at any given time. The server collects the log files of all active tablets. It also organizes them into sections for each group that contains a summary followed by the tablet logs in real time. This tool is also useful for post-processing log data to examine a user's trajectory through the activities.

G2

Current Activity: Workspace
Last Active: 02/04 02:10
Activities Explored: Workspace

G2 Date:1517677832746 Workspace Element:categoryclick OldVal:Add NewVal:Add Block: Workspace:
G2 Date:1517677832456 Workspace Element:categoryclick OldVal:Add NewVal:Add Block: Workspace:
G2 Date:1517677813502 Workspace Element:categoryclick OldVal:Add NewVal:Add Block: Workspace:
G2 Date:1517677809327 Workspace Element:categoryclick OldVal:Add NewVal:Add Block: Workspace:
G2 Date:1517677800297 Workspace Element:sound OldVal:null NewVal:delete Block: Workspace:
G2 Date:1517677800293 Workspace Element:selected OldVal:1_?L/Pg]cSO(,U8]s+bZ NewVal:null Block: Workspace:
G2 Date:1517677798929 Workspace Element:selected OldVal:null NewVal:1_?L/Pg]cSO(,U8]s+bZ Block: Workspace:m01_mystery ;
G2 Date:1517677797597 Workspace Element:selected OldVal:1_?L/Pg]cSO(,U8]s+bZ NewVal:null Block: Workspace:m01_mystery ;
G2 Date:1517677791992 Workspace Element:selected OldVal:null NewVal:1_?L/Pg]cSO(,U8]s+bZ Block: Workspace:m01_mystery ;
G2 Date:1517677785809 Workspace Element:categoryclick OldVal:Add NewVal:Add Block: Workspace:m01_mystery ;
G2 Date:1517677780590 Workspace Element:selected OldVal:1_?L/Pg]cSO(,U8]s+bZ NewVal:null Block: Workspace:m01_mystery ;
G2 Date:1517677777410 Workspace Element:categoryclick OldVal:Add NewVal:Add Block: Workspace:m01_mystery ;
G2 Date:1517677762564 Workspace Element:selected OldVal:null NewVal:1_?L/Pg]cSO(,U8]s+bZ Block: Workspace:m01_mystery ;
G2 Date:1517677761259 Workspace Element:selected OldVal:1_?L/Pg]cSO(,U8]s+bZ NewVal:null Block: Workspace:m01_mystery ;
G2 Date:1517677757541 Workspace Element:categoryclick OldVal:Look NewVal:Add Block: Workspace:m01_mystery ;
G2 Date:1517677337743 Workspace Element:stackclick OldVal:undefined NewVal:undefined Block: Workspace:m01_mystery ;

G3

Current Activity: Workspace
Last Active: 02/04 02:18
Activities Explored: Workspace RPS

Figure 10: Screenshot of tablet logging interface

4 The POP Curriculum: Design & Technical Development

This chapter describes how the activities in the POP platform were developed into a developmentally-appropriate curriculum to teach children three AI concepts: rule-based systems, supervised machine learning, and generative AI. Young children's ability to understand these concepts have not been previously explored. They were chosen because of their relative simplicity and relevance to the kinds of AI that children are usually exposed to through smart toys and media about intelligent machines. and entertain apps that children interact with. The following activities were used in the evaluative study to answer **Research Question 3**:

Can children learn about key AI concepts such as rule-based systems, supervised machine learning, and generative AI by coding and interacting with a social robot platform? How does this inform a developmentally appropriate curriculum for AI and its attributes?

4.1 Introduction to Programming

Children do a series of five mini-tasks (like the two in Fig. 11) to get an introduction to programming with PopBlocks. Each task has a limited number of blocks so that children can progressively learn the interface. On each page, the robot tells the child what task to complete. Pressing the blocks in the toolbox (dark blue part of the screen) will make the robot give a hint about how to use that block. The three AI concepts allow children to explore how machines learn, make inferences, and act creatively. They are briefly described in the following sections.
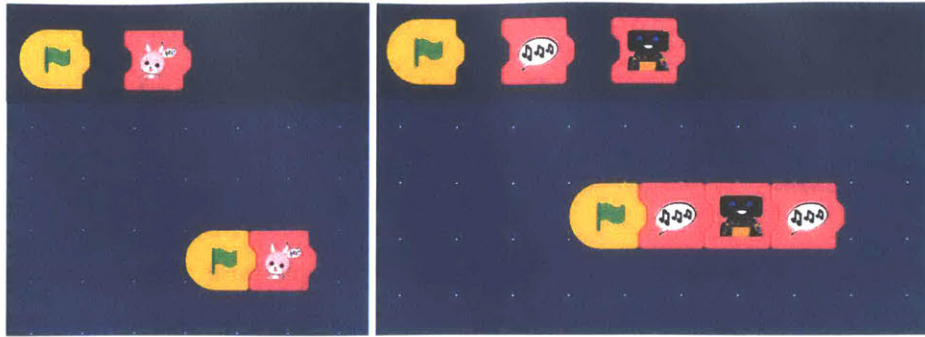
Figure 11: Screenshots of introduction to programming activity

Task	Prompt
Connect blocks together, learn green flag trigger	Can you teach me to say my name and introduce myself to others? Touch a block to learn about it. Drag it down onto the screen to program me.
Create a sequence of multiple blocks	I was invited to a party at the castle. I am very excited. Can you program me to be excited and sing my song?
Learn the color blocks, learn to turn on LEGO WeDo hub	I need to get ready for the party. Can you dress me up?
Learn the motor blocks and loop blocks	I'm already at the castle party. Now I'm ready to dance. Can you teach me how?
Use the programming interface with all of the different blocks available	Great! Now you know how to program me. Let's have some fun.

Table 3: Introduction to programming tasks

4.2 Rule-Based Systems

Rule-based systems (or expert systems) are a common form of AI. These systems contain two main components: a way to represent knowledge, and a way to act on that knowledge. In commercial smart devices, games such as tic-tac-toe tend to be implemented using rule-based systems. I chose this as an AI topic because it was a relatively small leap from programming a robot. Learning about rule-based systems allows children to see how robots can have persistent knowledge and use it to make decisions.
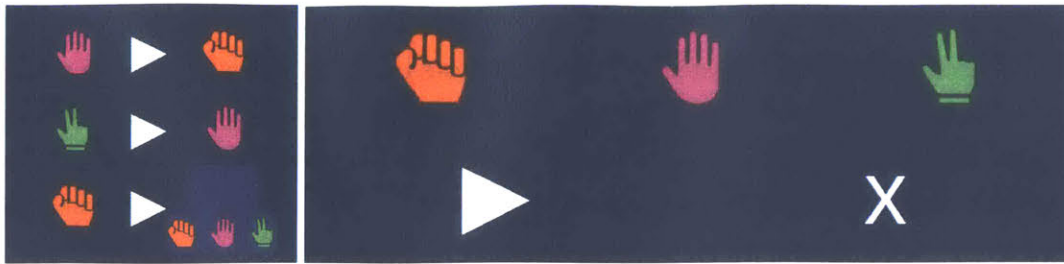
Figure 12: Screenshots of rule-based systems activity: rock, paper, scissors. Left, interface for teaching rules. Right, interface for playing game.

In the rule-based systems activity, the robot has a rule matrix and a probability transition matrix. It fills the probability transition matrix to predict the user's next move and decide its subsequent move based on what the player did. First, children train the robot's rule matrix with the three rules of rock, paper, scissors. The screen on the left of Fig. 12 shows how children input the rules of the game. The top line is read "paper beats rock". After children are done programming the rules, the robot reads the rules back to them. The interface allows for children to teach the robot any kind of rule, except for teaching it that a move can beat itself (e.g., scissors beats scissors). Next, children program the robot to react to winning or losing using the programming interface. As a group, we talk about being a good sportsman rather than a sore-winner or a sore-loser. However, just like in the case with teaching the robot the rules, children are free to create whatever kind of robot they would like. Finally, the child and the robot play the game against one another. As the child plays, the robot uses a state transition matrix for the game to predict the child's next move based on their last three moves. If the robot's guess for a move is greater than chance (33%) then the robot will say "I think you will put X, so I will put Y because Y beats X." Otherwise, the robot says "I'm not sure what you'll put, I'll just guess." As the game progresses, the robot encourages the child to keep playing to improve the robot's accuracy, "I'm getting good at this. The more you play me, the better I will get."

As a group, we talk about how the robot guesses what move will come next by using past examples and referencing the rules the child taught it. Children also get the opportunity to teach the robot the opposite rules and play against it then. If the child does not teach the robot all three rules, the robot will ask the child whether it won or not and then incorporate that new knowledge into its set of rules.

4.3 Supervised Machine Learning

Supervised machine learning is related to rule-based systems, but is more common in modern AI systems. It involves forming a knowledge base on the fly by learning from examples. For instance, this method is used in personalized recommender systems for media streaming applications. I chose this topic to introduce the idea of a training set to children. They can explore how different training sets improves the robot's accuracy as well as see how robots learn.
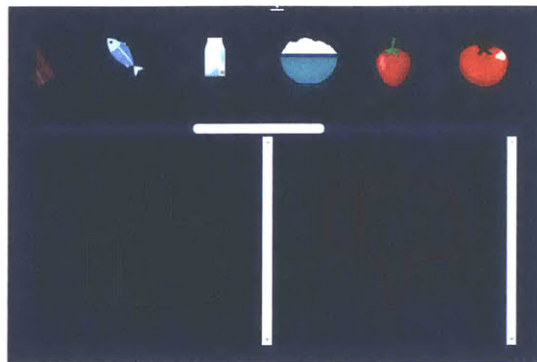
Figure 13: Screenshot of supervised machine learning activity: food sorting

Children train the robot to recognize healthy and unhealthy foods (indicated by the thumbs up and down in Fig. 13), and then program it to respond to the foods. Before beginning interaction with the robot, children sort stickers on the table into groups of "good food" and "bad food". Next, the researcher asks the children to talk about things that the groups have in common -- for example, healthy foods contain a lot of fruits and vegetables. As children begin teaching the robot foods, the robot will also begin to guess labels for the food by comparing it to known foods. Let's say that strawberries have been labeled and the child asks the robot about tomatoes. The robot would say, "Tomatoes are a lot like strawberries, so tomatoes go in the healthy group." Children experiment with the number and kinds of foods that are labeled to see when the robot does a good job guessing and when it does a bad job. Children can also try to confuse the robot by telling it the wrong foods are healthy or unhealthy. In all, children can play with 20 different foods shown below in Table 4.

Food	Food Group	Color	Number of Calories (per 100g)	Grams of Sugar (per 100g)
Banana	Fruit	Yellow	89	12
Bread	Grain	Brown	265	5
Broccoli	Vegetable	Green	34	2
Carrot	Vegetable	Orange	41	5
Cheese	Dairy	Yellow	371	2
Chicken	Meat	Brown	197	0
Chocolate	Fats and Sweets	Brown	488	56
Fish	Meat	Blue	105	0
Milk	Dairy	White	61	5
Rice	Grain	White	130	0
Strawberry	Fruit	Red	33	5
Tomato	Vegetable	Red	18	3
Soda	Fats and Sweets	Brown	38	15
Juice	Fruit	Orange	45	8
Butter	Fats and Sweets	Yellow	717	0
Cookie	Fats and Sweets	Brown	467	20
Potato Chips	Vegetable (also Fats and Sweets)	Yellow	536	0
Ice Cream	Dairy (also Fats and Sweets)	White	207	21
Corn	Vegetable	Yellow	365	0
Blueberries	Fruit	Blue	57	10

Table 4: Food attributes in PopBot database

The algorithm uses k-nearest neighbors with k=3. The robot knows about the color, food group, number of calories, and grams of sugar per serving of 20 different foods. Discrete features such as food group and color were mapped to numerical values. The mapping takes proximity into account.

For example, the fruit food group is closer to the vegetable food group than the dairy food group -- and red foods are closer to orange foods than white foods. The algorithm calculates the distance between a child's chosen food and all of the labelled food items. The distance between two foods is the sum of the normalized distance between the foods' features. When there are no labeled items, the robot says, "I don't know where anything goes yet." When the number of labeled items is less than or equal to 3, the algorithm chooses the one food that is the closest match. The robot says "X is a lot like Y, so X goes in the same group as Y." When the number of labeled items is greater than 3, the robot will simply say "X goes into the A group." If the child uses the help command to inquire why, the robot will explain "X is like <top 3 foods> and most of these are in the A group." The robot can further explain its reasoning down to the specific features that were similar, for example, "I think chocolate and bread are similar because they are both brown."

4.4 Generative AI

Generative AI is very different from the other two examples. Rather than learning rules, the system is given parameters and is allowed to generate responses within the bounds of those parameters. This can be seen in camera apps that use filters to change the appearance of original photos. I chose this activity to show children that robots do not always follow the rules, they can also be creative (in their own way).
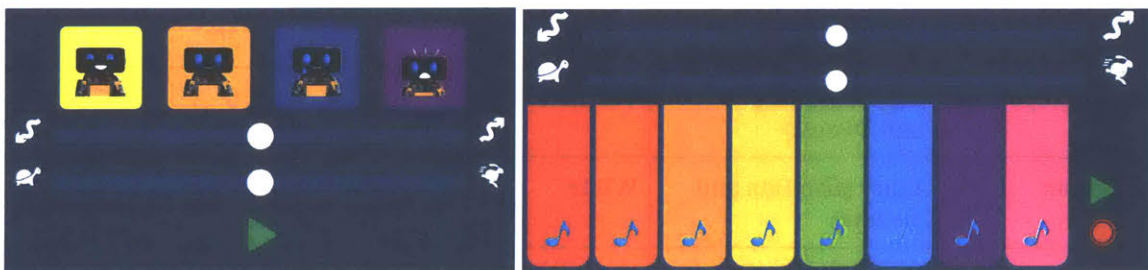
Figure 14: Screenshot of generative AI activity: music remix. Left, interface for programming emotions. Right, interface for creating new songs.

Children explore how a robot creates new music by changing the tempo and progressions of a given piece of music. First, we talk about how tempo and chord progressions (whether music goes up or down) translate to emotions in music. For example, happy songs have a faster tempo and go up. Children associate different combinations of tempo and progressions with emotions and teach

"musical emotions" to the robot. By hitting the play button on the left screen, the tablet will play a piece of music and send it to the robot to play back. Children can adjust the slide bars to change how the robot plays the music back by varying either the tone or the progression of the song. Then children can transform any musical input using these emotions. Finally, children can use the screen with the keyboard to remix their own creations.

Key: A	I	III	V	VIII
I			50%	50%
II			50%	50%
III			50%	50%
IV				100%
V		33%	33%	33%
VI	50%			50%
VII				100%
VIII			50%	50%

Table 5: Probability transition matrix for ascending progressions

Key: A	I	II	III	IV
I	100%			
II	50%	50%		
III	50%	50%		
IV		50%	50%	
V	33%	33%		33%
VI		50%		50%
VII	33%	33%	33%	
VIII	33%	33%	33%	

Table 6: Probability transition matrix for descending progressions

In changing the tempo of the song, the generative music algorithm uses simple rules. If the tempo needs to go faster, then it divides one long note into several rapid notes with evenly spaced breaks. To make the tempo slower, the algorithm increases the length of the notes and slightly decreases the space between them to make them flow into one another. The algorithm for changing the tune uses a probability transition matrix to choose a sequence of notes that goes up in a consonant interval or down in a dissonant one. This algorithm adds new, random notes to the middle and end of the song, keeping the rest intact so that it is clear that the robots output is based on the input. I developed the probability transition matrices in Table 5 and Table 6 by referencing musical theory about consonant and dissonant notes in major keys and through empirical evaluation.

4.5 Assessments

I developed a set of assessments for each activity. Each is delivered in a multiple choice format, (e.g., Fig. 15). In my protocol, the researcher reads the question and directs children to choose a response on the tablet. First, children take a set of pretests before engaging in the activites: a Theory of Mind Assessment, an AI Perception questionnaire, and a Engineering and Science Attitudes assessment. After, children work through a sequence of the activities described above. Immediately following each activity, children take an associated AI assessment comprised of 3--4 questions about the AI concept and algorithm introduced in that activity. For each algorithm, questions probed for children's understanding of the algorithm's basic functionality, edge cases, and initialization. The questions asked in each AI assessments for each activity is summarized in Table 7.
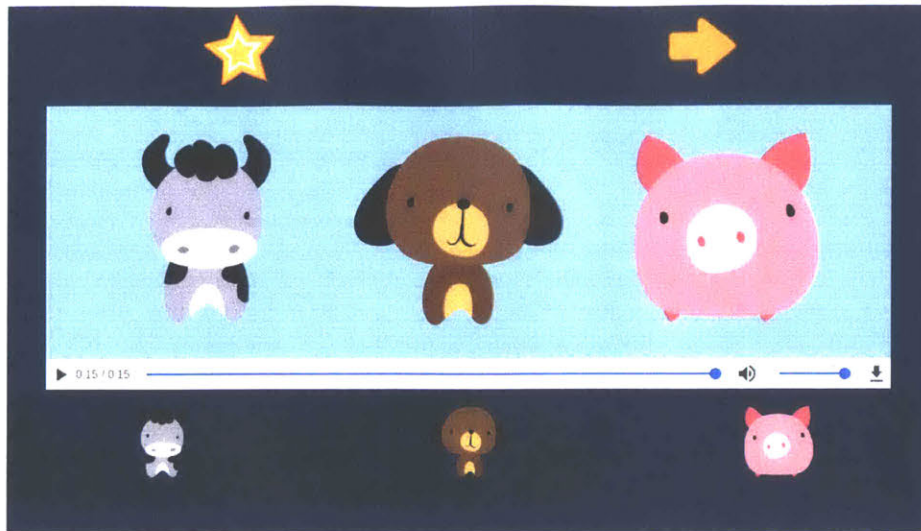
Figure 15: Screenshot of assessment question. Control question: which one of these animals makes a moo sound?

Question	Rationale
Rule-Based Systems	
1. Control: Which of these is rock? Rock, paper, or scissors?	The answer is rock. Control question to remind children of the format of the questions and how to respond.
2. We teach the robot the normal rules. Then, Sally plays rock and the robot plays paper, who does the robot think has won? Sally or the robot?	The answer is the robot. This is a basic functionality that asks the child to see that the robot decides who wins based on the response.
3. Sally plays paper five times. What does the robot think she will play next? Rock, paper, or scissors?	The answer is paper. This is a basic functionality question that asks the child to notice that the robot's predictions are based on the child's past behavior.
4. The robot thinks that Sally will play paper next. What will the robot play so that it can beat Sally? Rock, paper, or scissors?	The answer is scissors. This is an extended functionality question that requires the child to recognize that the robot's belief in Sally's next move leads it to behave in a certain way.
5. We changed the rules so that they are all opposite rules (paper beats scissors). Sally plays scissors and the robot plays paper. Who does the robot think has won? Sally or the robot?	The answer is the robot. This is an extended functionality question that requires the child to recognize that the robot only knows what rules you tell it, even if they are wrong. This requires false belief reasoning.

46

Supervised Machine Learning	
1. Control: Which one of these foods is bad for your teeth? Strawberry, ice cream, or corn?	The answer is ice cream. Control question to remind children of the format of the questions and how to respond.
2. You start the robot and put strawberries and tomatoes into the good group. Which group will the robot think chocolate goes in? The good group or the bad group?	The answer is the good group. This is an initialization question that requires the child to recognize that the robot does not know about any bad foods yet. Therefore, it will only predict the good group. This requires knowledge access reasoning.
3. What food does the robot think is most like a tomato? Strawberry, banana, or milk?	The answer is strawberry. This is a basic functionality question that requires the child to recognize that the robot knows about things like color and food group which makes the strawberry most similar.
4. You put ice cream in the good category and bananas in the bad category. What category will the robot put corn in? The good category or the bad category?	The answer is the bad category. This is an extended functionality question that requires the child to recognize that the robot will see the banana and corn as most similar and therefore group them together. This requires knowledge access and false belief reasoning.
Generative AI	
1. Control: Which one of these notes will make the robot's eyes go orange? Purple note, orange note, or green note?	The answer is the orange note. Control question to remind children of the format of the questions and how to respond.
2. Priya asks the robot to play back with the bars in the middle. Does the robot play the same song or a different song?	The answer is the same notes. This is a basic functionality question that requires children to understand that when the bars are in the middle the robot will neither change the speed nor the tune.
3. Priya asks the robot to play back with the bars to the right. Does the robot play the same song or a different song?	The answer is different notes. This is a basic functionality question that requires children to understand that the robot will change the speed and tune of the song when these options are set.
4. Does the robot's song have to have the same notes as the input?	The answer is yes. The robot's output song is always based on the input song.

Table 7: Assessment questions used in final study

5 Bringing POP into the Classroom: Iterative Design with Children and Teachers

In developing the activities and assessments for POP, I ran two pilots with children ages 3-7 from the Greater Boston Area. The initial pilot consisted of six children who did five, 2-hour-long private sessions (10 hours total for each child) interacting with the POP curriculum in the lab. The second pilot workshop was a one-hour long class with eleven children in a local STEM center. The goal of these pilot workshops was to iterate on the designs of the PopBot activities and assessments with stakeholders -- trying different teaching methods including using real-world examples and having children work in groups. The work of this chapter was performed to answer

Research Question 4:

What teaching methods are most effective for engaging and supporting young children with an early AI curriculum in classrooms?

First, I will discuss the initial pilot workshops. Six children (two only-children, two pairs of siblings) were invited to do five sessions each over the course of a few weeks. The only people in the room were the children, their parents, and the researcher. Each session was centered around one of the four main activities of the POP curriculum, introduction to programming, rock paper scissors, food classification, and generative AI. The last session was a character AI activity that did not make it into the final curriculum. During each session, first we would begin with a design activity, then we would look at a real-world application of our topic, next, we could do the activity from the POP curriculum, then discuss what we learned, and, finally, have free play where children could explore any of the POP materials. Free play gave children time to make their own PopBot creations and reinforce things they learned in previous sessions.

5.1 First Pilot Workshop: Explaining AI to Children

Every session included hands-on design activities and real-world AI activities to help children draw connections between their interests and what they were learning. The creative design activities also provided a lens into what kinds of things children wanted robots for and how their thinking about AI was developing. The real-world activities allowed children to connect what we were doing in the lab

with things they could continue to explore at home. With these activities I was also evaluating the extent to which children were able to make sense of less transparent AI artifacts.

Creative Design Activities

The design activities involved allowing children to first conceptualize and draw their own AI robots (sample drawings shown in Fig. 16). Next they were invited to build a robot body using the PopBot construction kit, and then program their PopBot to be a specific character.

The first day was an introduction to AI and robotics. Children and I discussed what the term "artificial intelligence" means, and how its ability to emulate human intelligence makes it unique. After showing children pictures of three examples of artificially intelligent technologies (self-driving cars, facial recognition software, a social robot named TEGA), children drew their own AI robot.
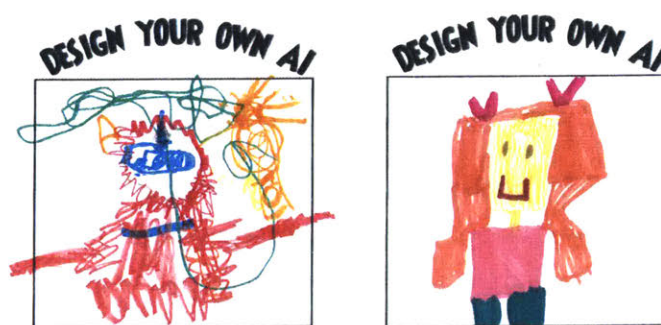
Figure 16: Children's AI designs

Six of the seven children in the pilot workshop finished their unique AI concept designs and drawings. A common theme throughout the designs was that the older children drew robots that could engage socially and that could assist others, while the younger children's robots were recreations of their toys. The youngest children (a three-year-old girl and a four-year-old boy) created robot designs that were similar to the AI examples that we discussed, a TEGA robot and a car; while the older children created their own designs. Two of these designs were humanoid robots that were, in fact, robot versions of the child him/herself. The purpose of these robots was to take the place of the child, especially in school, and to help the child learn more. The other two designs were more machine-like robots in appearance. One was a doctor and the other was a cleaning and cooking robot to assist the child's mother.

After this, children learned how to program the PopBots and had the freedom to design their own PopBot by building with LEGO blocks and craft materials (examples shown in Fig. 17). Children had the opportunity to design their own PopBots at every session during free. They designed everything from human-like characters to machines. Beyond making the robot move, one of the most popular functions was recording sounds for the robot to play. Children recorded songs or dialogue for the robot to speak to other characters. Then children would engage in pretend play with their creations. Some of the PopBot creations were machines such as a karate chopping robot or a swinging robot. Often times, children were interested in only building with LEGOs and spent less time programming the PopBot. Therefore, when running the final evaluation studies, I required children to use pre-built PopBots so that they would focus more on programming.
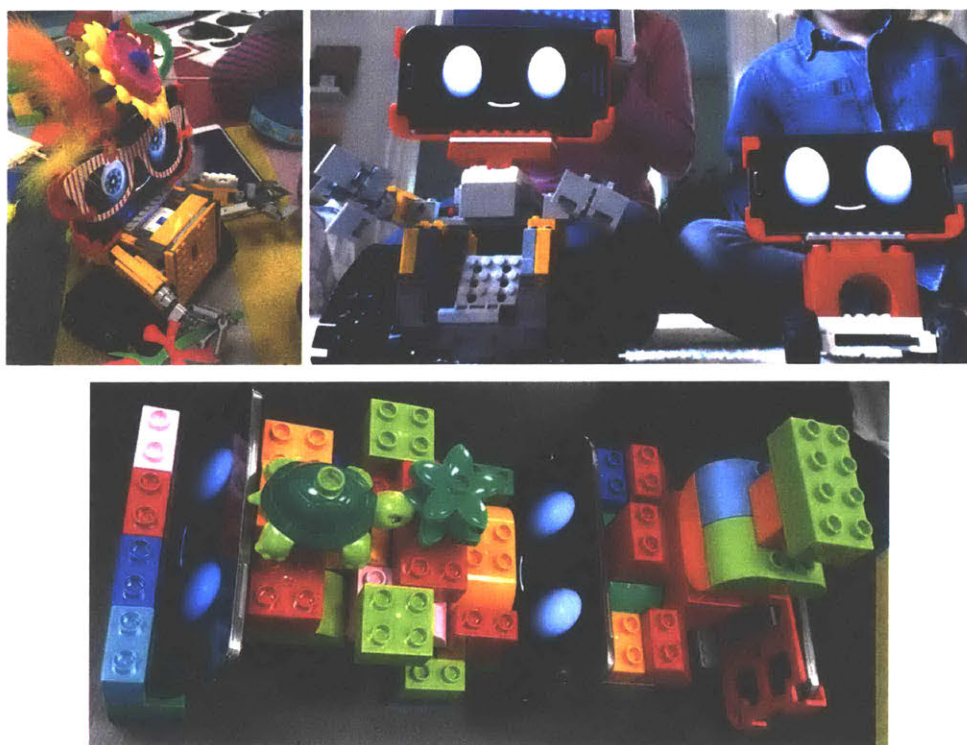
Figure 17: Children's PopBot designs

Exploring Real-World AI Activities

During the real-world activities, children interacted with commercial examples of AI algorithms (see Fig. 18) to connect the functionality of the PopBots to other use cases. To explore supervised machine learning, children used *Google Quick Draw*[5] to see how an AI system can learn to recognize

---

[5] Google Quick Draw https://quickdraw.withgoogle.com/

drawings from millions of examples. Google's *Quick Draw* app was developed to generate training data for Google's doodle database. The app prompts the user to draw a particular figure in 30 seconds or less and then tries to guess the picture. If the app's guess matches what the user was assigned to draw, the user gets a point and moves to the next drawing. In a noncompetitive version of the app, we had time to slow down and see what the robot's 20 best guesses were for a given drawing. Children and I discussed how the computer was able to make its guess by learning from a training set of hundreds of thousands of other drawings. The technology underlying Google *Quick Draw* is a neural net, so it is less transparent than other algorithms. A really useful activity was to think of ways that we could help the computer guess ("If we want the computer to recognize a foot, then we could draw toenails!"), as well as ways that we could trick the computer. Another helpful activity was playing pictionary, the child against the researcher. The child could then reflect on his or her own way of figuring out what a drawing represented ("Oh, I've never seen it drawn like that before." "I knew it was a shoe because it had the long shoe laces.") Children could then understand the computer's successes and mistakes a bit better. Older children definitely had an easier time with this than younger children. The most confusing part for children was why the robot needed so many examples. It was somewhat hard for the children to grasp that the computer had never seen the object before in real life and was learning from scratch.

