



Long-Term Personalization of an In-Home Socially Assistive Robot for Children With Autism Spectrum Disorders

Caitlyn Clabaugh^{1†}, Kartik Mahajan^{1*†}, Shomik Jain¹, Roxanna Pakkar¹, David Becerra¹, Zhonghao Shi¹, Eric Deng¹, Rhianna Lee¹, Gisele Ragusa² and Maja Mataric¹

¹ Interaction Lab, Department of Computer Science, Viterbi School of Engineering, University of Southern California, Los Angeles, CA, United States, ² STEM Education Research Group, Division of Engineering Education, Viterbi School of Engineering, University of Southern California, Los Angeles, CA, United States

OPEN ACCESS

Edited by:

Raquel Ros,
Universitat Ramon Llull, Spain

Reviewed by:

Tony Belpaeme,
University of Plymouth,
United Kingdom
Mehdi Khamassi,
Centre National de la Recherche
Scientifique (CNRS), France

*Correspondence:

Kartik Mahajan
kmahajan@usc.edu

[†]These authors have contributed
equally to this work

Specialty section:

This article was submitted to
Human-Robot Interaction,
a section of the journal
Frontiers in Robotics and AI

Received: 19 June 2019

Accepted: 15 October 2019

Published: 06 November 2019

Citation:

Clabaugh C, Mahajan K, Jain S,
Pakkar R, Becerra D, Shi Z, Deng E,
Lee R, Ragusa G and Mataric M
(2019) Long-Term Personalization of
an In-Home Socially Assistive Robot
for Children With Autism Spectrum
Disorders. *Front. Robot. AI* 6:110.
doi: 10.3389/frobt.2019.00110

Socially assistive robots (SAR) have shown great potential to augment the social and educational development of children with autism spectrum disorders (ASD). As SAR continues to substantiate itself as an effective enhancement to human intervention, researchers have sought to study its longitudinal impacts in real-world environments, including the home. Computational personalization stands out as a central computational challenge as it is necessary to enable SAR systems to adapt to each child's unique and changing needs. Toward that end, we formalized personalization as a hierarchical human robot learning framework (hHRL) consisting of five controllers (disclosure, promise, instruction, feedback, and inquiry) mediated by a meta-controller that utilized reinforcement learning to personalize instruction challenge levels and robot feedback based on each user's unique learning patterns. We instantiated and evaluated the approach in a study with 17 children with ASD, aged 3–7 years old, over month-long interventions in their homes. Our findings demonstrate that the fully autonomous SAR system was able to personalize its instruction and feedback over time to each child's proficiency. As a result, every child participant showed improvements in targeted skills and long-term retention of intervention content. Moreover, all child users were engaged for a majority of the intervention, and their families reported the SAR system to be useful and adaptable. In summary, our results show that autonomous, personalized SAR interventions are both feasible and effective in providing long-term in-home developmental support for children with diverse learning needs.

Keywords: long-term human-robot interaction, personalization, socially assistive robotics, reinforcement learning, home robot, autism spectrum disorders, early childhood

1. INTRODUCTION

Human development follows non-linear trajectories unique to each individual (Vygotsky, 1978). Therefore, socially assistive interventions need to be tailored toward the specific needs and preferences of each participant over time. In a long-term setting, this means interventions must continuously and rapidly adapt toward the user's unique personality. Given the complexity, unpredictability, and uniqueness of each user's progress, intervention strategies must be adapted

in situ via untrained human feedback. Creating autonomous long-term personalized adaptation poses many computational and engineering challenges.

Benefits of personalization are well-established across the domains of education (Bloom, 1984; Anderson et al., 2001) and healthcare (Artz and Armour-Thomas, 1992; Beevers and McGeary, 2012; Cesuroglu et al., 2012; Swan, 2012). While personalized services are paramount, they are neither universally nor equitably affordable. This provides motivation for human-machine interaction research that seeks to develop personalized assistance via socially assistive agents, whether disembodied, virtually embodied (DeVault et al., 2014), or physically embodied (Matarić, 2017).

Socially assistive robotics (SAR) combines robotics and computational methods to broaden access to personalized, socially situated, and physically co-present interventions (Feil-Seifer and Matarić, 2011). A large body of work has supported the importance of physical embodiment (Deng et al., 2019), including its role in increasing compliance (Bainbridge et al., 2008), social engagement (Lee et al., 2006; Wainer et al., 2006), and cognitive learning gains (Leyzberg et al., 2012). Correspondingly, there has been a significant body of work using various types of robots for children with autism spectrum disorders (ASD) in short-term studies (Diehl et al., 2012; Scassellati et al., 2012; Begum et al., 2016), and one long-term study (Scassellati et al., 2018).

The majority of past work with SAR for ASD has been related to social skills. However, it is well-established that learning in general is impacted by social factors; this is particularly important for young learners, because their learning is most often socially mediated (Durlak, 2011). Social difficulties often interfere with children's learning; therefore embedding social contexts in learning environments presents a developmentally appropriate practice that is preferable over isolating social behaviors from cognitive activities (Zins et al., 2004). Consequently, this work addresses the social and cognitive learning domains in tandem, in an intervention that is specifically designed for such learning by children with ASD (White et al., 2007; Guadalupe, 2016).

Personalizing the learning process is especially important in ASD. Given sufficient domain knowledge, personalization of SAR can be achieved through human-in-the-loop or Wizard of Oz (WoZ) frameworks, wherein intervention strategies are mapped to individuals *a priori* or *in situ* via human input (Riek, 2012). However, in practice, considering diverse individual needs and the noise of real-world environments, and the scale of need in ASD, non-autonomous personalization of SAR is infeasible. Reinforcement learning (RL) methods have been successfully applied to adapting to a user's learning habits over time, particularly in early child development studies (Ros et al., 2011). Moreover, recent long-term SAR studies have demonstrated success in maintaining persistent co-present support for educators, students, and caregivers (Bongaarts, 2004). There is therefore an opportunity to develop RL-based personalized long-term learning SAR systems, especially when teaching abstract concepts, such as mathematics (Clabaugh et al., 2015).

In this work, we propose a *personalized* SAR intervention framework that can provide accessible and effective long-term,

in-home support for children with ASD. To accommodate the variable nature of ASD, our framework personalizes to each user's individual needs. To that end, we introduce a hierarchical framework for Human Robot Learning (hHRL) that decomposes SAR interventions into computationally tractable state-action subspaces contained within a meta-controller. The meta-controller consists of disclosure, promise, instruction, feedback, and inquiry controllers that personalize instruction challenge levels and robot feedback based on each child's unique learning patterns. The framework is implemented and evaluated in a fully autonomous SAR system deployed in homes for session-based, single-subject interventions with 17 child participants diagnosed with ASD aged 3–7 years old. Using space-themed mathematics problems, the system combined tenets of educational robotics and computational personalization to maximize each child participant's cognitive gains. Our findings show that the SAR system successfully personalized its instruction and feedback to each participant over time. Furthermore, most families reported the SAR system to be useful and adaptable, and correspondingly, all users were engaged for a majority of the in-home intervention. As a result, all participants showed improvements in math skills and long-term retention of intervention content. These outcomes demonstrate that computational personalization methods can be successfully incorporated in long-term personalized SAR deployments to support children with diverse learning needs.

This paper is organized as follows. *Background* overviews SAR in the relevant contexts of learning, ASD, and personalization. *Formalizing Personalization in SAR* describes the hierarchical human robot learning framework, with a focus on personalization of the challenge level and robot feedback. *Personalized SAR Intervention Design* details the study design, data collection, and outcome measures. The *Results* section details the adaptation performance of the SAR system, its influence on user engagement, participating families' perspectives, and cognitive learning gains over the long-term interaction. *Discussion* and *Conclusion* summarize key insights and recommendations for future work.

2. BACKGROUND

Socially Assistive Robotics (SAR) lies at the intersection of socially interactive robotics and assistive robotics, and focuses on developing intelligent, socially interactive robots that provide assistance through social interaction, with measurable outcomes (Feil-Seifer and Matarić, 2005; Matarić and Scassellati, 2016). We review the relevant background in the main contribution areas of this work: SAR for *learning* (section 2.1) and SAR for *personalization* (section 2.2), both with a particular emphasis on the ASD context, given particular challenges and opportunities for SAR.

2.1. SAR for Learning

A large body of evidence across multiple disciplines supports personalized instruction as a means of positively impacting development and motivation of *individual* learners. Examples include personalized tutoring systems in human-computer interaction research (Wenger, 2014), personalized robot tutors

in HRI and SAR research (Leyzberg et al., 2014) and optimal challenge points (Guadagnoli and Lee, 2004), and the Zone of Proximal Development methodologies in education research (Chaiklin, 2003).

A significant body of SAR research has focused on user learning, with a specific focus on developing personalized robot tutors for young children (Clabaugh and Matarić, 2019). Many SAR and HRI studies have found a robot's embodiment to augment learning in a variety of settings (Gallese and Goldman, 1998; Lee et al., 2006; Gazzola et al., 2007; Wainer et al., 2007; Bainbridge et al., 2008; Leyzberg et al., 2012; Fridin and Belokopytov, 2014). Additionally, several studies on intelligent tutoring systems (ITS) have involved computational models of student learning patterns; however, in contrast to SAR, these works have predominately focused on university students in highly controlled environments (Anderson, 1985; Murray, 1999). From that body of past work, key principles about SAR for learning have been grounded in theories of embodied cognition, situated learning, and user engagement.

Embodied cognition research has shown that knowledge is directly tied to perceptual, somatosensory, and motoric experience, and that a robot's physical embodiment can help contextualize a user's ideas (Niedenthal, 2007; Deng et al., 2019). For example, SAR has helped participants develop motor (Goldin-Meadow and Beilock, 2010), behavioral (Fong et al., 2003), and cognitive skills (Toh et al., 2016). SAR has also shown success in helping users learn abstract concepts; for example, Clabaugh et al. (2015) implemented a SAR system that used deictic gestures to help preschoolers learn number concepts.

Situated learning refers to the importance of the social and physical environment on the learning process and outcomes (McLellan, 1996). Cognitive gains are dependent on context and are enhanced by social interaction (Anderson et al., 1996). Therefore, SAR intervention efficacy must be analyzed in real-world learning settings, involving user learning in various spatial and social contexts (Sabanovic et al., 2006). Environmental conditions impact the quality of SAR interactions and the resulting assistive outcomes. However, real-world scenarios are inherently noisier and less predictable, requiring more complex experimental designs and robust robot platforms (Ros et al., 2011).

User engagement is an important measure of SAR's effectiveness and is inherently tied to learning. In the context of HRI, engagement is widely accepted as a combination of behavioral, affective, and cognitive constructs. Specifically, engagement involves on-task behavior, interest in the robot and task at hand, and a willingness to remain focused (Scassellati et al., 2012). Rudovic et al. (2018) successfully modeled users' engagement with a personalized deep learning framework, however the model was developed *post-hoc*, not in real time. As discussed in Kidd (2008), maintaining user engagement in real time is a major challenge for real-world, long-term studies, as are overcoming technological difficulties and accounting for external human actors.

All of the challenges of SAR for learning are significantly amplified in the ASD context, but ASD is also the context where the success of SAR in supporting learning is especially promising.

ASD is a complex developmental disorder that is often marked by delays in language skills and social skills, including turn-taking, perspective-taking, and joint attention (White et al., 2007). Personalized therapeutic and learning interventions are critical for individuals with ASD, but the substantial time and financial resources required for such services make them inaccessible to many (Ospina et al., 2008; Lavelle et al., 2014), creating an opportunity for SAR support.

There is a large and growing body of research on using SAR for a variety of ASD interventions, as reviewed in Diehl et al. (2012), Scassellati et al. (2012), and Begum et al. (2016). SAR has been shown to help children with ASD develop behavioral and cognitive skills, specifically increased attention (Duquette et al., 2008), turn-taking (Baxter et al., 2013), social interaction (Robins et al., 2005), and many other skills. SAR's ability to perceive, respond, and adapt to user behavior is especially critical in the ASD context (Clabaugh and Matarić, 2019), as users with ASD vary greatly in symptoms and severities, underscoring the need for personalization, as our work also demonstrates.

2.2. Personalization in SAR

SAR systems have shown great potential for providing long-term situated support for meeting individual learning needs. Autonomous or computational personalization in SAR often seeks to maximize the participants' focus and performance, using rule-, model-, or goal-based approaches to personalization.

Rule-based approaches to personalization have been successful in both short-term and long-term SAR interventions. For example, Ramachandran et al. (2018) designed single session interventions where the robot encouraged participants to think out loud. Scenarios were presented based on whether a participant successfully answered a question, and this simple rule-based method resulted in learning gains across all users. Additional studies have expanded rule-based approaches for sequential interactions using hierarchical decision trees (Kidd and Breazeal, 2008; Reardon et al., 2015). Furthermore, in a study setup similar to ours, Scassellati et al. (2018) developed a personalized SAR system for month-long interventions with children with ASD. The system adapted the challenge level of activities using past performance and fixed thresholds. As a result, participants showed increased engagement to the robot and improved attention skills with adults when not in the presence of the robot. In contrast, this work personalizes feedback and challenge level using a goal-based approach, discussed below.

