

# Building a storytelling conversational agent through parent-AI collaboration

**Zheng Zhang**  
zzhang37@nd.edu  
University of Notre Dame

**Ying Xu**  
ying.xu@uci.edu  
UC Irvine

**Yanhao Wang**  
yanhaow@gatech.edu  
Georgia Tech

**Tongshuang Wu**  
wtshuang@cs.washington.edu  
University of Washington

**Bingsheng Yao**  
yaob@rpi.edu  
Rensselaer Polytechnic Institute

**Daniel Ritchie**  
drritchi@uci.edu  
UC Irvine

**Mo Yu**  
yum@us.ibm.com  
IBM Research

**Dakuo Wang**  
dakuo.wang@ibm.com  
IBM Research

**Toby Jia-Jun Li**  
toby.j.li@nd.edu  
University of Notre Dame

## Abstract

In this paper, we describe the design of StoryBuddy, a prototype system that allows parents to collaborate with AI in creating interactive storytelling experiences. To accommodate dynamic user needs, StoryBuddy supports two modes: a parent-child joint reading mode where parents involve in book reading process and a child independent reading mode where parents have minimal involvement. StoryBuddy also allows parents to configure question types and track child progress, and generates questions automatically. A preliminary user study suggests that parents and children found StoryBuddy useful, helpful, and likable.

## Introduction

Storytelling is a common parent-child activity that provides many educational benefits such as improving children’s language fluency, communication skills, cultural and emotional awareness, and other aspects of cognitive development (Peck 1989). Interactive storytelling in particular, where a storyteller asks questions relevant to story content and prompts a child to express their thoughts about the story, has been shown to maximize the educational benefits of storytelling (Kotaman 2020).

However, many parents experience barriers such as having difficulty in coming up with appropriate questions, multi-tasking, time conflict between storytelling and personal business, and challenges with keeping track of the child’s progress (Dietz et al. 2021). Although there are some existing digital interactive storytelling systems (Dietz et al. 2021), all those systems only support a limited set of storybooks for which the developers manually crafted a fixed set of questions prior deployment.

To address these limitations, we developed StoryBuddy, a system adopting a human-AI collaboration approach (Wang et al. 2019) that incorporates the expertise and preferences of

parents into the development of interactive storytelling experiences for their children on any storybooks. StoryBuddy has the following innovative features: (1) Personalized automatic question generation for any user-provided storybooks; (2) Two modes for supporting both parent presence and absence cases; (3) Dashboard that allows parents to assess children’s performance and track children’s question-answering progress; (4) An integrated conversational bot that provides interactive questioning-and-answering experience for children.

## Related work

Various digital tools have been introduced to help facilitate different aspects in interactive storytelling (Xu et al. 2021a) such as an embodied learning companion robot that can introduce new vocabulary words during storytelling (Kory and Breazeal 2014) and an in-home learning companion robot that can make pre-programmed comments in stories (Michaelis and Mutlu 2017). StoryCoder presents two storytelling games and four computational thinking games to support the use of interactive storytelling as a way to teach computational thinking concepts (Dietz et al. 2021).

For the automatic story question and answer generation (QAG), most works on QAG systems approach the problem from a pure machine learning perspective, trying to invent new rule-based (e.g., (Yao, Bouma, and Zhang 2012)) or neural-network-based (e.g., (Du, Shao, and Cardie 2017)) models that can generate “more accurate” questions and answers. While these systems perform well in generating correct and relevant questions. Their generated questions usually lack educational values and are ineffective in maintaining child engagement.

## Participatory Design

We launched participatory design sessions (Muller and Kuhn 1993) with four parents with children between 4 and 8 to uncover design goals and design strategies for StoryBuddy. In each participatory design session, we presented