As an example of generative AI, we used the *Prisma* app[6] to transform pictures that we took with a mobile device into versions with different styles. *Prisma* provides 20+ examples of different styles that one photo can be transformed into, and provides the ability to swipe between the original and the edited photo. The app takes a few seconds to transform a picture into a new creation, sometimes longer. The picture transformations change the colors, highlights, and sometimes the backgrounds of photos. Children enjoyed taking pictures with the app and choosing the best style for them. Often, they wanted to print out a lot of the pictures to keep. To understand how the app worked, we talked about what kinds of things were different between the original and the new pictures. *Prisma* also works with neural nets, so it was not possible for children to completely see the inner workings of the app or create their own styles on the fly. While Prisma was engaging, it was unclear that children were at all grasping how the algorithm worked. For example, children would say that the computer added two pictures together to make another picture. However, they could not identify that the computer selected particular features to change them. One way to make this more transparent would have been to have simpler transformation available. For example, transformations that would just change the color or the depth of lines. Then children could add these transformations together to incrementally move from the original picture to a new creation.

---

[6] Prisma AI https://prisma-ai.com/
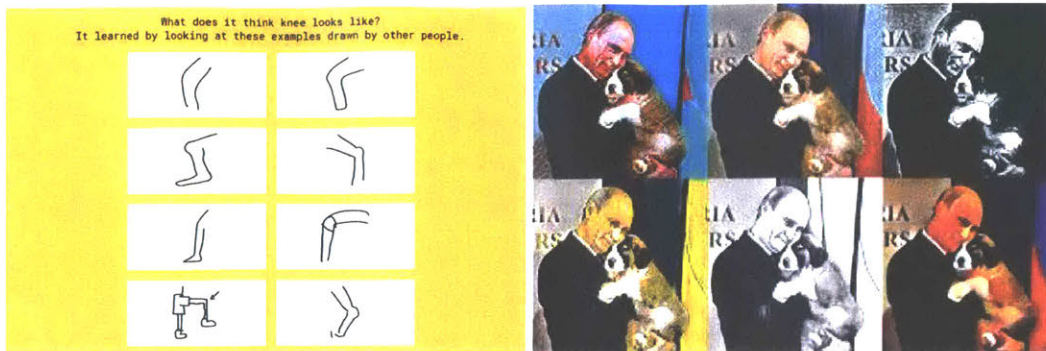
Figure 18: Real world examples of AI: Google Quick Draw and Prisma AI

Iterations on the POP Curriculum

One of the main goals of the pilot studies was to iterate on the designs of the interfaces and activities that were presented in Chapter 4. Fig. 19 shows what the first interfaces looked like. Children's feedback allowed me to make the interfaces simpler and more intuitive, change the icons to more understandable pictures, and expand or condense the activities as necessary.
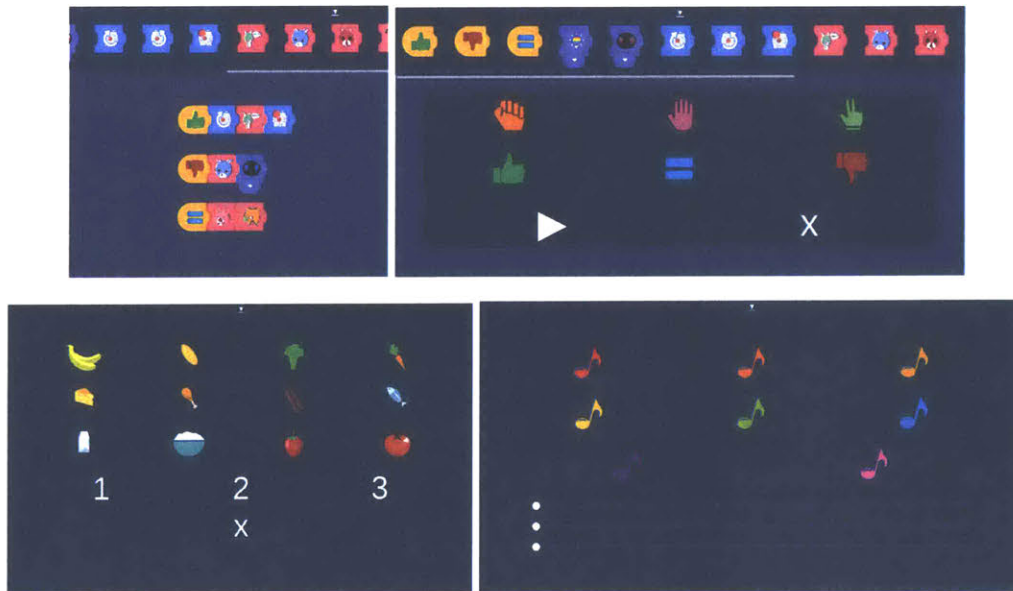
Figure 19: Early interfaces of the POP curriculum

Changes to interfaces after Pilot Studies:

- Removal of categories of blocks. Originally, blocks were sorted into six different categories (motion, music, animations, trigger, control, and looks). However, children would get stuck in

52

one category and not know where to find particular blocks. Having all of the blocks in a line was easier to navigate.

- Changing bunny icons and motor icons. The pink animation blocks in Fig. 19 (top) use bunnies to represent different emotions of the robot. This was done to remove abstract away the character that would do the animation. However, the icons were small and confusing and they were all changed to robots. Also, the motor icons originally were "start motor X" and "stop motor X." Children could place a wait block or animation in between to control how long the animation ran. However, children did not understand the icon and the wait block was difficult to introduce at the very beginning, especially for younger children.

- Separate pages for each activity. Fig. 19 (top, right) shows that the interface to play rock, paper, scissors was a pop-up window over the programming interface. Children understood what they were supposed to be doing more when every part of the activity had its own page.

- Making the food interface drag and drop. The original food interface required children to click a food then click a number (group 1, 2 or 3) to sort foods. However, younger children could not navigate this interface by themselves at all and the researcher had to do it for them.

- Making the music interface have a "record" function. Children would often want to play the same song repeatedly with different parameters, but children had a hard time remembering their song. Children were also more familiar with a music interface that used a keyboard rather than the buttons that are currently shown.

The AI activities and assessments discussed in the prior section, as well as the robot's behavior, were iteratively developed and refined during these first workshops. To evaluate children's conceptual understanding, I asked children to describe the activity in their own words. The discussion questions for each activity are listed below in Table 8.

Activity	Pilot Study Discussion Questions
Rule-Based Systems	How does the robot learn the rules of the game? How does the robot figure out what move to make every time?
Supervised Machine Learning	How did the robot sort things into groups? When was it easy and when did the robot get a lot of things wrong?
Generative AI	How did the robot decide which notes to play What does it mean for a robot to be creative? Is it creative the same way you are?

Table 8: Discussion questions for Pilot workshop activities

Out of all the activities, children in the pilot workshops grasped the rule-based systems activity through rock, paper, scissors most easily. In the pilot tests, children taught the robot the rules as they played. When the robot encountered a rule it did not know, it would ask the child, "Did I win?" and the child would use the tablet to respond yes, no, or it was a tie.

> Researcher: How does the robot learn the rules of the game?
> Lily (6-years-old): I taught the robot the rules of the game...it would learn as I go.
> Researcher: So who won more? You or the robot?
> Lily: The robot.
> Ivy (6-years-old): The robot. Well, at first me, but then the robot kept saying 'I think you will put rock' and I had put rock so it won. The robot got smarter the more we played.
> Researcher: So who's smarter now? You or the robot?
> Lily: Well...maybe the robot. But I taught it everything it knows. So actually I'm still smarter for now but I think the robot can get a lot smarter.

Children understood supervised machine learning somewhat, but needed a lot of hints from the robot to verbalize their understanding of how the algorithm worked. At first, children would just say that the robot knew because it was "smart." It took some prompting from the researcher to help the child dissect what the robot actually did and how it came to know where all of the foods belonged. One problem with this activity during the pilot sessions was that there were only 12 foods (there were 24 in the final study) which made it harder for children to see when the robot was grouping similar foods, such as all fruits, together into one group.

> Researcher: ...How did the robot figure out where the bread went?
> Jesse (four years old): Because he's so smart that he knowed it.

> Researcher: Do you think the robot just figured it out because it was
> smart?
> Jesse: It learned it.
> Researcher: It learned it by itself? How did it figure it out.
> Jesse: He figured out which food went to there. ...He figured it out
> because the first time he didn't do it but he got smarter and
> smarter and smarter until he found out where it went.
> Researcher: So the robot just learned the rules like in rock, paper
> scissors?
> Jesse: No the robot could guess it.

Finally, children understood the generative AI activity the least. In the pilot study, the generative AI activity just consisted of the robot playing back the same notes or throwing in new notes with increasing degrees of randomness. However, the randomness was too unrelated to the child's input, and children did not see how the robot's output song tied back to their own. In addition, what made this activity particularly difficult was just how much fun it was to make noise with the tablet and the robot. What was interesting in the children's descriptions was the agency that children gave the robot. The activity was set up as call-and-response, first the child would play and then the robot would. Even if the children played extremely long sequences of notes, they would wait for the robot to finish its turn before playing more notes. It is possible that they might have been mesmerized by watching the robot's eyes change colors as they played. Also, even though children did not see any sense in what the robot played, they still listened to its song. In their statements, they referred to the resulting music as though it were an intentional creative endeavour by the robot.

> Researcher: ...How did the robot make his song?
> Carly: Well when the bar is like this (indicates to the tablet, sets
> "randomness" variable to off) the robot just plays the same
> song. But sometimes (sets the "randomness" variable to 10) it
> makes its own music.
> Researcher: How does the robot make its own music?
> Carly: It just picks whatever notes it wants and plays them.
> Researcher: So the robot can just play whatever it wants like you?
> Carly: Yes, it just likes to play a lot of notes.

5.2 Second Pilot Workshop: Transforming the Activities for the Classroom

I did the second pilot workshop in a local STEM center. The workshop was a one-time, hour long class with 11 children (ages 4 to 6). The goal of the workshop was to transition the curriculum to a classroom context. The key difference between the first pilot and the second one was that children worked with the robot in a group, rather than each having their own robot. This allowed me to test the ability of the activities to engage small groups of children without an instructor always present and test the how well children could understand the robot's disclosures about the algorithm. Besides me, there were parents and two other teachers, employees at the STEM center, in the room. The

teachers had experience running technology education classes with young children. During this pilot, I recorded children's interactions with the tablet and robot, made observations about children's understanding, and took notes from the children, parents, and teachers about what they saw as successful or not successful.

The 11 children self-selected into groups of three or four and I assigned each group a robot. First, as an entire class we discussed what artificial intelligence means and how it relates to robots. Then, I explained the activity for the day -- teaching a robot the rules of rock, paper, scissors. The class discussed how to play rock, paper, scissors and what the best strategies are for winning. Children mentioned different strategies such as always playing one move over and over, but many were convinced that you just had to get lucky or read minds.

Next, the children played the rock, paper, scissors activity with the robot in their groups. Children would take turns choosing a move and keep score of who was winning more. Like in the first pilot study, rather than being explicitly taught the rules, the robot would learn the rules as it played. When the robot encountered a rule it did not know, it would say, "You played X and I played Y. Did I win?" Children would provide the answer and the robot would save the new rule. The robot would also remember children's past moves to guess their next move.

Children played with the robot for about 15 minutes. During this time, the other teachers and I spent a few minutes at each table guiding children and asking what the robot was learning and how. After this activity, we came back together as a class and I asked the discussion questions from the pilot activity: how does the robot learn the rules of the game and how did it guess what you would play next? Finally, children got the chance to program the robot for free play. Unlike the previous pilot workshops and the final study (discussed in Chapter 6), this pilot did not include exploring children's perceptions of AI. The main goal was to refine the curriculum, instruction protocol, and activities.

5.3 Lessons from Individual versus Group Work

In an early education classroom, every child does not have a dedicated instructor to guide them through activities. They also did not each get their own robot. Therefore, the POP curriculum needed to be flexible so that children could work in small groups without the need for constant oversight by an instructor. Changes to the curriculum included making the PopBots more engaging and building in more scaffolding. The following insights came directly from the feedback from instructors and parents at the local STEM center after the second pilot workshop.

Insights for Classroom Scaffolding

Children had a hard time paying attention to the words that the robot said unless their attention was called to it, which made it difficult for them to understand the activities on their own. The robot communicated how it made decisions through speech, for example, "I think you will put rock so I choose paper." Rather than focusing on the robot's speech, children focused instead on how many games they won or lost and on making sure that everyone had a fair turn. Therefore, when I asked the children how the robot played the game or why the robot was winning more often, they would not have an answer. I believe that this was due in part to the dynamics of having other children around and the robot not being as compelling as the other children. The robot's voice was expressionless, and it did not command attention as much as another person. This was not as much of a problem in the smaller studies.

The solution to this was to go through the entire activity once as one large group with the researcher pointing out all of the things that children should notice. Then, children would break into their groups to interact with the robots. Finally, the researcher would again ask children to recall things the robot said to explain its actions.

Children in groups sustained their interest in the activities for longer periods of time than children working individually. In groups, children formed mini-games to keep the activities interesting -- for example keeping score of who won against the robot the most. This meant that children got to experience more of the robot becoming really, really good at predicting their future moves. Children in individual studies often had to be coaxed to go through enough iterations to observe the robot's matured expertise.

Lastly, there were not enough robots such that every child could have a robot, and there were not enough adults present such that every child could be watched at all times. This made for situations where children would become distracted and wander off into another activity. During the time that it took to get a child to focus on the activity and complete it, other children would have finished already and begin to lose interest.

One way to prevent this from happening was to make sure that child had a role to contribute at all times. For example, one person could be figuring out the rules to teach the robot while the other was thinking about how to program the robot to win or lose. Another helpful tactic was to have several short status check breaks where children talk about what they observed so far. Children who had begun drifting would immediately re-engage when asked a question.

Another limitation of the classroom approach is that there are a limited number of instructors for each child. Therefore, it was best to have checkpoints in activities where children could complete one simple task before moving on. In the final curriculum, I added checkpoints for each activity on a separate web page so that children would not move on until everyone was ready. To make this less boring for children who were ahead, I also embedded open-ended activities into each checkpoint. In the final study, each activity had two or three checkpoints that took no more than 5 minutes to complete.

Assessing each child's understanding was difficult in the classroom because there was no time to have each child explain what the robot was doing in their own words. Therefore, I developed multiple choice assessments for the children to take. However, it is important to remember that many of the children cannot read and have very limited experience taking tests. This meant that the questions had to be extremely well-worded and straightforward, with responses chosen carefully to capture different lines of reasoning.

Developing Robot Cues for Scaffolding

After the pilot study, I created a state machine for each of the activities that would give the child occasional feedback about what the robot was thinking. For example, in rock, paper, scissors the robot would point out to the child when it had won several times in a row to help the child see that the robot was improving. The robot also would more actively encourage the child to keep playing after each game saying, "This is fun! Let's keep going" or "I want to play again."

In the supervised machine learning game, the important thing that child had to grasp was that the robot was making comparisons between foods that it knew and foods that it did not know. Therefore, this state machine included information like "I think corn is most like <names k=3 labelled foods>, so corn must go in <whichever group>." The child could use the tablet to further prompt why it thought things were similar, which would show the child what kinds of things the robot knows and what it does not. Also, the robot admitted to making mistakes, "Oops, I mean the <other> group" if its prediction did not match an input the user gave later.

Finally, in the generative AI activity the robot displayed curiosity to encourage the child to notice subtleties like which sliders controlled what aspects of the music and what happens when the sliders are both in the middle. The robot also declared with it was its turn and when it was done, to encourage more turn-taking with the music.

58

5.4 Summary

These pilot workshops revealed that children between the ages of 4-7 have the ability to complete the PopBots curriculum and explore AI concepts with the activities and materials. Children in the initial pilot workshops had the opportunity to see real-world examples of AI, interact with the PopBot, and internalize the knowledge by answering questions about how the robot worked. In the classroom pilot workshops, I observed that the Pop Curriculum designed for individuals was not effective with more children and fewer teachers. My insights from these workshops allowed me to revise the curriculum and add appropriate scaffolds for a classroom context. This included making each activity shorter with more intermediate goals as checkpoints. Also, the curriculum changed to allow a lot more interaction with the PopBot. These changes were important to keep children's attention and make the AI transparent enough for them to understand.

6 Evaluative Study Design

The main questions pursued in the final evaluations explore how and what children learned about AI from the Pop curriculum and what impact the curriculum had on their perceptions of AI and on identifying themselves as engineers. I predict that the amount that children learn from the curriculum is dependent on their grade and ToM development rather than engineering skills or any predisposition to mathematics. The curriculum focused on helping children build intuition by understanding the PopBot as a social and "thinking" being. In addition to the children's skills, my own skill in developing activities that make AI transparent are evaluated here. In cases where development is not the main factor for children's understanding, activity design is the next candidate. All activities were experimental and need revision. By analyzing how children came to grasp different concepts, I can begin to understand what kinds of activities will or will not work for this age group and developmental stage.

I suspect that children's perception of AI and themselves will also be affected by going through the POP curriculum. My curriculum gives young children a vocabulary and a set of examples on which to base their reasonings about AI. Whereas, before the curriculum, children's perceptions of AI may be based on a collection of memories about past experiences, how family and friends talk about AI, or how media portrayals of technology -- after completing the POP early AI curriculum, children will be able to reason about AI with a basic understanding of how the algorithms work. My hope and expectation is that this increased understanding, this hands-on engagement of coding, training and playing with an AI social robot, will lead to children feeling empowered to be creators of this technology.

In this final study, my goal was to first validate my ToM assessment, then evaluate the effectiveness of the POP curriculum in terms of children's performance on the AI assessment, analyze developmental differences in children's understanding of AI concepts, and analyze how children's perceptions of AI devices and perceptions of themselves change before and after completing the curriculum. For statistical analysis I set my significance level $\alpha=0.05$ and for effect size $\varphi=0.1$ is a small effect, 0.3 is a medium effect, and 0.5 is a large effect.

6.1 Participants

To evaluate the full curriculum, I recruited 80 four to six-year-old children from four schools in the Greater Boston Area. I worked with five different classrooms in all. Three schools were a part of the

Somerville Public School System and one was a Cambridge Private School. There were six children from classroom A, 16 children from classroom B, 19 from classroom C, 17 from classroom D1, and 22 from classroom D2. Children in classroom D2 had an advantage because they were already discussing robots in their classroom. None of the other children reported experience in robotics or computational thinking except for one student from classroom A who played Minecraft. Classroom D1 contained only Pre-K children, Classrooms A, C, and D2 contained only Kindergarten children, and classroom B was mixed with Pre-K and Kindergarten.

	Female	Male	
Pre-K	14	13	27
Kindergarten	25	28	53
	39	41	80

Table 9: Participants by gender and grade

6.2 Materials and Methods

The order of the curriculum remained consistent across the different classrooms. First, children completed the pretest assessments (AI perception and Engineering and Science Attitudes). The assessments were delivered on a tablet or on paper using the materials detailed in Appendix A. Then, children went through the POP curriculum:

1. Introduction to Programming with the PopBots
2. Rock, Paper, Scissors
3. Food Classification
4. Generative AI

The activities and assessments were exactly as described in Chapter 4. Also, see Appendix B for a more detailed protocol including the scripts the researcher used. Each activity was designed to teach the core material in 10-15 minutes, after which we would take the corresponding AI assessment. After all of the AI activities were done, children completed the post-tests (AI perception and Engineering and Science Attitudes assessment) on paper or on tablet.

From the assessments, I collected quantitative data about children's responses. To support these data, I also made observations while delivering the curriculum. In particular, I noted children's questions and observations, their understanding of the activities as they went through them, and

any reasonings about their responses that they could verbalize. Each session was also audio or video recorded for later analysis.

The exact procedure and timing for each classroom varied depending on the amount of time I had each day, the number of days that I could come to the class, and the number of children that I had to work with at once. In classroom A, I worked with all six children at once for an hour each session over six days. Given that the activities were meant to take 15 minutes, children in classroom A got a lot of free time to play with the robot. In classroom B, I worked with groups of three to four children at a time for 15 minutes each group over 7 days. I worked with all of the children in classroom C all at once, but divided them up into smaller groups of 3 or 4. Each session in classroom C went for one hour over the course of four days. On the first day we did the pretests and introduction to programming and on the last day we did generative AI and post-tests on the same day. Children in this classroom recorded their responses to the assessments on paper rather than tablets. In the D classrooms, I only had two hours in each classroom. Therefore, pretests and post-tests were conducted by the class's regular teachers. In D1 and D2, I worked with half of the class at a time for 30 minutes each both days. We did two activities each day, first the introduction and rock paper scissors then food classification and generative AI. Half of the D2 classroom did not complete generative AI because we ran out of time. Despite these differences in the amount of time children could spend doing the curriculum, I found no significant differences between schools on any of the assessments.

Given that the curriculum was run outside of school hours, it was somewhat common for children to leave early or miss a day. The curriculum activities do not build on each other so children could miss one and still be prepared for the next activity. Children also primarily worked together on the activities which meant that children who missed days could leverage the knowledge of their peers.

6.3 Assessments and Measures

To validate my ToM assessments, I compared children's scores on the three tasks (knowledge access, false belief, and explicit false belief) to the expected values determined by Wellman and Liu using Chi-square analysis. I also looked at differences in ToM development between the two grades and analyze them using a Chi-square test. I expected that my results for the ToM assessments will be similar to Wellman and Liu and that Kindergarten students will have slightly higher scores that Pre-K students.

To validate the POP AI assessments, my measure of success was the number of correct responses on the target questions; the control questions were not used in this analysis. I determined significance

using a Chi-square test to compare the number of children who responded correctly to the expected number of children who would respond correctly by chance (the inverse of the number of choices for that question, generally either 50% or 33%). I also looked at differences in assessment performance based on the ToM assessments and children's grade using Chi-square.

To analyze the responses of children on the AI perception test and Engineering and Science Attitudes assessment, I used a Chi Square test on three groups (children who responded Agree, children who said Neutral, and Children who responded Disagree to compare the distribution of children's answers against a chance distribution of Agreem Neutral, and Disagree being equally likely. Next, on the AI perception test, I compared children by grade and different ToM understandings using Chi-square tests and Cramer's φ. For children's attitudes towards engineering, I looked at differences between grade and gender again using Chi-square and Cramer's φ. Finally, I analyzed how children's answers changed from pre to post-test using McNemar Chi square tests.

7 Qualitative and Quantitative Results

In this chapter, we present data and analysis to address the following research questions of this thesis:

Research Question

1. *How do young children come to perceive and understand AI concepts? What factors influence this process?*
2. *What can children learn about key AI concepts such as rule-based systems, supervised machine learning, and generative AI by coding and interacting with a social robot platform? How does this inform a developmentally appropriate curriculum for AI and its attributes?*
3. *How does learning about AI change children's perception of artificially intelligent devices and themselves as engineers?*

7.1 Theory of Mind Assessment

The ToM assessment questions presented in Appendix A are from Wellman and Liu (2004). Children's scores were graded in the same way as Wellman and Liu (2004), where a question is only correct if a child gets both the target and control question correct. Data for this sample of 4-6-year-old children were quite consistent with Wellman and Liu's data for 3-5-year olds. The knowledge access question was the easiest for children to get correct, at 75%. The false belief questions proved to be more

difficult, with both at 55% of correct answers. I used a Chi-square test against the percent correct in this sample and in Wellman and Liu to determine if there was a significant difference between the two. I found that there was no significant difference for any of the components of Theory of Mind. This suggests that the ToM assessment delivered on the tablet is comparable to the one that Wellman and Liu developed.