Model-based approaches use models to evaluate the user's success and make optimal decisions. Bayesian Knowledge Tracing (BKT), a domain-specific form of Hidden Markov Models (HMMs), is a common *model-based* approach to personalization in SAR where the hidden state is based on the user's performance and loosely represents their knowledge (Desmarais and Baker, 2012; van De Sande, 2013). For example, BKT can assess how well a participant understands a concept, such as basic addition by examining the sequence of the user's correct and incorrect responses. To represent the variability present in most learning interactions, BKT uses two domain-specific parameters: the probability that a participant will slip and the probability that they will guess. These parameters are

dependent on the interaction context; students with ASD may have difficulty concentrating for extended periods and thus may slip more frequently than typically-developing users (Schiller, 1996). BKT has been successfully applied in SAR; Gordon and Breazeal (2015) and Schodde et al. (2017) used it to adapt to user age and experience, leading to increased learning gains. Leyzberg et al. (2014) applied BKT to training a SAR system to help users solve challenging puzzles more quickly. While outside of ASD, these studies demonstrate the value of BKT in adapting SAR to varying learner needs.

Goal-based methods help the SAR system to select actions that maximize the user's progress toward an assistive outcome. Reinforcement Learning (RL) is a popular goal-based approach, where each user action produces some reward representing progress toward the goal. Throughout the interaction, RL develops a unique, personalized strategy for each participant based on reward-favoring paths. Within HRI, RL has been used to maximize the user's affective state, leading to more effective interactions (Conn et al., 2008; Chan and Nejat, 2011; Castellano et al., 2012; Gordon et al., 2016). Prior studies have shown RL to require deep datasets given the noise of real-world environments. In a single-session context, Gordon et al. (2016) showed RL to successfully adapt in an average of three out of seven sessions. Castellano et al. (2012) also utilized RL to increase engagement; the model was trained on a 15 min interaction and was no better at adapting than a randomized empathetic policy. Conn et al. (2008) showed a RL which was able to adapt quickly, but simplified the robot state to three distinct behaviors. To enable a broader range of behaviors, Chan and Nejat (2011) implemented a hierarchical RL model to personalize feedback within a memory-based SAR interaction. As demonstrated by past work, long-term studies provide the datasets needed for effective RL-based personalization.

Related work has addressed improving social skills of children with ASD. To manage noisy environments and the unpredictable nature of ASD, two studies are particularly relevant as they used RL to parameterize action spaces and speed up robot learning. Velentzas and Khamassi (2018) used RL to personalize the robot's actions to maximize a child's engagement; the robot guided children through a Tower of Hanoi puzzle and used RL to effectively identify non-verbal cues and teach at the learning rate of the participant. That work parameterized the robot's action space to enable efficient decision making and learn single moves in the absence of traditional hierarchical models. Khamassi and Tzafestas (2018) utilized a parameterized action space to select the appropriate robot reaction that maximizes a child's engagement. That work also used RL to maximize the participant's engagement when interacting with a robot. By using the participant's gaze and past variations in engagement, their Q-Learning algorithm became more robust over time. The work used a parameterized environment to simultaneously explore a discrete action space (e.g., moving an object) and a continuous stream of movement features (e.g., expressivity, strength, velocity). These two studies provide insight into maximizing engagement in the absence of hierarchical models, especially when encouraging social interaction (e.g., talking, moving).

The work described in this paper is complementary but different from past work in that it analyzes a long-term SAR intervention for abstract concept learning, specifically helping children with ASD learn mathematics skills. As the next section details, a goal-based RL approach was developed to personalize the instruction and feedback provided to each child by the SAR system.

3. FORMALIZING PERSONALIZATION IN SAR

To address the challenge of long-term personalization in SAR in a principled way, we present a solution to the problem as a controller-based environment which we define as hierarchical human-robot learning (hHRL).

3.1. Human-Robot Learning

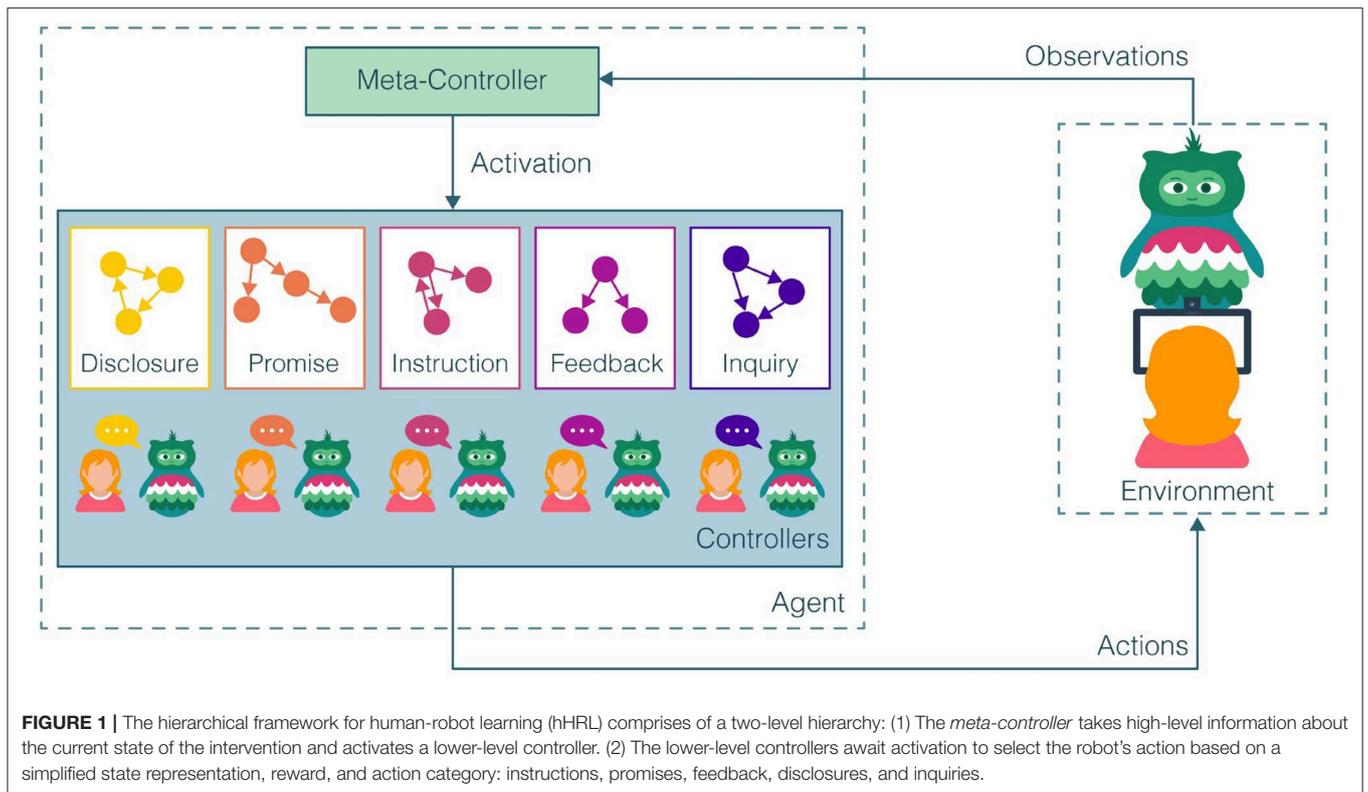
Past work has explored methods for computational personalization, with the objective of finding an optimal sequence of actions that steers the user toward a desired goal. While this problem has been studied in the contexts of user modeling in HCI (Fischer, 2001), machine teaching (Chen et al., 2018), as well as active (Cohn et al., 1996), and interactive machine learning (ML) (Amershi et al., 2014; Dudley and Kristensson, 2018), computational personalization is yet to be formalized in the context of SAR.

We define and formalize *Human-Robot Learning (HRL)* as the interactive and co-adaptive process of personalizing SAR. At the highest level, the quality of a SAR intervention can be assessed relative to some goal G . Since SAR contexts often involve long-term goals, success is better assessed via intermediate measures of progress toward G . Hence, it is important to design and represent SAR intervention interactions in a manner that maximizes observability. In this work, HRL is framed from the perspective of the robot, so that optimization is limited to the robot's actions and not those of the human, in contrast to human-robot collaboration and multi-agent learning (Littman, 1994; Nikolaidis et al., 2016).

3.2. Hierarchical HRL

We introduce a *hierarchical framework for HRL (hHRL)* as one that decomposes SAR interventions into computationally tractable state-action subspaces. Building on the work of Kulkarni et al. (2016), the hHRL framework is structured as a two-level hierarchy, shown in **Figure 1**. At the top level, a meta-controller considers high-level information about the intervention state and activates some lower-level controller. SAR-specific controllers wait for activation to select the robot's action based on a simplified state representation. The hHRL framework assumes that SAR interventions can be characterized by five abstract action categories: (1) instructions I , (2) promises P , (3) feedback F , (4) disclosures D , and (5) inquiries Q . Each category is modeled by a separate controller activated by the overarching meta-controller.

Each controller is responsible for a theoretical subset of SAR actions, henceforth referred to as *SAR acts*. In this work, SAR acts are formalized based on the directive, commissive, and representative illocutionary speech acts, or simply illocutions,



originally defined in linguistic semantics (Searle and Searle, 1969; Austin, 1975). Searle (1976) defined *illocutions* in terms of speaker, hearer, sincerity condition, psychological state, propositional content, and direction of fit. In the SAR context, the speaker is the robot, the hearer is the human user, and the direction of fit is either action-to-state, where the objective is to make the robot's action match the state of the intervention, or state-to-action, where the objective is to make the state of the intervention match what is expressed through the robot's action. Illocutions uniquely manifest themselves through other modes of communication, such as gestures (Mehrabian, 2017), pictures (Danesi, 2016), music (Kohn et al., 2004), and other multimodal signals (Horn, 1998; Forceville and Urios-Aparisi, 2009). These alternative signals are particularly relevant to SAR because robots have inherently expressive embodiments (Fong et al., 2003). Additionally, SAR interventions target special populations (Feil-Seifer and Matarić, 2005), such as linguistic minorities [e.g., American Sign Language (Stokoe et al., 1976)] or persons with disabilities that involve speech and language difficulties or delays [e.g., Dysarthria (Darley et al., 1969) or autism spectrum disorder (Kasari et al., 2012)]. Therefore, SAR acts are defined to be illocutions irrespective of communicative modality.

3.3. Abstract Controllers

Within the hierarchical model, *instructions* are defined as attempts by the robot to get the user to do something that might generate progress toward the intervention goal. Within the instruction controller, there may be some predefined or

learned ordering among instructions, such as the level of challenge or specificity.

Feedback is defined as beliefs expressed to the user by the robot about their past and current interactions. The *direction of fit* is action-to-state, the *sincerity condition* is belief B , and the *propositional content* is that some past or current state s had or has some property p . In this way, feedback F is defined as a specific form of representatives. *Representatives* were defined by Searle (1976) to commit the speaker to the truth of the expressed proposition. The *propositional content* is information about the state relative to some instruction or goal. The feedback controller is responsible for selecting the information or assistance given to the user by the robot. Feedback can be modeled in a variety of ways, the impacts of which have been studied in psychology and human-machine interaction. Specifically, feedback can be adapted to match individual proficiency or independence, as in scaffolded (Finn and Metcalfe, 2010) or graded cueing models (Feil-Seifer and Matarić, 2012; Greczek et al., 2013). It can also be modeled to increase self-efficacy, as in the growth mindset (O'Rourke et al., 2014; Park et al., 2017) and constructive feedback models (Ovando, 1994). Additionally, feedback timing has also been studied (Kulik and Kulik, 1988), such as feedback in response to help-seeking (Roll et al., 2011) and disengagement (Leite et al., 2015).

Extrinsic motivation is a well-studied driver of behavior, explored in educational (Vallerand et al., 1992), professional (Amabile, 1993), and personal settings (Sansone and Harackiewicz, 2000), as well as a common measure in evaluating the effectiveness of human-robot interaction (Breazeal, 1998;

Dautenhahn, 2007; Fasola and Mataric, 2012). In the hHRL framework, *promises* are defined as commitments made by the robot for performing future actions that aim to motivate the user through the promise controller. Promises also relay information critical to collaboration and transparency, expressed via verbal or non-verbal signals, such as gross motion (Dragan et al., 2013). Although they are not directly tied to quantitative measures, promises help to make the robot more personable and consistent over a long-term study period.

Disclosures are defined as beliefs expressed to the user by the robot about its past or current self. The disclosure controller selects internal information for the robot to share with the user as a means of fostering human-robot reciprocity and solidarity. Robot transparency has shown to increase trust (Hancock et al., 2011; Yagoda and Gillan, 2012), improve collaboration (Breazeal et al., 2005; Kim and Hinds, 2006), and build empathetic relationships (Leite et al., 2013). Past work has also shown that non-verbal signals can be particularly effective in disclosing internal states, such as emotion (Bruce et al., 2002).

