

The theoretical and methodological opportunities afforded by guided play with young children

Yue Yu<sup>1</sup>, Patrick Shafto<sup>1</sup>, Elizabeth Bonawitz<sup>1</sup>, Scott Cheng-Hsin Yang<sup>1</sup>, Roberta M. Golinkoff<sup>2</sup>,

Kathleen H. Corriveau<sup>3</sup>, Kathy Hirsh-Pasek<sup>4</sup>, & Fei Xu<sup>5</sup>

<sup>1</sup>Rutgers University – Newark

<sup>2</sup>University of Delaware

<sup>3</sup>Temple University

<sup>4</sup>Boston University

<sup>5</sup>University of California – Berkeley

#### Acknowledgements

This study is supported by NSF Award SMA-1640816 to P.S., E.B., R.G., K.H.C., K.H-P, & F.X.

Correspondence concerning this article should be addressed to Dr. Yue Yu, Department of Psychology, Rutgers University-Newark, NJ 07102. Email: [yue.yu@rutgers.edu](mailto:yue.yu@rutgers.edu).

Word count: 3937

## Abstract

For infants and young children, learning takes place all the time and everywhere. How children learn best both in and out of school has been a long-standing topic of debate in education, cognitive development, and cognitive science. Recently, *guided play* has been proposed as an integrative approach for thinking about learning as a child-led, adult-assisted playful activity. The interactive and dynamic nature of guided play presents theoretical and methodological challenges and opportunities. Drawing upon research from multiple disciplines, we discuss the integration of cutting-edge computational modeling and data science tools to address some of these challenges, and highlight avenues towards an empirically grounded, computationally precise and ecologically valid framework of guided play in early education.

Keywords: guided play; computational modeling; data science; direct instruction; free play

The theoretical and methodological opportunities afforded by guided play with young children

Learning in school is often characterized by structured courses and tasks with discrete and explicit objectives. Yet, learning is a continuous process that also takes place outside the classroom where explicit objectives are not always evident. This is especially true in early childhood interactions at home, where children often learn from everyday interactions with both the physical environment and with social partners (Bruner, 1961; Csibra & Gergely, 2009). How to best navigate between explicit, objective-directed learning and more flexibly driven exploration has been a longstanding topic of debate in education, developmental psychology, and cognitive science (Kirschner, Sweller, & Clark, 2006; Tobias & Duffy, 2009). This debate surfaces in a number of forms, as direct instruction vs. discovery learning or as work vs. play (Bonawitz, Shafto, et al., 2011; Clements & Sarama, 2014; Hirsh-Pasek & Golinkoff, 2011). Pitting these two interests against each other has neither optimized our understanding of learning, nor produced optimal methods of learning (Wise & O'Neill, 2009). Here, we discuss an integrated approach, *guided play*, that enables us to rethink learning as a child-led, adult-assisted activity (Weisberg, Hirsh-Pasek & Golinkoff, 2013; Weisberg, Hirsh-Pasek, Golinkoff & McCandliss, 2014; Weisberg, et al., 2016). Focusing on everyday interactions in early childhood, guided play is operationally defined as learning that is active and engaged, where the child takes initiative in a playful learning environment and the adult supports, rather than directs, the learning experience. Sitting between free play, where children explore by themselves, and direct instruction, where the interaction is led by an adult and children take a passive role, guided play takes advantage of the latest research in the science of learning.

Educational research indicates that student-led discovery learning that is facilitated by teachers outperforms both direct instruction and unassisted discovery (Honomichl & Chen, 2012;

Mayer, 2004). In a meta-analysis comparing explicit instruction, unassisted discovery, and assisted discovery (Alfieri, Brooks, Aldrich, & Tenenbaum, 2011), learning outcomes were more favorable for assisted discovery than for other forms of instruction. These results held for learners of different ages and across different learning domains. Similarly, developmental studies have shown an advantage of adult guidance over both direct instruction and free play, even before children start formal schooling (Fisher, Hirsh-Pasek, Newcombe, & Golinkoff, 2013; Haden, Cohen, Uttal, & Marcus, 2016; Han, Moore, Vukelich & Buell, 2010; Ridge, Weisberg, Ilgaz, Hirsh-Pasek, Golinkoff, 2015; Sim & Xu, 2017; Yu, Landrum, Bonawitz, & Shafto, under review). In both bodies of literature, “guidance” has referred to a variety of practices including modeling, questioning, encouragement, and feedback, and thus it is unclear what particular aspects of guidance are associated with learning (Honomichl & Chen, 2012; Wise & O'Neill, 2009).

In *guided play*, learning opportunities may be explicitly structured, but importantly the activity is child-led. Specifically, we define “*guidance*” as adults’ involvement that subtly channels the dyadic interactions to fulfill certain pedagogical objectives, while not interfering too much so that the activities remain child-led. The pedagogical objectives can be multi-level: They can focus on specific content knowledge, but can also focus on the emotional, motivational, and metacognitive aspects of the learning process, such as cultivating children’s love of learning, promoting their engagement, or making them aware of their own learning process (Weisberg, et al., 2014). Our concept of guidance is inspired by the Vygotskian concept of scaffolding (Vygotsky, 1934/1978; Wood, Bruner, & Ross, 1976; Fernández, Wegerif, Mercer, & Rojas-Drummond, 2001) and Barbara Rogoff’s theory of guided participation (Rogoff, et al., 1993). In addition to guidance being tailored to fit individual children’s needs and skill level (which is

similar to scaffolding), in guided play we also emphasize that guidance should never shift children away from controlling their own learning process. The pedagogical objectives of guidance are therefore broader—besides helping children to master particular knowledge or skills, guided play also aims to provide children with an opportunity to enjoy, control, and reflect upon their own learning process, which may facilitate independent inquiry and discovery in the future.