One of the main findings in Wellman and Liu's paper was that the development of various ToM components was highly correlated with age. Therefore, I looked at the differences between Pre-K and Kindergarten students in terms of the Theory of Mind assessment (Fig. 20). I found that there was little difference between the two grades. This is most likely due to the fact that Pre-K and Kindergarten classes both have 5-year-old children. This is an interesting finding because later any differences based on grade will suggest that factors other than the ToM components that I measured may be playing a part in differences between the grade levels.
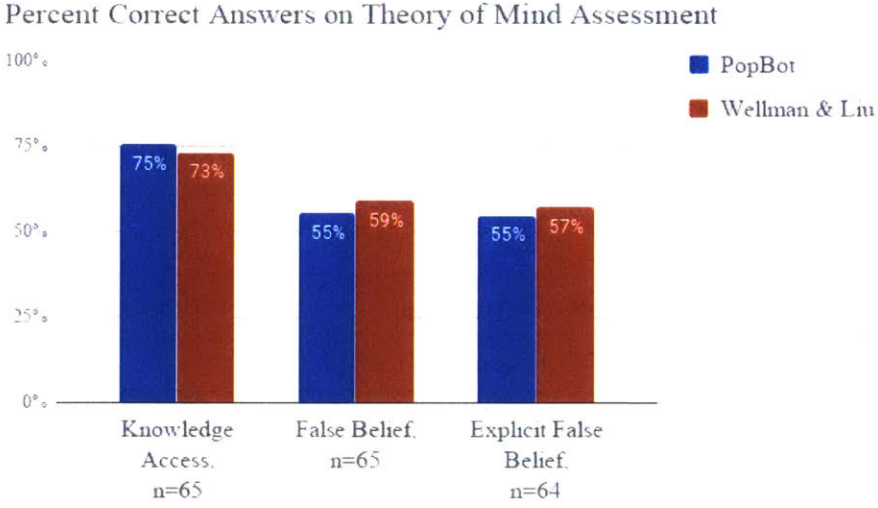
Figure 20: Percent of students who correctly answered Theory of Mind Assessments, this study versus Wellman & Liu, 2004

ToM Assessment	Number Correct	$X^2(df, N)$	p
Knowledge Access	49	$X^2(1, 65)=0.159$	0.690
False Belief	36	$X^2(1, 65)=0.125$	0.723
Explicit False Belief	35	$X^2(1, 64)=0.032$	0.859

Percent Correct Theory of Mind Answers by Grade
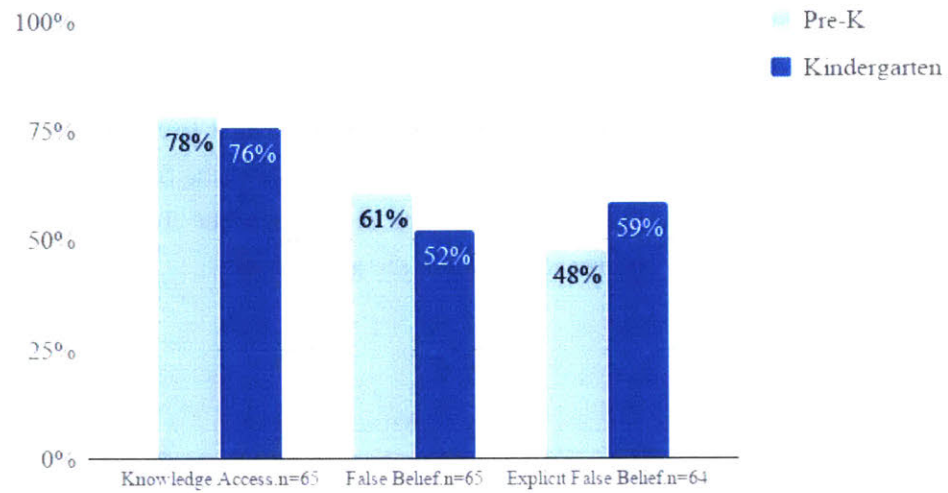
Figure 21: Percent correct answers on ToM assessment separated by grade

ToM Assessment	Number Correct		X^2(df, N)	p
	Pre-K	Kindergarten		
Knowledge Access	18	31	X^2(1, 65)=0.058	0.810
False Belief	14	22	X^2(1, 65)=0.433	0.510
Explicit False Belief	11	24	X^2(1, 64)=0.682	0.409

7.2 AI Assessment of POP Curriculum, and Relation to Grade and Developmental Factors

The POP curriculum AI assessment questions are in Appendix A, Section 1. Since this experiment had no control group, I used Chi-square tests to compare the number of questions that children got correct to the expected number they would get correct if they were just randomly guessing. Out the 10 target assessment questions (13 questions total, but three of them were control questions), statistical analysis shows that children performed better than chance on 7 of the assessment questions.

Rock Paper Scissors Assessment Performance

The most successful activity was *rock, paper, scissors*. The number of children who got the 4 assessment questions correct was significantly different from chance on every question (Fig. 22). Children understood how the robot used the rules to say who won or lost (83%), even when the rules were reversed (71%). They also understood that once the robot made a prediction, it would use that information to choose its next move (76%). The most difficult question was about the robot predicting what its opponent would play next, about 59% of children got this question correct. Understanding this concept required children to play multiple games with the robot to see the robot start to make predictions.

On one of the questions, grade (Pre-K or Kindergarten) was a significant factor in the number of children who got the question correct. Fig. 23 (left) shows that Kindergarten children were significantly more likely than pre-K children to understand how the robot chose its next move given a prediction 85% vs. 58% ($X^2(1, 65) = 5.96$, p=0.0146, φ=0.30). Both groups spent the same amount of time doing the activity, but more Kindergarten children took note of the robot's comments about the progression of the game. Before the robot made a move, the robot would try to guess what the child would play. More Kindergartners exclaimed that the robot had guessed correctly (or incorrectly) and was therefore cheating or somehow doing magic. The most common incorrect answer that Pre-K children gave was rock, almost 40% of the Pre-K children gave this answer. Perhaps children were confused about what the question was asking (e.g., they suggested a move that would make the robot lose rather than win).

In addition, children who, according to the ToM assessment, understood *knowledge access* were significantly more likely to correctly answer the question about who wins when the robot has been taught the normal rules of *rock, paper, scissors*. Fig. 23 (right) shows that 88% of children who

understood knowledge access got this question correct compared to almost 55% of children who did not understand knowledge access ($X^2(1, 54) = 6.64$, p=0.0097, φ=0.35). It makes sense that children without an understanding of knowledge access might have a difficult time drawing a connection between what the robot knows, and how it interprets the output of the game. However, there was no difference based on knowledge access skill for a similar question where the rules of the robot were reversed. A possible explanation for this is that seeing a similar question the second time gave children the opportunity to reflect on how the robot interprets knowledge that it is given.
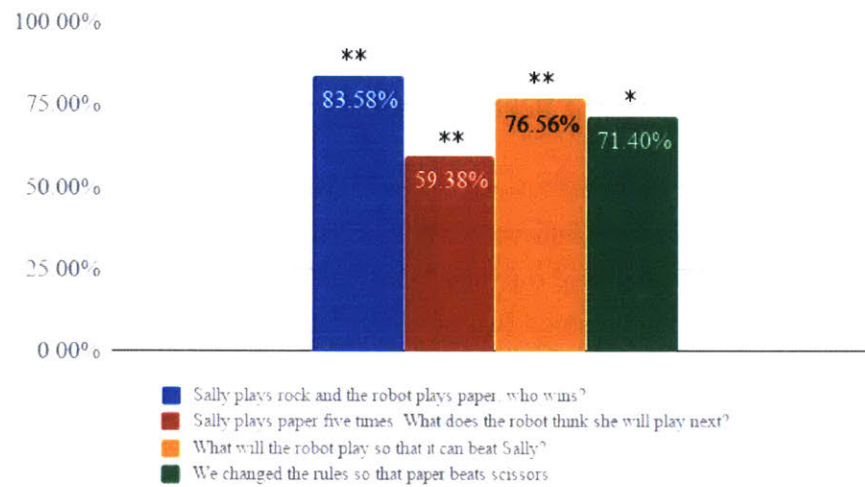
Figure 22: Percent of correct answers on the rock, paper, scissors assessment

Question	X^2(df, N)	p	Cramer's φ
Sally plays rock and the robot plays paper, who wins?	$X^2(1, 67) = 16.3778$	< 0.01 **	0.49
Sally plays paper five times. What does the robot think she will play next?	$X^2(1, 64) = 8.6557$	< 0.01**	0.37
What will the robot play so that it can beat Sally?	$X^2(1, 64) = 24.7172$	< 0.01**	0.62
We changed the rules so that paper beats scissors. Sally plays scissors and the robot plays paper. Who wins?	$X^2(1, 63) = 5.6438$	0.0175*	0.30
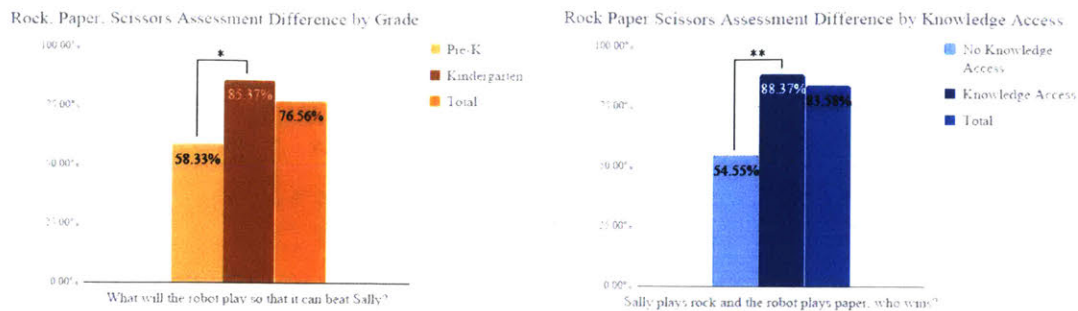
Figure 23: Differences in percent correct on rock paper scissors assessment by grade (left) and understanding of knowledge access (right). Table in Appendix C

Supervised Learning Assessment Performance

For the food classification activity, the number of correct answers was significantly greater than chance for two of the three questions (Fig. 24). On the question about nearest neighbor associations, 83.6% of children answered correctly. Then, on the question where the robot uses the nearest neighbor association to classify a food, 85% of children answered correctly. Children struggled with the question about where chocolate should be placed when tomato and strawberry are put in the good group (45%). Many children assumed that the robot knew that chocolate was unhealthy. However, in the robot's "mind" it only knew things like chocolate's color, food group and grams of sugar. Since the robot only had positive examples, it would guess that everything belonged in the healthy group. Children's confusion may have been a failure in recognizing that because the robot's database was limited, it had no prior intuitions about what was good or bad.

I found that grade played a role in children's understanding of this question. Kindergarten children were significantly more likely to get this question correct (97.6%) compared to Pre-K children (38.5%; $X^2(1, 65) = 25.38$, p<0.001, φ=0.68) (Fig. 25). A likely explanation for this strong difference is that kindergarten children were more equipped to understand how the robot's knowledge database worked. Before playing with the robot, children sorted foods into different groups and identified some of the common properties of the foods, for example a lot of the fruits and vegetables were in the healthy group. Kindergarten children found it easier to identify these common properties, suggesting that they had a stronger grasp of the robot's knowledge base.

68

Percent Correct Answers on Food Classification Assessment
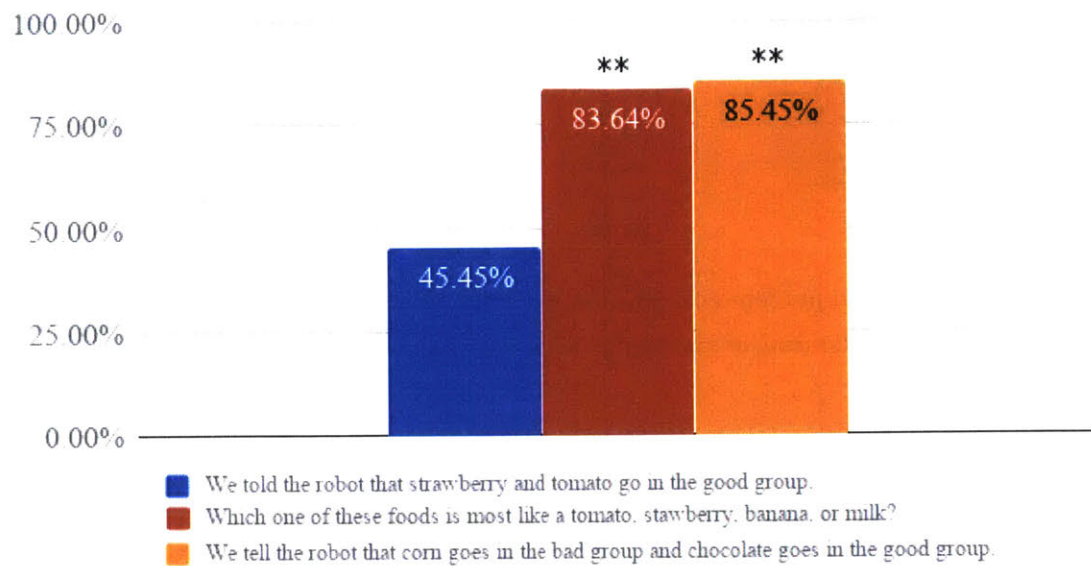
Figure 24: Percent of correct answers on food classification assessment

Question	X^2(df, N)	p	Cramer's φ
We told the robot that strawberry and tomato go in the good group. Where does the robot think chocolate will go? In the good group or the bad group?	X^2(1, 55) = 0.3277	0.567	–
Which one of these foods is most like a tomato, strawberry, banana, or milk?	X^2(1, 55) = 29.697	< 0.01**	0.73
We tell the robot that corn goes in the bad group and chocolate goes in the good group. Which group will the robot put bananas in? The good group or the bad group?	X^2(1, 55) = 15.1276	< 0.01**	0.52

Food Classification Assessment Difference by Grade

Which one of these foods is most like a tomato, stawberry, banana, or milk?
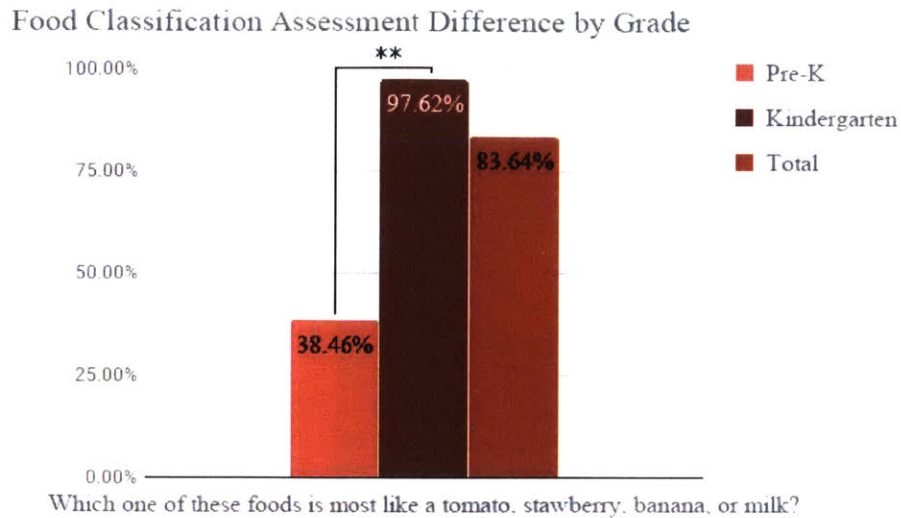
Figure 25: Percent of correct answers on food classification assessment separated by grade. Table in Appendix C

Generative Assessment Performance

Finally, the most challenging activity was the generative music activity (Fig. 26). Children were only significantly more likely than chance be correct on one question -- the one about the robot changing songs to sound different (83.3%). Children did not seem to register that the robot played the same notes as they did when the speed and tune sliders were set in the middle (57%). They also did not make the connection between the robot's song and their song input. On this particular question, children were significantly likely to get the question incorrect (14%). In this case, the activity did not seem to convey the AI concept clearly.

There was a difference between children who got the *explicit false belief question* on the ToM assessment correct and those who did not. Children who understood explicit false belief were significantly more likely to understand that when the robot's parameters were set to the middle of the scale (or 0 improvisation), then the robot would play the same song back to the child. The difference was 72.2% for children who understood explicit false belief and 35.7% for those who did not ($X^2(1, 32) = 4.27$, p=0.0389, φ=0.37 (Fig. 27)). If understanding explicit false belief was truly the factor that made the difference, then perhaps this is because children were more prepared to grasp the idea of parameters and how they led to different kinds of output that the robot could produce. However, as is the case with all of these questions, it is also possible that children who performed

70

better on the false belief task also had a higher level of language understanding and understood the question better than their peers.

Percent Correct Answers on Generative Music Assessment
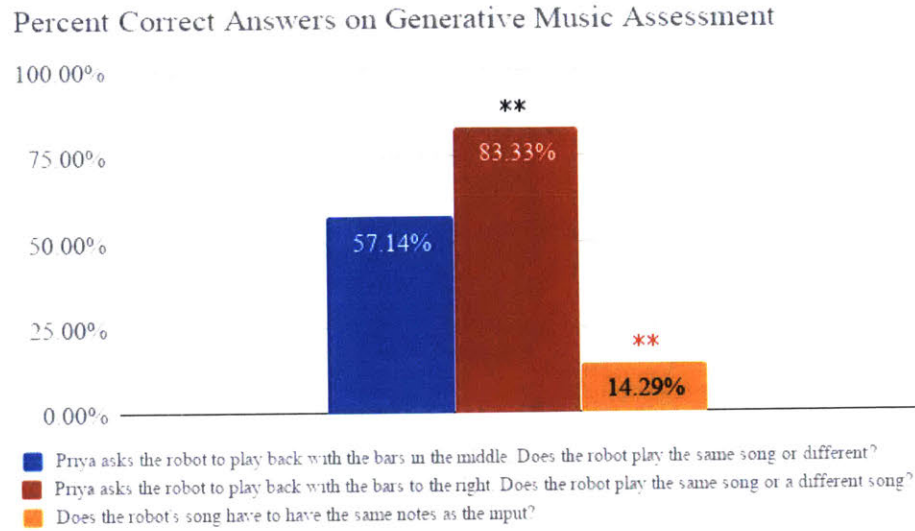
Figure 26: Percent of correct answer on generative music assessment

Question	$X^2(df, N)$	p	Cramer's φ
Priya asks the robot to play back with the bars in the middle. Does the robot play the same song or a different song?	$X^2(1, 42) = 0.4308$	0.5116	–
Priya asks the robot to play back with the bars to the right. Does the robot play the same song or a different song?	$X^2(1, 42) = 10.5$	< 0.01**	0.5
Does the robot's song have to have the same notes as the input?	$X^2(1, 42) = 12.2807$	< 0.01**	0.54

Generative Music Assessment Difference by Explicit False Belief

Legend:
- No Explicit False Belief
- Explicit False Belief
- Total

Bar values: 35.71%, 72.22%, 57.14%

X-axis label: Priya asks the robot to play back with the bars in the middle. Does the robot play the same song or different?
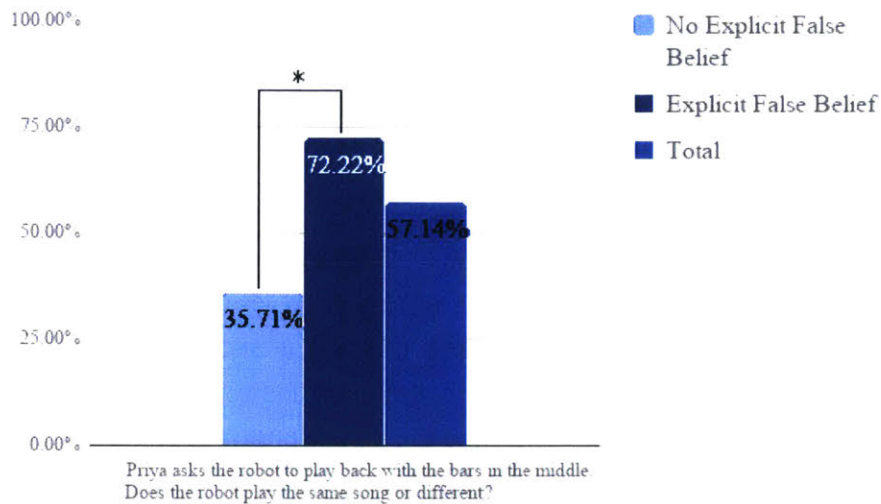
Figure 27: Percent of children's correct answers on generative music assessment separated by children's understanding of explicit false belief. Table in Appendix C

7.3 Children's Perception of AI Robots

The pre- and post-tests of children's perception investigated children's views of AI robots and how those views might have changed after interacting with the PopBots. Given that again, there was no control group, we cannot explicitly say that learning about AI alone led to changes.

The *AI Perception Questionnaire* probed how children think about the PopBot as a cognitive and social entity (Druga et al., 2017). The first two questions dealt with the robot as an intelligent being -- do robots always have to follow rules and can they learn? Cognitive beings may have rules to follow, but they may also be able to decide whether to follow those rules or not. The next question asked how children perceived the intelligence of the robot. In prior studies, children in this age range were quite indecisive about the intelligence of robots so it was interesting to see if learning about AI helped children decide one way or another (Druga et al., 2017). The last two questions asked how children saw the robot as a social being and what level of maturity they would assign to it. In these questions, children of this age group tend to already see robots as social beings, but it is not clear if children see them as having any particular age.

To analyze children's answers in the pre- and post-tests, I used a Chi-square goodness-of-fit test test to compare the distribution of children's responses to an expected distribution where the responses are split evenly across the three options, Agree, Disagree, and Neutral.

Pretest Results

Using a Chi-square goodness of fit test between the three groups, I found that children's responses on all pretest questions were significantly different from an even distribution. I further analyzed the data using Chi-square goodness of fit tests between two groups at a time with $\alpha=0.017$. Most children agreed that robots always have to follow the rules (62%). Very few children disagreed with this statement (3%), there was a significant difference between the number of children who disagreed and those who were neutral ($X^2(1,24)=15.04$, $p<0.0001$) or agreed ($X^2(1,61)=31.6$, $p<0.0001$). Most children also said that robots can learn (65%). On this question, there was a significant difference between children who agreed and those who disagreed ($X^2(1,50)=19.22$, $p<0.0001$) or were neutral ($X^2(1,53)=14.8$, $p=0.0001$). These results suggest that children believed that while the PopBot was able to expand its intelligence, it still chose or was required to operate within set rules.

Most children were neutral in their opinion of robots as smarter than them (64%) and in their perception of robots as people versus toys (62%). The children's indecisiveness about whether robots were smarter or not was also seen in an earlier study. Children in the younger age range, 4-6-years old, were less confident about the robot's intelligence than their 7-10-year-old counterparts (Druga et al, 2017). Finally, we also see that most children saw robots as children or something between adults and children. Very few children saw robots as adults (10%). In the post-hoc tests, there was a significant difference between the number of children who disagreed with this statement than children who agreed ($X^2(1,34)=12.98$, $p=0.0003$) or said neutral ($X^2(1,34)=12.98$, $p=0.0003$). When reasoning about this answer, children referred to their previous attributions. They believed that the robot was like an adult because adults follow rules. However, they also believed the robot was not like an adult because adults are usually smarter than they are.

Children's Perception of AI: Pretest
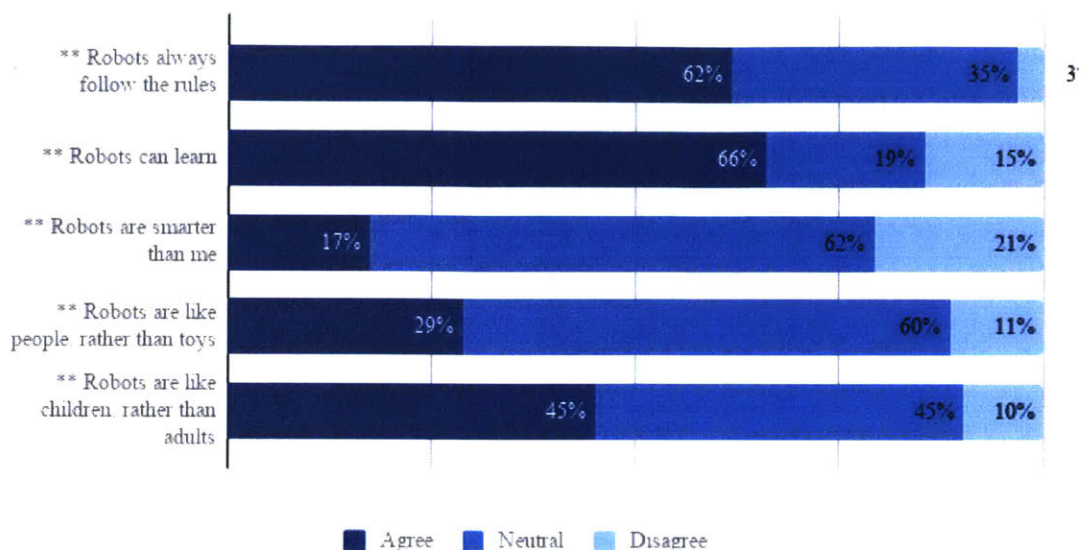
Figure 28: Children's perceptions of AI responses on pretest

Question	Agree	Neutral	Disagree	X^2, df=2	p
Robots always follow the rules	39	22	2	32.67	<0.0001 **
Robots can learn	41	12	9	30.23	<0.0001 **
Robots are smarter than me	11	39	13	23.24	<0.0001 **
Robots are more like people than toys	18	37	7	22.29	<0.0001 **
Robots are more like children than adults	28	28	6	15.61	0.0004 **

The number of responses is not equal for each question. Children were strongly encouraged, but not required to provide an answer of "Agree" "Neutral" (in the middle) or "Disagree." Also, some children were picked up early and could not complete the assessment.