Inquiries are defined as attempts by the robot to get the user to express some truth. The *inquiry controller* selects what information the robot should attempt to elicit. Inquires may be posed for a variety of interaction benefits, such as improving engagement, relationship, and trust (Hancock et al., 2011). Inquires may also be used to gather feedback about the robot or information about the user, as in interactive machine learning (Amershi et al., 2014).

3.4. Computational Personalization

To personalize the SAR system, the proposed hHRL controller was instantiated as a group of domain controllers, based on the abstract controllers defined in section 3.3. **Figure 2** represents how the abstract controllers were contained within a domain-specific *meta-controller*. The meta-controller activated one controller at a time. We used insights and data from our prior work, reported in Clabaugh et al. (2015), to inform the design of the controllers for SAR personalization. Our prior study collected data from 31 typically developing preschool children who interacted with a SAR tutor in a single session at their child development center preschool. The data were used to develop a model that predicted a child's performance on the game. We used this performance model to bootstrap the prediction of the children's performance in our study. Specifically, the instruction controller in the personalization framework optimized the level of challenge (pLoC) and the level of feedback (pLoF) to match each child's performance, as described next.

3.4.1. Personalization of the Level of Challenge

Personalization was partially accomplished within the instruction controller. Learning games g were randomly sampled without replacement from all games G and parameterized by some personalized *level of challenge* (LoC) $c \in [1, 5]$. The instruction controller was designed to optimize LoC to match individual proficiency. This optimization problem was based on the concepts of *optimal challenge* from the *Challenge Point Framework* by Guadagnoli and Lee (2004) and from the research on the *Zone of Proximal Development* (ZPD) by Chaiklin (2003);

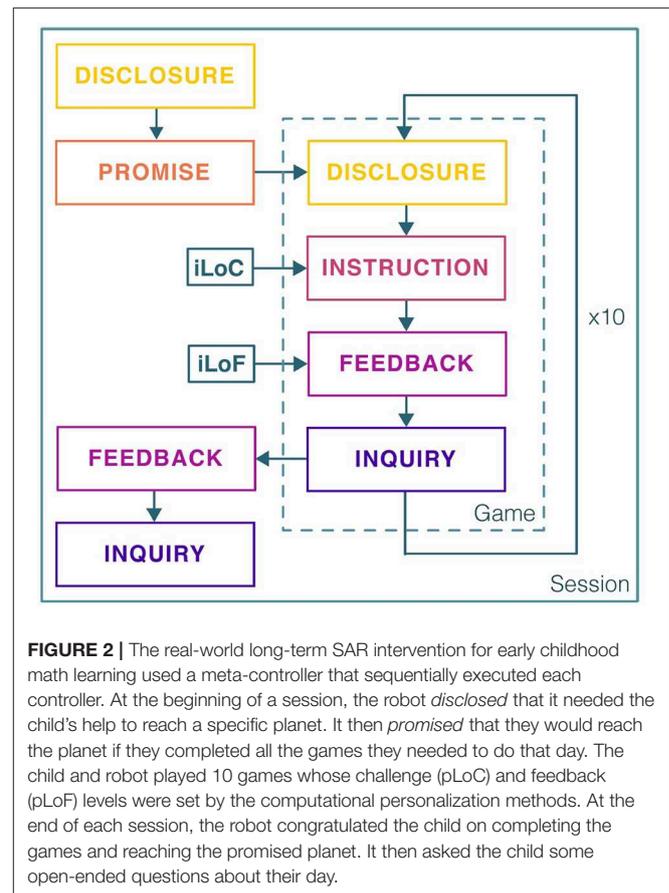


FIGURE 2 | The real-world long-term SAR intervention for early childhood math learning used a meta-controller that sequentially executed each controller. At the beginning of a session, the robot *disclosed* that it needed the child's help to reach a specific planet. It then *promised* that they would reach the planet if they completed all the games they needed to do that day. The child and robot played 10 games whose challenge (pLoC) and feedback (pLoF) levels were set by the computational personalization methods. At the end of each session, the robot congratulated the child on completing the games and reaching the promised planet. It then asked the child some open-ended questions about their day.

both define the goal as challenging individuals enough that they are presented with new information, but not so much that there is too much new information to interpret.

Since the goal is long-term adaptation, *personalized LoC* (pLoC) was framed as a RL problem, trained using Q-learning (Watkins and Dayan, 1992). Within the instruction controller, a reward function was used to quantify the intervention state and supply Q-learning. The intervention, at time t , was defined by:

1. the current game g_t ,
2. the current LoC c_t , and
3. the current number of mistakes m_t

More formally, the state space was defined as G and the action space was defined as C , for a total of $G \times C = 10 \times 5 = 50$ (state, action) couples. As previously explained, the next game g was randomly sampled without replacement from all games G . Therefore, the RL seeks to find a policy with the optimal LoC $c \in [1, 5]$ per game for the individual child.

Given the formulation above, the RL would select and evaluate different LoCs for each child. If some LoC in some game was too difficult or too easy for a child, the RL would learn to select a different LoC for that game, over time. This was accomplished through a reward function designed to maximize LoC without pushing the learner to make too many mistakes. Formally, at time t , let m_t be the number of

mistakes a learner has made and M be the pre-defined threshold of maximum mistakes [we used a threshold of five, based on our empirical findings from prior research (Clabaugh et al., 2015)]. The reward function $R(t)$ returns a value equivalent to the LoC c_t , unless the $m_t > M$; then, $R(t)$ returns the inverse of LoC.

$$R(t) = c_t \cdot MC(t), \quad (1)$$

where

$$MC(t) = \begin{cases} 1, & \text{if } m_t \leq M \\ -1, & \text{otherwise.} \end{cases} \quad (2)$$

3.4.2. Personalization of the Level of Feedback

Feil-Seifer and Matarić (2012) and Greczek et al. (2014) applied the concept of *graded cueing* to adapt feedback in the context of SAR interventions for children with ASD. A similar approach was taken in this work to instantiate the feedback controller, as mentioned in section 3.4. Analogous to the instruction controller, the feedback controller was modeled as a MDP, wherein the decision was to select one of five levels of feedback (LoF) $f \in [1, 5]$ to match individual need. The feedback actions were specific to early mathematics learning.

Personalized LoF (pLoF) was framed as a RL problem, trained using Q-learning (Watkins and Dayan, 1992) over many repeated interactions. Within the feedback controller, at time t , the intervention was represented by four parameters:

1. the current game g_t ,
2. the current LoF f_t ,
3. the current number of mistakes m_t , and
4. the current number of help requests h_t

Similar to LoC, the state space for the feedback controller consisted of the $G = 10$ game states. The action space consisted of the four LoFs $f \in [1, 4]$. The final LoF $f = 5$ was not included as part of the personalization problem. The final feedback level was selected if and only if the child made more than the five allotted mistakes, and the meta-controller would move on to the next interaction. Therefore, the feedback controller included a total of $G \times F = 10 \times 4 = 40$ (state, action) couples.

The reward function was designed to minimize LoF without pushing the learner to make too many M mistakes (where M was the predefined threshold of maximum mistakes) or penalizing them too heavily for making help requests.

$$R(t) = \left(-1 \cdot \frac{f_t}{m_t + h_t + 1} \right) + MC(t), \quad (3)$$

where

$$MC(t) = \begin{cases} 5, & \text{if } m_t \leq M \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

4. PERSONALIZED SAR INTERVENTION DESIGN

The SAR personalization framework was instantiated in a SAR systems designed for and evaluated in a month-long, in-home SAR intervention in the homes of children with ASD, and approved under USC IRB UP-16-00755. The details of the SAR system, study design, data collection, and outcomes measures are described next.

4.1. System Design

The physical robot was designed to be a near-peer learning assistant, intended to act as the child's companion rather than tutor. Toward that end, it was given a neutral, non-threatening character that presented educational games on a tablet and provided personalized feedback.

4.1.1. Physical Design

To enable long-term in-home deployments, including ensuring the protection of the system's sensitive components, we designed a self-contained and portable system, shown in **Figure 3**, consisting of the robot, and a container that encompassed the robot's power supply, speakers, and tablet. The container was approximately the same width as the robot to minimize the overall system footprint.

The robot platform we designed was modified the Stewart Platform Robot for Interactive Tabletop Engagement (SPRITE) with the Kiwi skin (Short et al., 2017). SPRITE used CoR-Dial, also known as the Co-Robot Dialogue system, the software stack that controls the robot's physical movements and virtual face. The SPRITE consists of a 3D printed base, housing electronic components and threaded rods that support a laser-cut platform with six degrees of freedom. Within the exterior skin, a small display was used to animate the robot's face that included two eyes, eyebrows, and a mouth, all of which were controlled using Facial Action Coding System (FACS) coding in CoR-Dial.

The Kiwi skin and character were designed to appeal to the target user population. Children with ASD are often overwhelmed by sensory input, so Kiwi was designed to be non-threatening and simple in its affective displays. It was also gender-neutral in its appearance, allowing each child to assign the robot's gender if and as desired.

4.1.2. Game Design

The design of the SAR intervention was conceptualized by our multidisciplinary team of researchers, leveraging established game design principles, including iterative prototyping (Adams, 2013). Through these processes, Clabaugh et al. (2018) designed an intervention that balanced the needs of the domain with the limitations of SAR technology. The initial game prototype was presented to a focus group of early childhood educators who served as subject matter experts and provided formative feedback on what was developmentally appropriate for children with ASD diagnoses. This informed the second generation of the game design, which was then piloted in a preschool classroom. Following these pilot studies, further adjustments were made

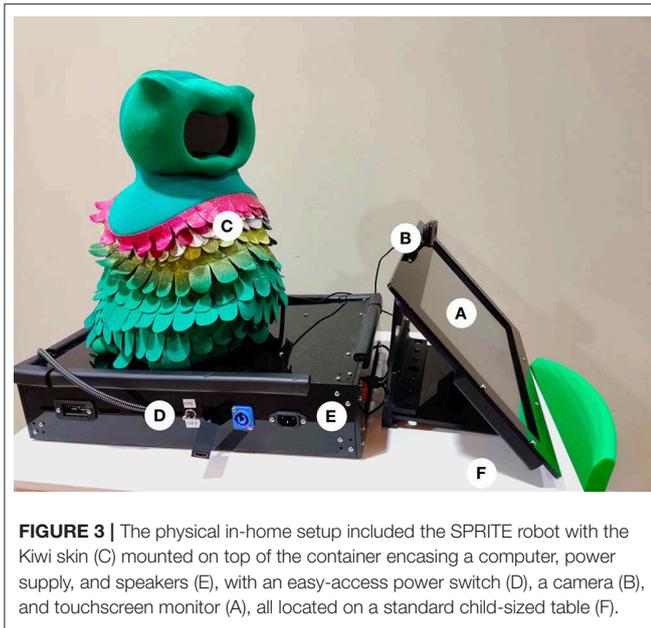


FIGURE 3 | The physical in-home setup included the SPRITE robot with the Kiwi skin (C) mounted on top of the container encasing a computer, power supply, and speakers (E), with an easy-access power switch (D), a camera (B), and touchscreen monitor (A), all located on a standard child-sized table (F).

to accommodate specific needs of children with ASD before the system was iteratively deployed and validated over multiple, long-term, in-home interventions for 17 children with ASD.

The game types within the system were tailored specifically for children with ASD, based on previous case studies, developmentally appropriate practices in working with children with ASD (Copple and Bredekamp, 2009), and standards recognized by the National Association of the Education of Young Children. More specifically, they were developed in concert with developmentally appropriate practices for young children ages 3–8 and informed by contemporary learning theory Omrod et al. (2017). Each game employed a scaffolded approach to gradually increasing difficulty level as the child navigated through successful completion of a particular game level (Sweller et al., 2007). Both the content and difficulty levels were also aligned both to best practices in child development standards of the National Association of Education of Young Children (NAEYC) and the National Common Core Mathematics Standards [for more advanced levels; CCMS (Copple and Bredekamp, 2009)].

The games were also aligned with the Wechsler Individual Achievement Test (WIAT), a developmental-level standard assessment (Wechsler, 2005), used as a pre-post measure of the impact of the game, as described in section 4.3.1. Numerical operation and math reasoning were selected as pre-academic content for the games because they are the early math skills needed by children in preschool and kindergarten. As a result, they are also areas that control for potential social biases found in many early childhood games.

Figure 4 illustrates an example of the different challenge levels of one of the games; game challenge levels were personalized to each child participant as described in section 3.4.1.