Because guided play requires seamless integration between the adult's objectives to support learning and child-led activity that can be highly fluid, characterizing appropriate guidance requires an understanding of the dynamic nature of an adult-child interaction in context. First, guided play is *interactive*. How well children can learn from a playful interaction depends on their mental state (Putnam, 1980) at the moment—including their level of knowledge, goal, attention, emotion, trust towards the play partner, etc. Therefore, effective guidance should take into account and be contingent upon the mental state of the child. This requires theories to consider the dyad as a system moving toward a joint objective (Fogel & Garvey, 2007; Heller & Rohlfing, 2017; Lavelli et al., 2015), and requires experimental designs and analytical tools that go beyond between-group comparisons to focus on individual dyads. Second, guided play is *dynamic*. Timing is critical for the guidance to be effective. Providing a label, for example, can be educational at a moment when a child is focusing on the target object, but can be confusing when the child is focusing on multiple objects (Pereira, Smith, & Yu, 2014). Similarly, demonstrating object functions when an infant is pointing to the object also supports learning (Begus, Gliga, & Southgate, 2014). For preschoolers, revealing causal features of objects right before, but not after, a demonstration of categorization facilitates children's category learning (Yu, & Kushnir, 2016). Existing theories, such as direct instruction and free

play, and methodological tools, such as standard statistical tests, are optimized for discrete interventions and are usually applied uniformly across groups of individuals. Characterizing the dynamic nature of guided play will require development of new theories and tools to capture interventions along a continuous timeline. In what follows, we detail these theoretical and methodological matters, the tools that may be used to address them, and the prospects for a theory of guided play.

### **Theoretical challenges and opportunities for guided play**

Free play and direct instruction have long been contrasted in education and cognitive development (Dewey, 1933; Hirsh-Pasek, Golinkoff, Berk, & Singer, 2008; Kirschner, Sweller, & Clark, 2006; Mayer, 2004), and existing mathematical and computational models for the two scenarios have likewise been developed separately because they typically focus on different aspects of learning (Nelson, 2005; Shafto, Goodman, & Griffiths, 2014). Free play is based on the constructivist views of learning, which portrays learning as an active process during which the learner repeatedly intervenes on their environment, and update their beliefs based on information gathered from these experiences (Piaget, 1952). Correspondingly, computational models of free play have largely focused on how to sequentially choose evidence during learning (Markant & Gureckis, 2014; Nelson, 2005; Settles, 2010; McCormack, et al., 2016). These models generate predictions about how the optimal next step will depend on the current state and are therefore *dynamic*. However, such models are inadequate to capture the *interactive* aspect of guided play because they do not usually simulate a social partner whose behavior is contingent on the learner.

In contrast, direct instruction emphasizes the necessity of outside instructions for learners to successfully navigate a learning task (Kirschner, Sweller, & Clark, 2006), and focuses on what

content should be delivered by instruction (Mayer, 2004). Correspondingly, computational models of direct instruction have focused on the evidence teachers should select to lead learners to the correct answer, given the learner's current beliefs (Rafferty, Brunskill, Griffiths, & Shafto, 2015; Shafto & Goodman, 2008; Shafto, Goodman, & Griffiths, 2014; Shafto, Goodman, Frank, 2012; Zhu, 2015; Frank, 2014). Some of these models simulate the *interactive* nature of teaching and learning through modeling the teacher and the learner's reasoning about the other's knowledge levels and objectives (Shafto & Goodman, 2008; Shafto, et al., 2012, 2014). However, these models are not *dynamic*; they select evidence with the immediate goal of the learner arriving at the correct inference. When dynamic extensions have been proposed, they encounter significant computational challenges that render the models of limited use for modeling real-life scenarios (Rafferty, et al., 2015; Yang & Shafto, 2017).

Theories and models of epistemic trust may inform modeling of dynamic interactions between a teacher and a learner. The literature on epistemic trust has investigated the dynamics of reasoning, focusing on a learner's sensitivity to both a teacher's prior knowledge in a given domain (Pasquini, Corriveau, Koenig, & Harris, 2007; Sobel & Corriveau, 2010) as well as her social group membership when making decisions about whom to trust (Chen, Corriveau, & Harris, 2013; Kinzler, Corriveau, & Harris, 2011). Models of epistemic trust (Shafto, Eaves, Navarro, & Perfors, 2012; Eaves & Shafto, 2012; 2017) tend to build upon aforementioned models of direct instruction. Although both of these bodies of work make the prediction that children's epistemic and social evaluation of a teacher should influence their trust in her (and therefore, their sensitivity to her guidance), to date, both the experimental and computational work has focused on the dynamics of trust, but not learning.

Finally, ecological psychology and dynamic systems approaches have been applied to analyze dynamic interactions between adults and children (Bronfenbrenner, 1986; Fogel & Garvey, 2007; Thelen & Smith, 1996). These approaches were foundational in emphasizing the need to view adult-child interactions as a system that evolves through time, as well as the need to situate these interactions in the immediate environment. They also provided invaluable computational tools to analyze patterns of co-activities that emerges along time. Because formal dynamic systems models often focus on overt behavior, applying these models to guided play may require an extension which takes into account the mental state and inferential capacities of both learners and guiding adults.