In separating the data by grade and ToM understanding, Chi-square analyses reveal significantly different answers on some questions with regard to grade, but not with regard to ToM understanding. Kindergarten children were significantly more likely to agree that robots are more like people than toys (46.2%) compared to the Pre-K children (0%, $X^2(1, 62)=5.89e-5$, p=0.0061). Pre-K children were significantly more likely to agree that robots are more like children than adults (77.3%)

compared to Kindergarteners (27.5%, $X^2(1, 62)=1.19e-5$, p=0.00275). Children's reasoning on this question did not make it apparent why these differences existed. One possible explanation, is that in previous studies we saw that as children get older they tend to be more sensitive to an object's intentionality, drawing bigger distinctions between a device that is acting of its own volition versus being propelled or controlled in some way (Jipson & Gelman, 2007; Druga et al., 2018). Therefore, perhaps Kindergarten children, who tended to be older than the Pre-K children, had a higher regard for the autonomy of many robots and saw them more as people. Conversely, Pre-K children saw robots as toys and associated toys more with children. However, on both of these questions, half of the children answered neutrally so perhaps most children would be unsure unless some other experience caused them to lean in a particular direction.
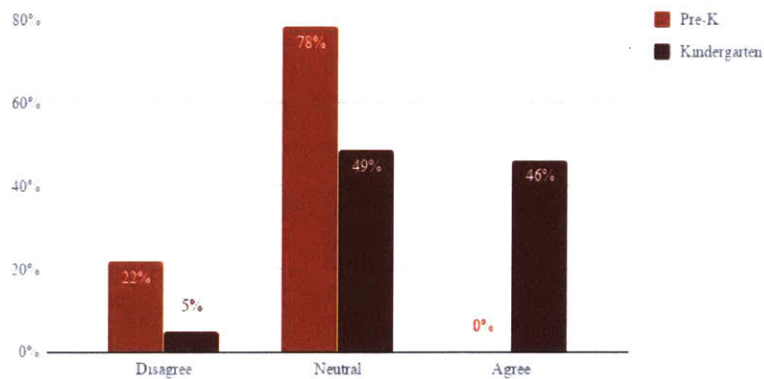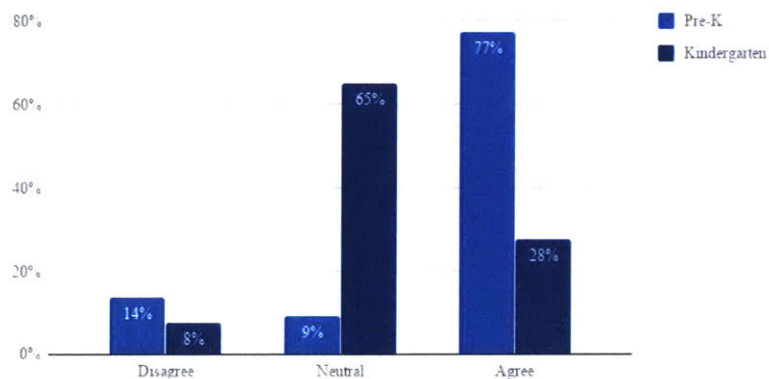
Figure 29: Differences in grade for children's responses to AI pretest

Post-Test Results: Changes in Perception of AI

Children's Perception of AI: Post-test

	Agree	Neutral	Disagree
** Robots always follow the rules	47%	39%	14%
** Robots can learn	75%	12%	13%
Robots are smarter than me	26%	38%	36%
Robots are like people, rather than toys	35%	33%	33%
Robots are like children, rather than adults	25%	49%	26%
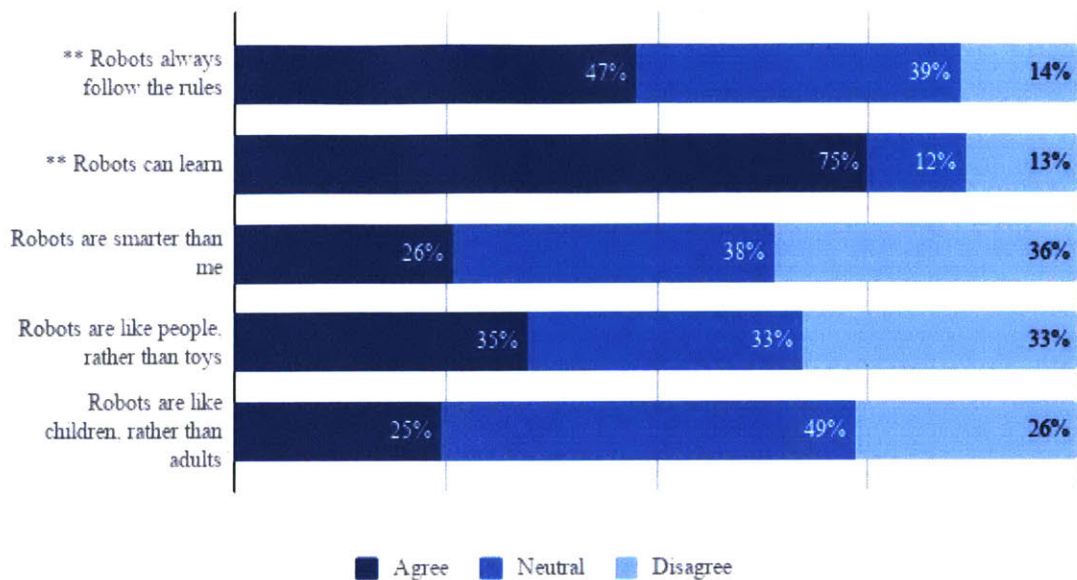
Figure 30: Children's perceptions of AI responses on post-test

Question	Agree	Neutral	Disagree	X^2, df=2	p
Robots always follow the rules	27	22	8	10.21	0.0061 **
Robots can learn	39	6	7	40.65	<0.0001 **
Robots are smarter than me	15	22	21	1.48	0.4771
Robots are more like people than toys	19	18	18	0.04	0.9802
Robots are more like children than adults	13	26	14	2.98	0.2254

The sample of children on the posttest is smaller than on the pretest. Most of the studies occurred in an after school program where children did not attend every day and sometimes were picked up early. In classroom D1, nine children were not present for the post-test. All children in the post-test participated in two or more of the AI activities.

Analysis with a Chi Square test across three groups found that children's post-test responses on whether robots had to follow the rules and whether robots could learn were significantly different from chance. As seen in the pretest, very few children disagreed that the robot always had to follow the rules (14%). Post-hoc analysis revealed that there was a significant difference between children who agreed that the robot had to follow the rules and those that disagreed ($X^2(1,49)=9.26$, p=0.0023). However, more children were neutral saying that sometimes the robot followed rules, but sometimes it did not. Almost every child agreed that robots can learn on the post-test. Post-hoc analysis revealed a significant difference between the number of children who agreed and the number who disagreed ($X^2(1,45)=20.9$, p<0.0001) and said neutral ($X^2(1,46)=22.76$, p<0.0001). Children's reasoning was very clearly that they knew that robots could learn because they had taught one. There was an extremely large difference between Pre-K and Kindergarten on the question about whether the robot was more like a toy or a person. Eighty percent of Pre-K children said that the robot was like a toy compared to 15% of Kindergarten children ($X^2(1, 55) = 9.5e-7$, p<0.001). Even after becoming more familiar with the intellect of the robot, Pre-K children still overwhelming viewed the robot as a toy. This surprising finding seems to warrant further investigation, particularly into how children behaved with the robot. For example, if Pre-K children were overall less interested in the robot's utterances this may have contributed to them seeing the robot as more of a toy. It is important to note that the sample of Pre-K children included a class of mixed Pre-K and Kindergarten children, so it is unlikely that different classrooms or different approaches to instruction caused the difference seen here.

Changes in Answers to Perception of AI: Pretest to Posttest
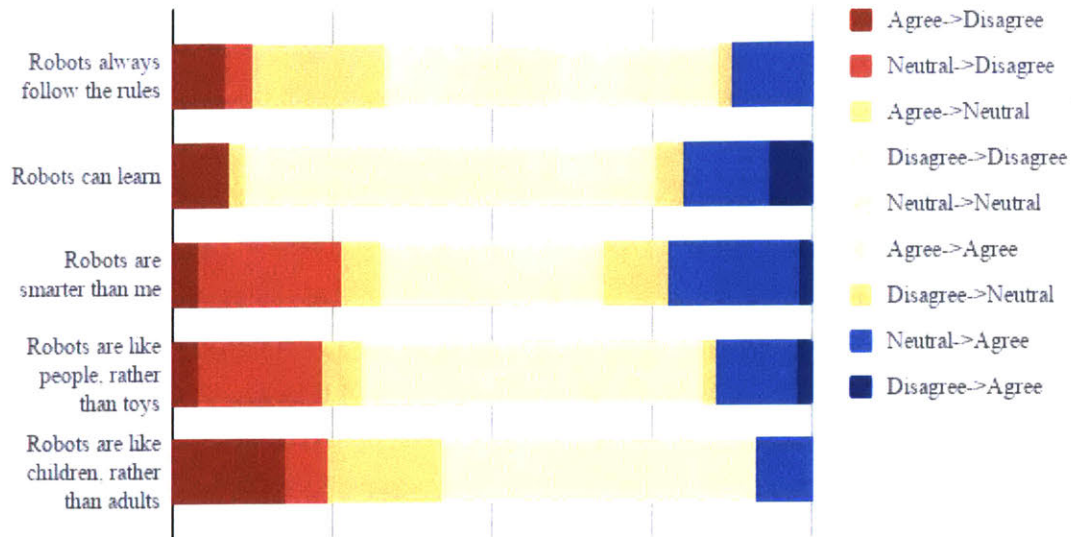
Figure 31: Changes in perceptions of AI responses, pretest -> post-test. Red signifies that the child moved towards disagreement with the statement, blue shows that the child moved towards agreement, light yellow indicates no change.

Question	Agree, Neutral to Disagree	Disagree, Neutral to Agree	n	Z	p
Robots always follow the rules	6	6	48	-2.086	0.039*
Robots can learn	4	9	45	-0.569	0.5687
Robots are smarter than me	13	11	49	0.14	0.889
Robots are more like people than toys	11	7	47	1.14	0.254
Robots are more like children than adults	11	4	45	3.22	0.0013 **

Fig. 31 shows how children's responses changed from pretest to post-test. To analyze how children's answers changed from pretest to post-test, I used a Wilcoxon signed-rank test on each question. On the question about whether robots were more like children or adults, pretest ranks were significantly different from post-test ranks Z=3.22, p=0.0013. Eleven children, 7 of them Pre-K

78

children, who initially agreed or were neutral about seeing the robot as a person moved towards seeing them as toys.

> Researcher: Are robots toys or are they like people?
> Robert (5-years-old): Both because sometimes you can go in it.
> Jane (6-years-old): You can't even go in it.
> Robert: Sometimes there's a ladder and you can climb and go in it.
> Jane: No because it's made out of glass. And the robot's not even real. So you can't go in it. They're just toys and you can only play with them. The only real robots are in movies.
> Researcher: Well I have robots in my lab.
> Jane: But they're not real robots. Robots like tell you what to do and [say] whatever you say.

In their discussion, Jane and Robert were discussing two robots that they had seen. Robert had played on a robot-shaped jungle gym and Jane had seen a glass robot in a movie. Although both children had already completed the curriculum at this point, some of their ideas about robots were still colored by previous experiences they had. In Robert and Jane's case, they both thought that the PopBots were like toys, but they had prior experience with larger robots that they considered to be real.
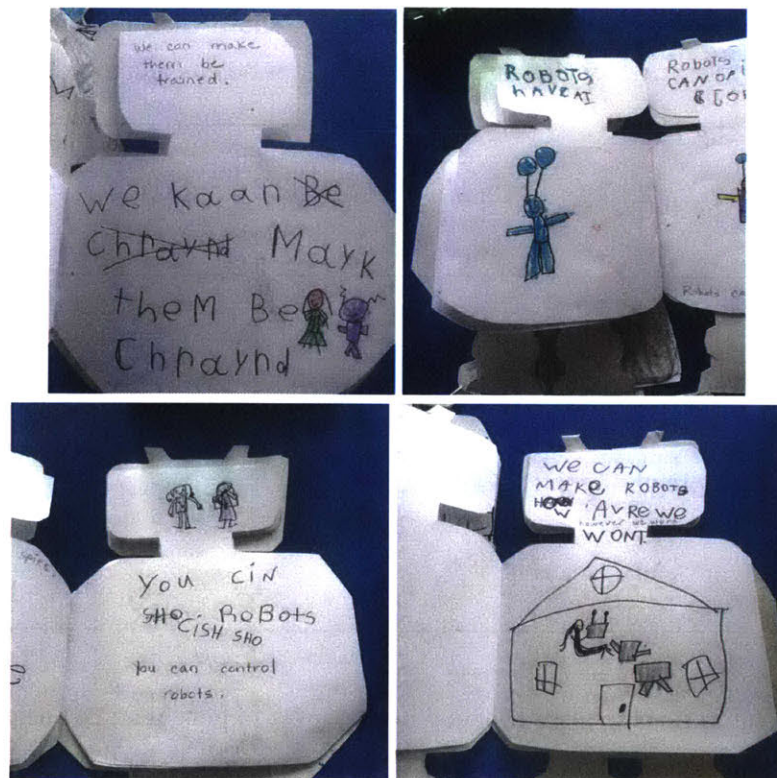
Figure 32: Children's drawings in their "Robot Facts" booklets. (a) We can make [robots] be trained. (b) Some robots have AI. (c) You can control robots. (d) We can make robots however we want.

79

The Kindergarten classroom that was studying robotics for a week created booklets for children to write facts that they learned about robots. From here, we can see other things that children internalized about robots. Some children talked about how robots can be trained and taught. Others talked about how robots can be built and controlled (possibly by programming). In describing what robots could do, lots of children suggested things like helping others or playing games. This was very similar to the children in the pilot study, who also gravitated toward designing assistive and entertaining robots.

7.4 Attitudes Towards Engineering

The Engineering and Science Attitudes assessment (see Appendix A4) centered around children's desires to be engineers, work in engineering related jobs, and their perception of the relevance of such work to their everyday lives. The questionnaire is adapted from the Engineering and Science Attitudes Assessment used by Engineering is Elementary, but shortened to five questions that focus more on engineering than science (Cunningham & Lachapelle, 2010).

Pretest Results

The pretest showed that children believed that science had a lot to do with real life ($X^2(2,39)=51.95$, $p<0.0001$), that they enjoyed figuring things out ($X^2(2,36)=36.5$, $p<0.0001$) and that they *did not* know what an engineer was ($X^2(2,36)=16.17$, $p=0.0003$). Children's familiarity with science was probably the result of the recent focus on STEM in pre-K and K classrooms. In all of the classrooms I worked with, children had some knowledge of the scientific method and had experience proposing a hypothesis and running experiments to test them. However, underscoring the argument that technology and engineering education often go underemphasized in the early classroom, only 25% of the children claimed that they knew what an engineer was. Children would report facts like "My dad (or other family member) is an engineer!" However, their definitions for what an engineer was would range from people who build bridges to people who drove trains. This makes it especially interesting that so many children reported wanting to be an engineer. It is possible that children chose this answer because the researcher identified herself as an engineer right before giving the assessment.

Children's Self-Identification as Engineers: Pretest
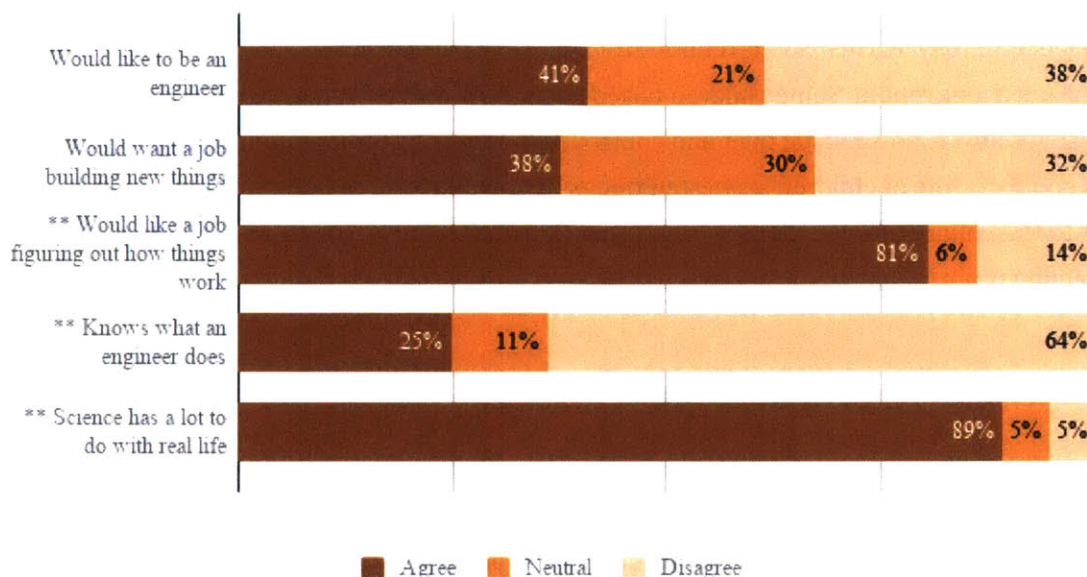
Figure 33: Engineering and Science Attitudes assessment responses on pretest

Question	Agree	Neutral	Disagree	X^2, df=2	p
Would like to be an engineer	16	8	15	2.92	0.2322
Would want a job building new things	14	11	12	0.38	0.827
Would like a job figuring out how things work	29	2	5	36.5	<0.0001 **
Knows what an engineer does	9	4	23	16.17	0.0003 **
Science has a lot to do with real life	33	2	2	51.95	<0.0001 **

There was a disconnect between children wanting to figure things out and wanting to build things, although these questions are related. All of the classrooms included building materials like LEGO, K'nex, and Magna-tiles that the children regularly played with. However, children's experiences and pleasure building with these toys did not align with the responses they gave. It is difficult to assess

81

whether this result is typical. Another study that asked children these questions grouped children's answers about building with their answers about enjoying figuring out how things work (Sullivan, 2016). In this study, children's responses to this question leaned towards agreement. Another possibility is these children did not understand the question and it should be revised in the future.

Post-Test Results: Changes in Attitudes Towards Engineering

Children's Self-Identification as Engineers: Posttest
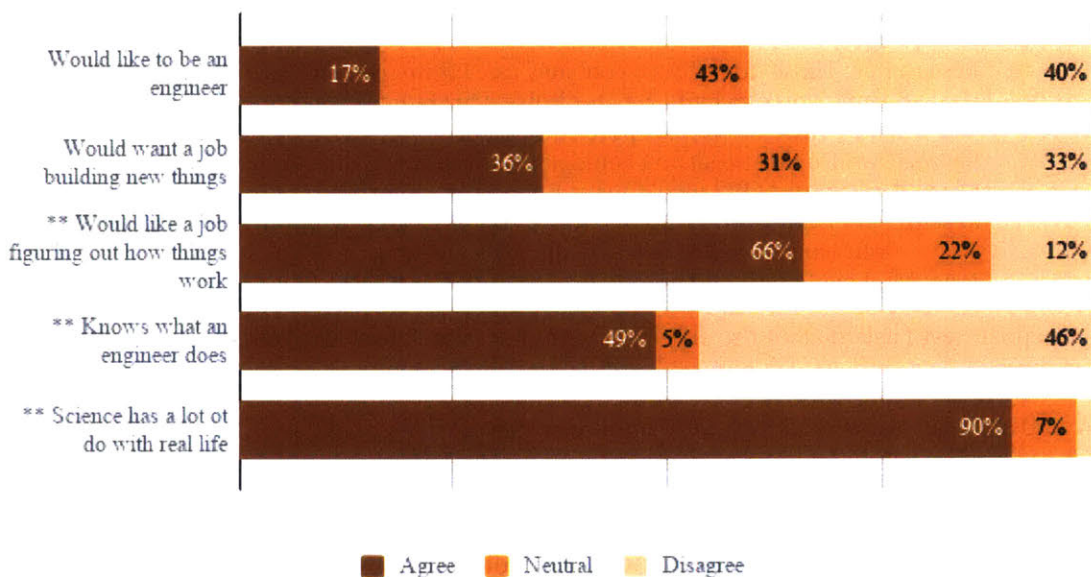
Figure 34: Engineering and Science Attitudes assessment responses on post-test

Question	Agree	Neutral	Disagree	X^2, df=2	p
Would like to be an engineer	7	18	17	5.29	0.071
Would want a job building new things	15	13	14	0.14	0.9324
Would like a job figuring out how things work	27	9	5	20.1	<0.0001 **
Knows what an engineer does	20	2	19	13.35	0.0013 **
Science has a lot to do with real life	37	3	1	59.9	<0.0001 **

During the post-test (Fig. 34), Chi square analysis showed that children were still significantly more likely than the expected to say that they would like a job figuring out how things worked ($X^2(2,41)$=20.1, $p<0.0001$), that science had a lot to do with real life ($X^2(2,41)$=59.9, $p<0.0001$). More children said they knew what an engineer was on the posttest than on the pretest. However, fewer children said that they would like to be an engineer. Initially this seemed discouraging, but children's reasoning for this answer suggested something different.

> Researcher: None of you want to be engineers? What would you like to be then?
> Children: A soccer player. An art teacher. A dentist because my dentist is really, really nice.
> Researcher: Those are all very cool jobs too. I think you can do any of them. Do you think you can build robots? Raise your hand if you think you could build a robot.
> *All children raise hands except for one.*
> Researcher: Do you think you could build the robot with some help?
> Oswald (5-years-old): I would want my parents to help me.
> Researcher: That's fair. What about the rest of you. Do you think you could build it, but you would need help from your parents?
> Jane: No, I could build it by myself. I don't need any help.

During the post-test I asked all of the children if they thought they could build robots on their own, and all answered "yes" except for three who said that they would want their parents to help. Given this, I would say that overall children were empowered by the AI curriculum, or at the very least they were not at all discouraged by it. One possible explanation for so few wanting to be engineers during the post-test is that they still did not know what an engineer was. Children clearly enjoyed working with the robots but, like their other building toys, they did not view the curriculum as an engineering task. Another possibility is that the word "engineer" did not have enough positive reinforcement in children's lives. The words "robot" or "scientist" immediately spark joy and excitement in children, but "engineer" did not. After the curriculum, one teacher said that she would intentionally begin creating a more positive climate around the word "engineer." Finally, each classroom only spent 2-6 hours with the curriculum. This short amount of time may not have been enough for children to change their opinions on about being an engineer.

In sum, children did not necessarily all want to become engineers after the POP curriculum, but they were still empowered to feel that they could build robots themselves or with some help. Ultimately, the goal of any experience in the early classroom is to open children's minds to different opportunities. Before the curriculum, a total of 23 children out of 80 had experience with computational thinking and one child had experience with hardware through a LEGO WeDo class. After the curriculum, children were able to program and build PopBots themselves, and expressed interest in wanting to learn more. Therefore, I can confidently say that programming and building

robots has been added to the toolbox of children who did the curriculum. Later in life, children will likely reflect on this positive, early experience in STEM (Aschbacher, Li, & Roth, 2010).

Changes in Answers to Perception of Self as Engineer: Pretest to Posttest
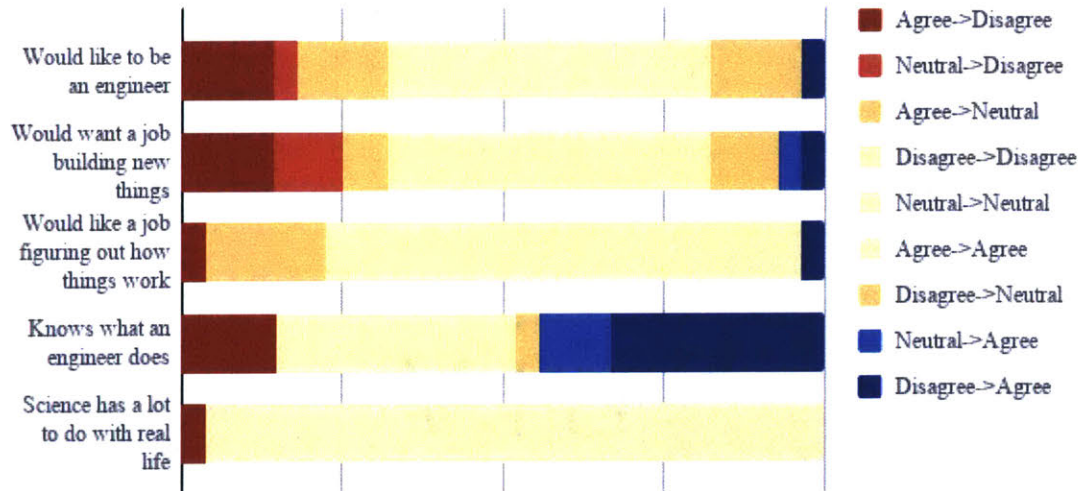
Figure 35: Changes in attitudes towards engineering assessment responses, pretest -> post-test. Red signifies that the child moved towards disagreement with the statement, blue shows that the child moved towards agreement with the statement, yellow indicates no change.

Question	Agree, Neutral to Disagree	Disagree, Neutral to Agree	n	Z	p
Would like to be an engineer	5	1	28	1.27	0.204
Would want a job building new things	7	2	28	1.27	0.204
Would like a job figuring out how things work	1	1	27	-	-
Knows what an engineer does	4	12	27	-1.53	0.126
Science has a lot to do with real life	1	0	27	-	-

A Wilcoxon signed-rank test did not indicate any significant changes between pretest and post-test, likely because of the small sample of children who completed both the pretest and post-test for this assessment. However, according to Fig. 35, the largest changes were in children saying that they knew what an engineer was and in children saying they would not like to be an engineer or have a job building new things. The POP curriculum did not explicitly focus on children understanding what engineers do, but many children seemed to grasp the concept over the course anyways. However, their lack of interest in building things or pursuing engineering still stands in contrast with their anecdotal statements about enjoying programming and building robots. Children's reasoning around these questions seemed to be that engineers only built specific things and this misunderstanding may have led to many children saying that they did not want that kind of job.

8 Discussion

8.1 Preschool Children's Understanding of AI

Platforms exist that teach children about a range of computational thinking concepts from sequencing to debugging; this thesis shows that the *Preschool-Oriented Programming Curriculum* can help children understand three AI topics: rule-based systems, supervised machine learning and, to a lesser extent, generative AI. Through building, training, and programming a social robot, children encountered these AI concepts. Then, they demonstrated their comprehension through their performance on an AI assessment.

The easiest concept for children to grasp was rule-based systems. These algorithms involve a knowledge base that the robot can act on. This concept is very straight forward with many parallels to conditionals in programming. The next easiest topic was supervised machine learning. In this activity, the robot used knowledge about different properties of foods (e.g., color, food group, calories) to guess which foods should be grouped together. In the pilot workshops, children somewhat understood how the robot made guesses, but they were missing how the robot figured out what foods were similar based on their properties. In the evaluative study, the robot explicitly stated its reasoning every time it made a decision. The amount of feedback given by the robot was most likely the factor that led to children understanding better. Finally, the activity that children understood the least was generative AI. The robot combined many rules to transform a child's song into a new creation. The output was different enough that some children did not recognize their original songs. This disconnect led to them not understanding how the robot's output was based on their input. From observations of how children interacted with the POP curriculum and analysis of the relationship between development and what children learned, we can gain insights into how children learned these AI concepts.

Show and Tell and Show Again

Children need to see the concepts demonstrated and reinforced several times for the concepts to stick. When ideas were lightly reinforced, children had a more difficulty internalizing them. This was seen on the rock, paper, scissors question about how the robot tracked previous moves to guess the next move. The robot had to play a couple of rounds before it could predict children's moves.

Predictions did not occur often enough for every child to notice that it was happening. If children had some mechanism to track their previous moves it would have been easier for them to see their own patterns and how the robot took advantage of them. Another addition to the activity could be having children watch a video or observe their peers play with a move tracker. That way, there are more opportunities for children to grasp the idea. In a session longer than 15 minutes, we might have been able to explore this topic with children.