4.1.3. Child-Robot Interaction

The Kiwi character described itself as a space explorer and a peer to the child user that continuously needed help from the

child in order to return to its home planet. Users were told they could help Kiwi by playing the provided tablet-based games. The games tested a variety of preschool and kindergarten math skills, including addition, counting, and pattern matching. The SAR system offered ten different types of games based on five different levels of challenge (LoC). Child participants were encouraged to play at least one game during each interaction; the games involved the user performing the following on-screen tasks:

1. Pack Moon-Rocks: Drag 1–10 moon-rocks into a box.
2. Select Galaxy: Select the galaxy with more or fewer stars.
3. Select Planet: Select the planet with a particular number.
4. Feed Space Pets: Evenly divide a set of stars between two “alien pets.”
5. Pets on a Spaceship: Drag numbered “alien pets” into a spaceship in increasing or decreasing order.
6. Organize Moon-Rocks: Separate and organize moon-rocks based on sprite and number.
7. Organize Space Objects: Separate and organize various space-themed objects based on sprite and number.
8. Pattern Completion: Complete a pattern with the provided space objects.
9. Identify Alien Emotion: Determine the emotion of one or more “alien friends” based on their facial expressions.

The graphics in the game used an age-appropriate comic book design style, with colorful aliens guiding the user through the games. Each game allowed up to five mistakes; every mistake was followed by a verbal hint delivered by Kiwi paired with child-like body movements that signaled whether the user was struggling or excelling. The feedback actions were specific to the game context of early mathematics learning. For example, if the child was presented with the instruction “Put five energy crystals into a box” but used too few crystals, the feedback controller executed one of the following actions:

1. “We need to have a total of five energy crystals inside the box.”
2. “Try counting out loud as you drag each crystal one by one.”
3. “You have too few energy crystals. Try adding some to the box.”
4. “We currently have three energy crystals. So we need two more energy crystals. Can you drag two more crystals into the box?”
5. “Let’s try something else.”

4.2. Study Participants

Seventeen children with ASD were included in this research, hereafter referred to as P1–P17. Families were recruited through regional centers within the state’s Department of Developmental Services and through local school districts. Together, these two recruitment venues provide services for >10,000 children and adults with ASD, with ~1/3 of the population under the age of ten.

Study recruitment flyers were provided to service coordinators, school district administrators, and family research center coordinators who are employed in regional centers and the schools. Families who were interested in participating in the research contacted our research team and provided written information about their child with ASD. A licensed psychologist



on our team reviewed each child's developmental and health information for a match with the study's inclusion criteria:

1. Age between 3 and 8 years old
2. Stable physical, sensory (hearing, vision), and medical health
3. English as a primary language spoken in the family
4. Clinical diagnosis of ASD in mild to moderate ranges as described in the Diagnostic and Statistical Manual of Mental Disorders–Version 5 (Van Bourgondien et al., 1992; Baird et al., 2003; Dover and Le Couteur, 2007; Kanne et al., 2008).

Of the 17 children in the study, 2 were female and 15 male. They were between 3 years, 4 months and 7 years, 8 months of age. Additionally there were 3 sets of sibling pairs (P3 and P4, P5 and P6, and P16 and P17). More information about each participants living situation, education level, and age can be found in the **Supplementary Materials**.

Due to the challenges of ASD and real-world studies, there were some exceptions among the participants. Specifically, there are no personalization data for P1 and P2, as the system was not yet fully developed for those first two deployments. Additionally, P3 did not complete the post-study assessments for personal reasons, but did participate in the study for over a month and provided all other study data. Besides these exceptions, the rest of the participants participated in the entire study.

There is no control condition in this study, as is common in ASD studies, because individuals on the autism spectrum present an extremely broad range of symptoms, symptom combinations, and symptom severities. Consequently, work with ASD participants typically follows a single-case study model rather than the randomized trial model. The single-case study model relies on pre/post-comparisons, as was done in this paper (Lobo et al., 2017). The pre/post-WIAT Interventions in

section 4.3.1 serve as a sample baseline to evaluate participant improvement over the course of the study.

All child participants in the study were enrolled in full-time educational and therapeutic interventions that were consistent with the state's educational and developmental services standards and statutes. These services varied based on child needs and family preferences. All child participants had intelligence scores within "normal" limits levels (scores >70) based either on the Leiter International Performance Scale-3 (Roid et al., 2013) or the by the Differential Ability Scales (Elliott, 2012).

The child participants' ASD diagnoses were obtained via clinical best estimate (CBE) by trained psychologists or psychiatrists who had >10 years of experience in diagnosing children with ASD and other developmental disabilities. The tools used to diagnose ASD varied across clinician and referring agency. In each case, multiple measures were used to determine the diagnosis and level of ASD. Common measures used for the ASD diagnoses were the Autism Diagnostic Interview-Revised (Wing et al., 2002; Tadevosyan-Leyfer et al., 2003), Autism Diagnostic Observation Schedule (Lord et al., 2000; Gotham et al., 2008), and Child Autism Rating Scale (Van Bourgondien et al., 1992). All children with ASD diagnoses in the study had diagnoses in the mild to moderate range (Van Bourgondien et al., 1992).

4.3. Procedure

The SAR intervention was deployed in the home of each participating family for at least 30 days. The duration of each deployment was determined by when the minimum number of 20 child-robot interactions was completed; the average duration of deployment was 41 days, with a standard deviation of 5.92 days. On the day of deployment for each family, all

system equipment was provided and assembled by the research team; the only requirement from participating families was a power outlet and sufficient space. During system setup, child participants were assessed by an educational psychologist using the measures described in section 4.2. After the system was set up, the research team conducted a system tutorial with the child participant and family.

To capture natural in-home interactions, the SAR system was fully autonomous and could be turned on and off whenever the family desired. The child participants were encouraged but not required to complete five sessions per week. Similarly, during each session, they were encouraged but not required to play each of the 10 games at least once.

4.3.1. Objective Measures

A large corpus of multi-modal data was collected, including video, audio, and performance on the games. The USB camera mounted at the top of the game tablet recorded a front view of the child participant. A second camera recorded the child-robot interaction from a side view. All interactions with the tablet were recorded, including help requests and answers to game questions.

User engagement was annotated by analyzing the camera data. A participant was considered to be engaged when paying full attention to the interaction, immediately responding to the robot's prompts, or seeking further guidance or feedback from others in the room.

Due to numerous technological challenges common in noisy real-world studies, we were able to analyze sufficient video and audio data from seven participants (P5, P7, P9, P11, P12, P16, P17). A primary expert coder annotated whether a participant was engaged or disengaged for those seven participants. To verify the absence of bias, two additional annotators independently annotated 20% of data for each participant; inter-rater reliability was measured using Fleiss' kappa, and a reliability of $k = 0.84$ was achieved between the primary and verifying annotators.

The primary quantitative measure of cognitive skills gained throughout the study were the pre- and post-assessments, inspired by the standardized Wechsler Individual Achievement Test (WIAT II) (Wechsler, 2005) used to assess the academic achievement of children, adolescents, college students, and adults, aged 4–85. The test evaluates a broad range of academics skills using four basic scales: Reading, Math, Writing, and Oral Language. Within those, there are nine subtest scores, including two math subtests, *numerical operations* (NO) and *math reasoning* (MR), which were the most relevant to the SAR intervention content. For young children, NO refers to early math calculations, number discrimination, and related skills; MR refers to concepts of quantity and order, early word problems, patterning, and other skills that require reasoning to solve problems. WIAT II was selected over the WIAT III because the timing of math fluency in version III presents a potential bias for children with ASD diagnoses.

The WIAT II provides raw and composite scores. Standard scores and percentile ranking are computed by comparing an individual assessment to large national samples of typically developing individuals aged 3 to adult (i.e., 2015 US normative sample $N = 2,950$). A standard score of 100–110 is considered

an “average achievement score” by national standards. The percentile ranking indicates how an individual compares to the national sample on which the tests were normed. The WIAT-II was used as a pre-post comparison measure to determine achievement gains over the SAR intervention. Procedurally, the pre-assessment was conducted during the first few days of the intervention and the post-assessment was conducted at the end of the intervention for each child.

4.3.2. Subjective Measures

We conducted biweekly interviews with participating families throughout the deployments to evaluate the system in terms of its usefulness and relationship with the participating child, rating responses on a 7-point Likert scale with 1 being least likable and 7 being most likable. Given the variable nature of in-home studies and different degrees of ASD across the participants, the surveys used a single-subject design (Horner et al., 2005) wherein each child served as their own unique baseline. The semi-structured interviews contained similar questions, each tailored for a specific evaluation criterion, as follows.

Based on prior work by Moon and Kim (2001), these were the questions about Kiwi's usefulness:

- Does Kiwi help your child do better on the tasks? Why or why not?
- How could Kiwi be more useful?
- How involved do you have to be while your child is playing with Kiwi?

Based on prior work by Lee et al. (2005) and Rau et al. (2009), these were the questions about the child-robot relationship:

- Do you think Kiwi is your friend?
- Do you think Kiwi listens to you?
- Do you feel like Kiwi knows you?

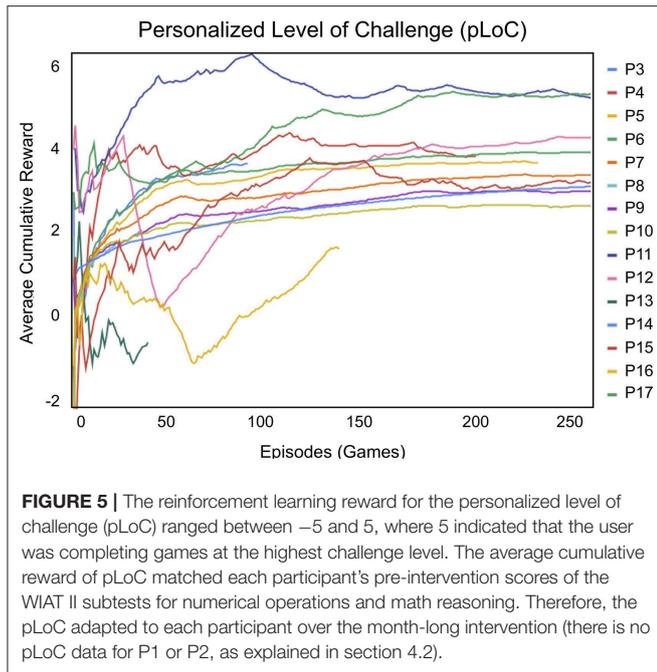
5. RESULTS

The presented month-long in-home deployments produced a large set of results. Sections 5.1 and 5.2 describe the patterns and quantitative results, respectively, of the hHRL framework instantiation. Section 5.4 discusses how the adaptive system influenced the engagement of the child participants. Section 5.5 reports on how the adaptive SAR system influenced cognitive skills gains across all participants, as measured by the pre-post intervention assessments.

5.1. Personalized Level of Challenge

As illustrated by the learning curve in **Figure 5**, the personalized level of challenge (pLoC) changed over time and varied by participant. Since the goal of the adaptation was to find the optimal LoC for each participant, this learning curve cannot be interpreted in a traditional sense. For instance, if a child was not proficient at math, the learning system may not have been able to reach higher reward values, because the reward is based on both LoC and child performance.

Therefore, other factors must also be considered in interpreting the pLoC results. First, more than 100 episodes or



games were required for pLoC to begin to converge. For example, for participants P3 and P8, the pLoC curve did not have a chance to converge over the few games these participants played. On the other hand, for P6, P11, and P15, the system was able to smoothly adapt given the long interaction periods.

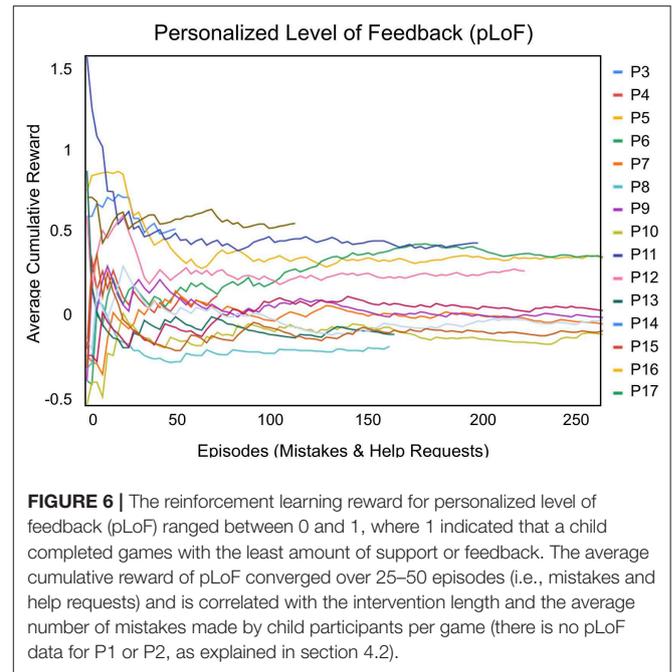
If a child played 10–20 games in a session, consistent with the 13.27 average of the study, then 10 sessions were required before the pLoC began to converge, totalling to ~ 132 episodes. This is a reasonable requirement given that the participants completed an average of 14.10 sessions with the robot. Excluding participants P3 and P8, who, as noted above, played significantly fewer games per session, we find an average of 17.57 sessions with the robot.

Consequently, we can conclude that the SAR system was able to adapt and personalize to each child over time. Specifically, the pLoC implementation of the instruction controller did personalize to each child, but required a minimum number of episodes and interaction consistency to do so.