A unified theory of guided play must combine strengths from previous research to capture the interactive and dynamic nature of learning. A key challenge for proposing such a theory is the development of theoretical frameworks that avoid simulating every possible mental state of the teacher and the learner, which would create intractable computational problems. Even the simplest learning situations involve many potential choices by both learners and guiding adults. For example, when an adult guides a child to learn the name of an object, the adult could choose from a variety of actions (e.g., pointing to the object, holding it, looking at the child, or looking at the object) as well as utterances (e.g., naming the object, or asking a question), and the child could also respond in a variety of ways (e.g., reaching for the object, repeating the word, or displaying a puzzled face). Adults and children nevertheless navigate such situations, making choices while balancing short- and long-term objectives. To simulate these capacities, one approach is to adopt simplified computational models similar to those employed in the educational technology literature. One example is Bayesian knowledge tracing, which instead of modeling the learner's full belief state, focuses on whether the learner has the correct

concepts (Corbett, Anderson, 1995; Yudelson, Koedinger, & Gordon, 2013). A second approach is to use task-specific information to limit the set of relevant actions. For example, an approach that pairs observation of naturalistic adult-child interaction during a task with an experiment that measures the learning outcome of that task could help to identify the task-relevant subset of information (Yu, Bonawitz, & Shafto, 2017). Subsequent experimental studies could then test predictions of the model on this reduced set of relevant information rather than the whole set of logical possibilities.

### **Methodological challenges and opportunities for guided play**

The interactive and dynamic properties of guided play also pose questions for experimental design and analysis that may require modifications of existing tools and the development of new ones. One source of methodological challenges arises from variations in the effectiveness of guidance based on individual characteristics of the child. Guidance *content* that is effective for one child may not be effective for a different child. For example, two children may have different misconceptions about what constitutes a triangle (Fisher et al., 2013): One may think a triangle needs to have the point at the top, whereas the other may think a triangle needs to have all acute angles. In this case, different examples should be presented to guide these two children away from their respective misconceptions: It would be more effective to show the first child a real triangle with point in the bottom, and show the second child an obtuse triangle. This intuition is supported by research: Research in category learning has shown that a set of evidence that is effective in facilitating one person's learning may be less effective when presented to another person (Markant & Gureckis, 2014; Sim, Tanner, Alpert, & Xu, 2015). In addition, individual differences in children's background knowledge, cognitive style, and experiences with different sociocultural practices can all influence the effectiveness of presenting

certain content to them (Gutiérrez & Rogoff, 2003; McNamara, Kintsch, E., Butler-Songer, & Kintsch, W. 1996; Price, 2004). Individual differences remain an important topic for further research.

The *timing* of guidance is also important: Well-timed guidance that is contingent upon the child's prior actions may impact child learning outcomes differently than if the same guidance is not well-timed (Pereira, Smith, & Yu, 2014). Such variability in guidance content and timing poses challenges to typical random-assignment controlled experiments, as uniform interventions applied to groups of randomly assigned individuals do not necessarily test the interactive and dynamic predictions of guided play. Yet observational designs are insufficient to tease apart the causal relations between components of guided play and children's learning outcomes. Therefore, new methods and analytical tools are required to select the content and timing of guidance to maximally inform our understanding of the mechanisms involved in guided play.

Advances in data science and technology may provide tools for addressing some of these challenges by providing an opportunity for real-time analysis and feedback, as well as (semi-)automatic analysis of large amounts of time series data. For example, in word-learning scenarios, children look at the experimenter more when they are uncertain about an object label (Hembacher, deMayo, & Frank, 2017). Thus, an overt behavior, here eye gaze, reveals important information about the learner's mental state, and could represent opportunities for guidance. Technological advances in eye-tracking equipment and data sharing mechanisms have allowed for the collection and sharing of large-scale, live-stream video data from naturalistic adult-child interactions (Databrary, 2012; Franchak, Kretch, Soska, & Adolph, 2011). However, coding and analysis of children's looks are usually conducted manually, which restricts the amount of data

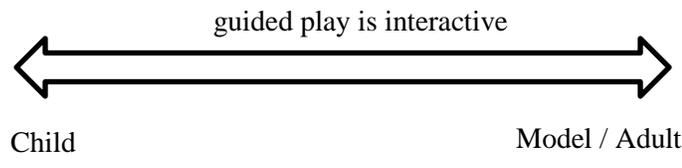
that can be utilized and precludes real-time feedback during the interaction. Applying tools of automatic decoding of eye movements and looking, such as those used in vision research (Borji & Itti, 2014; Duc, Bays, Husain, 2008; Gottlieb, Oudeyer, Lopes, & Baranes, 2013), may allow for the online recognition of the referent associated with the child's gaze, which, in turn, may help to nominate a range of appropriate guidance "moves" that are contingent upon the child's attention and mental state. Indeed, research in social robotics has implemented gaze and action detection in robot *learners* to infer human teachers' pedagogical intent based on their gaze and actions, and to react in a contingent way (e.g., when the teacher showed an object with pedagogical cues, the robot turned head to the same object; then when the teacher looked back at the robot's eyes and labeled the object, the robot looked at the teacher and smiled). Human teachers were more engaged and more likely to attribute human-like traits to the robot when the robot displayed these contingent reactions (Lohan, et al., 2012). Similar algorithms may also support *teachers* who provide guidance contingent on the learner's behavior.