Another way to help children reinforce AI concepts is to use as many sensory examples as possible. While children enjoyed playing music in the generative AI activity, they were not able to grasp how the robot changed their input into a new output by their sense of hearing alone. I believe that the activity could have been much more successful with the aid of a visualization, like the one used in Google's AI Duet[7]. In AI Duet, the notes and duration of the user's input are superimposed with the computers output. The added benefit of such a visualization is that children will be creating musical and visual art with the robot. The fact that children rely on lots of feedback to understand AI concepts is consistent with other studies where children gravitated towards programming blocks with lots of feedback, for example motion and music blocks, in the *ScratchJr.* and *KIBO* platforms (Leidl et al., 2017; Sullivan, 2016).

Grade and Theory of Mind Understanding (Kind Of) Matter

On two of the 10 assessment questions, there was a significant difference between Pre-K and Kindergarten students' performance. The first question was in rock, paper, scissors, "The robot thinks that Sally will play paper. What should the robot play to beat Sally, rock, paper, or scissors?" and the second was in food classification, "Which one of these foods is most like a tomato: strawberry, banana, or milk?" With both of these questions, children who were incorrect almost all chose the same answer, rock on the former question and banana on the latter. The rock, paper, scissors question may have been too complex for Pre-K children. It required children to understand what the robot believed, reference the rules of rock, paper, scissors, and put themselves in the robot's shoes to figure out what it would decide. Perhaps along the way children became confused or missed a step. Although differences in understanding of computational thinking skills based on grade were seen in other platforms, the ones seen in this study seemed to occur for different reasons. With the *KIBO* platform, Kindergarteners did not perform as well as first and second graders on a triggering, sequencing, and repeats assessment (Sullivan, 2016). However, it was also reported that Kindergarteners wrote shorter programs and did not use triggers or repeats as often as the older children. They were less familiar with the concepts. The difference with this study is

---

[7]Google's AI Duet by Yotam Mann https://experiments.withgoogle.com/ai-duet .

that children were familiar with the concepts, so their incorrect answers were most likely the result of the question being more complex. Then, on the food question, one child mentioned that tomatoes and bananas both had the seeds on the inside, not the outside like the strawberry. Although this initially seemed like a rare response, other children at the same school chose the same answer. Therefore, it is possible that some experience led all of the children to the same answer.

Children's understanding of knowledge access and explicit false belief was a factor on two assessment questions. On the rock, paper, question "Sally plays rock and the robot plays paper. Who wins?" there was a significant difference between children who had an understanding of knowledge access and those who did not. This question requires children to understand the rules of rock, paper, scissors to determine who would win. But does not directly require children to understand what knowledge the robot has access too. Then, the generative music question "Priya asks the robot to play a song back when the (speed and rising/falling chord progression) bars are in the middle. Does the robot play the same song?" revealed a significant difference between children who understood explicit false belief and those who did not. This question also did not seem to require explicit false belief skills. On questions where I expected ToM understanding or age to make a difference, it did not, but on these questions it did.

From these results, it seems that children's ability to understand AI may not be related to ToM understanding or grade as measured by the assessments I chose. On one hand, this is a promising result: children's ability to understand AI with a social robot may not be limited by their cognitive perspective-taking abilities. The PopBot and POP platform were able to communicate in such a way that it helped young children understand. On the other hand, this leads me to wonder what the hidden factor is that caused children to not understand some concepts. The previous results show that age and gender were not factors, and grade and ToM understanding only occasionally made a difference.

One possible answer is that children's self-regulation ability when learning makes a difference. The ability to learn and teach is strongly intertwined with children's ToM abilities, it is something that develops with age and practice (Strauss et al., 2002). When teaching and learning, students have to constantly evaluate the gap between what the learners know and what they need to know. Children in this study were teaching the robot while also self-regulating their own learning about AI. Many of the assessment questions asked children to reflect on what the robot learned. Therefore, children with stronger teaching skills would be better equipped to answer these questions. Furthermore, children were engaged in exploratory learning. Although the researcher did somewhat scaffold the activities and call children's attention to certain things, it was up to the children to concretize the knowledge. For example, in the generative music activity many children just assumed that the robot

always played a different song. Others observed that sometimes the robot played the same song and sometimes it played the same song just a bit faster or slower. In future studies, I would like to try some of Strauss's assessments that assess children's understanding other's knowledge gaps. I would also like to develop or find more assessments for measuring children's self-regulation in learning. I predict that children with higher sensitivity to other's knowledge gaps and stronger self-regulation skills will be more successful learning AI with the robot.

8.2 Preschool Children's Perception of AI

Prior studies of children's perceptions of intelligent machines have shown that children's understandings of machines develops gradually with age and experience (Scaife & van Duuren, 1995; van Duuren & Scaife, 1995; Severson & Carlson., 2010; Kahn et al., 2012; Druga et al., 2017). In our recent study with children and smart toys, children ages 4 to 6 were the most indecisive about rating AI-enabled devices as intelligent, humanlike, or adultlike (Druga et al., 2017). In the AI perception pretest, I saw the same results -- children were split on whether the robot was smarter, more like a toy or person, and more like a child or adult. Children more clearly labelled robots as things that had to follow the rules and things that could learn. When children described robots in this way, they would reference previous experiences and understandings of robots from movies, books, and things "they just knew." Children's reasoning about their answers on these questions was in line with our study on children's intuitive attributions of intelligence to mice and robots (Druga et al., 2018). In this study, we saw that children and their parents were biased by previous exposure toward thinking that robots were more intelligent than animals. We saw that older children (7-9 years old) mirrored their parents' reasoning. This study suggests that younger children's attributions of intelligence are influenced by media representations of robots and the others' perceptions (such as parents, but also peers and others).

After the posttest, children's perceptions of robots changed a bit, but they still reasoned about their answers by referencing previous experiences. This underscores the power of outside influences on children's understanding of technology. It is extremely important that parents and the media communicate about robots, especially toys and robots for the home, in a way that is informative. Children are strongly affected by these outside views.

In previous research with children, learning to program robots led to more psychological understandings of the robots (Mioduser & Levy, 2010). This was seen through the increase in children who perceived the robots more as adults than children. Children said that the robots were more like adults because they listened like adults. I think that children were also affected by the

language and perfect pronunciation of the robot voice. However, not all children seemed to have a deeper psychological understanding of the robot. Pre-K children saw the robots more as toys than people after the curriculum. In this case, I wonder if children were affected by the LEGO form of the robot. Perhaps doing the same curriculum with a different robot would not have the same effect.

One important finding in this study is that after the curriculum, almost all of the children believed that the robot could learn. In a previous study, we saw that older children see voice personal assistants as smarter than them because of their know-it-all access to information (Druga et al., 2017). It is important for children to develop a sense that they can teach and train these devices as well. It is also important for children to gain a healthy respect for the training sets beneath the intelligent machines they interact with. During the first pilot workshop, we talked about this explicitly and children did seem to grasp the idea that the behavior of real-world intelligent systems was determined by the people who programmed them.

8.3 Helping Children Develop an Engineering Identity

At face value, children's responses towards the Engineering and Science Attitudes assessment showed that fewer children were interested in being engineers after the curriculum. This is inconsistent with all studies across ages done in the past (Sullivan, 2016; Cunningham & Lachapelle, 2010). I believe this was the result of few children in this sample knowing about engineers. Even after the curriculum, children's definitions of engineers were limited to driving trains and building bridges. Unfortunately, a few hours with the AI curriculum did not make a difference. In the future, I will be more intentional about helping children understand what an engineer is and how the POP curriculum relates to being an engineer.

That said, this curriculum did produce quite a few children who were interested in building robots. Fewer than a quarter of the children in the sample had ever tried programming or building robots before. After a few hours all of them felt excited and prepared to continue building robots. In previous studies, gender roles strongly impacted children's beliefs about their engineering abilities (Sullivan, 2016; Cunningham & Lachapelle, 2010; Cvencek & Meltzoff, 2010). However, there were no differences in gender seen in this study. Girls and boys alike enjoyed programming, building, and interacting with the robot.

8.4 Design Considerations for Helping Children Understand AI

Currently, the POP curriculum explains only three AI concepts out of all of the different ones that exist. I believe that there are other hands-on ways to make AI concepts like planning, perception, reasoning, and controls accessible to young children. Understanding these concepts will help young children better understand more complex mechatronics, self-driving cars, and chatbots. This would be powerful for children because it would allow them to start thinking about new ways to reason about the world by creating things that navigate complex problems in physical and social spaces.

Designing Future AI Activities

First, every activity should include a connection to real world AI, psychological reasoning, and ideas bigger than the AI concepts that are taught. Next, activities should be divided into many short steps that lead to the complete system. In the rock, paper, scissors activity during the pilot workshop, the robot learned the game at the same time as it was using the rules to reason. Having both happening at the same time made it difficult for children to discern that they had taught the robot something. Even for high school students learning AI, it was sometimes difficult for them to understand the difference between training and testing (Kumar, 2004). This needs to be explicit for young children so that they can understand when the robot is learning versus being programmed versus reasoning. Separating the full system into smaller parts also allows children to explore the nuances of each piece more fully. Children can see how changing one part affects the rest of the system. This is especially important for algorithms with many parameters like generative music. If children can explore every nuance they can more effectively construct their understanding. When working with a classroom, using smaller parts gives everyone a chance to do each part in small groups then come back to the class and report what they learned. Intermittent discussions about each of these small steps helps keep everyone on the same page.

Another important aspect of activity design is keeping the activities hands-on and personalizable. Outside of the POP curriculum, I observed many examples of hands-on and multisensory learning in the classrooms that I visited. Through taste, smell and touch children explored cooking; through touch and sight they explored the gardening, building animal habitats, and creating art. Therefore, an AI curriculum for these children should be more like building with blocks than playing with a computer. In the case of POP, tangible blocks limit how much can be explored with the activities. But, the PopBot exists in the physical world, therefore it can, and should interact with children through that medium. Children programmed the robot on the tablet, but they could touch the robot

or show it things to test their programs. Also, children need to be able to adjust the activities to suit their interests. The activities in the current POP curriculum should be expanded so that children have many different ways to approach learning any given topic.

Advancing Children's Understanding of AI

Outside of POP, this work explored other ways that facilitators can help children develop a healthy understanding of AI. The first step is to grasp what the child believes about AI and why. This study showed that prior experiences shape each child's perception uniquely. Therefore, adults need to work to understand each child's unique perspective. One great way to do that is to use design activities where children create their own AI. Having children do design activities before the curriculum showed that Kindergarten children cared about using AI to help people and that Pre-K children wanted to use AI to make entertaining devices or toys. During and after the curriculum, children's designs showed how their understanding of the potential of AI grew. Most of them continued to want assistive or entertaining robots, but they saw many more areas where AI could be applied.

Real-world examples of AI were also very useful. Children could see that intelligent technology was all around them and that it manifested itself in many different ways (functional, entertaining, creative, assistive). It's important that children not only see these examples, but have an opportunity to interact and tinker with them as well. Example of more tinkerable examples include Prisma, Google Draw, and Google's A.I. Duet. If an example is not tinkerable, then children can use their own selves to run AI algorithms the same way that the technology does. In the future, AI devices should be more transparent and trainable so that children can relate their own cognition to them more. For example, voice personal assistants should try to explain why they could not help (e.g., did not hear correctly, did not find any results) and have a feedback loop for children to teach and personalize them. Smart toys should make it clear if and when they are recording and sending information to the Internet. These kinds of changes will give children more agency when interacting with smart devices.

Helping Children Develop an Identity in Engineering

In the movement towards more STEM in early classroom education, it is important that technology and engineering are not forgotten. Children had extremely positive attitudes toward science, but they had little knowledge of engineering and programming. Toolkits for young engineers like POP are extremely important because they give children exposure to an exciting field in a hands-on way.

Barriers to including this kind of technology include price and facilitator training. As for price, POP can use old smartphones, perhaps donated by parents, and LEGOs or other building materials that most classrooms already have. As for teacher and parent training, I am working on materials for POP that can give facilitators enough background on AI that they can learn with their children.

Children are listening to the way that people talk about science and engineering. If people talk about the people who build robots (engineers) as excitedly and as much as they talk about robots, then children may begin to get excited about engineering as well.

9 Conclusion and Contributions

This thesis is the first in-depth design, development and assessment of a computational toolkit and curriculum for young children to explore and understand AI concepts in a developmentally informed way. The main products of this work are the PopBots platform (i.e., PopBot construction kit, PopBlocks app and supporting technologies), and the POP curriculum (i.e., the hands-on activities, workshop protocol, and assessments). These materials were iteratively developed and evaluated with numerous workshops and a final study with 80 children that yields quantitative and qualitative results.

9.1 Answers to Research Questions

How do young children come to perceive and understand AI concepts? What factors influence this process? (Chapter 2 & 7)

We found that children's perceptions of AI were shaped by their previous encounters with the technology. Children learn about intelligent machines like robots from television, movies, and conversations with others. All of these help them develop an opinion of the nature of these devices. Their personal interactions with the devices also shape their opinions. Children learned AI concepts in the POP curriculum after encountering them in their interactions with the PopBots.

Is a social robot platform an effective way to deliver developmentally appropriate curriculum for young children to learn about AI concepts. What features support children's exploration, perception and understanding? (Chapter 3, 5 & 7)

The POP platform and POP curriculum was effective in helping children learn about rule-based systems, supervised machine learning and, to a lesser extent, generative AI. The social robot in the POP platform learned from children and gave continuous feedback on its actions and reasoning. This helped reinforce the AI concepts presented in the curriculum and led to children's comprehension of them. Repetition, tinkerability, and clear presentation of various concepts were key to facilitating understanding.

What can children learn about key AI concepts such as rule-based systems, supervised machine learning, and generative AI by coding and interacting with a social robot platform? How does this inform a developmentally appropriate curriculum for AI and its attributes? (Chapter 4, 6 & 7)

Even after just 15 minutes of interaction, children demonstrated comprehension of rule-based systems and supervised machine learning with a k-nearest neighbors algorithm. However, it did not seem that children understood generative AI. Analysis of children's performance with regard to grade and ToM understanding revealed that these were loosely correlated with children's performance on the assessments. Most likely, grade and ToM understanding are not the developmental factors that have the greatest effect on what children can understand. I hypothesize that children's ability to self-regulate their learning and understand the robot's knowledge gaps make a bigger difference. In future studies I will seek to show that this is the case.

What teaching methods are most effective for engaging and supporting young children with an early AI curriculum in classrooms? (Chapter 5 & 7)

When teaching the POP curriculum, it was important to mix moments of independent work with group discussions. This created a balance between children building their own understanding and having enough time to reflect on what they saw. In addition to showing children examples of real-world AI, it was important that the activities were hands-on and multisensory so that children could explore robotics themselves.

How does learning about AI change children's perception of artificially intelligent devices and themselves as engineers? (Chapter 7)

After learning about AI, children came to perceive robots as machines that could learn and, specifically, as things that they could build and train themselves. Children wanted to build robots to help and provide entertainment. Working with the POP platform resulted in children believing that they could build their own robots to do these things. Unfortunately, children did not know what engineers were and they did not connect this interest to wanting to be an engineer. However, through the POP platform more children had an idea of what an engineer was.

9.2 Future Work

There is also a lot more work to be done in understanding what children think about AI. The preschool to kindergarten population is one of the most interesting ones to study this topic since their ideas are so fluid and they are not yet hindered by perceptions of intelligence that arise when you begin school. It is worth exploring what kinds of prior experiences influence children's perceptions of intelligent device because of how much children's perceptions were affected by outside experiences. Especially since these perceptions even persisted after working with the curriculum.

Working with PopBots allowed children to see how robots learned. Next it would be interesting to look at what new tools and new assessments might be used to further understand the nature of children's relationships with this technology. For example, having the robot ask questions to children directly or having children program and teach robots that present more as adults or humans.

The next steps for the POP Curriculum are expansion and release as an open source project. One of the biggest limitations to releasing this curriculum into the real world is that parents and teachers have not necessarily had training to work with these kinds of platforms. Working with teachers during my study, I saw an interest in wanting to engage and learn about the technology, but also some apprehension about trying this new subject. Therefore, in addition to open sourcing the toolkit, I believe that it is also important to create these guides to support facilitators.

--

As intelligent machines become more social and more informative, it is important that children develop both an understanding of how machines work and a sense that they can build their own intelligent machines. My hope is that this work is useful to educators and parents, and well as for companies building AI-enabled products for family and children's use.

References

1. Apperly, I. A. (2012, 05). What is "theory of mind"? Concepts, cognitive processes and individual differences. Quarterly Journal of Experimental Psychology, 65(5), 825-839. doi:10.1080/17470218.2012.676055

2. Aschbacher, P. R., Ing, M., & Tsai, S. M. (2014, 07). Is Science Me? Exploring Middle School Students' STEM Career Aspirations. Journal of Science Education and Technology, 23(6), 735-743. doi:10.1007/s10956-014-9504-x

3. Ashcraft, C., McLain, B., & Eger, E. (2016). Women in tech the facts. Boulder: National Center for Women & Technology (NCWIT).

4. Barnes-Holmes, Y., McHugh, L., & Barnes-Holmes, D. (2004). Perspective-taking and Theory of Mind: A relational frame account. The Behavior Analyst Today, 5(1), 15.

5. Beran, T. N., Ramirez-Serrano, A., Kuzyk, R., Fior, M., & Nugent, S. (2011). Understanding how children understand robots: Perceived animism in child–robot interaction. International Journal of Human-Computer Studies, 69(7-8), 539-550. http://dx.doi.org/10.1016/j.ijhcs.2011.04.003

6. Bernstein, D., & Crowley, K. (2008, 04). Searching for Signs of Intelligent Life: An Investigation of Young Children's Beliefs About Robot Intelligence. Journal of the Learning Sciences, 17(2), 225-247. doi:10.1080/10508400801986116

7. Bers, M. U. (2008). Blocks to robots: Learning with technology in the early childhood classroom. Teachers College Press.

8. Bers, M. U., & Horn, M. S. (2009). Tangible programming in early childhood: Revisiting developmental assumptions through new technologies: Childhood in a digital world. In I. R. Berson, & M. J. Berson (Eds.), High-tech tots: Childhood in a digital world Greenwich, CT: Information Age Publishing.

9. Bers, M. U. (2018). Coding as a playground: Programming and computational thinking in the early childhood classroom. Routledge.

10. Buolamwini, J. & Gebru, T.. (2018). Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. Proceedings of the 1st Conference on Fairness, Accountability and Transparency, in PMLR 81:77-91

11. Bolukbasi, T., Chang, K., Zou, J., Saligrama, V., & Kalai, A. (2016). Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings. Retrieved May 5, 2018, from https://arxiv.org/abs/1607.06520.

12. Borenstein, J., & Pearson, Y. (2013). Companion robots and the emotional development of children. Law, Innovation and Technology, 5(2), 172-189.

13. Breazeal, C., & Scassellati, B. (1999). How to build robots that make friends and influence people. Proceedings 1999 IEEE/RSJ International Conference on Intelligent Robots and Systems. Human and Environment Friendly Robots with High Intelligence and Emotional Quotients (Cat. No.99CH36289). doi:10.1109/iros.1999.812787

14. Breazeal, C. L. (2004). Designing sociable robots. The MIT Press.

15. Breazeal, C., Harris, P. L., Desteno, D., Westlund, J. M., Dickens, L., & Jeong, S. (2016, 03). Young Children Treat Robots as Informants. Topics in Cognitive Science, 8(2), 481-491. doi:10.1111/tops.12192

16. Burgsteiner, H., Kandlhofer, M., & Steinbauer, G. (2016, February). IRobot: Teaching the Basics of Artificial Intelligence in High Schools. In AAAI (pp. 4126-4127).

17. Campaign for a Commercial Free Childhood. (2017). Don't let Mattel's new 'digital nanny' trade children's privacy for profit. Retrieved from http://www.commercialfreechildhood.org/action/dont-let-mattels-new-digital-nanny-trade-childrens-privacy-profit

18. Chen, S. (2017, May 03). How to Explain AI to Kids – The Eliza Effect – Medium. Retrieved from https://medium.com/eliza-effect/science-fiction-movie-trailers-and-youtube-videos-i-use-to-help-kids-understand-artificial-38a6c08d4652

19. Clements, D. H., & Gullo, D. F. (1984). Effects of computer programming on young children's cognition. Journal of Educational Psychology, 76(6), 1051-1058. doi:10.1037//0022-0663.76.6.1051

20. Commercial Free Childhood (2017). Don't let Mattel's new 'digital nanny' trade children's privacy for profit. Retrieved May 7, 2018, from http://www.commercialfreechildhood.org/action/dont-let-mattels-new-digital-nanny-trade-childrens-privacy-profit#experts

21. Cunningham, C., & Lachapelle, C. (2010). The impact of Engineering is Elementary (EiE) on students' attitudes toward engineering and science. ASEE Annual Conference and Exposition.

22. Cvencek, D., & Meltzoff, A. N. (2011). Math–Gender Stereotypes in Elementary School Children. Encyclopedia of Diversity in Education. doi:10.4135/9781452218533.n463

23. Druga, S., Williams, R., Breazeal, C., & Resnick, M. (2017). "Hey Google is it OK if I eat you?" Proceedings of the 2017 Conference on Interaction Design and Children - IDC '17. doi:10.1145/3078072.3084330

24. Druga, S., Williams, R., & Park, H.W., Breazeal, C. (2018). "How smart are the smart toys?" Children's and parents' attributions of intelligence to computational objects." Proceedings of the 2018 Conference on Interaction Design and Children - IDC '18.

25. Druga, S. (2018). Cognimates. Retrieved May 12, 2018, from https://mitmedialab.github.io/cognimates-website/about/

26. eMarketer. (2017). Awareness of Artificial Intelligence Technology According to Internet Users Worldwide, March 2017 (% of respondents). Retrieved from http://www.emarketer.com/Chart/Awareness-of-Artificial-Intelligence-Technology-Accordin g-Internet-Users-Worldwide-March-2017-of-respondents/207331

27. Elkin, M., Sullivan, A., & Bers, M. U. (2016, 07). Programming with the KIBO Robotics Kit in Preschool Classrooms. Computers in the Schools, 33(3), 169-186. doi:10.1080/07380569.2016.1216251

28. Fior, M., Nugent, S., Beran, T. N., Ramirez-Serrano, A., & Kuzyk, R. (2010). Children's relationships with robots: robot is child's new friend.

29. Flannery, L. P., Silverman, B., Kazakoff, E. R., Bers, M. U., Bontá, P., & Resnick, M. (2013). Designing ScratchJr. Proceedings of the 12th International Conference on Interaction Design and Children - IDC '13. doi:10.1145/2485760.2485785

30. Fraser, N. (2013). Blockly: A visual programming editor. URL: https://code. google. com/p/blockly.

31. Gee, B., & Peck, D. (2016). The illusion of asian success: Scant progress for minorities cracking the glass ceiling from 2007-2015. Ascend.

32. Gelsomini, M., Park, H. W., Lee, J. J., & Breazeal, C. (2017). Engaging Children as a Storyteller. Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction - HRI '17. doi:10.1145/3029798.3036651

33. Gordon, G., Spaulding, S., Westlund, J., Lee, J., Plummer, L., Martinez, Das, M., Breazeal, C. (2016). Affective personalization of a social robot tutor for children's second language skills. AAAI'16 Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, 3951-3957.

34. Gordon, G., Breazeal, C., & Engel, S. (2015). Can Children Catch Curiosity from a Social Robot? Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction - HRI '15. doi:10.1145/2696454.2696469

35. Gordon, M., Ackermann, E., & Breazeal, C. (2015). Social Robot Toolkit. Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction Extended Abstracts - HRI'15 Extended Abstracts. doi:10.1145/2701973.2702001

36. Hwong, V., & Verto Analytics. (2017, July 24). How are personal assistant apps shaping digital consumer habits? Retrieved from https://www.vertoanalytics.com/personal-assistant-apps-shaping-digital-consumer-habits/

37. Imberman, S. P. (2003, 01). Teaching neural networks using LEGO handy board robots in an artificial intelligence course. ACM SIGCSE Bulletin, 35(1), 312. doi:10.1145/792548.611995

38. Jipson, J. L., & Gelman, S. A. (2007, 11). Robots and Rodents: Children's Inferences About Living and Nonliving Kinds. Child Development, 78(6), 1675-1688. doi:10.1111/j.1467-8624.2007.01095.x

39. Jung, M. F., Martelaro, N., Hoster, H., & Nass, C. (2014). Participatory materials. Proceedings of the 2014 Conference on Designing Interactive Systems - DIS '14. doi:10.1145/2598510.2598591

40. Kahn, P. H., Reichert, A. L., Gary, H. E., Kanda, T., Ishiguro, H., Shen, S., Gill, B. (2011). The new ontological category hypothesis in human-robot interaction. Proceedings of the 6th International Conference on Human-robot Interaction - HRI '11. doi:10.1145/1957656.1957710

41. Kahn, P. H., Kanda, T., Ishiguro, H., Freier, N. G., Severson, R. L., Gill, B. T., Shen, S. (2012). "Robovie, you'll have to go into the closet now": Children's social and moral relationships with a humanoid robot. Developmental Psychology, 48(2), 303-314. doi:10.1037/a0027033

42. Kennedy, J., Baxter, P., Senft, E., & Belpaeme, T. (2016, 03). Social robot tutoring for child second language learning. 2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI). doi:10.1109/hri.2016.7451757

43. Keren, G., & Fridin, M. (2014, 06). Kindergarten Social Assistive Robot (KindSAR) for children's geometric thinking and metacognitive development in preschool education: A pilot study. Computers in Human Behavior, 35, 400-412. doi:10.1016/j.chb.2014.03.009

44. Kory, J., & Breazeal, C. (2014, 08). Storytelling with robots: Learning companions for preschool children's language development. The 23rd IEEE International Symposium on Robot and Human Interactive Communication. doi:10.1109/roman.2014.6926325

45. Koski, M., Kurhila, J., & Pasanen, T. A. (2008). Why using robots to teach computer science can be successful theoretical reflection to andragogy and minimalism. Proceedings of the 8th International Conference on Computing Education Research - Koli '08. doi:10.1145/1595356.1595363

46. Kumar, A., Kumar, D., & Russell, I. (2006, 03). Non-traditional projects in the undergraduate AI course. ACM SIGCSE Bulletin, 38(1), 479. doi:10.1145/1124706.1121491

47. Kumar, A. N. (2004, 09). Three years of using robots in an artificial intelligence course. Journal on Educational Resources in Computing, 4(3), 1-15. doi:10.1145/1083310.1083311

48. Lee, J. J., Breazeal, C., & Desteno, D. (2017, 10). Role of Speaker Cues in Attention Inference. Frontiers in Robotics and AI, 4. doi:10.3389/frobt.2017.00047

49. Levy, S. T., & Mioduser, D. (2010, 04). Approaching Complexity Through Planful Play: Kindergarten Children's Strategies in Constructing an Autonomous Robot's Behavior. International Journal of Computers for Mathematical Learning, 15(1), 21-43. doi:10.1007/s10758-010-9159-5

50. Levy, S. T., & Mioduser, D. (2007, 04). Does it "want" or "was it programmed to..."? Kindergarten children's explanations of an autonomous robot's adaptive functioning. International Journal of Technology and Design Education, 18(4), 337-359. doi:10.1007/s10798-007-9032-6

51. Lane, D. (2016). Machine Learning for Kids. Retrieved from https://machinelearningforkids.co.uk/#!/worksheets

52. Lovato, S., & Marie Piper, A. (2015). "Siri, is this you?": Understanding young children's interactions with voice input systems. In Proceedings of IDC 2015: The 14th International Conference on Interaction Design and Children (pp. 335-338). [2771910] Association for Computing Machinery, Inc. DOI: 10.1145/2771839.2771910

53. Martin, F. G. (2007, March). Real Robots Don't Drive Straight. In AAAI Spring Symposium: Semantic Scientific Knowledge Integration (pp. 90-94).