5.2. Personalized Level of Feedback

The learning curve for the personalized level of feedback (pLoF) model, shown in **Figure 6**, adapted the level of feedback to each participant more rapidly than pLoC. Analogous to pLoC, the pLoF learning curve cannot be interpreted in the traditional sense of simply maximizing cumulative reward; it is meant to match each child's need.

Participants with high mistake totals and long interventions usually had the longest feedback curves and, subsequently, allowed the system to adapt to their needs. The pLoF tail was longest for P10, who had the third highest mistake average, balanced with overall intervention length. Although P3 and P8 had the highest mistake averages, they also had the shortest interventions. This can be compared to the pLoF curves for P5 and P6, who had the lowest mistake averages and longest



intervention lengths; the cumulative reward is higher and tails are shorter for both P5 and P6 compared to those of P10. Subsequently, P10 stands out as the longest and flattest among the three, demonstrating the value of longer interactions. Overall, the pLoF model successfully adapted to each child participant over time.

5.3. Participant SAR Evaluation

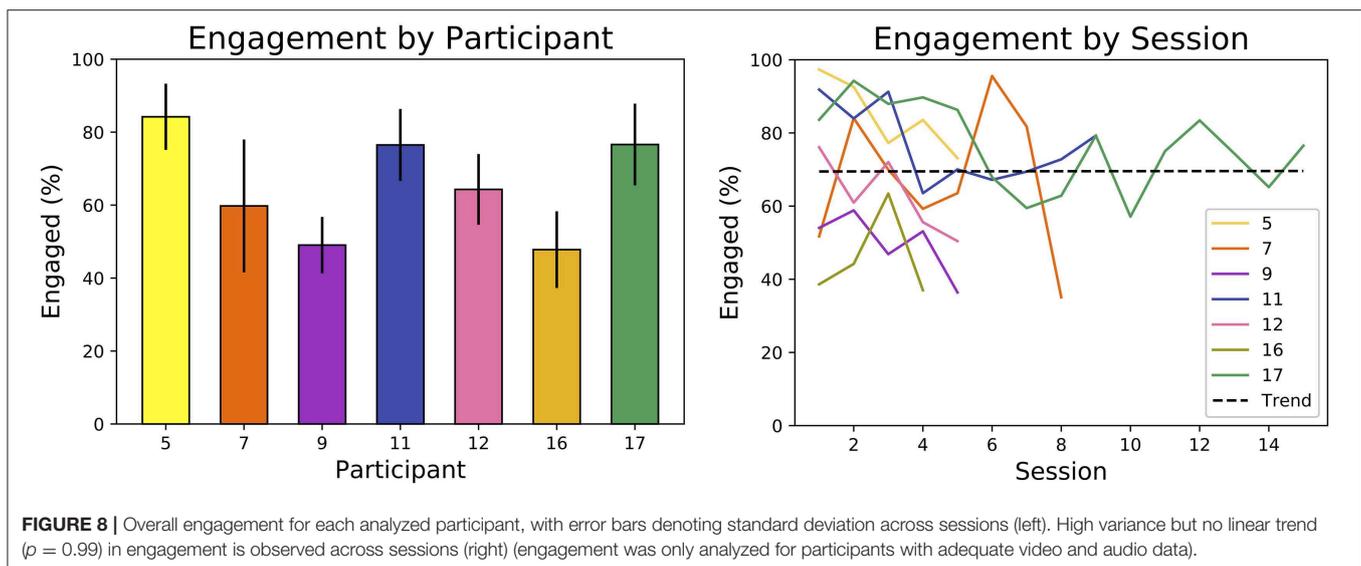
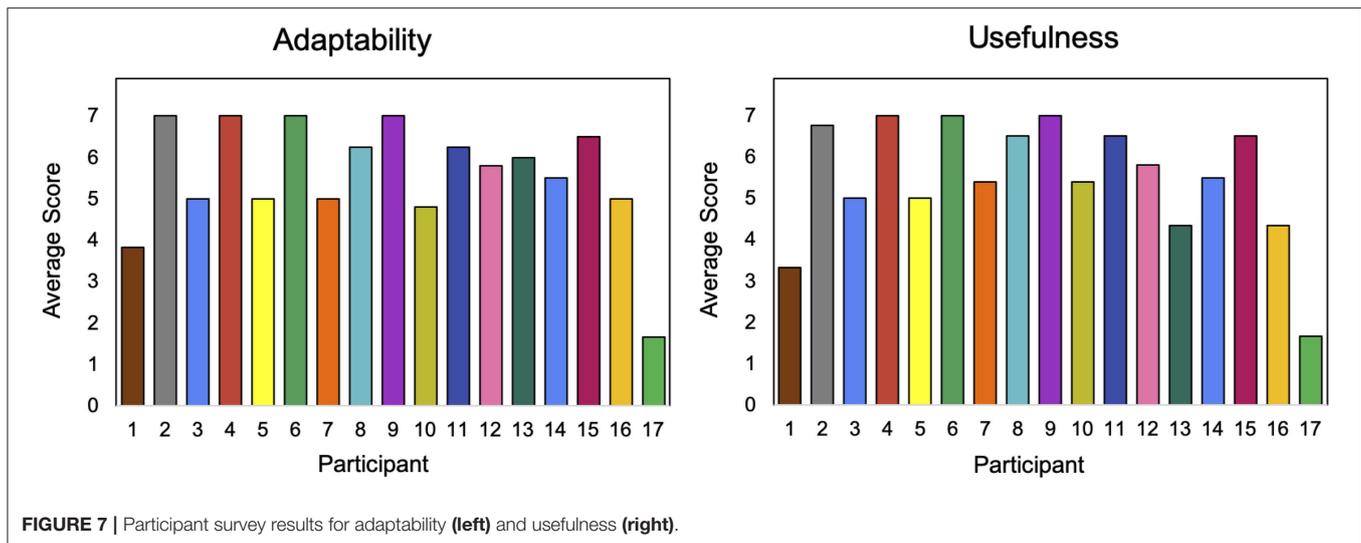
SAR survey results, utilizing a seven-point Likert scale, assessed the average adaptability and usefulness of the system throughout the study.

Participants' responses about the adaptability of the SAR system, seen in **Figure 7**, correlate with the challenges encountered in adapting to the individual needs of each participant. P17 reported the lowest average score for adaptability and usefulness. The result for P17 is likely due to the participant's age, as this was the second oldest and highest performing student in the study, so even the maximum difficulty was too easy for the participant.

The reported scores for usefulness were similar to those of adaptability, as seen in **Figure 7**, since the two measures are related: a system that is more adaptive to a participant is more useful. P1 and P17 were once again outliers with the lowest reported scores for usefulness.

P1 and P7 had personal similarities: they were less than a year apart in age and had parents with the same levels of education (high school). Consequently, one would expect the system to adapt relatively similarly to both participants. Their reported scores for usefulness (5 vs. 3) and adaptability (5 vs. 4) were similar, thus supporting the consistency of the system across participants.

Sibling pairs (P3 and P4, P5 and P6, and P16 and P17) showed discrepancies that can be explained by the fact that the system was



better suited to the needs of one sibling than the other, likely due to their age. For example, P3 was younger than P4, and therefore was not able to engage with the games as well, resulting in the lower adaptability and usefulness scores. Similarly, P6, the older sibling, reported higher scores for adaptability and usefulness than P5. The higher scores mean that the child liked the robot more and found it more adaptable and useful.

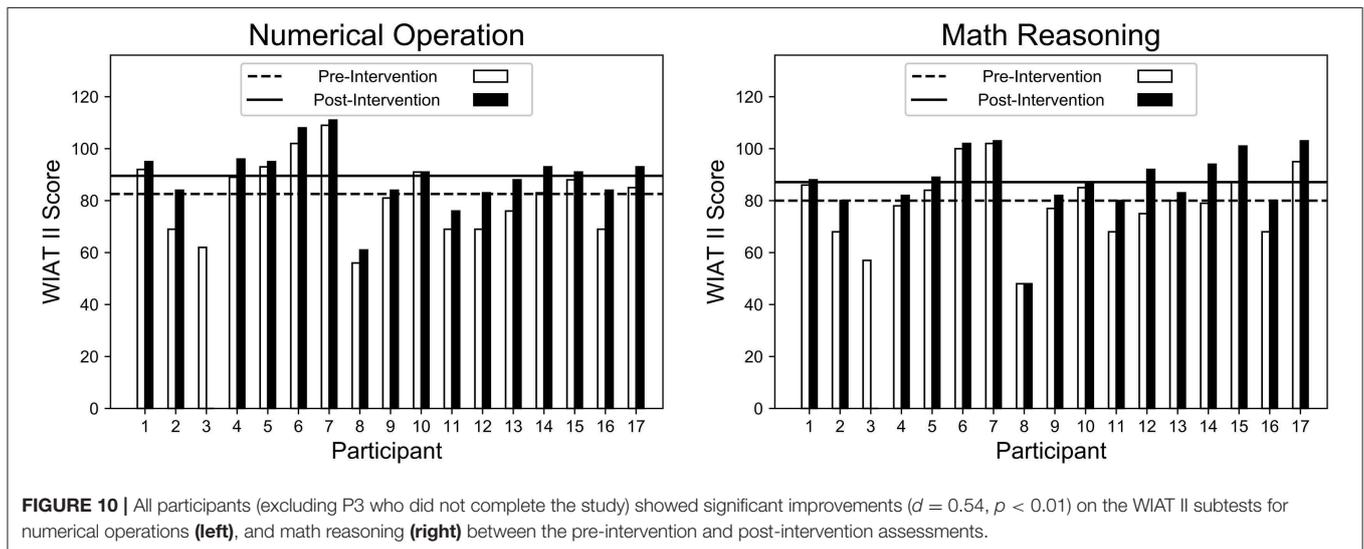
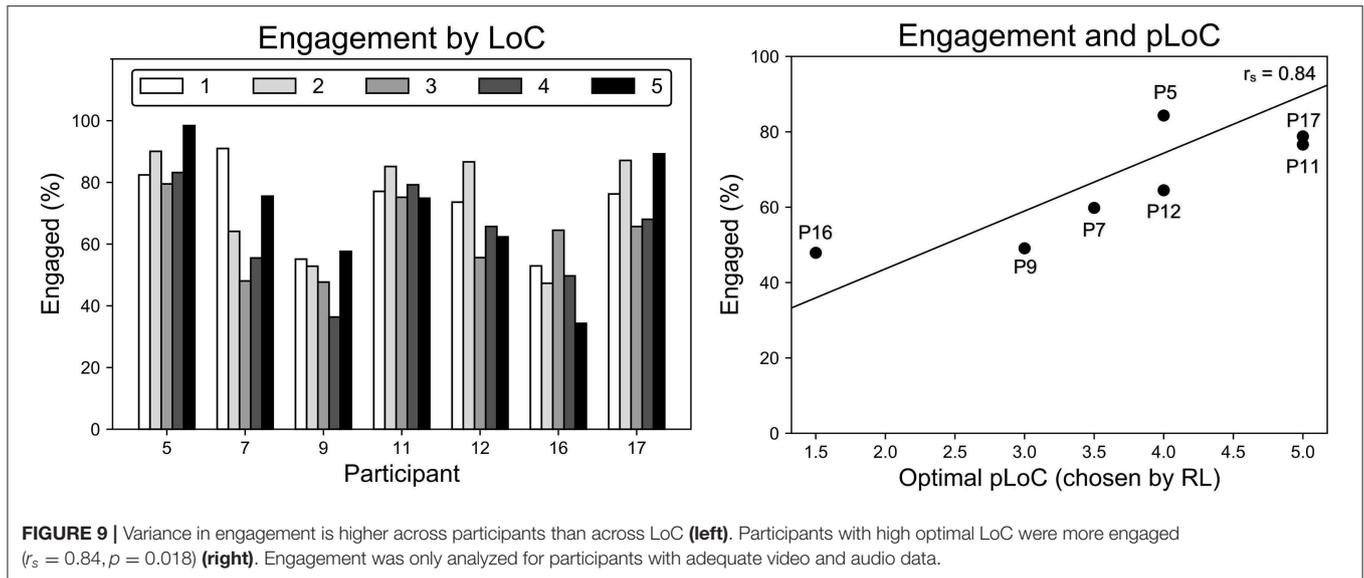
5.4. Effect of Personalization on Engagement

We found that our SAR system elicited and maintained participants' engagement throughout the month-long intervention, an important measure of effectiveness. As mentioned in section 4.3.1, we analyzed the seven participants (P5, P7, P9, P11, P12, P16, P17) with adequate video and audio data to analyze measures of engagement.

5.4.1. Short-Term and Long-Term Engagement

The SAR system maintained reasonable levels of participant engagement during individual sessions and over the month-long intervention. As shown in **Figure 8**, all participants were engaged on average 65% of the intervention. Across sessions, participants had an average engagement range of 32% and standard deviation of 11%. However, there was no statistically significant ($p = 0.99$) increase or decrease in engagement over the study, as determined by a regression t -test and shown by the plotted trend line. In addition, the median duration of continuous engagement over all participants was higher than the median duration of continuous disengagement: 13–5 s on average, respectively.