Similarly, the learner's affect and engagement plays an important role (Rader & Hughes, 2005; Greene & Noice, 1988). In guided play, the joy that accompanies play helps to sustain motivation, interest, and excitement, which should be associated with enhanced learning outcomes (Hirsh-Pasek & Golinkoff, 2003; Weisberg, et al., 2016). Unfortunately, given the time-intensive nature of affect coding, the evidence relating affective states to improved learning outcomes is less extensive. Data science tools may be used to automatically identify affect and engagement in real-time video streams for analysis, and to time guidance to foster affect that predict positive short term and long term learning (Baker, LoBue, Bonawitz, & Shafto, 2017; Littlewort, et al., 2006; Yao, Shao, Ma, & Chen, 2015). Such analytical tools would allow for direct tests of guided play predictions related to the timing of learning, while employing

experimental designs that are similar to those typically used in the developmental and educational literature.

### **Coupling computational models and data science tools**

A more ambitious possibility is to couple models and data science tools to create experiments highlighting times when interventions may yield the strongest test of the theory. Attempts at interactive, dynamic approaches to teaching can be found in the literature of social robotics and intelligent tutoring systems (Anderson, Boyle, & Reiser, 1985; Breazeal, 2002; Clement, Roy, Oudeyer, & Lopes, 2015; Lohan, et al., 2012; Nguyen & Oudeyer, 2014; Thomaz & Breazeal, 2008; Vollmer, et al., 2014), in which data from expert teachers have been used to train algorithms to learn the contingencies between learner's behavior and teachers' appropriate response (Ruvolo, Whitehill, Virnes, & Movellan, 2008). Such data-driven approaches can serve as a first step for identifying patterns in guided-play interactions. However, to understand characteristics of effective guidance, we also need theory-driven computational models that can represent children's mental states based on their behavior. Such models differ from existing intelligent tutoring systems in that instead of teaching knowledge in specific domains, they are designed to understand the general principles of effective guidance in a wide range of child-led activities that may or may not have an explicit learning goal. Coupling such models with empirical data could inform an algorithm that predicts appropriate guidance based on children's behavior, which could in turn be used in experiments to verify the effect of guidance on children's learning. These experiments would have significant advantages relative to classic training studies, as the intervention is based on an online algorithm which would adapt based on children's moment-by-moment behavior.



model trained with existing guided play interactions and theories to identify relevant behavior and hypotheses for a specific task  
*[setting task-relevant behavior to child's and adult's look, point, facial expression and language, setting hypotheses to "point at the top", "all acute angles", "closed shape with three edges"]*

displaying task-relevant behavior  
*["Is this a triangle because the point is at the top?"]*



inferring child's mental state based on behavior  
*[assigning high probability on the wrong hypothesis "point at the top"]*



identifying optimal guidance  
*[an up-side-down triangle is most effective to shift away from wrong hypothesis]*



guiding in a way consistent with the principles of guided play  
*["Does this one have a point at the top? Is it a real triangle?"]*



updating belief  
*["Hmm... maybe it's not about where the point is"]*



displaying further task-relevant behavior  
*[Looking doubtfully at an obtuse triangle]*



updating child's mental state based on behavior  
*[assigning uncertainty on the wrong hypothesis "all acute angles"]*

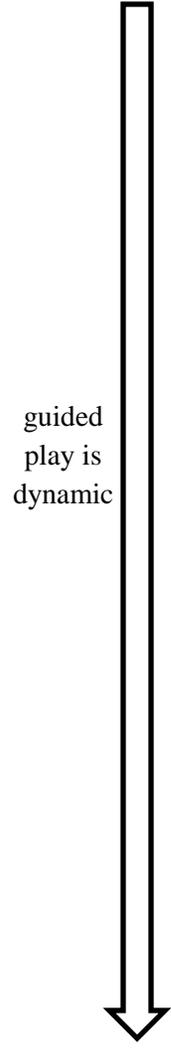


Figure 1. We propose a framework that integrates computational modeling and data science to address challenges brought by the interactive and dynamic nature of guided play. By modeling children's moment-to-moment mental state from their task-relevant behavior, the proposed framework identifies guidance that are optimized in terms of timing and form, with the objective of sustaining the children's interests toward the learning goal. The *italic text* provides an example of learning geometric shapes (Fisher et al., 2013) to show how the framework could be implemented to a specific guided play interaction. This framework can facilitate research of guided play by identifying key aspects of guidance within the dynamic and complex interactions children experience in their everyday environment.

Consider how such computational models could be applied to a recent study of guided play (Fisher et al., 2013). This study examined different pedagogical methods on preschoolers' learning of geometric shapes, with increased learning in guided play as compared to didactic instruction and free play. In the guided play condition, the experimenter presented two typical examples (e.g., upright triangles) and two atypical examples (e.g., inverted triangles) in a playful manner, and asked children to determine what makes them the same shape. During children's active exploration the experimenter used questions, encouragement, and feedback to guide them towards the correct answer. Yet, because the interaction was dynamic, the manner and timing of adult guidance were not prespecified in the experimental design, which makes it difficult to pinpoint what aspects of guidance resulted in the enhanced learning outcomes.