54. McReynolds, E., Hubbard, S., Lau, T., Saraf, A., Cakmak, M., & Roesner, F. (2017). Toys that Listen. Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17. doi:10.1145/3025453.3025735

55. Melson, G. F., Kahn, P. H., Beck, A., Friedman, B., Roberts, T., Garrett, E., & Gill, B. T. (2009, 03). Children's behavior toward and understanding of robotic and living dogs. Journal of Applied Developmental Psychology, 30(2), 92-102. doi:10.1016/j.appdev.2008.10.011

56. McNally, M. F., & Klassner, F. (2007). Demonstrating the Capabilities of MindStorms NXT for the AI Curriculum. In AAAI Spring Symposium: Semantic Scientific Knowledge Integration (pp. 103-104).

57. Milford, M. (2018). The Complete Guide to Artificial Intelligence for Kids. Retrieved from https://www.kickstarter.com/projects/1740916372/the-complete-guide-to-artificial-intelligence-for

58. Minsky, M. (2009, January 22). OLPC Memo 5: Education and Psychology. Retrieved May 2, 2018, from OLPC Memo 5: Education and Psychology

59. Mioduser, D., & Levy, S. T. (2010, 06). Making Sense by Building Sense: Kindergarten Children's Construction and Understanding of Adaptive Robot Behaviors. International Journal of Computers for Mathematical Learning, 15(2), 99-127. doi:10.1007/s10758-010-9163-9

60. Nomura, T., Kanda, T., Kidokoro, H., Suehiro, Y., & Yamada, S. (2017). Why do children abuse robots?. Interaction Studies, 17(3), 347-369.

61. O'Neil, C. (2017). Weapons of Math Destruction. Penguin Books.

62. Papert, S. (1980). Mindstorms; Children, Computers, and Powerful Ideas. BASIC.

63. Park, H. W., Gelsomini, M., Lee, J. J., & Breazeal, C. (2017). Telling Stories to Robots. Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction - HRI '17. doi:10.1145/2909824.3020245

64. Perez, S. (2018, January 16). 39 million Americans now own a smart speaker, report claims. Retrieved from https://techcrunch.com/2018/01/12/39-million-americans-now-own-a-smart-speaker-report-claims/

65. Perlman, R. (1974). TORTIS: Toddler's Own Recursive Turtle Interpreter System. Retrieved from http://hdl.handle.net/1721.1/6224

66. Portelance, D. J., & Bers, M. U. (2015, June). Code and Tell: Assessing young children's learning of computational thinking using peer video interviews with ScratchJr. In Proceedings of the 14th International Conference on Interaction Design and Children (pp. 271-274). ACM.

67. Portelance, D. J., Strawhacker, A. L., & Bers, M. U. (2015, 08). Constructing the ScratchJr programming language in the early childhood classroom. International Journal of Technology and Design Education, 26(4), 489-504. doi:10.1007/s10798-015-9325-0

68. Premack, D., & Woodruff, G. (1978, 12). Does the chimpanzee have a theory of mind? Behavioral and Brain Sciences, 1(04), 515. doi:10.1017/s0140525x00076512

69. Resnick, M., Martin, F., Berg, R., Borovoy, R., Colella, V., Kramer, K., & Silverman, B. (1998). Digital manipulatives. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '98. doi:10.1145/274644.274684

70. Resnick, M., Silverman, B., Kafai, Y., Maloney, J., Monroy-Hernández, A., Rusk, N., Silver, J. (2009, 11). Scratch. Communications of the ACM, 52(11), 60. doi:10.1145/1592761.1592779

71. Resnick, M., & Robinson, K. (2017). Lifelong Kindergarten: Cultivating creativity through projects, passion, peers, and play. The MIT Press.

72. Scaife, M. and Duuren, M. (1995), Do computers have brains? What children believe about intelligent artifacts. British Journal of Developmental Psychology, 13: 367-377. doi:10.1111/j.2044-835X.1995.tb00686.x

73. Severson, R. L., & Carlson, S. M. (2010). Behaving as or behaving as if? Children's conceptions of personified robots and the emergence of a new ontological category. Neural Networks, 23(8-9), 1099-1103.

74. Short, E., Swift-Spong, K., Greczek, J., Ramachandran, A., Litoiu, A., Grigore, E. C., Scassellati, B. (2014, 08). How to train your DragonBot: Socially assistive robots for teaching children about nutrition through play. The 23rd IEEE International Symposium on Robot and Human Interactive Communication. doi:10.1109/roman.2014.6926371

75. Slavin, T. (2016, February 01). What is Artificial Intelligence?. Retrieved from https://www.kidscodecs.com/what-is-artificial-intelligence/

76. Snow, J. (2018, February 16). "We're in a diversity crisis": Cofounder of Black in AI on what's poisoning algorithms in our lives. Retrieved from https://www.technologyreview.com/s/610192/were-in-a-diversity-crisis-black-in-ais-founder-on-whats-poisoning-the-algorithms-in-our/

77. Strauss, S., Ziv, M., & Stein, A. (2002, 09). Teaching as a natural cognition and its relations to preschoolers' developing theory of mind. Cognitive Development, 17(3-4), 1473-1487. doi:10.1016/s0885-2014(02)00128-4

78. Strawhacker, A., Lee, M., Caine, C., & Bers, M. (2015). ScratchJr demo. Proceedings of the 14th International Conference on Interaction Design and Children - IDC '15. doi:10.1145/2771839.2771867

79. Strawhacker, A., Lee, M., & Bers, M. U. (2017, 02). Teaching tools, teachers' rules: Exploring the impact of teaching styles on young children's programming knowledge in ScratchJr. International Journal of Technology and Design Education. doi:10.1007/s10798-017-9400-9

80. Sullivan, A., Elkin, M., & Bers, M. U. (2015). KIBO robot demo. Proceedings of the 14th International Conference on Interaction Design and Children - IDC '15. doi:10.1145/2771839.2771868

81. Talaga, P., & Oh, J. C. (2009). Combining AIMA and LEGO mindstorms in an artificial intelligence course to build real world robots. Journal of Computing Sciences in Colleges, 24(3), 56-64.

82. Tencent. (2017). 2017 Global Artificial Intelligence White paper. Retrieved from http://www.tisi.org/Public/Uploads/file/20171201/20171201151555_24517.pdf

83. Tractica. (2015). Global unit shipments of consumer market robots (household, toy & educational) from 2015 to 2020 (in 1,000s)*. In Statista - The Statistics Portal. Retrieved May 1, 2018, from https://www.statista.com/statistics/608676/consumer-robotics-shipments-worldwide/.

84. Turkle, S. (2017). Alone together: Why we expect more from technology and less from each other. Basic Books.

85. Turkle, S., Taggart, W., Kidd, C. D., & Dasté, O. (2006, 12). Relational artifacts with children and elders: The complexities of cyber companionship. Connection Science, 18(4), 347-361. doi:10.1080/09540090600868912

86. Turkle, S. (2012). Sherry Turkle: Connected, but alone? PsycEXTRA Dataset. doi:10.1037/e547002012-001

87. Turkle, S. (2005). The second self computers and the human spirit. MIT Press.

88. van Duuren, M., & Scaife, M. (1996). "Because a robot's brain hasn't got a brain, it just controls itself"—Children's attributions of brain related behaviour to intelligent artefacts. European Journal of Psychology of Education, 11(4), 365.

89. Van Kruistum, C, van Steensel, R. (2016). Young children (0-8) and digital technology. National report - The Netherlands, Free University Amsterdam and Erasmus University Rotterdam. Available at: https://www.academia.edu/22009256/Young_children_0-8_and_digital_technology_Dutch _national_report_ Google Scholar

90. Wang, Z. (2015, 02). Theory of mind and children's understanding of teaching and learning during early childhood. Cogent Education, 2(1). doi:10.1080/2331186x.2015.1011973

91. Wellman, H. M., & Liu, D. (2004). Scaling of theory-of-mind tasks. Child development, 75(2), 523-541.

92. Westlund, J. M., Jeong, S., Park, H. W., Ronfard, S., Adhikari, A., Harris, P. L., Breazeal, C. L. (2017, 06). Flat vs. Expressive Storytelling: Young Children's Learning and Retention of a Social Robot's Narrative. Frontiers in Human Neuroscience, 11. doi:10.3389/fnhum.2017.00295

93. Wolber, D., Abelson, H., Spertus, E., & Looney, L. (2014). App Inventor 2: Create your own Android apps. Beijing: OReilly.

94. Wyeth, P., & Wyeth, G. (2008). Robot Building for Preschoolers. RoboCup 2007: Robot Soccer World Cup XI Lecture Notes in Computer Science, 124-135. doi:10.1007/978-3-540-68847-1_11

Appendix A: Assessments

1. Theory of Mind Assessments

The following questions are selected from Wellman & Liu (2004). It is used as a pre-assessment in this thesis. There are two main areas explored: *knowledge access* and *false belief.*

1.1 Knowledge Access

Present children with an image of a plastic box with a drawer containing a small toy dog inside the closed drawer. "Here's a drawer. What do you think is inside the drawer?" (The child can give any answer he or she likes or indicate that he or she does not know).

Next, the drawer is opened and the child is shown the content of the drawer. "Let's see ... it's really a dog inside!"

Close the drawer: "Okay, what is in the drawer?"

Then, a girl is produced: "Polly has never seen inside this drawer. Now here comes Polly.

So, does Polly know what is in the drawer? (the target question)

"Did Polly see inside this drawer?" (the memory question).

To be correct, the child must answer the target question "no" and the memory control question "no". In Wellman and Liu, 73% of 3-5 year olds answered this correctly.

1.2 Content False Belief

The child sees a clearly identifiable Band-Aid box with a toy pig inside the closed Band-Aid box. "Here's a Band-Aid box. What do you think is inside the Band-Aid box?"
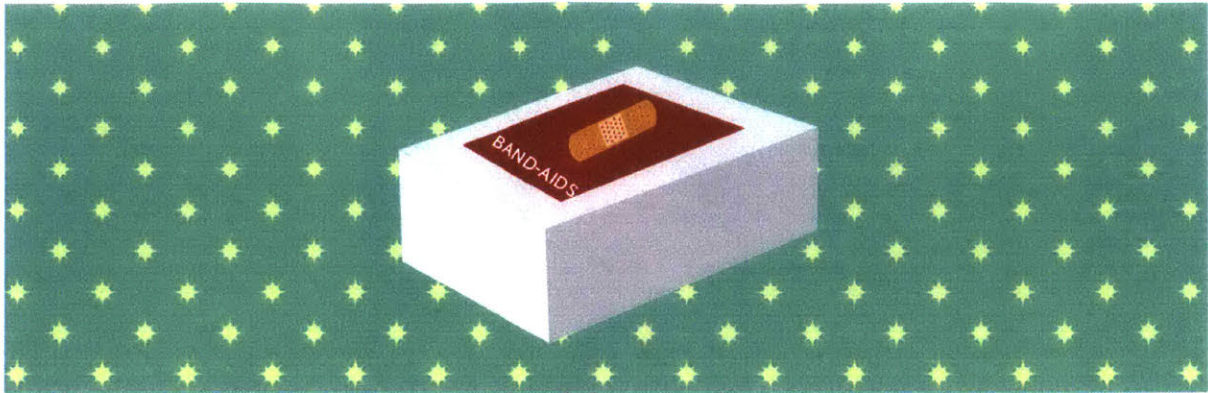
Next, the Band-Aid box is opened: "Let's see ... it's really a pig inside!"

The Band-Aid box is closed: "Okay, what is in the Band-Aid box?"
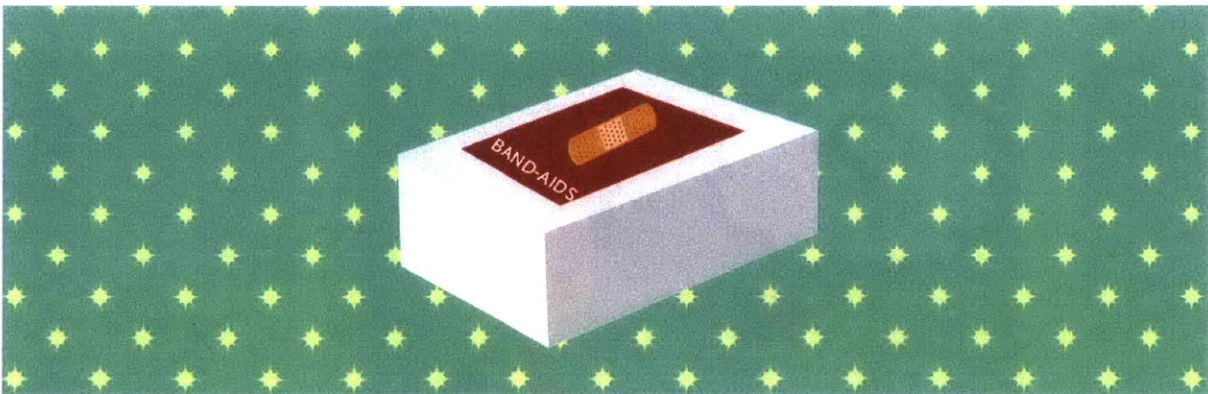
Then a boy is produced: "Peter has never ever seen inside this Band-Aid box.

Now here comes Peter. So, what does Peter think is in the box? Band-Aids or a pig (the target question)

"Did Peter see inside this box? (the memory question).

To be correct the child must answer the target question "Band-Aids" and answer the memory question "no." In Wellman and Liu, 59% of children answered this correctly.

1.3 Explicit False Belief

Children see boy and a picture with a backpack and a closet drawn on it. "Here's Scott. Scott wants to find his mittens.

His mittens might be in his backpack or they might be in the closet. Really, Scott's mittens are in his backpack. But Scott thinks his mittens are in the closet.

So where will Scott look for his mittens? In his backpack or in the closet?" (the target question)

"Where are Scott's mittens really? In his backpack or in the closet?" (the reality question).

To be correct the child must answer the target question "closet" and answer the reality question "backpack." In Wellman and Liu, 57% of children answered this correctly.

2. Understanding Target AI Concepts

(10 Total Probe Questions)

I developed a set of questions that were administered after each associated POP activity to assess children's understanding of key AI concepts.

2.1 Rule-Based Systems Assessment

(5 Questions)

1. Control: Which of these is rock? Rock, paper, or scissors?

2. We teach the robot the normal rules. Then, Sally plays rock and the robot plays paper, who does the robot think has won? Sally or the robot?
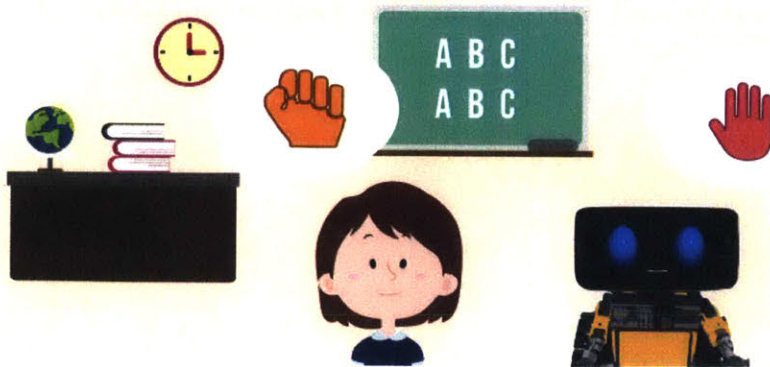
3. Sally plays paper five times. What does the robot think she will play next? Rock, paper, or scissors?
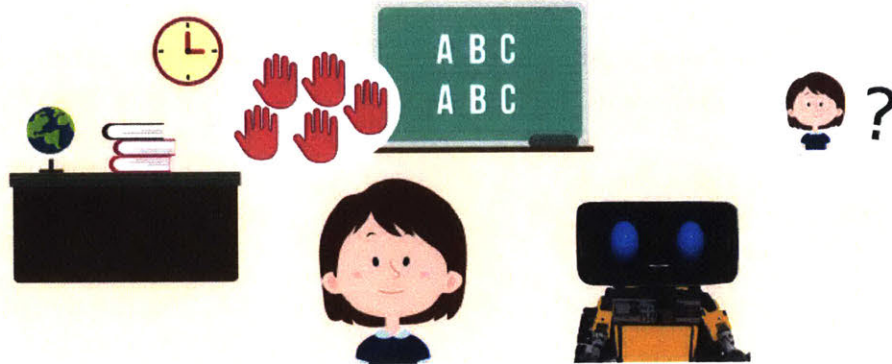
4. The robot thinks that Sally will play paper next. What will the robot play so that it can beat Sally? Rock, paper, or scissors?
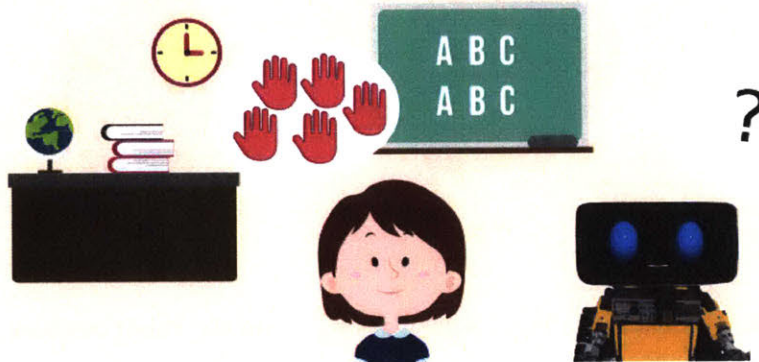
5. We changed the rules so that they are all opposite rules (paper beats scissors). Sally plays scissors and the robot plays paper. Who does the robot think has won? Sally or the robot?
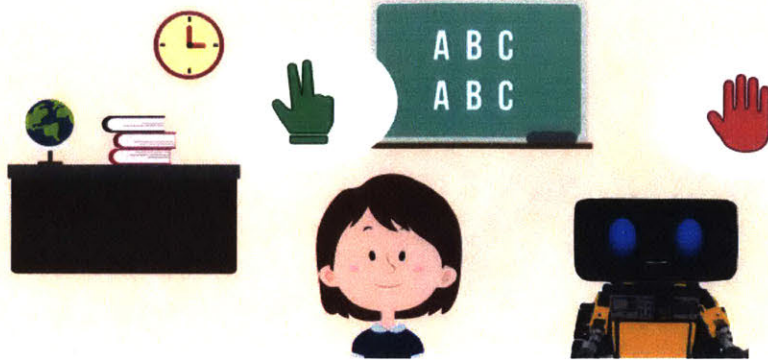
2.2 Supervised Machine Learning

(4 Questions)

1. Control: Which one of these foods is bad for your teeth? Strawberry, ice cream, or corn?
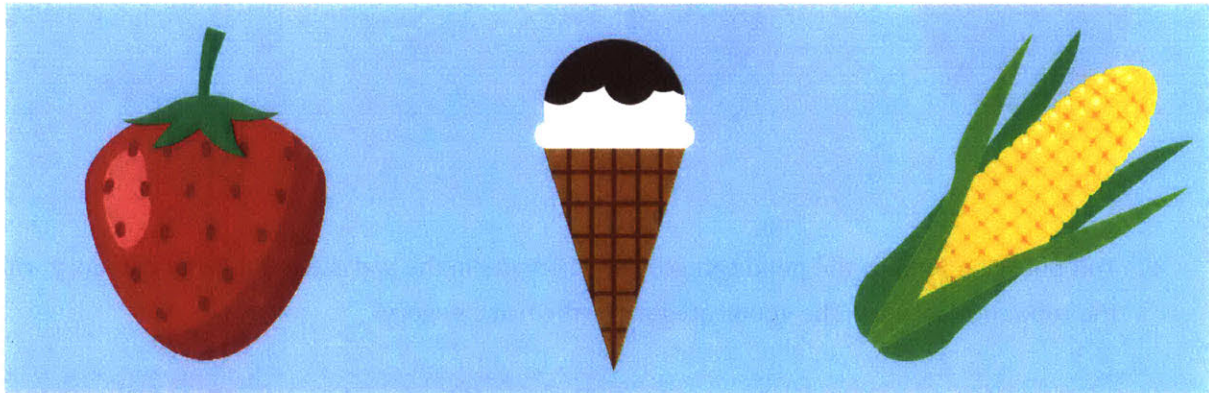
2. You start the robot and put strawberries and tomatoes into the good group. Which group will the robot think chocolate goes in? The good group or the bad group?
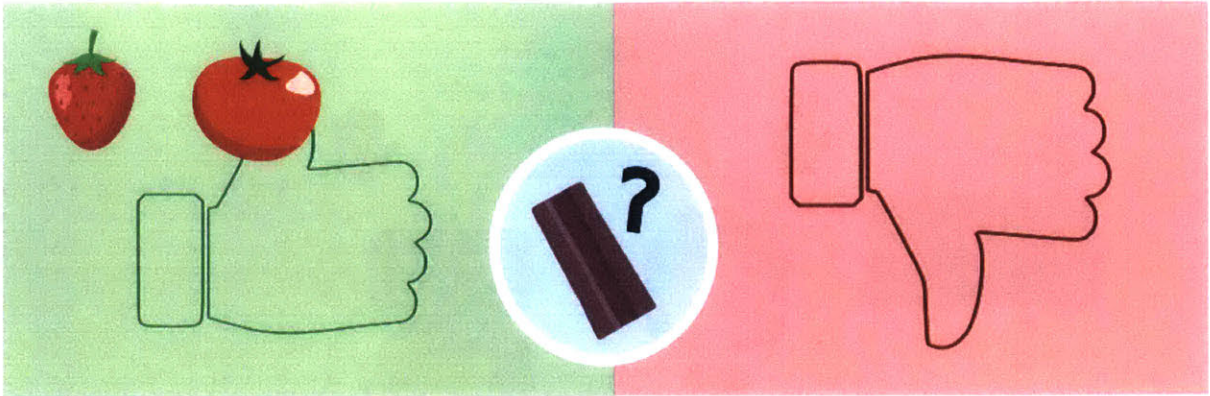
3. What food does the robot think is most like a tomato? Strawberry, banana, or milk?
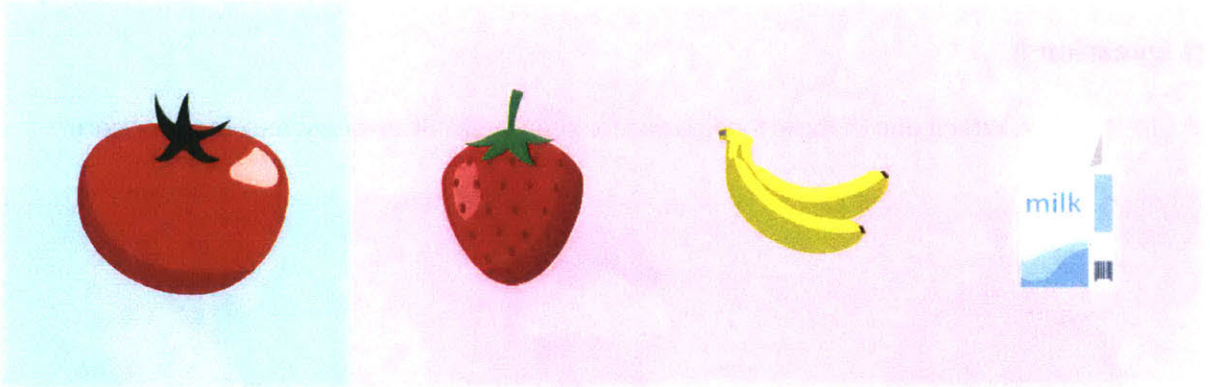
4. You put ice cream in the good category and bananas in the bad category. What category will the robot put corn in? The good category or the bad category?
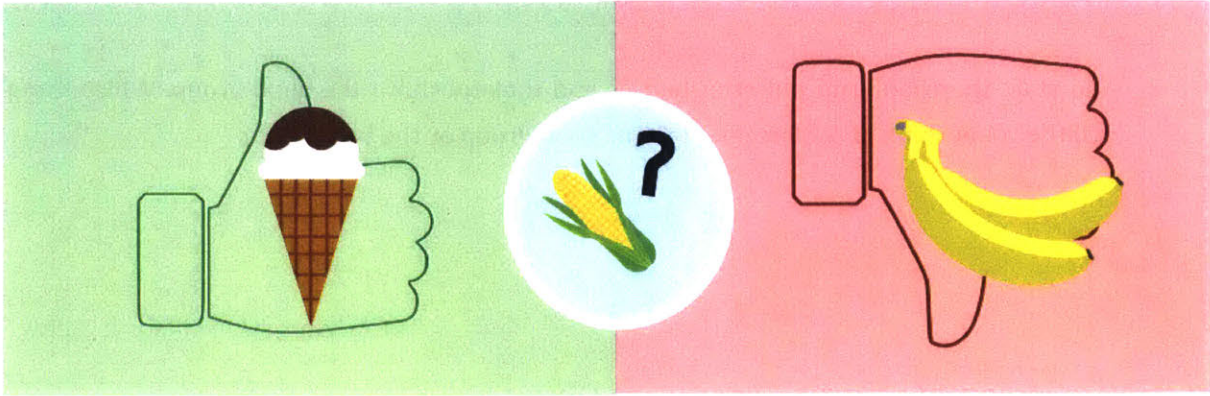
2.3 Generative AI Assessment