Furthermore, the robot was able to elicit and maintain user engagement during each game. Engagement was higher shortly after the robot had spoken; participants were engaged about 70% of the time when the robot had spoken in the previous minute, but <50% of the time when the robot had not spoken for over



a minute. Participants also remained engaged after 5 min of starting a game nearly 60% of the time.

5.4.2. Engagement and the Level of Challenge

Engagement varied significantly across participants and their level of challenge (LoC), as shown in Figure 9. A two-way analysis of variance (ANOVA) showed ($p < 0.01$) that average engagement for each participant varied significantly and that average engagement under each LoC also varied significantly. The variance across participants accounted for 91% of the total variance, indicating the importance of personalization in SAR.

The personalized level of challenge (pLoC) did not necessarily maximize engagement. As discussed above, pLoC eventually converged to an optimal LoC for each participant. But, as shown in Figure 9, participants whose optimal LoC was low were less engaged ($r_s = 0.84, p = 0.018$). We hypothesize that this effect

is due to the time required for the learning system to adapt to each user; it took >100 games for the pLoC to begin to converge, and thus participants with a lower LoC were presented with many games of higher challenge level before convergence. This further supports the importance of personalization for increasing engagement, especially with a sufficiently fast convergence rate.

5.5. Impact on Math Learning

Overall, this study observed positive gains in math learning for all participants, excluding P3 who did not complete the study. As seen in Figure 10, participants' pre- and post-intervention scores on the WIAT II subtests increased significantly for numerical operations (NO) ($p = 0.002$) and math reasoning (MR) ($p < 0.001$), as determined by a t -test. In addition, both NO and MR scores had a significant effect size of $d = 0.53$ and $d = 0.54$, respectively, as calculated using Cohen's d .

The result reveal that NO and MR both increased even when there was a large discrepancy between the initial assessment of certain participants. For example, P17 scored much higher on MR than on NO on the pre-assessment and even with such different starting points, both NO and MR increased at the post-assessment. On the other hand, P11 started with the same MR and NO scores, and both scores improved after the intervention.

When observing total cognitive gains, it is important to consider developmental factors: the age and subsequent skill level of each participant. Where older students generally had smaller net gains, they started near or above average. On the other hand, younger students started far below average, and thus had much room to improve. P8's pre-intervention scores (MR = 48; NO = 56) were significantly below the national average. Given P8's age (3.75 years), the scores are cautiously computed in terms of what they represent nationally. In another case, P16's pre-intervention scores (MR = 68; NO = 69) were far below the national average. P16 was the youngest participant (3.11 years) and still made significant progress, improving by over 10 points in both categories (MR = 80; NO = 84). On the other hand, P17 was tied second oldest (7.2 years) and only made marginal gains, despite making few mistakes and performing at the highest challenge level.

6. DISCUSSION

The results of the long-term in-home deployment provide several insights for personalization in SAR.

We found that both the personalized levels of challenge (pLoC) and feedback (pLoF) converged for almost all participants. After ~100 games, the feedback and challenge curves stabilized, showing that the system adapted to an appropriate LoC for each student. Therefore, the long-term nature of the study was important for successful personalization. The participants with the longest episodes in the pLoC were P6, P16, and P11, with 715, 592, and 520 episodes (games played), respectively. In contrast, pLoF interacted most with for P10, P7, and P9, who had 353, 237, and 228 episodes (mistakes and help requests), respectively. The SAR system adapted to the participants in both cases.

Regardless of the difference in sessions, participants who yielded a consistent score by the end of the interaction in the pLoC had similar success with pLoF, and vice versa. This happened for participants who interacted with equal or above average 113.4 and 302.5 episodes for pLoF and pLoC, respectively. On the other hand, P16 illustrated the negative impact of minimal interaction, as both the pLoC and pLoF failed to standardize given only 129 total interactions both on pLoC and pLoF, reaching over 171 episodes below average for pLoC. Within the interaction, the pLoC reward for P16 fluctuated by 2.6 points between the 61st and 129th episode. For reference, the second highest fluctuation in this same interval was 1.02 points by P16, whose system ultimately converged after 592 games.

Overall, the pLoF and pLoC demonstrate the ability to adapt to each user's preferences given their willingness to interact with the robot and provide the system opportunities to learn. The

participant surveys support this conclusion and provide the user's perspective on the SAR's ability to adapt.

P6, P7, P9, and P11, who both the pLoC and pLoF had adapted to, reported in their post-interaction interviews an average rating of 6.25, showing a shared appreciation for the system's adaptiveness. The only study participant who believed the system did not adapt was P17, who likely felt this way because of limited success with the feedback model; P17 had only 112 feedback episodes over 298 total games. Aside from this outlier, the survey results supported the effectiveness of pLoC and pLoFs. P9 was an ideal participant, who believed the system adapted and had above average episodes while stabilizing both pLoC and pLoF.

Usefulness questionnaire data provide additional insights into the value of creating an adaptive system. All participants reported very similar scores for usefulness and adaptiveness, implying that the usefulness of the system is related to its adaptiveness. The pre-post assessments supported this finding while providing quantitative data about the learning gains of each participant as a result of SAR personalization.

Participants whose optimal LoC was lower were less engaged, as shown in **Figure 9**. For example, the system converged to the lowest pLoC for P16, who also had the second lowest engagement. This is likely because P17 was presented with games of higher challenge before the system began to converge to an optimal LoC. When also considering that P16 had a below average number of episodes, it is likely the robot failed to adapt quickly ultimately discouraging the participant from interacting further.

The analysis of the objective and subjective outcome measures supports the success of the system as a whole, with all participants improving in math skills over the course of the long-term in-home interaction. Regardless of whether the system was able to adapt to an optimal LoC, all participants demonstrated cognitive gains. The participants gained an average of 7.0 points on numerical operation (NO) and 7.125 points on math reasoning (MR). Although for P16 the system was unable to adapt both in pLoF and pLoC personalization, that participant was still in the top five in both NO and MO gains, with an increase of 15 and 12 points, respectively. This is likely due to the participant's initially low scores that allowed much room for improvement. All participants who had at or above average number of episodes (either in pLoC and pLoF) showed strong positive gains. P8 illustrated the disadvantages of insufficient interaction time, being the participant with the least episodes in both pLoF and pLoC and resulting with below average gains in both NO and MR assessments.

7. CONCLUSION

Socially assistive robotics (SAR) has demonstrated tremendous potential for use in high impact domains, such as personalized learning for special needs populations. This work considered the problem of computational personalization in the context of long-term real-world SAR interventions. At the intersection of HRI and machine learning, computational personalization seeks to

autonomously adapt robot interaction to meet the unique needs and preferences of individual users, providing a foundation for personalization.

This work presented a formalized framework for human-robot learning as a hierarchical decision-making problem (hHRL) that decomposes a SAR intervention for tractable computational personalization, and utilized a reinforcement learning approach to personalize the level of challenge and feedback for each user. The approach was instantiated within the interactive games and tested in month-long in-home deployments with children with ASD. The SAR system was able to personalize to the children with ASD who demonstrated cognitive gains, supporting the effectiveness of the approach.

The body of results of the presented study demonstrate that the hHRL framework and its instantiation can engage and adapt to children with diverse needs in math learning over multiple weeks. These findings highlight the tremendous potential of in-home personalized SAR interventions.

DATA AVAILABILITY STATEMENT

The dataset analyzed in this study includes identifiable video and audio data of children with autism spectrum disorders and their families, along with video information and images representing their homes. Consequently, the University IRB prohibits distribution of the dataset to protect the privacy of the research participants.

ETHICS STATEMENT

The studies involving human participants were reviewed and approved by Institutional Review Board (IRB). Written informed consent to participate in this study was provided by the participants' legal guardian/next of kin.

AUTHOR CONTRIBUTIONS

CC developed the learning framework and its instantiation and parts of the system software, oversaw the implementation

of the rest, most of the deployments, data annotation, and provided source text for the paper. KM outlined, wrote, and edited the paper while managing various stakeholders throughout the writing process. SJ designed and modeled the engagement process, and also detailed the Background and Engagement Results. RP facilitated the bi-weekly surveys and edited the Results section. DB contributed to the game performance model development, and led the first set of in-home deployments and data collections. ZS designed the graphs used throughout the paper and served as a secondary editor. ED contributed to the design of the Kiwi robot hardware and the overall SAR system for time-extended in-home deployments. RL rated the survey responses and transcribed all survey questions. GR provided domain expertise in autism and early child learning, assessment methods and tools, lead the participant recruitment, and administered the pre- and post-study assessments and interviews. MM was the project lead. she advised all students, coordinated the robot and study designs, oversaw data analysis, and extensively edited the paper.

FUNDING

This research was supported by the National Science Foundation Expedition in Computing Grant NSF IIS-1139148.

ACKNOWLEDGMENTS

The authors thank National Science Foundation Expedition Grant for funding this study along with the valuable contributions of various annotators including Balasubramanian Thiagarajan, Kun Peng, Leena Mathur, and Julianna Keller.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/frobt.2019.00110/full#supplementary-material>

REFERENCES

- Adams, E. (2013). *Fundamentals of Game Design. 3rd Edn.* Pearson, GA: New Riders.
- Amabile, T. M. (1993). Motivational synergy: toward new conceptualizations of intrinsic and extrinsic motivation in the workplace. *Hum. Resour. Manage. Rev.* 3, 185–201.
- Amershi, S., Cakmak, M., Knox, W. B., and Kulesza, T. (2014). Power to the people: the role of humans in interactive machine learning. *AI Mag.* 35, 105–120. doi: 10.1609/aimag.v35i4.2513
- Anderson, J. R. (1985). *Cognitive Psychology and Its Implications.* New York, NY: WH Freeman; Times Books; Henry Holt & Co.
- Anderson, J. R., Reder, L. M., and Simon, H. A. (1996). Situated learning and education. *Educ. Res.* 25, 5–11.
- Anderson, L. W., Krathwohl, D. R., Airasian, P., Cruikshank, K., Mayer, R., Pintrich, P., et al. (2001). A taxonomy for learning, teaching and assessing: a revision of bloom's taxonomy. New York, NY: Longman Publishing.
- Artz, A. F., and Armour-Thomas, E. (1992). Development of a cognitive-metacognitive framework for protocol analysis of mathematical problem solving in small groups. *Cogn. Instr.* 9, 137–175. doi: 10.1207/s1532690xci0902_3
- Austin, J. L. (1975). *How To Do Things With Words, Vol. 88.* London: Oxford University Press.
- Bainbridge, W. A., Hart, J., Kim, E. S., and Scassellati, B. (2008). "The effect of presence on human-robot interaction," in *The 17th IEEE International Symposium on Robot and Human Interactive Communication, 2008. RO-MAN 2008* (Munich: IEEE), 701–706.
- Baird, G., Cass, H., and Slonims, V. (2003). Diagnosis of autism. *BMJ* 327, 488–493. doi: 10.1136/bmj.327.7413.488
- Baxter, P., Kennedy, J., Belpaeme, T., Wood, R., Baroni, I., and Nalin, M. (2013). "Emergence of turn-taking in unstructured child-robot social interactions," in *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (Christchurch: IEEE), 77–78.