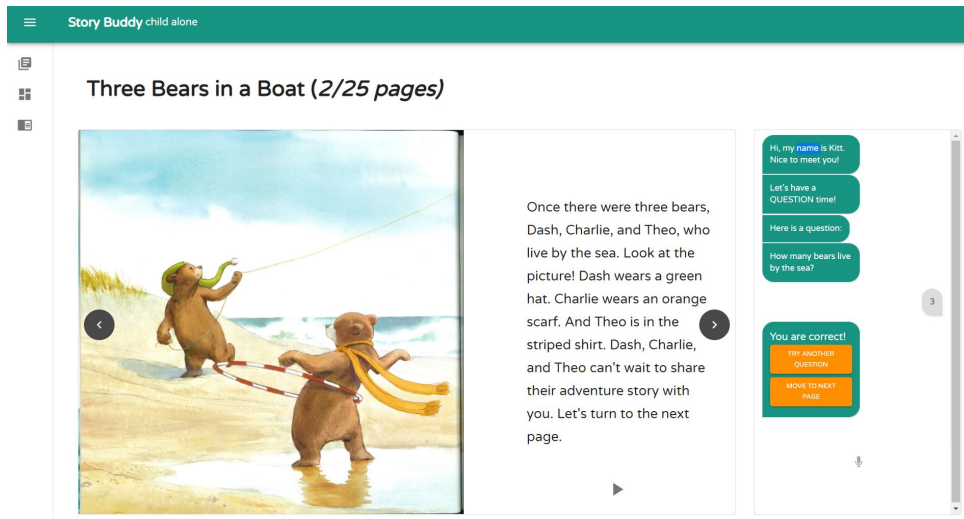


Figure 1: StoryBuddy's story reading interface in the automated bot-reading mode

the participant with four variations of storyboards in random order. These four variations differed from each other in two aspects: (1) whether the parent is present at the storytelling; and (2) whether the system runs on a tablet or a smart speaker.

The storyboards set scenarios for users to reflect on their needs, constraints, and practices (Bodker 1999). In the storyboards, we purposely deemphasized the details in the interface designs of the system by avoiding directly showing screen contents (if showing screens was necessary, our storyboards used low-fidelity sketches). Instead, the storyboards focused on illustrating the parent motivations and goals in the scenarios, the constraints in time, attention, and cognitive capabilities, and the interaction dynamic among the parent, child, and agent in different scenarios depending on whether the parent is present and which type of device is used. The main goal was to elicit the feelings and emotions of participants towards different design decisions and to build empathy with them. After validating user needs illustrated in the storyboards, we asked participants to think about how their personal experience with parent-child storytelling (or the lack of it) can connect to the storyboard scenarios. Through this process, each parent identified things they liked, things they disliked, and their concerns about the different paradigms of AI involvement in the parent-child interactive storytelling process. Lastly, we asked them to think of and propose new ideas on design features, interfaces, or interaction techniques for (1) addressing the issues they identified in the storyboards; and (2) bridging the gaps between the scenarios presented in the storyboards and their own personal scenarios.

In the participatory design sessions, the users came up with several design strategies for improving the varieties of interaction patterns to optimize child engagement in storytelling, accommodating parents' flexible learning and developmental goals of storytelling, and supporting flexible parent engagement to balance the conflict between the parent's

desire to be present, play an active role and build parent-child relationship; and the constraint that sometimes they are not available. Based on the insights from the participatory design process, We developed the following six design goals:

1. **DG1:** Maintain child attention and optimize for child engagement through varying the interaction means among the parent, the child, and the agent
2. **DG2:** Assist parents with facilitating question-answering in joint-reading through recommending questions
3. **DG3:** Support different parental goals on whether to focus on the assessment objective in interactive storytelling
4. **DG4:** Generate follow-up questions to deepen the questioning-and-answering session
5. **DG5:** Accommodate varied parent preferences of question generation
6. **DG6:** Support storytelling in both parent present and absent environment

## StoryBuddy Prototype System

Based on insights from the participatory design sessions, we developed a prototype of StoryBuddy—an AI-enabled interactive tool for configuring, augmenting, and automating interactive storytelling with children. The prototype was implemented as a web app that can run on desktops, laptops, tablets, or smartphones. StoryBuddy presents several features that allow flexible parent involvement in both the configuration and delivery of interactive story contents while supporting diverse parent needs in children's skill development, progress assessment, and engagement. Below we summarize several key features of StoryBuddy. More detail on them as well as other features such as an interactive dashboard for child progress tracking and performance assess-

ment can be found in the demo video<sup>1</sup>.

### Question generation

StoryBuddy uses an automated question-answer generation(QAG) model trained on the FairytaleQA dataset (Xu et al. 2021b). This QAG model can automatically generate high-quality QA pairs from any children’s storybooks. The questions generated by the QAG model are designed to mimic the style “as if a teacher or parent is to think of a question to improve children’s language comprehension ability while reading a story to them (Yao et al. 2021).” This QAG model also supports generating questions and the corresponding answers from a specific question type. Based on this QAG model, StoryBuddy is able to generate multiple QA pairs given the story section that users are reading.