(4 Questions)

1. Control: Which one of these notes will make the robot's eyes go orange? Purple note, orange note, or green note?
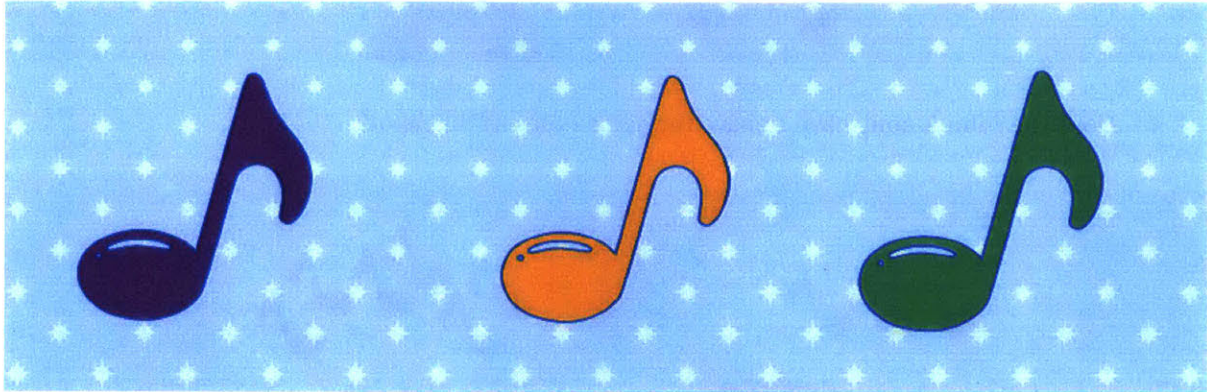
2. Priya asks the robot to play back with the bars in the middle. Does the robot play the same song or a different song?
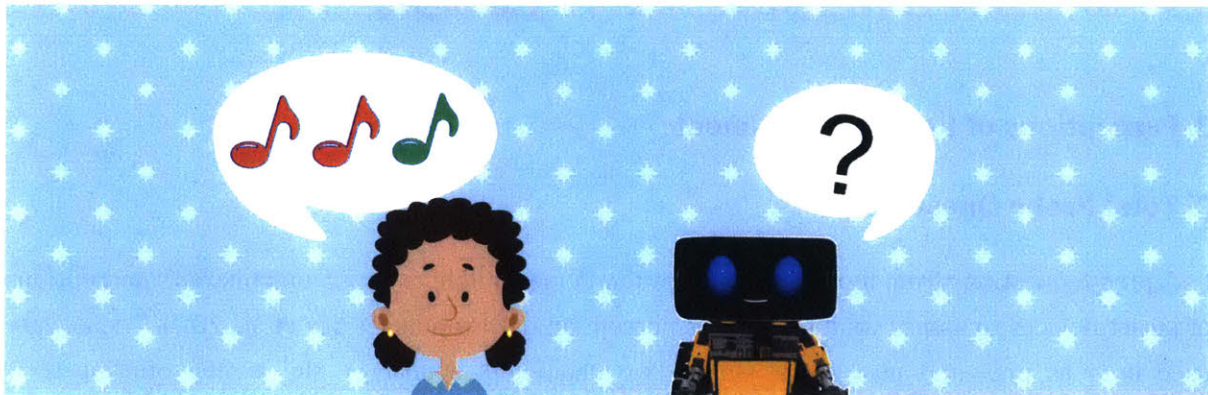
3. Priya asks the robot to play back with the bars to the right. Does the robot play the same song or a different song?
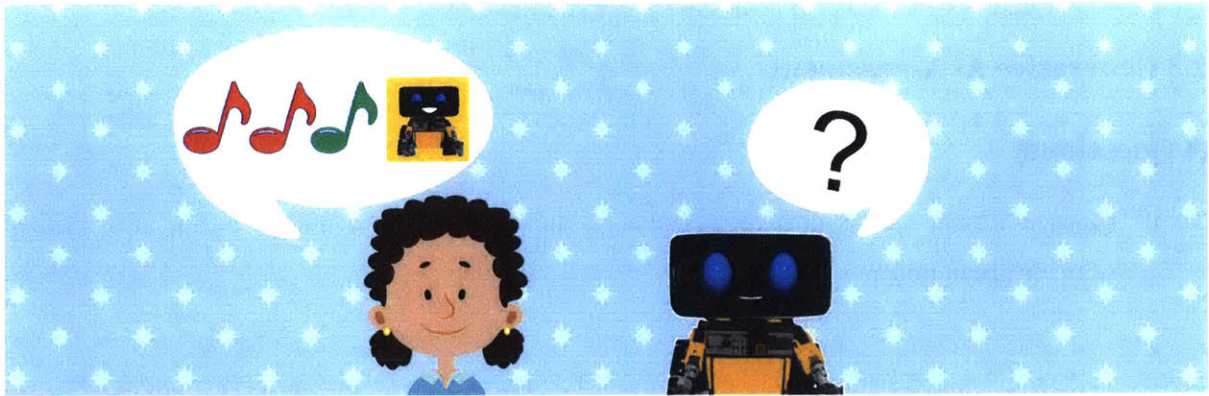
4. Does the robot's song have to have the same notes as the input?

3. Perceptions of Robots Assessment

(5 Total Probe Questions)

I adapted 5 questions from a prior study from the Personal Robots Group on children's perceptions of smart devices to assess young children's perception of AI robots (Druga et al., 2017). These were used both as pre- and post-assessments. Two characters appear on the screen offering their differing (often opposite) opinions on the target question. The child is asked which character they agree with more, or if their opinion is somewhere in between the two.

Which view do you agree with more, or are you somewhere in the middle?

1. Do you agree with Bluey or Greeney more or are you somewhere in the middle?
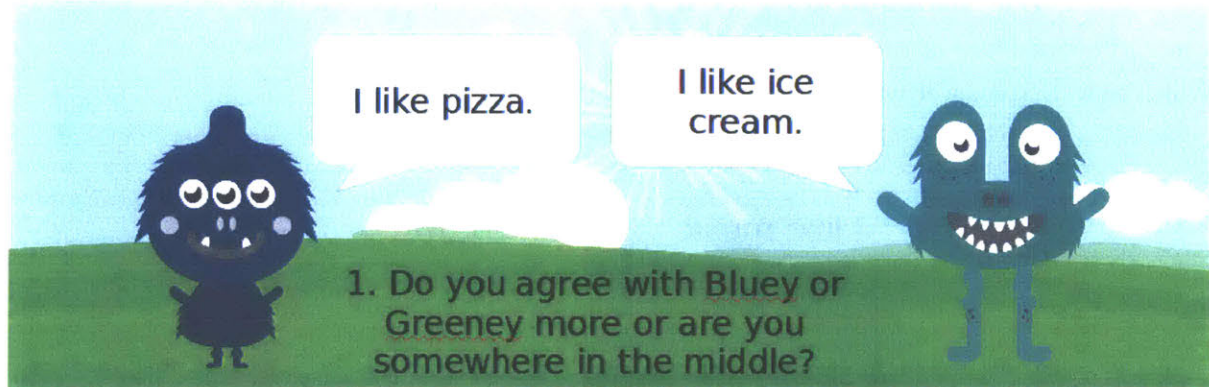
    A. Control: I love pizza / I love ice cream.
    B. Robots follow rules / robots do not follow rules.
    C. I am smarter than robots / robots are smarter than me.
    D. Robots are like toys / robots are like people.
    E. Robots cannot learn new things / robots can learn new things.
    F. Robots are like friends / robots are like adults.

4. Engineering and Science Attitudes Assessment

(6 Total Probe Questions)

The following questions were taken from the Engineering and Science Attitudes Assessment that was used by Amanda Sullivan to assess children's attitudes towards being an engineer before and after doing the KIBO curriculum (Sullivan, 2016). I used these assessments before and after children finished the POP curriculum. To make the assessment more appropriate for young children, I used the Monster Game from a prior study on children's perceptions of AI to ask the questions (Druga et al., 2017). Two characters appear on the screen offering their differing (often opposite) opinions on the target question. The child is asked which character they agree with more, or if their opinion is somewhere in between the two.

Which view do you agree with more, or are you somewhere in the middle?

I like music

I like drawing

1. Do you agree with Orangey or Yellowy more or are you somewhere in the middle?
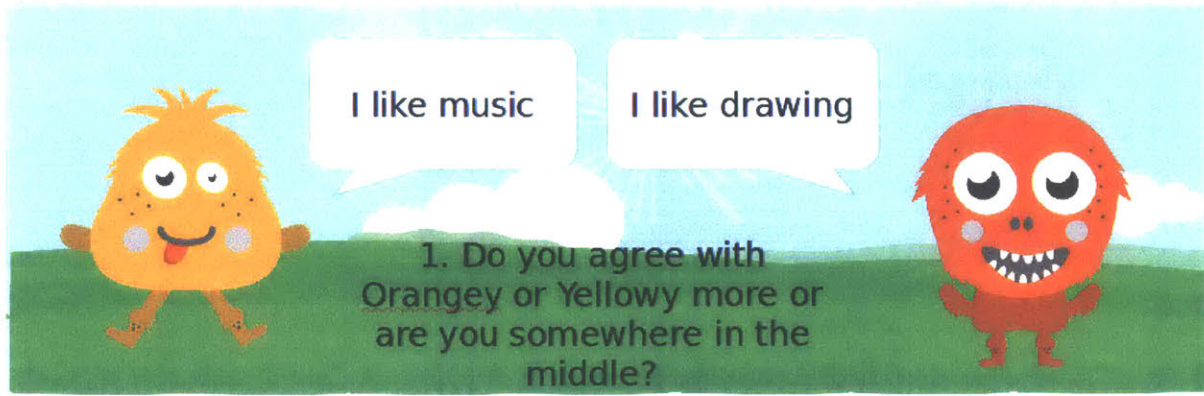
A. I like music / I like drawing.

B. I would like to be an engineer / I would not like to be an engineer.

C. I would not want a job building new things / I would want a job building new things.

D. Science has nothing to do with real life / science has a lot to do with real life.

E. I would like a job figuring out how things work / I would not like a job figuring out how things work.

F. I do not know what engineers do / I do know what engineers do.

Appendix B: Detailed Evaluative Study Protocol

1 Session One: Pretests

Hello there! My name is Randi and I am an engineer, I build robots. **Who can tell me what a robot is?**

For the next few weeks, we are going to play with a very special robot. This robot is one that you can build and teach to make it smarter. **Has anyone built or programmed a robot before?**

(Hand out tablets.) *First, I have some questions for you all. I have the questions here on your tablets. We are going to use the tablets to program the robots so let's get familiar with them.*

First, we have a video to watch. We will watch the videos all together. Here's an example. (Play control question video). *Now, you have to choose which answer you think is best by clicking one of the pictures at the bottom. So which one do you choose?* (Make sure that all children choose the cow.) *Great job, when you're ready for the next question, you click the yellow arrow. That's how all of the questions are. Now let's continue.*

Assessments (15-20 minutes): Theory of Mind Test, AI perception questionnaire, Engineering and Science Attitudes assessment

2 Session Two: Introduction to PopBots

Learning Goals: have children program their own robot and become familiar with the interface

How to program your robot (15-20 minutes - it's OK if they don't get to the end)

(Take out one robot.) *These are the robots that we are going to build and program. You can see that they have eyes and a mouth. Do you think the robots can see? Can they speak? The robots don't have ears, so they cannot hear. If the robot cannot hear, then how do we talk to it? We use the tablet to talk to the robot.*

Last time, we talked about whether robots are like toys or like people. **How is this robot kind of like a toy or like a person?**

You each get your own robot and your own tablet to program the robot. The robot will talk to you and help you. If you need help you can press the question mark at the bottom.

Children will go through a 5-part story that guide children through programming their robot. The story focuses on core parts of the PopBots: turning on the motors, playing an animation, using the motors, and changing colors. The story is framed as open-ended challenges (e.g., teach your robot how to dance) where children will learn to use these blocks. Also help children learn to assist one another by showing others things they figured out and using their words, not their hands.

Bonus: Allow children to freely program their robots or try more complex blocks like loops or recording audio.

3 Session Three: Teach Your Bot the Rules

Learning Goals: have children train their robot to play rock, paper, scissors.

Part 1: Train your robot (5 minutes)

Last time we had a lot of fun programming the robot. You figured out how to make the robot speak, change colors, and drive around. You built something new. That means that you're all engineers. Congratulations!

I asked before if robots always have to follow the rules. **Do you think robots always have to follow the rules?** *Today we're going to find out how robots follow the rules by teaching our robots to play a game: Rock, Paper Scissors.*

(Take out the rock, paper, scissors cards and poster.) *Do you know the rules to rock paper scissors? Let's put them on the poster. There are three rules: rock beats scissors, scissors beat paper, and paper beats rock. Now we're going to teach the robot these rules using our tablets. On the first page tap each square and put the shapes in the square. When you're done, hit the train button and your robot will speak out the rules it knows.*

Part 2: Program your robot (10-15 minutes)

Now the robot is almost ready to play the game but we need to teach it one more important thing. Do you know what it is? We need to teach the robot how to be a good sport. That means the robot needs to say and do nice things whether it wins or loses. We're going to use these triggers to program the robot to react to when it wins or loses. (Explain what the three game triggers mean for the robot.)

Part 3: Play your robot (7 minutes)

Finally, children will be able to play against their robot. Let them play for a few rounds.

Questions for discussion and reflection: *Is the robot following the rules you taught it? Does it always follow the rules? What if you changed the rules, would it still follow them? Did you notice that the robot tries to guess what you will play next. Why does the robot do that and how does he guess? The robot gets better at playing the game with practice. Do you think the robot can get smarter than you at the game? If yes, but you taught the robot all of the rules so does that mean you're still smarter? If no, then why not?*

Now I'm going to pause the robots so that we can answer some questions about this game.

Assessment: Rock Paper Scissors Assessment (See Appendix A)

Bonus: Allow children to retrain or reprogram their robots. See if children notice that the robot tries to predict the child's next move to become better at playing.

4 Session Four: Train Your Bot

Learning Goals: have children see that robots can learn to guess by using a food sorting game where children give the robots some examples of food and the robot learns about them.

Part 1: Yummy for my tummy (5 minutes)

I have here a bunch of different foods. (Take out picture cards of different foods). *Let's sort these foods into two group: foods that are good for you and food that are bad for you.*

Now, let's look at the two groups we've made. What do some of the foods in the good group have in common (food group, color, sweetness, etc.)? Now what what the bad group? We can see that a lot of the foods in each group have things in common.

Part 2: Train your robot (5-10 minutes)

We have to teach our robot about what foods are good and bad. But unlike yesterday with rock, paper, scissors, there are more than three rules. Actually, how many foods are there in the world? That's right. It would take forever to teach the robot about every food one by one, so we're going to give the robot a few examples to help it learn.

Give each child some of the foods from the healthy list and unhealthy list to train their robot. Every robot will only get some of the foods in their training examples, but not all.

Part 3: Test your robot (10+ minutes)

Great, so now our robot knows about a few foods, let's ask it to guess where this new food goes. (The robot will give an answer and explain its reasoning e.g., I put the tomato in the unhealthy group because a lot of unhealthy foods are red). *Why did the robot make that guess?* (Make sure children are paying attention to robot's reasoning).

Let's test the robot with all of the foods and see how many does it gets right or wrong? Which ones seem to be hardest for the robot?

Bonus: Train the robot in some new way: foods you like and foods you don't like or opposites.

Assessment: Food Classification Assessment (See Appendix A)

5 Session Five: Create with Your Bot

Learning Goals: have children learn about music and emotions and how robots do not always follow the rules, they can be creative too.

Part 1: It makes me want to dance (5 minutes)

Let's listen to some music. (Hum three songs, one fast and bright, one fast and somber, one slow and somber). *Which one of those songs sounded happy to you? Why? Which one sounded sad? What about the other song? Some ways that we can make music sound like different emotions is by changing the speed and whether the song goes up or down.*

Part 2: Train your robot (5 minutes)

Now we're going to teach the robot about music. On the tablet we can tell the robot to play music that goes up or down and fast or slow. The robot will play back music after we play to it. Let's see what happens when we play music with the bars (on the tablet) in the middle. See, the robot plays back the same song. Now let's try with the speed all the way down. Was that the same song just a bit slower? OK, let's try the other bar. Was that the same song?

Have each child teach their robot which tempos and which chord progressions make what emotions. Train the robot with four emotions: happy, sad, scared and excited.

Part 3: Make some noise (10-15 minutes)

Now, let's play whatever song we would like and have the robot change it. Allow each child use the piano on the tablet to make and record a song for their robot to play. Then, use the music emotions to have the robot remix the song.

Bonus: Play and record a robot symphony where each robot plays their robot songs one at a time and has their robot dance to it.

Assessment: Generative AI Assessment (See Appendix A)

6 Session Six: Closing and Post-Tests

Learning Goals: allow children to have time to ask any lingering questions and reflect on what they learned in a group setting

Have everyone talk about their one favorite thing that they learned about robots.

Assessments (15-20 minutes): AI perception questionnaire, Engineering and Science Attitudes Assessment (See Appendix A).

Appendix C: Data Tables

1. Biographical data of participants and Theory of Mind Assessment

2. Pretest: AI Perception

3. Post-test: AI Perception

4. Pretest: Engineering and Science Attitudes

5. Post-test: Engineering and Science Attitudes

6. AI Assessments

Biographical data of participants and Theory of Mind Assessment

	Gender	Grade	Age	Knowledge Access	Knowledge Access Control	False Belief	False Belief Control	Explicit False Belief	Explicit False Belief Control
001	M	K	6	Correct	Correct	Incorrect	Correct	Incorrect	Correct
002	F	K	6	Correct	Correct	Incorrect	Correct	Correct	Correct
003	M	K	6	Correct	Correct	Correct	Correct	Incorrect	Correct
004	M	K	5	Correct	Correct	Incorrect	Correct	Correct	Correct
005	F	K	5						
006	M	K	5	Correct	Correct	Incorrect	Correct	Correct	Correct
101	M	Pre-K	5	Correct	Correct	Incorrect	Correct	Incorrect	Correct
102	F	Pre-K	5	Correct	Correct	Correct	Correct	Correct	Correct
103	F	Pre-K	5						
104	F	Pre-K	5	Incorrect	Correct	Correct	Incorrect	Correct	Correct
105	M	K	5	Correct	Correct	Correct	Correct	Correct	Correct
106	M	Pre-K	4	Incorrect	Correct	Correct	Incorrect	Correct	Correct
107	M	Pre-K	5	Incorrect	Correct	Correct	Correct	Incorrect	Correct
108	M	Pre-K	5	Correct	Correct	Incorrect	Incorrect	Correct	Correct
109	M	K	6	Correct	Correct	Incorrect	Correct	Correct	Correct
110	F	Pre-K	5	Incorrect	Correct	Correct	Correct	Correct	Correct
111	M	K	6	Correct	Incorrect	Correct	Correct	Correct	Correct
112	F	K	5	Incorrect	Correct	Correct	Correct	Incorrect	Correct
113	M	K	5	Correct	Correct	Correct	Correct	Correct	Correct
114	M	Pre-K	4						
115	F	Pre-K	4						
116	M	K	6						
201	F	K	5	Incorrect	Incorrect	Incorrect	Incorrect	Incorrect	Correct
202	F	K	6	Correct		Incorrect	Correct	Incorrect	
203	F	K	5	Correct	Correct	Incorrect	Incorrect	Incorrect	Correct
204	M	K	6	Correct		Incorrect	Correct	Correct	
205	F	K	6	Incorrect	Incorrect	Incorrect	Correct	Incorrect	Correct
206	M	K	6	Correct	Correct	Incorrect	Correct	Incorrect	Correct
207	F	K	5	Correct	Correct	Correct	Correct	Incorrect	Correct
208	M	K	5	Correct	Correct	Incorrect	Correct	Correct	Correct
209	F	K	5	Incorrect	Incorrect	Incorrect	Correct	Incorrect	Incorrect
210	F	K	6	Correct	Incorrect	Incorrect	Incorrect	Incorrect	Correct
211	F	K	5	Incorrect	Incorrect	Incorrect	Incorrect	Incorrect	Correct
212	M	K	5						
213	M	K	6						
214	F	K	6						
215	M	K	5						
216	M	K	5						
217	F	K	5						
218	M	K	5						
219	F	K	5						
301	F	K	6	Correct	Correct	Incorrect	Correct	Incorrect	Correct
302	F	K	6	Correct	Correct	Correct	Correct	Incorrect	Correct
303	M	K	5	Correct	Correct	Correct	Correct	Incorrect	Correct
304	F	K	6	Correct	Correct	Incorrect	Correct	Incorrect	Correct
305	M	K	5	Correct	Correct	Incorrect	Incorrect	Incorrect	Correct
306	F	K	6	Correct	Correct	Correct	Correct	Correct	Correct
307	M	K	6	Correct	Correct	Incorrect	Correct	Incorrect	Correct
308	M	K	6	Correct		Incorrect	Correct	Incorrect	Correct
309	F	K	5	Correct	Correct	Correct	Correct	Correct	Correct

310 F	K	5 Correct	Correct		Correct	Correct	Correct
311 M	K	5 Correct	Correct	Correct	Correct	Correct	Correct
312 F	K	5 Correct	Correct	Correct	Correct		
313 F	K	5 Incorrect	Correct	Correct	Correct	Incorrect	Correct
314 F	K	6 Correct	Correct	Incorrect	Correct	Incorrect	Correct
315 F	K	6 Correct	Correct	Incorrect	Correct	Incorrect	Correct
316 M	K	6 Correct	Correct	Incorrect	Correct	Incorrect	Correct
317 M	K	5 Correct	Correct	Incorrect	Correct	Incorrect	Correct
318 M	K	6 Correct	Correct	Incorrect	Correct	Correct	Correct
319 F	K	5 Correct	Correct	Incorrect	Correct	Incorrect	Correct
320 M	K	5 Correct	Correct	Incorrect	Correct	Correct	Correct
321 M	K	6 Correct	Correct	Correct	Correct	Incorrect	Correct
322 M	K	5					
401 M	Pre-K	5 Correct	Correct	Incorrect	Correct	Incorrect	Correct
402 M	Pre-K	5					
403 M	Pre-K	5 Correct	Correct	Correct	Correct	Incorrect	Correct
404 F	Pre-K	4 Correct	Correct	Incorrect	Correct	Incorrect	Correct
405 F	Pre-K	4 Correct	Correct	Correct	Correct	Incorrect	Correct
406 F	Pre-K	5 Correct	Correct	Incorrect	Correct	Correct	Incorrect
407 M	Pre-K	5 Correct	Correct	Incorrect	Correct	Correct	Incorrect
408 M	Pre-K	4 Correct	Correct	Incorrect	Correct	Correct	Incorrect
409 F	Pre-K	4 Correct	Correct	Incorrect	Correct	Correct	Incorrect
410 M	Pre-K	5 Correct	Correct	Correct	Correct	Incorrect	Correct
411 M	Pre-K	5 Correct	Correct	Incorrect	Correct	Incorrect	Correct
412 F	Pre-K	4 Correct	Correct	Incorrect	Correct	Incorrect	Correct
413 M	Pre-K	4 Correct	Correct	Incorrect	Correct	Correct	Incorrect
414 F	Pre-K	5 Correct	Correct	Incorrect	Correct	Correct	Incorrect
415 F	Pre-K	4 Correct	Correct	Incorrect	Correct	Incorrect	Correct
416 F	Pre-K	5 Correct	Correct	Incorrect	Correct	Incorrect	Correct
417 F	Pre-K	5 Incorrect	Correct	Incorrect	Correct	Correct	Incorrect

Pretest: AI Perception

	Gender	Grade	Age	Pretest: Rules	Pretest: Smarter	Pretest: Toy	Pretest: Learn	Pretest: Child
001	M	K	6	Agree	Agree	Neutral	Agree	Neutral
002	F	K	6	Agree	Neutral		Agree	Neutral
003	M	K	6	Agree	Neutral	Neutral		Neutral
004	M	K	5	Neutral	Agree	Agree	Agree	Agree
005	F	K	5					
006	M	K	5	Neutral	Disagree	Agree	Agree	Agree
101	M	Pre-K	5	Neutral	Neutral	Neutral	Disagree	Agree
102	F	Pre-K	5	Neutral	Neutral	Disagree	Agree	Agree
103	F	Pre-K	5					
104	F	Pre-K	5	Neutral	Disagree	Disagree	Disagree	Agree
105	M	K	5	Neutral	Neutral	Neutral	Agree	Agree
106	M	Pre-K	4	Neutral	Disagree	Disagree	Disagree	
107	M	Pre-K	5	Agree	Agree	Neutral	Neutral	Disagree
108	M	Pre-K	5	Agree	Neutral	Neutral	Neutral	Agree
109	M	K	6	Agree	Neutral	Agree	Neutral	Neutral
110	F	Pre-K	5	Neutral	Neutral	Neutral	Disagree	Agree
111	M	K	6					
112	F	K	5	Agree	Disagree	Agree	Disagree	Agree
113	M	K	5	Agree	Neutral	Agree	Disagree	Neutral
114	M	Pre-K	4					
115	F	Pre-K	4					
116	M	K	6					
201	F	K	5	Agree	Agree	Disagree	Agree	Agree
202	F	K	6	Neutral	Neutral	Neutral	Neutral	Neutral
203	F	K	5	Neutral	Neutral	Agree	Agree	Disagree
204	M	K	6	Disagree	Neutral	Agree	Agree	Disagree
205	F	K	6	Neutral	Neutral	Neutral	Neutral	Neutral
206	M	K	6	Neutral	Neutral	Agree	Agree	Agree
207	F	K	5	Agree	Agree	Neutral	Agree	Neutral
208	M	K	5	Neutral	Neutral	Neutral	Neutral	Neutral
209	F	K	5	Agree	Neutral	Neutral	Agree	Neutral
210	F	K	6					
211	F	K	5	Agree	Agree	Disagree	Agree	Agree
212	M	K	5					
213	M	K	6					
214	F	K	6					
215	M	K	5					
216	M	K	5					
217	F	K	5					
218	M	K	5					
219	F	K	5					
301	F	K	6	Agree	Disagree	Neutral	Agree	Neutral
302	F	K	6	Agree	Disagree	Agree	Agree	Neutral
303	M	K	5	Neutral	Neutral	Neutral	Neutral	Neutral
304	F	K	6	Agree	Disagree	Agree	Agree	Neutral
305	M	K	5	Agree	Disagree	Agree	Agree	Neutral
306	F	K	6	Agree	Neutral	Neutral	Neutral	Neutral
307	M	K	6	Neutral	Neutral	Neutral	Agree	Neutral
308	M	K	6	Neutral	Neutral	Neutral	Agree	Neutral
309	F	K	5	Agree	Agree	Neutral	Disagree	Neutral