Following the aforementioned framework, existing videos of guided play interactions could be used to train a computational model of learning geometric shapes in four steps: First, data science tools can identify a set of common task-relevant behavior during children's active

exploration, and cluster behaviors into categories (e.g., children's looking and pointing may be categorized as seeking guidance from the experimenter; their emotion as confident vs. doubtful; their language as statements or questions). Tools of this stage could build upon advances in (semi-)automatic recognition of eye gaze (e.g., Lohan, et al., 2012; Smith, Yu, Yoshida, & Fausey, 2015), emotion (e.g., Baker, LoBue, Bonawitz, & Shafto, 2017; Nojavanasghari, Huang, & Khan, 2018), natural language including information-seeking questions (e.g., Rothe, Lake, & Gureckis, 2016), among others.

Second, a computational model can be used to simulate children's moment-to-moment beliefs about geometric shapes based on these behavioral patterns. For example, if children point to an upright triangle, look doubtfully at the experimenter, and ask "Is this a triangle because the point is at the top?", their presumed belief about triangles would shift towards the wrong hypothesis of "point at the top", with a flat distribution indicating uncertainty. The model at this stage could be built upon existing work that links behavior with mental states on a microgenetic scale, including those that model shifting hypotheses (e.g., Bonawitz, Denison, Gopnik, & Griffiths, 2014), epistemic trust (e.g., Eaves & Shafto, 2017), and automatic goal inference (inverse reinforcement learning; e.g. Baker, Saxe, & Tenenbaum, 2009).

Third, a model of guidance can identify the most effective intervention given children's current belief. For example, in the aforementioned scenario, to shift children's belief away from the wrong hypothesis and towards the correct hypothesis, the best example to show may be a real triangle with the point at the bottom. Existing models of teaching, such as the model presented in Rafferty, et al., 2015, has used partially observable Markov decision process to optimize teaching actions given the learner's observed behaviors as well as previous teaching actions. Similar approaches could be used to build models that optimize guidance based on children's

current belief. Importantly, the model is not intended to immediately lead the child to the correct hypothesis as in direct instruction (e.g., “Triangles are shapes bounded by three edges and three vertices”), rather it optimizes the child’s interest to guide them toward the correct hypothesis. In this way, guided play remains child-led.

Finally, the recommended intervention can be carried out by the experimenter in a way that is consistent with the principles of guided play (e.g., through questions like “What about this one [pointing to the inverted triangle]? Does it have point at the top? Is it a real triangle?”).

Once trained, this model will significantly advance our understanding of 1) how individual children grasp concepts of geometric shapes; 2) common misconceptions along the way; and 3) optimal interventions. The resulting model-based interventions allow for guidance tailored to the learner’s moment-by-moment belief states.

### **Prospects and directions for a theory of guided play**

For children, learning takes place everywhere, all the time, and often involves interactions by the learner with more knowledgeable individuals. This ubiquity of learning opportunities can be exploited by providing subtle guidance that is contingent on the environment and children’s current mental state (Ridge, et al, 2015). Although research has highlighted the advantage of guided play, as compared to direct instruction or free play for facilitating learning (Alfieri, et al., 2011; Fisher, et al., 2013; Haden, et al., 2016; Sim & Xu, 2017; Yu, Landrum, Bonawitz, & Shafto, under review), pinpointing the optimum content and timing of guidance requires an understanding of the interactive and dynamic nature of an adult-child interaction.

We suggest that integrating computational models and data science tools may help lay out an avenue towards an empirically grounded and computationally precise framework of

guided play. *By modeling children's moment-to-moment mental state and the responsive behavior from adults, the proposed model has the potential to identify different components of guided play from dynamic and individualized interactions, and recommend model-based interventions that are optimized in terms of timing and form, with the objective of sustaining the child's interests toward the learning goal.* The resulting theory of guided play could identify key aspects of guidance that makes guided play effective in a particular context, while maintaining the complexity and ecological validity that comes with the interactive and dynamic nature of the theory. The goal is to use this framework to understand how learning proceeds and when it succeeds, which will also depend on the cultural context and individual learner. Future work could further extend the framework from one-on-one interactions in early childhood to more complex learning scenarios and topics, such as those in a classroom setting. We hope such a framework will shed light on principles of optimal environments and practices to facilitate children's learning, and present an example of using new approaches to studying cognitive development.

## References

Alfieri, L., Brooks, P. J., Aldrich, N. J., & Tenenbaum, H. R. (2011). Does discovery-based instruction enhance learning? *Journal of Educational Psychology, 103*, 1–18.

doi:10.1037/a0021017

Anderson, J. R., Boyle, C. F., & Reiser, B. J. (1985). Intelligent tutoring systems. *Science, 228*, 456-462.

Baker, L., LoBue, V., Bonawitz, E.B. & Shafto, P. (2017). Towards automated classification of emotional facial expressions. *Proceedings of the 39th annual conference of the Cognitive Science Society*.

Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition, 113*(3), 329-349. doi: 10.1016/j.cognition.2009.07.005

Begus, K., Gliga, T., & Southgate, V. (2014). Infants learn what they want to learn: Responding to infant pointing leads to superior learning. *PloS one, 9*(10), e108817.

doi:10.1371/journal.pone.0108817

Bonawitz, E., Denison, S., Gopnik, A., & Griffiths, T. L. (2014). Win-Stay, Lose-Sample: A simple sequential algorithm for approximating Bayesian inference. *Cognitive psychology, 74*, 35-65. doi: 10.1016/j.cogpsych.2014.06.003

Bonawitz, E., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery.