### Question preference configuration

StoryBuddy provides a preference configuration panel, where the parent can choose the preferred types of generated questions for the back-end model. From the participatory design insights, we learned that some parents wish to have controls of the generated questions at a finer granularity. The preference configuration panel allows them to customize the generation of questions to better align with the learning and development goals they have for their children. The use of this panel is optional.

### Conversational agent

StoryBuddy has a built-in conversational bot that can lead a questioning-and-answering conversation with children. At storytelling time, children can see the book page including text and illustrations as shown in Figure 1. On the first page, the agent first greets the child. The agent will then read the story text on each page, say “OK, here is a question” and then ask the question as configured by the parent. The child can answer the question to the bot by clicking on the microphone icon, the transcript of their speech will automatically appear in the dialog. After receiving the child’s answer, the agent will judge the correctness of the answer. If the answer is correct, the agent will say “You are correct! Good job!”, the child can choose either “move to next page” or “try another question” (if another question is available); if the answer is wrong, the child will see an additional “try again” option for them to retry the same question.

### Modes for parent presence and absence

StoryBuddy provides two distinct modes to reconcile the parent’s desire and availability to be present. In the parent-present mode, StoryBuddy assists the parent by helping them identify opportunities to ask questions and recommend questions to use in order to help reduce their cognitive load. Besides, StoryBuddy supports dynamic interaction patterns—the parent and the child can switch between different ways of interaction including parent-child reading, child-bot reading, parent-child questioning-and-answering, and child-bot questioning-and-answering. This capability to

switch between different interaction patterns was shown to be effective in maintaining the child’s engagement in our preliminary user study.

In contrast, a child-alone mode allows parents to preset the story and questions, after the setting StoryBuddy will launch a child-bot joint reading without requiring the parent’s intervention. The bot will lead the questioning-and-answering as described in the conversational agent section.

## Preliminary User Study

We conducted a preliminary user study with 12 parent-child pairs over Zoom to assess the usefulness and usability of StoryBuddy. All 12 pairs successfully completed the two 20-minute experiment sessions: assigned parent-AI co-reading and automated bot-reading. Our participants generally found StoryBuddy easy and enjoyable to use, despite that a small number of participants were less enthusiastic. User frustration reported in the study mostly originated from technical limitations, such as the inaccuracy in speech recognition or answer detection in the child-agent interaction session via voice interface.

Parents praised the usefulness of Storybuddy, believed that automatically generated questions could reduce their burden of making questions by themselves. Parents also praised StoryBuddy’s engagement design for keeping their children entertained through the storytelling process. Several parents envisioned its value for tracking child progress, assessing their performance assessment, and help with language skill development. Such benefit can be more significant in non-native-English-speaking families.

## Design Implication

### Supporting Flexibility in Parent Involvement

An important design consideration for parent involvement is how we can help parents stay involved when they are unavailable for synchronous parent-child interactions. The conflict between (1) parents’ desire for fulfilling their children’s storytelling needs and stay involved in the process; and (2) parents’ limited and constrained time availability was a recurring theme in our studies. This problem is aggravated by the increasingly common work-from-home arrangement for parents. In work-from-home situations, parents may *seem available* to children for interaction since they are physically home when they are, in fact, unavailable. An interactive storytelling agent that can keep children engaged without requiring parental intervention would be particularly useful in such scenarios. To address this issue, the core strategy that StoryBuddy used is to turn synchronous involvement into asynchronous involvement—While the delivery of the story content and the interaction for question-answering is facilitated by the StoryBuddy agent, the parent has (1) control over the content of the interaction through the configuration *before* a storytelling session; and (2) knowledge on the progress and the performance of the child through viewing the dashboard *after* a storytelling session.

<sup>1</sup>[https://youtu.be/vz5AYO\\_eJ54](https://youtu.be/vz5AYO_eJ54)

## AI as a Partner in Parent-Child Interaction

In the parent-AI co-reading mode, StoryBuddy introduced an AI agent into an *existing* activity that used to be only the parent and the child. This existing activity of parent-child joint reading is complex as it usually fulfills multiple goals at the same time: relationship building, skill development and assessment, and entertainment. Therefore, it is critical to consider the roles that the AI system will play in every aspect of this activity, and how its roles contribute (positively or negatively) to the user's goals.