- Beevers, C. G., and McGeary, J. E. (2012). Therapygenetics: moving towards personalized psychotherapy treatment. *Trends Cogn. Sci.* 16, 11–12. doi: 10.1016/j.tics.2011.11.004
- Begum, M., Serna, R. W., and Yanco, H. A. (2016). Are robots ready to deliver autism interventions? A comprehensive review. *Int. J. Soc. Robot.* 8, 157–181. doi: 10.1007/s12369-016-0346-y
- Bloom, B. S. (1984). The 2 sigma problem: the search for methods of group instruction as effective as one-to-one tutoring. *Educ. Res.* 13, 4–16.
- Bongaarts, J. (2004). Population aging and the rising cost of public pensions. *Popul. Dev. Rev.* 30, 1–23. doi: 10.31899/pgy1.1012
- Breazeal, C. (1998). “Early experiments using motivations to regulate human-robot interaction,” in *AAAI Fall Symposium on Emotional and Intelligent: The Tangled Knot of Cognition, Technical Report FS-98-03* (Cambridge, MA), 31–36.
- Breazeal, C., Kidd, C. D., Thomaz, A. L., Hoffman, G., and Berlin, M. (2005). “Effects of nonverbal communication on efficiency and robustness in human-robot teamwork,” in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2005 (IROS 2005)* (Cambridge, MA: IEEE), 708–713.
- Bruce, A., Nourbakhsh, I., and Simmons, R. (2002). “The role of expressiveness and attention in human-robot interaction,” in *Proceedings ICRA’02. IEEE International Conference on Robotics and Automation, 2002, Vol. 4* (Pittsburgh, PA: IEEE), 4138–4142.
- Castellano, G., Martinho, C., Paiva, A., Leite, I., and Pereira, A. (2012). “Modelling empathic behaviour in a robotic game companion for children: an ethnographic study in real-world settings,” in *International Conference on Human-Robot Interaction (HRI)* (Boston, MA: ACM), 367–374.
- Cesuroglu, T., van Ommen, B., Malats, N., Sudbrak, R., Lehrach, H., and Brand, A. (2012). Public health perspective: from personalized medicine to personal health. *Pers. Med.* 9, 115–119. doi: 10.2217/pme.12.16
- Chaiklin, S. (2003). The zone of proximal development in vygotsky’s analysis of learning and instruction. *Vygotskys Educ. Theory Cult. Context* 1, 39–64. doi: 10.1017/CBO9780511840975.004
- Chan, J., and Nejat, G. (2011). “A learning-based control architecture for an assistive robot providing social engagement during cognitively stimulating activities,” in *International Conference on Robotics and Automation (ICRA)* (Shanghai: IEEE), 3928–3933.
- Chen, Y., Singla, A., Mac Aodha, O., Perona, P., and Yue, Y. (2018). Understanding the role of adaptivity in machine teaching: The case of version space learners. *arXiv:1802.05190*.
- Clabaugh, C., Jain, S., Thiagarajan, B., Shi, Z., Mathur, L., Mahajan, K., et al. (2018). “Attentiveness of children with diverse needs to a socially assistive robot in the home,” in *2018 International Symposium on Experimental Robotics* (Buenos Aires: University of Southern California).
- Clabaugh, C., and Mataric, M. (2019). Escaping oz: Autonomy in socially assistive robotics. *Annu. Rev. Control Robot. Auton. Syst.* 2, 33–61. doi: 10.1146/annurev-control-060117-104911
- Clabaugh, C., Ragusa, G., Sha, F., and Mataric, M. (2015). “Designing a socially assistive robot for personalized number concepts learning in preschool children,” in *2015 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)* (Seattle, WA: IEEE), 314–319.
- Cohn, D. A., Ghahramani, Z., and Jordan, M. I. (1996). Active learning with statistical models. *J. Artif. Intell. Res.* 4, 129–145.
- Conn, K., Liu, C., Sarkar, N., Stone, W., and Warren, Z. (2008). “Affect-sensitive assistive intervention technologies for children with autism: an individual-specific approach,” in *International Conference on Robot and Human Interactive Communication (RO-MAN)* (Munich: IEEE), 442–447.
- Copple, C., and Bredekamp, S. (2009). *Developmentally Appropriate Practice in Early Childhood Programs: Serving Children From Birth Through Age 8*. Washington, DC: National Association for the Education of Young Children, 3.
- Danesi, M. (2016). *The Semiotics of Emoji: The Rise of Visual Language in the Age of the Internet*. London: Bloomsbury Publishing.
- Darley, F. L., Aronson, A. E., and Brown, J. R. (1969). Differential diagnostic patterns of dysarthria. *J. Speech Lang. Hear. Res.* 12, 246–269.
- Dautenhahn, K. (2007). Socially intelligent robots: dimensions of human-robot interaction. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 362, 679–704. doi: 10.1098/rstb.2006.2004
- Deng, E., Mutlu, B., and Mataric, M. J. (2019). Embodiment in socially interactive robots. *Found. Trends Robot.* 7, 251–356. doi: 10.1561/23000000056
- Desmarais, M. C., and Baker, R. S. (2012). A review of recent advances in learner and skill modeling in intelligent learning environments. *User Model. User Adapt. Interact.* 22, 9–38. doi: 10.1007/s11257-011-9106-8
- DeVault, D., Artstein, R., Benn, G., Dey, T., Fast, E., Gainer, A., et al. (2014). “Simsensei kiosk: a virtual human interviewer for healthcare decision support,” in *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems* (Paris: International Foundation for Autonomous Agents and Multiagent Systems), 1061–1068.
- Diehl, J. J., Schmitt, L. M., Villano, M., and Crowell, C. R. (2012). The clinical use of robots for individuals with autism spectrum disorders: a critical review. *Res. Autism Spectr. Disord.* 6, 249–262. doi: 10.1016/j.rasd.2011.05.006
- Dover, C. J., and Le Couteur, A. (2007). How to diagnose autism. *Arch. Dis. Child.* 92, 540–545. doi: 10.1136/adc.2005.086280
- Dragan, A. D., Lee, K. C., and Srinivasa, S. S. (2013). “Legibility and predictability of robot motion,” in *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (Budapest: IEEE), 301–308.
- Dudley, J. J., and Kristensson, P. O. (2018). A review of user interface design for interactive machine learning. *ACM Trans. Interact. Intell. Syst.* 8:8. doi: 10.1145/3185517
- Duquette, A., Michaud, F., and Mercier, H. (2008). Exploring the use of a mobile robot as an imitation agent with children with low-functioning autism. *Auton. Robots* 24, 147–157. doi: 10.1007/s10514-007-9056-5
- Durlak, J. (2011). The impact of enhancing students’ social and emotional learning: a meta-analysis of school-based universal interventions: social and emotional learning. *Child Dev.* 82, 405–432. doi: 10.1111/j.1467-8624.2010.01564.x
- Elliott, C. D. (2012). “The differential ability scales,” in *Contemporary Intellectual Assessment: Theories, Tests, and Issues. 2nd Edn.* Chapter 13, eds D. P. Flanagan, and P. L. Harrison (New York, NY: Guilford Press), 336–356.
- Fasola, J., and Mataric, M. J. (2012). Using socially assistive human-robot interaction to motivate physical exercise for older adults. *Proc. IEEE* 100, 2512–2526. doi: 10.1109/JPROC.2012.2200539
- Feil-Seifer, D., and Mataric, M. (2012). “A simon-says robot providing autonomous imitation feedback using graded cueing,” in *Poster Paper in International Meeting for Autism Research (IMFAR)* (Toronto, ON).
- Feil-Seifer, D., and Mataric, M. J. (2005). “Defining socially assistive robotics,” in *9th International Conference on Rehabilitation Robotics, 2005. ICORR 2005* (Chicago, IL: IEEE), 465–468.
- Feil-Seifer, D., and Mataric, M. J. (2011). Socially assistive robotics. *IEEE Robot. Autom. Mag.* 18, 24–31. doi: 10.1109/MRA.2010.940150
- Finn, B., and Metcalfe, J. (2010). Scaffolding feedback to maximize long-term error correction. *Mem. Cogn.* 38, 951–961. doi: 10.3758/MC.38.7.951
- Fischer, G. (2001). User modeling in human-computer interaction. *User Model. User Adapt. Interact.* 11, 65–86. doi: 10.1023/A:1011145532042
- Fong, T., Nourbakhsh, I., and Dautenhahn, K. (2003). A survey of socially interactive robots. *Robot. Auton. Syst.* 42, 143–166. doi: 10.1016/S0921-8890(02)00372-X
- Forceville, C., and Urios-Aparisi, E. (2009). *Multimodal Metaphor, Vol. 11*. Berlin: Walter de Gruyter.
- Fridin, M., and Belokopytov, M. (2014). Embodied robot versus virtual agent: involvement of preschool children in motor task performance. *Int. J. Hum. Comput. Interact.* 30, 459–469. doi: 10.1080/10447318.2014.888500
- Gallese, V., and Goldman, A. (1998). Mirror neurons and the simulation theory of mind-reading. *Trends Cogn. Sci.* 2, 493–501.
- Gazzola, V., Rizzolatti, G., Wicker, B., and Keysers, C. (2007). The anthropomorphic brain: the mirror neuron system responds to human and robotic actions. *Neuroimage* 35, 1674–1684. doi: 10.1016/j.neuroimage.2007.02.003
- Goldin-Meadow, S., and Beilock, S. L. (2010). Action’s influence on thought: the case of gesture. *Perspect. Psychol. Sci.* 5, 664–674. doi: 10.1177/1745691610388764
- Gordon, G., and Breazeal, C. (2015). “Bayesian active learning-based robot tutor for children’s word-reading skills,” in *AAAI* (Portland, OR), 1343–1349.
- Gordon, G., Spaulding, S., Westlund, J. K., Lee, J. J., Plummer, L., Martinez, M., et al. (2016). “Affective personalization of a social robot tutor for children’s second language skills,” in *AAAI* (Phoenix, AR), 3951–3957.
- Gotham, K., Risi, S., Dawson, G., Tager-Flusberg, H., Joseph, R., Carter, A., et al. (2008). A replication of the autism diagnostic observation schedule