*Cognition, 120*(3), 322-330. doi:10.1016/j.cognition.2010.10.001

Borji, A., & Itti, L. (2014). Defending Yarbus: Eye movements reveal observers' task. *Journal of Vision, 14*(3), 29-29. doi:10.1167/14.3.29

Breazeal, C. L. (2002). *Designing sociable robots (Intelligent robotics and autonomous agents)*. Cambridge: MIT Press.

Bronfenbrenner, U. (1986). Ecology of the family as a context for human development: Research perspectives. *Developmental psychology*, 22(6), 723.

Bruner, J. S. (1961). The act of discovery. *Harvard educational review*.

Chen, E. E., Corriveau, K. H., & Harris, P. L. (2013). Children trust a consensus composed of outgroup members—but do not retain that trust. *Child Development*, 84(1), 269-282. doi: 10.1111/j.1467-8624.2012.01850.x

Clement, B., Roy, D., Oudeyer, P.-Y., & Lopes, M. (2015). Multi-Armed Bandits for Intelligent Tutoring Systems. *Journal of Educational Data Mining (JEDM)*, 7(2), 20-48.

Clements, D. H., & Sarama, J. (2014). *Learning and teaching early math: The learning trajectories approach*. New York: Routledge

Corbett, A. T.; Anderson, J. R. (1995). Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction*, 4(4): 253–278.

Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences*, 13(4), 148-153. doi:10.1016/j.tics.2009.01.005

Databrary. (2012). *The Databrary Project: A video data library for developmental science*. New York: New York University. Retrieved from <http://databrary.org>

Dewey, J. (1933). *How we think: A restatement of the relation of reflective thinking to the educational process*. Lexington, MA: Heath, 35, 64.

Duc, A. H., Bays, P., & Husain, M. (2008). Eye movements as a probe of attention. *Progress in Brain Research*, 171, 403-411. doi: 10.1016/S0079-6123(08)00659-6

Eaves, B. S., & Shafto, P. (2012). Unifying pedagogical reasoning and epistemic trust. *Advances in Child Development and Behavior*, 43, 295-319. doi: 10.1016/B978-0-12-397919-3.00011-3

Eaves, B. S., & Shafto, P. (2017). Parameterizing developmental changes in epistemic trust. *Psychonomic Bulletin & Review*, 24(2), 277-306. doi: 10.3758/s13423-016-1082-x

Fernández, M., Wegerif, R., Mercer, N., & Rojas-Drummond, S. (2001). Re-conceptualizing " scaffolding" and the zone of proximal development in the context of symmetrical collaborative learning. *The Journal of Classroom Interaction*, 36/37, 40-54.

Fisher, K. R., Hirsh-Pasek, K., Newcombe, N., & Golinkoff, R. M. (2013). Taking shape: Supporting preschoolers' acquisition of geometric knowledge through guided play. *Child Development*, 84(6), 1872-1878. doi: 10.1111/cdev.12091

Fogel, A., & Garvey, A. (2007). Alive communication. *Infant Behavior and Development*, 30(2), 251-257. doi:10.1016/j.infbeh.2007.02.00

Franchak, J. M., Kretch, K. S., Soska, K. C., & Adolph, K. E. (2011). Head-mounted eye tracking: A new method to describe infant looking. *Child Development*, 82(6), 1738-1750. doi: 10.1111/j.1467-8624.2011.01670.x

Frank, M. (2014). Modeling the dynamics of classroom education using teaching games. *Proceedings of the 36th annual conference of the Cognitive Science Society*.

Gottlieb, J., Oudeyer, P.-Y., Lopes, M., & Baranes, A. (2013). Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends in Cognitive Sciences*, 17(11), 585-593. doi:10.1016/j.tics.2013.09.001

Greene, T. R., & Noice, H. (1988). Influence of positive affect upon creative thinking and problem solving in children. *Psychological Reports*, 63(3), 895-898. doi:

10.2466/pr0.1988.63.3.895

Gutiérrez, K. D., & Rogoff, B. (2003). Cultural ways of learning: Individual traits or repertoires of practice. *Educational researcher*, 32(5), 19-25. doi:10.3102/0013189X032005019

Haden, C. A., Cohen, T., Uttal, D., & Marcus, M. (2016). Building learning: Narrating and transferring experiences in a children's museum. In D. Sobel & J. Jipson (Eds.), *Cognitive Development in Museum Settings: Relating Research and Practice*. (p. 84 103). New York, NY: Psychology Press.

Han, M., Moore, N., Vukelich, C., & Buell, M. (2010). Does play make a difference? How play intervention affects the vocabulary learning of at-risk preschoolers. *American Journal of Play*, 3(1), 82-105.

Heller, V., & Rohlfing, K. J. (2017). Reference as an interactive achievement: Sequential and longitudinal analyses of labeling interactions in shared book reading and free play. *Frontiers in psychology*, 8, 139.

Hembacher, E., & Frank, M. C. (2017). Children's social referencing reflects sensitivity to graded uncertainty. *Proceedings of the 39th annual conference of the Cognitive Science Society*.

Hirsh-Pasek, K. & Golinkoff, R. (2003). *Einstein never used flash cards*. Emmaus, PA: Rodale Press.

Hirsh-Pasek, K., & Golinkoff, R. M. (2011). The great balancing act: Optimizing core curricula through playful learning. *The pre-K debates: Current controversies and issues*, 110-115.

