DreamTales: Reframing AI's Role in Parent-Child Storytelling Through Pre-Post Reflection

JIANING YU*, Harvard Graduate School of Design, USA
HISAKO NOMOTO*, Harvard University, USA
SHOHREH GHORBANI*, MIT Media Lab, USA

Existing AI co-reading systems often act as in-situ mediators, risking reduced parental agency and disrupted emotional exchange. DreamTales redefines AI's role as a Reflective Companion via a pre-post reflection workflow. Leveraging GPT-4o and DALL·E 3, it generates personalized, context-aware stories, passively analyzes parent-child vocal/linguistic synchrony during reading, and provides a non-judgmental Reflective Dashboard post-session. In a pilot study (N=2), parents expressed a strong preference for the reflective workflow over live AI mediation, citing improved preservation of the parent–child dyadic bond. Notably, personalized story generation emerged as a potential catalyst for emotional synchrony. Together, this work highlights the Reflective Companion model as a promising design direction and sheds light on opportunities for further exploration and validation of personalized storytelling as a mechanism for supporting parent–child emotional bonding and parental reflective awareness.

CCS Concepts: • **Human-centered computing** → **Emotional computing**; **Natural language interfaces**; *Mobile devices.*

Additional Key Words and Phrases: AI Storytelling, Emotional Synchrony, Parental Reflective Functioning, HCI for Caregiving, Attachment Theory

ACM Reference Format:
Jianing Yu, Hisako Nomoto, and Shohreh Ghorbani. 2026. DreamTales: Reframing AI's Role in Parent-Child Storytelling Through Pre-Post Reflection. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '26).* ACM, New York, NY, USA, 7 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

Shared storytelling is foundational to parent–child bonding, fostering emotional co-regulation, attachment security, and social-emotional skill development—relying on caregivers' ability to sensitively respond to children's affective cues (attunement). However, parents often struggle with adjustment amid external stressors, emotional fatigue, or uncertainty in feelings-based conversations. Recent AI co-reading systems (e.g., ContextQ [8], StoryBuddy [14]) typically act as in-situ mediators, offering real-time prompts to boost short-term engagement. However, inserting AI into live emotional exchanges disrupts the parent-child dyad: it reshapes authority and intimacy, diminishes parental agency and reflective capacity, and risks "delegated intimacy" [5]. Parents also describe such real-time interventions as intrusive to the natural rhythm of shared storytelling.

*These authors contributed equally to this work.

Authors' Contact Information: Jianing Yu, nomy_yu@gsd.harvard.edu, Harvard Graduate School of Design, Cambridge, MA, USA; Hisako Nomoto, hisako_nomoto@gse.harvard.edu, Harvard University, Cambridge, MA, USA; Shohreh Ghorbani, ghorbani@mit.edu, MIT Media Lab, Cambridge, MA, USA.

Manuscript submitted to ACM

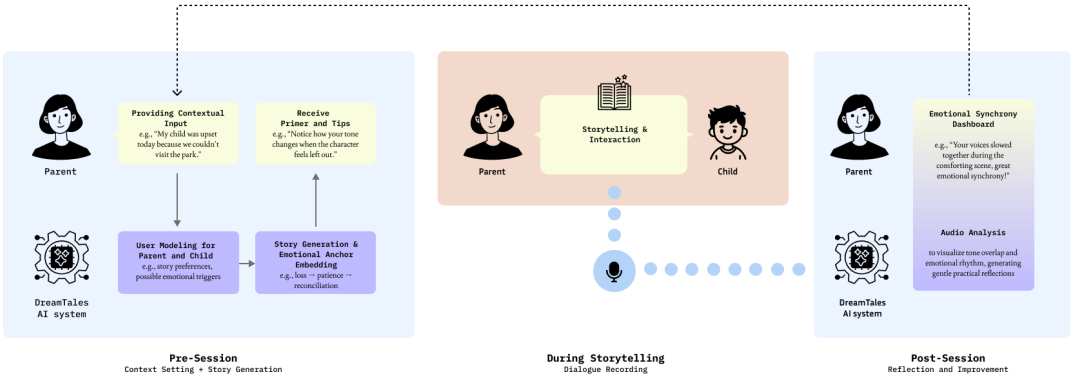

Fig. 1. DreamTales reframes AI's role in parent–child storytelling by shifting support away from real-time intervention toward preparation and reflection.

To address these tensions, we reframe AI as a **Reflective Companion**, shifting support to pre- and post-session moments rather than live intervention. Grounded in parental reflective functioning theory [12], this approach preserves the dyadic bond while nurturing caregivers' ability to interpret children's internal states. We embody this in **DreamTales**, an AI-supported workflow for bedtime storytelling: pre-session, it generates personalized stories (integrating the child's interests/experiences) and non-intrusive reflection primers; during reading, it remains passively silent to capture interaction data; post-session, a Reflective Dashboard visualizes emotional synchrony with non-judgmental cues (see Figure 1).