For example, in our design, StoryBuddy assists the parent by helping them identify opportunities for asking questions, recommending questions to use, and proposing follow-up questions. These assistance approaches help parents come up with better questions for fulfilling their skill development and assessment goal, reduce their cognitive load so that they can allocate more attention to the interaction with their child for the relationship-building goal, and keep their child engaged and entertained through the occasional interaction between the child and the agent. However, all these forms of AI involvement do not lessen the central role of the parent in their interactions with the child. It is a human-AI collaboration (Wang et al. 2020) instead of a human-AI competition.

## Conclusion

This paper presented StoryBuddy, a new system that allowed parents to collaborate with an AI system in creating storytelling experiences with interactive questioning-answering. Informed by results from a formative study and a participatory design study, we designed two distinct modes and several dynamic interaction paradigms that supported flexible degrees of parent involvement. A user study with 12 pairs of parents and children found that StoryBuddy was effective in providing parents with desired levels of control and involvement while maintaining children's engagement in the storytelling process. Parents and children considered StoryBuddy useful, helpful, and likable.

## References

- Bodker, S. 1999. Scenarios in user-centred design-setting the stage for reflection and action. In *Proceedings of the 32nd Annual Hawaii International Conference on Systems Sciences. 1999. HICSS-32. Abstracts and CD-ROM of Full Papers*, 11–pp. IEEE.
- Dietz, G.; Le, J. K.; Tamer, N.; Han, J.; Gweon, H.; Murnane, E. L.; and Landay, J. A. 2021. *StoryCoder: Teaching Computational Thinking Concepts Through Storytelling in a Voice-Guided App for Children*. New York, NY, USA: Association for Computing Machinery.
- Du, X.; Shao, J.; and Cardie, C. 2017. Learning to ask: Neural question generation for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1342–1352. Vancouver, Canada: ACL.
- Kory, J., and Breazeal, C. 2014. Storytelling with robots: Learning companions for preschool children's language development. In *The 23rd IEEE international symposium on robot and human interactive communication*, 643–648. IEEE.
- Kotaman, H. 2020. Impacts of dialogical storybook reading on young children's reading attitudes and vocabulary development. *Reading Improvement* 57(1):40–45.
- Michaelis, J. E., and Mutlu, B. 2017. *Someone to Read with: Design of and Experiences with an In-Home Learning Companion Robot for Reading*. New York, NY, USA: ACM. 301–312.
- Muller, M. J., and Kuhn, S. 1993. Participatory design. *Commun. ACM* 36(6):24–28.
- Peck, J. 1989. Using storytelling to promote language and literacy development. *The Reading Teacher* 43(2):138–141.
- Wang, D.; Weisz, J. D.; Muller, M.; Ram, P.; Geyer, W.; Dugan, C.; Tausczik, Y.; Samulowitz, H.; and Gray, A. 2019. Human-ai collaboration in data science: Exploring data scientists' perceptions of automated ai. *Proceedings of the ACM on Human-Computer Interaction* 3(CSCW):1–24.
- Wang, D.; Churchill, E.; Maes, P.; Fan, X.; Shneiderman, B.; Shi, Y.; and Wang, Q. 2020. From human-human collaboration to human-ai collaboration: Designing ai systems that can work together with people. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI EA '20, 1–6. New York, NY, USA: Association for Computing Machinery.
- Xu, Y.; Wang, D.; Collins, P.; Lee, H.; and Warschauer, M. 2021a. Same benefits, different communication patterns: Comparing children's reading with a conversational agent vs. a human partner. *Computers & Education* 161:104059.
- Xu, Y.; Wang, D.; Yu, M.; Ritchie, D.; Yao, B.; Tongshuang Wu, Zheng Zhang, T. J.-J. L.; Bradford, N.; Sun, B.; Hoang, T.; Yisi Sang, Y. H.; Ma, X.; Yang, D.; Peng, N.; Yu, Z.; and Warschauer, M. 2021b. Fantastic questions and where to find them: Fairytaleqa: An authentic dataset for narrative comprehension.
- Yao, B.; Wang, D.; Wu, T.; Hoang, T.; Sun, B.; Li, T. J.-J.; Yu, M.; and Xu, Y. 2021. It is ai's turn to ask human a question: Question and answer pair generation for children storybooks in fairytaleqa dataset.
- Yao, X.; Bouma, G.; and Zhang, Y. 2012. Semantics-based question generation and implementation. *Dialogue & Discourse* 3(2):11–42.