Hirsh-Pasek, K., Golinkoff, R., Berk, L., & Singer, D. (2008). *A Manifesto for playful learning in preschool: Presenting the scientific evidence*. New York: Oxford University Press.

Honomichl, R. D., & Chen, Z. (2012). The role of guidance in children's discovery learning. *Wiley Interdisciplinary Reviews: Cognitive Science*, 3(6), 615-622. doi: 10.1002/wcs.1199

Kinzler, K. D., Corriveau, K. H., & Harris, P. L. (2011). Children's selective trust in native-accented speakers. *Developmental Science*, 14(1), 106-111. doi: 10.1111/j.1467-7687.2010.00965.x

Kirschner, P. A., Sweller, J., & Clark, R. E. (2006). Why minimal guidance during instruction does not work: An analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching. *Educational Psychologist*, 41(2), 75-86. doi: 10.1207/s15326985ep4102\_1

Lavelli M., Barachetti., & Florit, E. (2015). Gesture and speech during shared book reading with preschoolers with specific language impairment. *Journal of Child Language*, 42(6), 1191-1218. doi:10.1017/S0305000914000762

Littlewort, G., Bartlett, M. S., Fasel, I., Susskind, J., & Movellan, J. (2006). Dynamics of facial expression extracted automatically from video. *Image and Vision Computing*, 24(6), 615-625. doi: 10.1016/j.imavis.2005.09.011

Lohan, K. S., Rohlfing, K. J., Pitsch, K., Saunders, J., Lehmann, H., Nehaniv, C. L., . . . Wrede, B. (2012). Tutor spotter: Proposing a feature set and evaluating it in a robotic system. *International Journal of Social Robotics*, 4(2), 131-146. doi:10.1007/s12369-011-0125-8

Markant, D. B., & Gureckis, T. M. (2014). Is it better to select or to receive? Learning via active and passive hypothesis testing. *Journal of Experimental Psychology: General*, *143*(1), 94. doi: 10.1037/a0032108

McNamara, D.S., Kintsch, E., Butler-Songer, N., & Kintsch, W. (1996). Are good texts always better? Interactions of text coherence, background knowledge, and levels of understanding in learning from text. *Cognition and Instruction*, *14*, 1–43. doi:10.1207/s1532690xci1401\_1

Mayer, R. E. (2004). Should there be a three-strikes rule against pure discovery learning?. *American psychologist*, *59*(1), 14. doi: 10.1037/0003-066X.59.1.14

McCormack, T., Bramley, N., Frosch, C., Patrick, F., & Lagnado, D. (2016). Children's use of interventions to learn causal structure. *Journal of Experimental Child Psychology*, *141*, 1-22. doi: 10.1016/j.jecp.2015.06.017

Nelson, J. D. (2005). Finding useful questions: On Bayesian diagnosticity, probability, impact, and information gain. *Psychological Review*, *112*(4), 979-999. doi:10.1037/0033-295x.112.4.979

Nguyen, S. M., & Oudeyer, P.-Y. (2014). Socially guided intrinsic motivation for robot learning of motor skills. *Autonomous Robots*, *36*(3), 273-294. doi:10.1007/s10514-013-9339-y

Nojavanasghari, B., Huang, Y., & Khan, S. (2018). Interactive generative adversarial networks for facial expression generation in dyadic interactions. *arXiv preprint*. arxiv:1801.09092.

Pasquini, E. S., Corriveau, K. H., Koenig, M., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental Psychology*, *43*(5), 1216. doi: 10.1037/0012-1649.43.5.1216

Pereira, A. F., Smith, L. B., & Yu, C. (2014). A bottom-up view of toddler word learning. *Psychonomic Bulletin & Review*, 21(1), 178-185. doi:10.3758/s13423-013-0466-4

Price, L. (2004). Individual differences in learning: Cognitive control, cognitive style, and learning style. *Educational Psychology*, 24(5), 681-698. doi:10.1080/0144341042000262971

Putnam, H. (1980). The nature of mental states. *Readings in philosophy of psychology*, 1, 223-231.

Rader, N., & Hughes, E. (2005). The influence of affective state on the performance of a block design task in 6- and 7-year-old children. *Cognition & Emotion*, 19(1), 143-150. doi: 10.1080/02699930441000049

Rafferty, A. N., Brunskill, E., Griffiths, T. L., & Shafto, P. (2016). Faster teaching via pomdp planning. *Cognitive Science*, 40(6), 1290-1332. doi: 10.1111/cogs.12290

Ridge, K. E., Weisberg, D. S., Ilgaz, H., Hirsh-Pasek, K. A., & Golinkoff, R. M. (2015). Supermarket speak: Increasing talk among low-socioeconomic status families. *Mind, Brain, and Education*, 9(3), 127-135. doi: 10.1111/mbe.12081

Rogoff, B., Mistry, J., Artin, G., Mosier, C., Chavajay, P., & Heath, S. B. (1993). Guided participation in cultural activity by toddlers and caregivers. *Monographs of the Society for Research in Child Development*, 58(8), i-179. doi:10.2307/1166109

Rothe, A., Lake, B. M., & Gureckis, T. M. (2016). Asking and evaluating natural language questions. *Proceedings of the 38th annual conference of the cognitive science society*.

Ruvolo, P., Whitehill, J., Virnes, M., & Movellan, J. R. (2008). Building a more effective teaching robot using apprenticeship learning. *Proceedings of the 7th IEEE International Conference on Development and Learning*. doi: 10.1109/devlrm.2008.4640831

Settles, B. (2010). *Active learning literature survey*. University of Wisconsin, Madison, 52(55-66), 11.