- (ADOS) revised algorithms. *J. Am. Acad. Child Adolesc. Psychiatry* 47, 642–651. doi: 10.1097/CHI.0b013e31816bfb7
- Greczek, J., Atrash, A., and Matarić, M. (2013). “A computational model of graded cueing: robots encouraging behavior change,” in *International Conference on Human-Computer Interaction* (Las Vegas, NV: Springer), 582–586.
- Greczek, J., Kaszubski, E., Atrash, A., and Matarić, M. (2014). “Graded cueing feedback in robot-mediated imitation practice for children with autism spectrum disorders,” in *2014 RO-MAN: The 23rd IEEE International Symposium on Robot and Human Interactive Communication* (Edinburgh: IEEE), 561–566.
- Guadagnoli, M. A., and Lee, T. D. (2004). Challenge point: a framework for conceptualizing the effects of various practice conditions in motor learning. *J. Motor Behav.* 36, 212–224. doi: 10.3200/JMBR.36.2.212-224
- Guadalupe, F. (2016). Social skills training for autistic children: a comparison study between inclusion and mainstreaming education. *Acad. J. Pediatr. Neonatol.* 1:555. doi: 10.19080/AJPN.2016.01.555564
- Hancock, P. A., Billings, D. R., Schaefer, K. E., Chen, J. Y., De Visser, E. J., and Parasuraman, R. (2011). A meta-analysis of factors affecting trust in human-robot interaction. *Hum. Factors* 53, 517–527. doi: 10.1177/0018720811417254
- Horn, R. E. (1998). *Visual Language*. Washington, DC: MacroVu Inc.
- Horner, R. H., Carr, E. G., Halle, J., McGee, G., Odom, S., and Wolery, M. (2005). The use of single-subject research to identify evidence-based practice in special education. *Except. Children* 71, 165–179. doi: 10.1177/001440290507100203
- Kanne, S. M., Randolph, J. K., and Farmer, J. E. (2008). Diagnostic and assessment findings: a bridge to academic planning for children with autism spectrum disorders. *Neuropsychol. Rev.* 18, 367–384. doi: 10.1007/s11065-008-9072-z
- Kasari, C., Rotheram-Fuller, E., Locke, J., and Gulsrud, A. (2012). Making the connection: randomized controlled trial of social skills at school for children with autism spectrum disorders. *J. Child Psychol. Psychiatry* 53, 431–439. doi: 10.1111/j.1469-7610.2011.02493.x
- Khamassi, M., V. G. T. T., and Tzafestas, C. (2018). Robot fast adaptation to changes in human engagement during simulated dynamic social interaction with active exploration in parameterized reinforcement learning. *IEEE Trans. Cogn. Dev. Syst.* 10, 881–893. doi: 10.1109/TCDS.2018.2843122
- Kidd, C. D. (2008). *Designing for long-term human-robot interaction and application to weight loss* (Ph.D. thesis), AAI0819995, Cambridge, MA, United States.
- Kidd, C. D., and Breazeal, C. (2008). “Robots at home: understanding long-term human-robot interaction,” in *International Conference on Intelligent Robots and Systems (IROS)* (Nice: IEEE), 3230–3235.
- Kim, T., and Hinds, P. (2006). “Who should I blame? Effects of autonomy and transparency on attributions in human-robot interaction,” in *The 15th IEEE International Symposium on Robot and Human Interactive Communication, 2006. ROMAN 2006* (Munich: IEEE), 80–85.
- Kohn, R., Saxena, S., Levav, I., and Saraceno, B. (2004). The treatment gap in mental health care. *Bull. World Health Organ.* 82, 858–866. doi: 10.1590/S0042-96862004001100011
- Kulik, J. A., and Kulik, C.-L. C. (1988). Timing of feedback and verbal learning. *Rev. Educ. Res.* 58, 79–97.
- Kulkarni, T. D., Narasimhan, K., Saeedi, A., and Tenenbaum, J. (2016). “Hierarchical deep reinforcement learning: integrating temporal abstraction and intrinsic motivation,” in *Advances in Neural Information Processing Systems* (Barcelona), 3675–3683.
- Lavelle, T. A., Weinstein, M. C., Newhouse, J. P., Munir, K., Kuhlthau, K. A., and Prosser, L. A. (2014). Economic burden of childhood autism spectrum disorders. *Pediatrics* 133, e520–e529. doi: 10.1542/peds.2013-0763
- Lee, K. M., Jung, Y., Kim, J., and Kim, S. R. (2006). Are physically embodied social agents better than disembodied social agents? The effects of physical embodiment, tactile interaction, and people’s loneliness in human-robot interaction. *Int. J. Hum. Comput. Stud.* 64, 962–973. doi: 10.1016/j.ijhcs.2006.05.002
- Lee, K. M., Park, N., and Song, H. (2005). Can a robot be perceived as a developing creature? *Hum. Commun. Res.* 31, 538–563. doi: 10.1111/j.1468-2958.2005.tb00882.x
- Leite, I., McCoy, M., Ullman, D., Salomons, N., and Scassellati, B. (2015). “Comparing models of disengagement in individual and group interactions,” in *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction* (Portland, OR: ACM), 99–105.
- Leite, I., Pereira, A., Mascarenhas, S., Martinho, C., Prada, R., and Paiva, A. (2013). The influence of empathy in human-robot relations. *Int. J. Hum. Comput. Stud.* 71, 250–260. doi: 10.1016/j.ijhcs.2012.09.005
- Leyzberg, D., Spaulding, S., and Scassellati, B. (2014b). “Personalizing robot tutors to individuals’ learning differences,” in *Proceedings of the 2014 ACM/IEEE International Conference on Human-Robot Interaction* (New York, NY: ACM), 423–430.
- Leyzberg, D., Spaulding, S., Toneva, M., and Scassellati, B. (2012). “The physical presence of a robot tutor increases cognitive learning gains,” in *Proceedings of the Cognitive Science Society, Vol. 34* (Sapporo).
- Littman, M. L. (1994). “Markov games as a framework for multi-agent reinforcement learning,” in *Machine Learning Proceedings 1994* (New Brunswick, NJ: Elsevier), 157–163.
- Lobo, M., Moeyaert, M., Baraldi, C. A., and Babik, I. (2017). Single-case design, analysis, and quality assessment for intervention research. *J. Neurol. Phys. Ther.* 31:187. doi: 10.1097/NPT.0000000000000187
- Lord, C., Risi, S., Lambrecht, L., Cook, E. H., Leventhal, B. L., DiLavore, P. C., et al. (2000). The autism diagnostic observation schedule-generic: A standard measure of social and communication deficits associated with the spectrum of autism. *J. Autism Dev. Disord.* 30, 205–223. doi: 10.1023/A:1005592401947
- Matarić, M. J. (2017). Socially assistive robotics: human augmentation versus automation. *Sci. Robot.* 2:eaam5410. doi: 10.1126/scirobotics.aam5410
- Matarić, M. J., and Scassellati, B. (2016). *Socially Assistive Robotics*, Chapter 73. Cham: Springer International Publishing, 1973–1994.
- McLellan, H. (1996). *Situated Learning Perspectives*. Englewood, NJ: Educational Technology.
- Mehrabian, A. (2017). *Nonverbal Communication*. New York, NY: Routledge.
- Moon, J.-W., and Kim, Y.-G. (2001). Extending the tam for a world-wide-web context. *Inform. Manage.* 38, 217–230. doi: 10.1016/S0378-7206(00)00061-6
- Murray, T. (1999). Authoring intelligent tutoring systems: an analysis of the state of the art. *Int. J. Artif. Intell. Educ.* 10, 98–129.
- Niedenthal, P. M. (2007). Embodying emotion. *science* 316, 1002–1005. doi: 10.1126/science.1136930
- Nikolaidis, S., Kuznetsov, A., Hsu, D., and Srinivasa, S. (2016). “Formalizing human-robot mutual adaptation: a bounded memory model,” in *The Eleventh ACM/IEEE International Conference on Human Robot Interaction* (New York, NY: IEEE Press), 75–82.
- Omrod, J., Anderman, E., and L.H., A. (2017). *Educational Psychology, Developing Learners* (Upper Saddle River, NJ: Merrill Prentice Hall), 4.
- O’Rourke, E., Haimovitz, K., Ballweber, C., Dweck, C., and Popović, Z. (2014). “Brain points: a growth mindset incentive structure boosts persistence in an educational game,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, ON: ACM), 3339–3348.
- Ospina, M. B., Seida, J. K., Clark, B., Karkhaneh, M., Hartling, L., Tjosvold, L., et al. (2008). Behavioural and developmental interventions for autism spectrum disorder: a clinical systematic review. *PLoS ONE* 3:e3755. doi: 10.1371/journal.pone.0003755
- Ovando, M. N. (1994). Constructive feedback: a key to successful teaching and learning. *Int. J. Educ. Manage.* 8, 19–22.
- Park, H. W., Rosenberg-Kima, R., Rosenberg, M., Gordon, G., and Breazeal, C. (2017). “Growing growth mindset with a social robot peer,” in *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction* (Vienna: ACM), 137–145.
- Ramachandran, A., Huang, C.-M., Gärtland, E., Scassellati, B., and Scas, B. (2018). “Thinking aloud with a tutoring robot to enhance learning,” in *International Conference on Human-Robot Interaction (HRI)* (Chicago, IL: ACM), 59–68.
- Rau, P. P., Li, Y., and Li, D. (2009). Effects of communication style and culture on ability to accept recommendations from robots. *Comput. Hum. Behav.* 25, 587–595. doi: 10.1016/j.chb.2008.12.025
- Reardon, C., Zhang, H., Wright, R., and Parker, L. E. (2015). “Response prompting for intelligent robot instruction of students with intellectual disabilities,” in *International Conference on Robot and Human Interactive Communication (RO-MAN)* (Kobe), 784–790.

- Riek, L. D. (2012). Wizard of Oz studies in HRI: a systematic review and new reporting guidelines. *J. Hum. Robot Interact.* 1, 119–136. doi: 10.5898/JHRI.1.1.Riek
- Robins, B., Dautenhahn, K., Te Boekhorst, R., and Billard, A. (2005). Robotic assistants in therapy and education of children with autism: can a small humanoid robot help encourage social interaction skills? *Univ. Access Inform. Soc.* 4, 105–120. doi: 10.1007/s10209-005-0116-3
- Roid, G. H., Miller, L. J., Pomplun, M., and Koch, C. (2013). *Leiter International Performance Scale (Leiter-3)*. Los Angeles, CA: Western Psychological Services.
- Roll, I., Aleven, V., McLaren, B. M., and Koedinger, K. R. (2011). Improving students' help-seeking skills using metacognitive feedback in an intelligent tutoring system. *Learn. Instr.* 21, 267–280. doi: 10.1016/j.learninstruc.2010.07.004
- Ros, R., Nalin, M., Wood, R., Baxter, P., Looije, R., Demiris, Y., et al. (2011). "Child-robot interaction in the wild: advice to the aspiring experimenter," in *Proceedings of the 13th International Conference on Multimodal Interfaces (Alicante: ACM)*, 335–342.
- Rudovic, O., Lee, J., Dai, M., Schuller, B., and Picard, R. W. (2018). Personalized machine learning for robot perception of affect and engagement in autism therapy. *Sci. Robot.* 3:eaa06760. doi: 10.1126/scirobotics.aa06760
- Sabanovic, S., Michalowski, M. P., and Simmons, R. (2006). "Robots in the wild: observing human-robot social interaction outside the lab," in *9th IEEE International Workshop on Advanced Motion Control, 2006 (Istanbul: IEEE)*, 596–601.
- Sansone, C., and Harackiewicz, J. M. (2000). *Intrinsic and Extrinsic Motivation: The Search for Optimal Motivation and Performance*. Woodbine, NJ: Elsevier.
- Scassellati, B., Admoni, H., and Mataric, M. (2012). Robots for use in autism research. *Annu. Rev. Biomed. Eng.* 14, 275–294. doi: 10.1146/annurev-bioeng-071811-150036
- Scassellati, B., Boccanfuso, L., Huang, C.-M., Mademtz, M., Qin, M., Salomons, N., et al. (2018). Improving social skills in children with ASD using a long-term, in-home social robot. *Sci. Robot.* 3:eaat7544. doi: 10.1126/scirobotics.aat7544
- Schiller, E. (1996). Educating children with attention deficit disorder. *Our Children* 22, 32–33.
- Schodde, T., Bergmann, K., and Kopp, S. (2017). "Adaptive robot language tutoring based on Bayesian knowledge tracing and predictive decision-making," in *International Conference on Human-Robot Interaction (HRI) (Vienna: ACM)*, 128–136.
- Searle, J. R. (1976). A classification of illocutionary acts. *Lang. Soc.* 5, 1–23.
- Searle, J. R., and Searle, J. R. (1969). *Speech Acts: An Essay in the Philosophy of Language, Vol. 626*. Cambridge, UK: Cambridge University Press.
- Short, Y. F. E., Short, D., and Mataric, M. J. (2017). *Sprite: Stewart Platform Robot for Interactive Tabletop Engagement*. Department of Computer Science, University of Southern California, Tech Report.
- Stokoe, W. C., Casterline, D. C., and Croneberg, C. G. (1976). *A Dictionary of American Sign Language on Linguistic Principles*. Washington, DC: Linstok Press.
- Swan, M. (2012). Health 2050: the realization of personalized medicine through crowdsourcing, the quantified self, and the participatory biocitizen. *J. Pers. Med.* 2, 93–118. doi: 10.3390/jpm2030093
- Sweller, J., Kirschner, P., and Clark, R. (2007). Why minimally guided teaching techniques do not work: a reply to commentaries. *Educ. Psychol.* 42, 115–121. doi: 10.1080/00461520701263426
- Tadevosyan-Leyfer, O., Dowd, M., Mankoski, R., Winklosky, B., Putnam, S., McGrath, L., et al. (2003). A principal components analysis of the autism diagnostic interview-revised. *J. Am. Acad. Child Adolesc. Psychiatry* 42, 864–872. doi: 10.1097/01.CHI.0000046870.56865.90
- Toh, L. P. E., Causo, A., Tzuo, P.-W., Chen, I., Yeo, S. H. (2016). A review on the use of robots in education and young children. *J. Educ. Technol. Soc.* 19:148.
- Vallerand, R. J., Pelletier, L. G., Blais, M. R., Briere, N. M., Senecal, C., and Vallieres, E. F. (1992). The academic motivation scale: a measure of intrinsic, extrinsic, and amotivation in education. *Educ. Psychol. Meas.* 52, 1003–1017.
- Van Bourgondien, M. E., Marcus, L. M., and Schopler, E. (1992). Comparison of dsm-iii-r and childhood autism rating scale diagnoses of autism. *J. Autism Dev. Disord.* 22, 493–506.
- van De Sande, B. (2013). Properties of the bayesian knowledge tracing model. *J. Educ. Data Mining* 5:1.
- Velentzas, G., Tsitsimis, T., Rañó, I., Tzafestas, C., and Khamassi, M. (2018). Adaptive reinforcement learning with active state-specific exploration for engagement maximization during simulated child-robot interaction. *Paladyn J. Behav. Robot.* 9, 235–253. doi: 10.1515/pjbr-2018-0016
- Vygotsky, L. (1978). Interaction between learning and development. *Read. Dev. Children* 23, 34–41.
- Wainer, J., Feil-Seifer, D. J., Shell, D. A., and Mataric, M. J. (2006). "The role of physical embodiment in human-robot interaction," in *The 15th IEEE International Symposium on Robot and Human Interactive Communication, 2006. ROMAN 2006 (Hatfield: IEEE)*, 117–122.
- Wainer, J., Feil-Seifer, D. J., Shell, D. A., and Mataric, M. J. (2007). "Embodiment and human-robot interaction: a task-based perspective," in *The 16th IEEE International Symposium on Robot and Human Interactive Communication, 2007. RO-MAN 2007 (Jeju: IEEE)*, 872–877.
- Watkins, C. J., and Dayan, P. (1992). Q-learning. *Mach. Learn.* 8, 279–292.
- Wechsler, D. (2005). *Wechsler Individual Achievement Test. 2nd Edn. (WIAT II)*. London: The Psychological Corporation.
- Wenger, E. (2014). *Artificial Intelligence and Tutoring Systems: Computational and Cognitive Approaches to the Communication of Knowledge*. San Francisco, CA: Morgan Kaufmann.
- White, S. W., Keonig, K., and Scahill, L. (2007). Social skills development in children with autism spectrum disorders: a review of the intervention research. *J. Autism Dev. Disord.* 37, 1858–1868. doi: 10.1007/s10803-006-0320-x
- Wing, L., Leekam, S. R., Libby, S. J., Gould, J., and Larcombe, M. (2002). The diagnostic interview for social and communication disorders: background, inter-rater reliability and clinical use. *J. Child Psychol. Psychiatry* 43, 307–325. doi: 10.1111/1469-7610.00023
- Yagoda, R. E., and Gillan, D. J. (2012). You want me to trust a robot? The development of a human-robot interaction trust scale. *Int. J. Soc. Robot.* 4, 235–248. doi: 10.1007/s12369-012-0144-0
- Zins, J. E., Weissberg, R. P., Wang, M. C., and Walberg, H. (2004). *Building Academic Success on Social and Emotional Learning: What Does the Research Say?* New York, NY: Teachers College Press.

Conflict of Interest: MM was co-founder of Embodied, Inc., but was no longer involved with the company. CC was a full-time employee of Embodied, Inc., but was not involved with the company while the reported work was done.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2019 Clabaugh, Mahajan, Jain, Pakkar, Becerra, Shi, Deng, Lee, Ragusa and Mataric. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.