310 F	K	5 Neutral	Disagree	Agree	Agree	Agree
311 M	K	5 Neutral	Disagree	Agree	Agree	Agree
312 F	K	5 Neutral	Neutral	Neutral	Neutral	Agree
313 F	K	5 Agree	Neutral	Neutral	Agree	Neutral
314 F	K	6 Agree	Neutral	Neutral	Agree	Neutral
315 F	K	6 Agree	Neutral	Agree	Agree	Neutral
316 M	K	6 Agree	Disagree	Agree	Agree	Neutral
317 M	K	5 Agree	Disagree	Agree	Agree	Neutral
318 M	K	6 Agree	Agree	Agree	Agree	Neutral
319 F	K	5 Neutral	Disagree	Neutral	Agree	Neutral
320 M	K	5 Disagree	Neutral	Neutral	Agree	Disagree
321 M	K	6 Agree	Neutral	Agree	Neutral	Agree
322 M	K	5				
401 M	Pre-K	5 Agree	Neutral	Neutral	Agree	Agree
402 M	Pre-K	5				
403 M	Pre-K	5 Agree	Neutral	Neutral	Agree	Agree
404 F	Pre-K	4 Agree	Neutral	Neutral	Agree	Agree
405 F	Pre-K	4 Neutral	Neutral	Neutral	Agree	Agree
406 F	Pre-K	5 Agree	Neutral	Disagree	Agree	Agree
407 M	Pre-K	5 Agree	Neutral	Disagree	Agree	Disagree
408 M	Pre-K	4 Agree	Agree	Neutral	Disagree	Agree
409 F	Pre-K	4 Agree	Neutral	Neutral	Neutral	Neutral
410 M	Pre-K	5 Agree	Agree	Neutral	Agree	Agree
411 M	Pre-K	5 Agree	Agree	Neutral	Agree	Agree
412 F	Pre-K	4 Agree	Neutral	Neutral	Agree	Agree
413 M	Pre-K	4 Agree	Neutral	Neutral	Disagree	Agree
414 F	Pre-K	5 Agree	Neutral	Neutral	Agree	Disagree
415 F	Pre-K	4 Agree	Neutral	Neutral	Agree	Agree
416 F	Pre-K	5 Agree	Neutral	Neutral	Agree	Agree
417 F	Pre-K	5 Neutral	Neutral	Neutral	Neutral	Neutral

Post-test: AI Perception

	Gender	Grade	Age	Post-test: Rules	Post-test: Smarter	Post-test: Toy	Post-test: Learn	Post-test: Child
001	M	K	6	Agree	Neutral	Agree		Disagree
002	F	K	6	Neutral	Neutral	Agree		Disagree
003	M	K	6		Agree			
004	M	K	5					
005	F	K	5	Agree	Neutral			Disagree
006	M	K	5					Disagree
101	M	Pre-K	5	Neutral	Neutral	Neutral	Neutral	Neutral
102	F	Pre-K	5	Neutral	Disagree	Disagree	Agree	Agree
103	F	Pre-K	5					
104	F	Pre-K	5	Disagree	Agree	Disagree	Agree	Disagree
105	M	K	5	Neutral	Disagree	Neutral	Neutral	Neutral
106	M	Pre-K	4					
107	M	Pre-K	5	Neutral	Neutral	Disagree		
108	M	Pre-K	5	Disagree	Disagree	Neutral	Neutral	Neutral
109	M	K	6	Neutral	Disagree	Disagree	Agree	Neutral
110	F	Pre-K	5	Neutral	Neutral	Disagree	Disagree	Agree
111	M	K	6					
112	F	K	5	Disagree	Neutral	Neutral	Agree	Disagree
113	M	K	5	Neutral	Neutral	Agree	Neutral	Agree
114	M	Pre-K	4	Neutral	Neutral	Disagree	Disagree	Agree
115	F	Pre-K	4					
116	M	K	6	Disagree	Disagree	Disagree	Disagree	Agree
201	F	K	5	Agree	Disagree	Agree	Agree	Disagree
202	F	K	6	Agree	Agree	Disagree	Agree	Agree
203	F	K	5	Agree	Agree	Agree	Agree	
204	M	K	6					
205	F	K	6					
206	M	K	6	Disagree	Neutral	Agree	Agree	Disagree
207	F	K	5	Disagree	Neutral	Neutral	Disagree	
208	M	K	5	Neutral	Agree	Neutral	Agree	
209	F	K	5	Agree	Disagree	Neutral	Agree	Agree
210	F	K	6					
211	F	K	5	Neutral	Agree	Neutral	Disagree	Disagree
212	M	K	5	Disagree	Neutral	Neutral	Neutral	Neutral
213	M	K	6	Agree	Neutral	Neutral	Agree	Neutral
214	F	K	6					
215	M	K	5					
216	M	K	5	Neutral	Neutral			
217	F	K	5	Agree	Agree	Disagree	Agree	Agree
218	M	K	5					
219	F	K	5	Agree	Disagree	Neutral	Neutral	Agree
301	F	K	6	Agree	Disagree	Disagree	Agree	Neutral
302	F	K	6	Agree	Disagree	Disagree	Agree	Neutral
303	M	K	5	Agree	Disagree	Agree	Agree	Neutral
304	F	K	6	Agree	Disagree	Agree	Agree	Neutral
305	M	K	5	Agree	Disagree	Agree	Agree	Neutral
306	F	K	6	Disagree	Disagree	Neutral	Agree	Neutral
307	M	K	6	Neutral	Disagree	Agree	Agree	Neutral
308	M	K	6	Agree	Disagree	Neutral	Agree	Neutral
309	F	K	5	Agree	Disagree	Neutral	Agree	Neutral

310 F	K	5	Agree	Disagree	Neutral	Agree	Neutral
311 M	K	5	Neutral	Neutral	Agree	Agree	Agree
312 F	K	5					
313 F	K	5	Neutral	Neutral	Agree	Agree	Neutral
314 F	K	6	Neutral	Neutral	Neutral	Agree	Neutral
315 F	K	6	Agree	Disagree	Neutral	Agree	Neutral
316 M	K	6	Neutral	Neutral	Agree	Agree	Neutral
317 M	K	5	Neutral	Neutral	Agree	Agree	Disagree
318 M	K	6	Neutral	Agree	Agree	Agree	Neutral
319 F	K	5	Neutral	Neutral	Agree	Agree	Agree
320 M	K	5	Neutral	Neutral	Neutral	Agree	Disagree
321 M	K	6	Agree	Neutral	Agree	Agree	Agree
322 M	K	5	Neutral	Disagree	Agree	Agree	Neutral
401 M	Pre-K	5	Agree	Disagree	Disagree	Disagree	Agree
402 M	Pre-K	5					
403 M	Pre-K	5	Agree	Agree	Agree	Agree	Neutral
404 F	Pre-K	4	Agree	Agree	Disagree	Agree	Neutral
405 F	Pre-K	4	Agree	Agree	Disagree	Agree	Disagree
406 F	Pre-K	5	Agree	Agree	Disagree	Disagree	Disagree
407 M	Pre-K	5	Agree	Agree	Disagree	Agree	Disagree
408 M	Pre-K	4					
409 F	Pre-K	4					
410 M	Pre-K	5	Agree	Agree	Disagree	Agree	Neutral
411 M	Pre-K	5					
412 F	Pre-K	4	Agree	Agree	Disagree	Agree	Neutral
413 M	Pre-K	4					
414 F	Pre-K	5					
415 F	Pre-K	4					
416 F	Pre-K	5					
417 F	Pre-K	5					

Pretest: Engineering and Science Attitudes

	Gender	Grade	Age	Pretest Be Engineer	Pretest Build	Pretest Science Real Life	Pretest Figure Out	Pretest Know Engineer
001	M	K	6	Agree	Disagree	Agree	Disagree	Agree
002	F	K	6	Agree	Agree	Agree	Disagree	Disagree
003	M	K	6	Agree			Disagree	Disagree
004	M	K	5		Disagree	Disagree		
005	F	K	5					
006	M	K	5		Disagree	Disagree		
101	M	Pre-K	5	Agree	Agree	Agree	Disagree	Disagree
102	F	Pre-K	5	Agree	Disagree	Neutral	Agree	Disagree
103	F	Pre-K	5					
104	F	Pre-K	5					
105	M	K	5	Disagree				
106	M	Pre-K	4					
107	M	Pre-K	5					
108	M	Pre-K	5					
109	M	K	6	Agree				
110	F	Pre-K	5	Disagree				
111	M	K	6					
112	F	K	5					
113	M	K	5					
114	M	Pre-K	4					
115	F	Pre-K	4					
116	M	K	6					
201	F	K	5	Disagree	Disagree	Agree	Agree	Disagree
202	F	K	6	Neutral	Neutral	Neutral	Agree	Disagree
203	F	K	5	Agree	Agree	Agree	Agree	Disagree
204	M	K	6	Disagree	Agree	Agree	Agree	Disagree
205	F	K	6	Neutral	Neutral	Agree	Agree	Disagree
206	M	K	6	Disagree	Neutral	Agree	Agree	Disagree
207	F	K	5	Disagree	Agree	Agree	Agree	Disagree
208	M	K	5	Neutral	Neutral	Agree	Agree	Neutral
209	F	K	5	Agree	Agree	Agree	Agree	Disagree
210	F	K	6					
211	F	K	5	Disagree	Disagree	Agree	Disagree	Disagree
212	M	K	5					
213	M	K	6					
214	F	K	6					
215	M	K	5					
216	M	K	5					
217	F	K	5					
218	M	K	5					
219	F	K	5					
301	F	K	6	Disagree	Agree	Agree	Agree	Agree
302	F	K	6	Agree	Agree	Agree	Agree	Disagree
303	M	K	5	Neutral	Agree	Agree	Agree	Neutral
304	F	K	6	Agree	Agree	Agree	Agree	Agree
305	M	K	5	Agree	Agree	Agree	Agree	Agree
306	F	K	6	Disagree	Disagree	Agree	Agree	Agree
307	M	K	6	Agree	Neutral	Agree	Agree	Disagree
308	M	K	6	Disagree	Neutral	Agree	Agree	Disagree
309	F	K	5	Agree	Neutral	Agree	Agree	Agree

310 F	K	5	Disagree	Disagree	Agree	Agree	Agree
311 M	K	5	Disagree	Agree	Agree	Agree	Agree
312 F	K	5	Neutral	Disagree	Agree	Agree	Agree
313 F	K	5	Neutral	Neutral	Agree	Neutral	Neutral
314 F	K	6	Neutral	Neutral	Agree	Neutral	Disagree
315 F	K	6	Neutral	Agree	Agree	Agree	Disagree
316 M	K	6	Disagree	Disagree	Agree	Agree	Disagree
317 M	K	5	Disagree	Disagree	Agree	Agree	Disagree
318 M	K	6	Disagree	Neutral	Agree	Agree	Neutral
319 F	K	5	Agree	Neutral	Agree	Agree	Disagree
320 M	K	5	Agree	Agree	Agree	Agree	Disagree
321 M	K	6	Agree	Disagree	Agree	Agree	Disagree
322 M	K	5					
401 M	Pre-K	5					
402 M	Pre-K	5					
403 M	Pre-K	5					
404 F	Pre-K	4					
405 F	Pre-K	4					
406 F	Pre-K	5					
407 M	Pre-K	5					
408 M	Pre-K	4					
409 F	Pre-K	4					
410 M	Pre-K	5					
411 M	Pre-K	5					
412 F	Pre-K	4					
413 M	Pre-K	4					
414 F	Pre-K	5					
415 F	Pre-K	4					
416 F	Pre-K	5					
417 F	Pre-K	5					

Post-test: Engineering and Science Attitudes

	Gender	Grade	Age	Post-test: Be Engineer	Post-test: Build	Post-test Science Real Life	Post-test: Figure Out	Post-test: Know Engineer
001	M	K	6					
002	F	K	6					
003	M	K	6					
004	M	K	5					
005	F	K	5					
006	M	K	5					
101	M	Pre-K	5					
102	F	Pre-K	5					
103	F	Pre-K	5					
104	F	Pre-K	5					
105	M	K	5					
106	M	Pre-K	4					
107	M	Pre-K	5					
108	M	Pre-K	5					
109	M	K	6					
110	F	Pre-K	5					
111	M	K	6					
112	F	K	5					
113	M	K	5					
114	M	Pre-K	4					
115	F	Pre-K	4					
116	M	K	6					
201	F	K	5	Disagree	Disagree	Agree	Agree	Disagree
202	F	K	6	Neutral	Neutral	Neutral	Neutral	Neutral
203	F	K	5	Agree	Agree	Agree	Agree	Agree
204	M	K	6					
205	F	K	6					
206	M	K	6	Disagree	Disagree	Agree	Agree	Disagree
207	F	K	5	Disagree	Agree	Agree	Agree	Disagree
208	M	K	5	Neutral	Agree			
209	F	K	5	Agree	Disagree	Agree	Agree	Disagree
210	F	K	6					
211	F	K	5	Disagree	Disagree	Agree	Agree	Agree
212	M	K	5	Neutral	Disagree	Agree	Agree	Disagree
213	M	K	6	Disagree	Disagree	Agree	Agree	Disagree
214	F	K	6					
215	M	K	5					
216	M	K	5	Agree	Agree	Agree	Agree	Agree
217	F	K	5	Neutral	Neutral	Neutral	Neutral	Neutral
218	M	K	5					
219	F	K	5	Disagree	Disagree	Agree	Agree	Agree
301	F	K	6	Neutral	Disagree	Agree	Agree	Agree
302	F	K	6	Neutral	Disagree	Agree	Agree	Agree
303	M	K	5	Neutral	Agree	Agree	Agree	Agree
304	F	K	6	Neutral	Agree	Agree	Agree	Disagree
305	M	K	5	Neutral	Agree	Agree	Agree	Disagree
306	F	K	6	Agree	Agree	Disagree	Disagree	Disagree
307	M	K	6	Neutral	Neutral	Agree	Neutral	Agree
308	M	K	6	Disagree	Disagree	Agree	Agree	Disagree
309	F	K	5	Disagree	Disagree	Agree	Agree	Disagree

310 F	K	5	Disagree	Disagree	Agree	Agree	Agree
311 M	K	5	Disagree	Neutral	Agree	Agree	Agree
312 F	K	5					
313 F	K	5	Neutral	Neutral	Agree	Neutral	Agree
314 F	K	6	Neutral	Neutral	Agree	Neutral	Disagree
315 F	K	6	Disagree	Disagree	Agree	Agree	Agree
316 M	K	6	Neutral	Neutral	Agree	Neutral	Agree
317 M	K	5	Neutral	Neutral	Agree	Agree	Agree
318 M	K	6	Neutral	Neutral	Agree	Neutral	Agree
319 F	K	5	Disagree	Neutral	Agree	Neutral	Disagree
320 M	K	5	Disagree	Neutral	Agree	Agree	Agree
321 M	K	6	Disagree	Neutral	Agree	Agree	Agree
322 M	K	5	Neutral	Neutral	Agree	Neutral	Agree
401 M	Pre-K	5	Disagree	Disagree	Neutral	Disagree	Disagree
402 M	Pre-K	5					
403 M	Pre-K	5	Agree	Agree	Agree	Agree	Disagree
404 F	Pre-K	4	Neutral	Agree	Agree	Agree	Agree
405 F	Pre-K	4	Disagree	Agree	Agree	Disagree	Disagree
406 F	Pre-K	5	Agree	Agree	Agree	Disagree	Disagree
407 M	Pre-K	5	Disagree	Agree	Agree	Disagree	Disagree
408 M	Pre-K	4					
409 F	Pre-K	4					
410 M	Pre-K	5	Agree	Agree	Agree	Agree	Disagree
411 M	Pre-K	5					
412 F	Pre-K	4	Neutral	Agree	Agree	Agree	Agree
413 M	Pre-K	4					
414 F	Pre-K	5					
415 F	Pre-K	4					
416 F	Pre-K	5					
417 F	Pre-K	5					

AI Assessments 1

	Gender	Grade	Age	RPS1	RPS2	RPS3	RPS4
001	M	K	6	Correct	Correct	Correct	Incorrect
002	F	K	6				
003	M	K	6	Correct	Correct	Correct	Incorrect
004	M	K	5	Correct	Correct	Correct	Incorrect
005	F	K	5	Correct	Correct	Correct	Incorrect
006	M	K	5				
101	M	Pre-K	5	Correct	Correct	Correct	Correct
102	F	Pre-K	5	Correct	Correct		
103	F	Pre-K	5	Correct	Incorrect	Correct	Correct
104	F	Pre-K	5				
105	M	K	5	Correct	Incorrect	Correct	Correct
106	M	Pre-K	4	Correct	Incorrect	Incorrect	Correct
107	M	Pre-K	5	Incorrect	Correct	Correct	Incorrect
108	M	Pre-K	5	Correct	Correct	Correct	Correct
109	M	K	6	Correct	Incorrect	Incorrect	Incorrect
110	F	Pre-K	5	Incorrect	Incorrect	Correct	Correct
111	M	K	6	Correct	Correct		
112	F	K	5	Incorrect	Incorrect	Correct	Correct
113	M	K	5	Correct	Incorrect	Correct	Correct
114	M	Pre-K	4	Correct	Incorrect	Correct	Correct
115	F	Pre-K	4	Correct	Incorrect	Incorrect	
116	M	K	6	Correct	Incorrect	Correct	Incorrect
201	F	K	5				
202	F	K	6				
203	F	K	5				
204	M	K	6				
205	F	K	6	Correct	Incorrect	Correct	Correct
206	M	K	6	Correct	Correct	Correct	Correct
207	F	K	5	Correct	Correct	Correct	Incorrect
208	M	K	5				
209	F	K	5	Correct	Correct	Incorrect	Correct
210	F	K	6				
211	F	K	5	Correct	Correct	Incorrect	Correct
212	M	K	5	Correct	Correct	Correct	Correct
213	M	K	6	Correct	Correct	Correct	Correct
214	F	K	6	Correct	Correct	Incorrect	Correct
215	M	K	5	Correct	Correct	Correct	Correct
216	M	K	5	Correct	Correct	Incorrect	Incorrect
217	F	K	5	Correct	Correct	Correct	Correct
218	M	K	5	Correct	Correct	Correct	Correct
219	F	K	5				
301	F	K	6	Incorrect	Incorrect	Correct	Correct
302	F	K	6	Incorrect	Incorrect	Correct	Correct
303	M	K	5	Incorrect	Incorrect	Correct	Correct
304	F	K	6	Correct	Incorrect	Correct	Correct
305	M	K	5	Correct	Incorrect	Incorrect	Incorrect
306	F	K	6	Correct	Correct	Correct	Incorrect
307	M	K	6	Incorrect	Incorrect	Correct	Correct
308	M	K	6	Incorrect	Incorrect	Correct	Correct
309	F	K	5	Correct	Incorrect	Correct	Correct
310	F	K	5				
311	M	K	5	Correct	Correct	Correct	Correct

312 F	K	5 Incorrect	Incorrect	Correct	Correct
313 F	K	5 Correct	Correct	Correct	Correct
314 F	K	6 Correct	Correct	Correct	Incorrect
315 F	K	6 Correct	Correct	Correct	Correct
316 M	K	6 Correct	Correct	Correct	Correct
317 M	K	5 Correct	Correct	Correct	Correct
318 M	K	6 Correct	Correct	Correct	Correct
319 F	K	5 Correct	Correct	Correct	Correct
320 M	K	5 Correct	Correct	Correct	Correct
321 M	K	6 Correct	Correct	Correct	Correct
322 M	K	5			
401 M	Pre-K	5 Correct	Incorrect	Incorrect	Correct
402 M	Pre-K	5 Incorrect	Correct	Correct	Incorrect
403 M	Pre-K	5 Correct	Correct	Correct	Correct
404 F	Pre-K	4 Correct	Incorrect	Correct	Correct
405 F	Pre-K	4 Correct	Incorrect	Incorrect	Correct
406 F	Pre-K	5 Correct	Incorrect	Incorrect	Correct
407 M	Pre-K	5 Correct	Correct	Correct	Incorrect
408 M	Pre-K	4 Correct	Incorrect	Incorrect	Correct
409 F	Pre-K	4			
410 M	Pre-K	5 Correct	Incorrect	Correct	Incorrect
411 M	Pre-K	5 Correct	Incorrect	Incorrect	Correct
412 F	Pre-K	4 Correct	Correct	Correct	Incorrect
413 M	Pre-K	4 Correct	Correct	Incorrect	Incorrect
414 F	Pre-K	5 Correct	Correct	Correct	Incorrect
415 F	Pre-K	4 Correct	Correct	Correct	Incorrect
416 F	Pre-K	5 Correct	Incorrect	Incorrect	Correct
417 F	Pre-K	5 Incorrect	Incorrect	Incorrect	Correct

AI Assessment 2

	Gender	Grade	Age	FOOD1	FOOD2	FOOD3	MUSIC1	MUSIC2	MUSIC3
001	M	K	6	Incorrect	Correct	Correct	Correct	Incorrect	Incorrect
002	F	K	6	Correct	Correct	Correct	Incorrect	Correct	Correct
003	M	K	6	Incorrect	Correct	Incorrect	Incorrect	Correct	Correct
004	M	K	5	Correct	Correct	Incorrect			
005	F	K	5	Correct	Correct	Correct	Correct	Incorrect	Incorrect
006	M	K	5	Incorrect	Correct	Correct	Incorrect	Correct	Correct
101	M	Pre-K	5	Incorrect	Correct	Correct	Correct	Incorrect	Incorrect
102	F	Pre-K	5				Correct	Correct	Incorrect
103	F	Pre-K	5	Incorrect	Correct	Correct			
104	F	Pre-K	5				Incorrect	Incorrect	Incorrect
105	M	K	5				Incorrect	Correct	Incorrect
106	M	Pre-K	4	Incorrect	Correct	Correct			
107	M	Pre-K	5	Incorrect	Correct	Incorrect	Incorrect	Incorrect	Incorrect
108	M	Pre-K	5	Incorrect	Correct	Incorrect	Correct	Correct	Incorrect
109	M	K	6	Incorrect	Correct	Correct	Incorrect	Correct	Incorrect
110	F	Pre-K	5				Correct	Correct	Correct
111	M	K	6						
112	F	K	5				Incorrect	Correct	Incorrect
113	M	K	5	Incorrect	Correct	Correct	Correct	Incorrect	Incorrect
114	M	Pre-K	4				Correct	Correct	Incorrect
115	F	Pre-K	4						
116	M	K	6	Incorrect	Correct	Correct	Correct	Incorrect	Incorrect
201	F	K	5	Correct	Correct	Correct			
202	F	K	6	Incorrect	Correct	Correct			
203	F	K	5	Correct	Correct	Correct			
204	M	K	6						
205	F	K	6	Incorrect	Correct	Correct	Correct	Correct	Incorrect
206	M	K	6	Incorrect	Correct	Correct	Incorrect	Correct	Incorrect
207	F	K	5				Incorrect	Correct	Incorrect
208	M	K	5	Incorrect	Incorrect	Incorrect			
209	F	K	5	Correct	Correct	Correct	Incorrect	Correct	Incorrect
210	F	K	6						
211	F	K	5	Incorrect	Correct	Correct	Correct	Correct	Incorrect
212	M	K	5	Correct	Correct	Correct	Incorrect	Correct	Incorrect
213	M	K	6	Incorrect	Correct	Correct	Correct	Correct	Incorrect
214	F	K	6				Incorrect	Correct	Incorrect
215	M	K	5				Correct	Correct	Incorrect
216	M	K	5	Correct	Correct	Correct	Correct	Correct	Incorrect
217	F	K	5	Correct	Correct	Correct	Incorrect	Correct	Incorrect
218	M	K	5				Incorrect	Correct	Incorrect
219	F	K	5	Correct	Correct	Correct			
301	F	K	6	Incorrect	Correct	Correct			
302	F	K	6	Incorrect	Correct	Correct			
303	M	K	5	Incorrect	Correct	Correct			
304	F	K	6	Incorrect	Correct	Incorrect			
305	M	K	5	Incorrect	Correct	Incorrect			
306	F	K	6	Incorrect	Correct	Correct			
307	M	K	6	Incorrect	Correct	Correct			
308	M	K	6	Incorrect	Correct	Correct			
309	F	K	5	Incorrect	Correct	Correct			
310	F	K	5	Incorrect	Correct	Correct			
311	M	K	5	Correct	Correct	Correct	Correct	Correct	Incorrect

ID	Sex	Grade	Num						
312	F	K	5						
313	F	K	5	Incorrect	Correct	Correct			
314	F	K	6	Correct	Correct	Correct	Correct	Correct	Correct
315	F	K	6	Incorrect	Correct	Correct			
316	M	K	6	Incorrect	Correct	Correct			
317	M	K	5	Correct	Correct	Correct	Correct	Correct	Incorrect
318	M	K	6	Correct	Correct	Correct			
319	F	K	5	Correct	Correct	Incorrect	Correct	Correct	Incorrect
320	M	K	5						
321	M	K	6	Correct	Correct	Correct	Correct	Correct	Incorrect
322	M	K	5	Correct	Correct	Correct			
401	M	Pre-K	5				Correct	Correct	Incorrect
402	M	Pre-K	5	Correct	Incorrect	Correct			
403	M	Pre-K	5	Correct	Incorrect	Correct			
404	F	Pre-K	4				Correct	Correct	Incorrect
405	F	Pre-K	4				Correct	Correct	Incorrect
406	F	Pre-K	5				Incorrect	Correct	Incorrect
407	M	Pre-K	5	Correct	Incorrect	Correct			
408	M	Pre-K	4				Incorrect	Correct	Incorrect
409	F	Pre-K	4						
410	M	Pre-K	5	Correct	Incorrect	Correct			
411	M	Pre-K	5				Correct	Correct	Correct
412	F	Pre-K	4	Correct	Incorrect	Correct			
413	M	Pre-K	4	Correct	Incorrect	Correct			
414	F	Pre-K	5	Correct	Incorrect	Correct			
415	F	Pre-K	4	Correct	Incorrect	Correct			
416	F	Pre-K	5				Correct	Correct	Incorrect
417	F	Pre-K	5				Incorrect	Correct	Incorrect