Shafto, P., & Goodman, N. (2008). Teaching games: Statistical sampling assumptions for learning in pedagogical situations. *Proceedings of the 30th annual conference of the Cognitive Science Society*.

Shafto, P., Eaves, B., Navarro, D. J., & Perfors, A. (2012). Epistemic trust: Modeling children's reasoning about others' knowledge and intent. *Developmental Science*, 15(3), 436-447. doi: 10.1111/j.1467-7687.2012.01135.x

Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014). A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive psychology*, 71, 55-89. doi: 10.1016/j.cogpsych.2013.12.004

Shafto, P., Goodman, N. D., & Frank, M. C. (2012). Learning from others: The consequences of psychological reasoning for human learning. *Perspectives on Psychological Science*, 7(4), 341-351. doi: 10.1177/1745691612448481

Sim, Z. L., Tanner, M. M., Alpert, N. Y., & Xu, F. (2015). Children learn better when they select their own data. In *Proceedings of the 37th annual conference of the Cognitive Science Society*.

Sim, Z. & Xu, F. (2017) Learning higher-order generalizations through free play: Evidence from 2-and 3-year-old children. *Developmental psychology*, 53(4), 642.

Smith, L. B., Yu, C., Yoshida, H., & Fausey, C. M. (2015). Contributions of head-mounted cameras to studying the visual environments of infants and young children. *Journal of Cognition and Development*, 16(3), 407-419. doi: 10.1080/15248372.2014.933430

Sobel, D. M., & Corriveau, K. H. (2010). Children monitor individuals' expertise for word learning. *Child Development*, *81*(2), 669-679. DOI: 10.1111/j.1467-8624.2009.01422.x

Thelen, E., & Smith, L. B. (1996). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT press.

Thomaz, A. L., & Breazeal, C. (2008). Teachable robots: Understanding human teaching behavior to build more effective robot learners. *Artificial Intelligence*, *172*(6-7), 716-737. doi: 10.1016/j.artint.2007.09.009

Tobias, S., & Duffy, T. M. (2009). *Constructivist instruction: Success or failure?*. New York: Routledge.

Vollmer, A.-L., Mühlig, M., Steil, J. J., Pitsch, K., Fritsch, J., Rohlfing, K. J. & Wrede, B. (2014). Robots show us how to teach them: Feedback from robots shapes tutoring behavior during action learning. *PLOS ONE* *9*(3): e91349. doi: 10.1371/journal.pone.0091349

Vygotsky, L. S. (1934/1987). Thinking and speech. In *L. S. Vygotsky, Collected works* (vol. 1, pp. 39–285) (R. Rieber & A. Carton, Eds; N. Minick, Trans.). New York: Plenum.

Weisberg, D. S., Hirsh-Pasek, K., & Golinkoff, R. M. (2013). Guided play: Where curricular goals meet a playful pedagogy. *Mind, Brain, and Education*, *7*(2), 104-112. doi: 10.1111/mbe.12015

Weisberg, D. S., Hirsh-Pasek, K., Golinkoff, R. M., & McCandliss, B. D. (2014). Mise en place: Setting the stage for thought and action. *Trends in Cognitive Sciences*, *18*(6), 276-278. doi: 10.1016/j.tics.2014.02.012

Weisberg, D. S., Hirsh-Pasek, K., Golinkoff, R. M., Kittredge, A. K., & Klahr, D. (2016). Guided play: Principles and practices. *Current Directions in Psychological Science*, *25*(3), 177-182.

Wise, A. F., & O'Neill, K. (2009). Beyond more versus less: A reframing of the debate on instructional guidance. In S. Tobias & T. M. Duffy (Eds.), *Constructivist instruction: Success or failure?* (pp. 82-105). New York: Routledge/Taylor & Francis Group.

Wood, D., Bruner, J. S., & Ross, G. (1976). The role of tutoring in problem solving. *Journal of child psychology and psychiatry*, 17(2), 89-100.

Yang, S.C-H. & Shafto, P. (2017). Teaching versus active learning: A computational analysis of conditions that affect learning. *Proceedings of the 39th annual conference of the Cognitive Science Society*.

Yao, A., Shao, J., Ma, N., & Chen, Y. (2015). Capturing au-aware facial features and their latent relations for emotion recognition in the wild. *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction*.

Yu, Y., Bonawitz, E., & Shafto, P. (2017). Inconvenient samples: Modeling the effects of non-consent by coupling observational and experimental results. *Proceedings of the 39th Annual Conference of the Cognitive Science Society*.

Yu, Y., & Kushnir, T. (2016). When what's inside counts: Sequence of demonstrated actions affects preschooler's categorization by nonobvious properties. *Developmental Psychology*, 52(3), 400-410. doi:10.1037/dev0000088

Yu, Y., Landrum, A., Bonawitz, E. & Shafto, P. (under review). Questioning supports effective transmission of knowledge and increased exploratory learning in pre-kindergarten children.

Yudelson, M. V., Koedinger, K. R., & Gordon, G. J. (2013). Individualized bayesian knowledge tracing models. *International Conference on Artificial Intelligence in Education*.

Zhu, X. (2015). Machine teaching: An inverse problem to machine learning and an approach toward optimal education. In *AAAI Conference on Artificial Intelligence*.