A formative pilot study (N=2 parent-child dyads) revealed strong parental preference for this reflective model over real-time mediation, citing preserved emotional flow and agency. Personalized storytelling also emerged as a catalyst for emotional synchrony. This work makes three core contributions: (1) introducing the Reflective Companion concept for parent-child interactions; (2) presenting the DreamTales prototype integrating LLMs and multimodal analysis; (3) offering preliminary empirical insights into non-intrusive AI's potential to support emotional bonding and parental reflection.

2 Related Work

Shared storytelling is a cornerstone of parent-child emotional co-regulation and attachment formation [3, 4], with parental attunement—sensitive responsiveness to a child's affective cues—playing a pivotal role in fostering empathy and emotional understanding [1, 9]. A randomized trial confirmed that while AI-guided reading boosts behavioral engagement, parent-led interactions yield stronger affective connection [13], highlighting the irreplaceable value of dyadic intimacy. Computational research has further linked parent-child behavioral and vocal synchrony to relationship quality [2, 10], yet existing technologies often prioritize behavioral routines over relational attunement.

To support shared reading, HCI community has developed AI systems primarily as in-situ mediators: ContextQ generates real-time dialogic prompts [8], while StoryBuddy positions AI as a co-storyteller [14]. These tools increase conversational turns and vocabulary diversity but reshape interaction dynamics—shifting parental roles from autonomous caregivers to AI facilitators, diminishing emotional bonding and sense of authorship. Such real-time intervention risks "delegated intimacy," displacing parental emotional labor and eroding relational authenticity [5, 7].

Emerging HCI paradigms prioritize reflective support over real-time automation, positioning AI as a mirror for self-awareness rather than a manager [6]. Systems like eaSEL offer post-experience summaries to avoid interrupting parent-child interaction [11], aligning with research showing passive detection of emotional synchrony via tone and motion cues [2, 10]. Grounded in parental reflective functioning theory—emphasizing caregivers' ability to interpret internal states [9]—DreamTales advances this direction by integrating pre-story preparation and post-story reflection, separating observation from intervention to preserve dyadic integrity while fostering long-term sensitivity.

3 Methodology

We adopted a research-through-design approach to explore alternative AI roles in parent-child storytelling, consisting of three interconnected stages: formative design exploration, DreamTales prototype development, and a small-scale pilot study.

3.1 Formative Design Exploration

We conducted semi-structured interviews with two parent-child dyads (a mother with a 4-year-old daughter and a father with a 3-year-old son) to understand existing co-reading practices, relational goals, and parental perceptions of AI roles. Interviews focused on parents' definitions of high-quality co-reading, stressors during bedtime routines, emotionally meaningful moments, and comparative views of two hypothetical AI models: In-Situ Mediator vs Reflective Companion. Key insights included parental concerns about delegating emotional labor and a desire to preserve culturally grounded, personalized narratives, which guided subsequent design decisions.

3.2 Prototype Design: DreamTales

Guided by insights from the formative exploration, we designed and implemented **DreamTales**, an AI-supported pre-post reflection workflow structured to separate observation from intervention, positioning AI as a reflective companion for parents. The prototype operates in three interconnected stages, with a user-centric design that balances technical functionality and natural caregiving dynamics.

3.2.1 ***Pre-Session: Personalization and Preparation.*** As shown in Figure 2, the parent initiates the process by inputting the child's context, including basic information such as name and age, as well as more personal and reflective questions such as his current interests ("loves race cars") and a recent emotional struggle ("was sad we couldn't go to the park"). The personalized story generation pipeline involves both GPT-4o and DALL-E 3 for text and illustration generation:

1) Story Engine (GPT-4o): Based on the parent's input, GPT-4o generates a bespoke story text. The prompt is meticulously engineered to enforce constraints crucial for the target age group (3–6 years): simple, low-complexity vocabulary; short sentence structure; and the integration of the child's context as a key emotional narrative element.

2) Visual Asset Pipeline (DALL-E 3)): To support the visual aspect of the storybook, DALL-E 3 generates consistent, simple illustrations based on character descriptions generated by the LLM.

The generated 5-6 page illustration storybook is followed by tailored tips and suggestions. We used a specialized lightweight model (GPT-4o-mini) to analyze the generated story content to produce a single, concise pre-reading tip (the Reflection Primer). For example, if the story involves a moment of disappointment, the primer might be: "Before starting, notice where the disappointment appears in your voice..." This primes the parent's emotional awareness without being overly prescriptive.

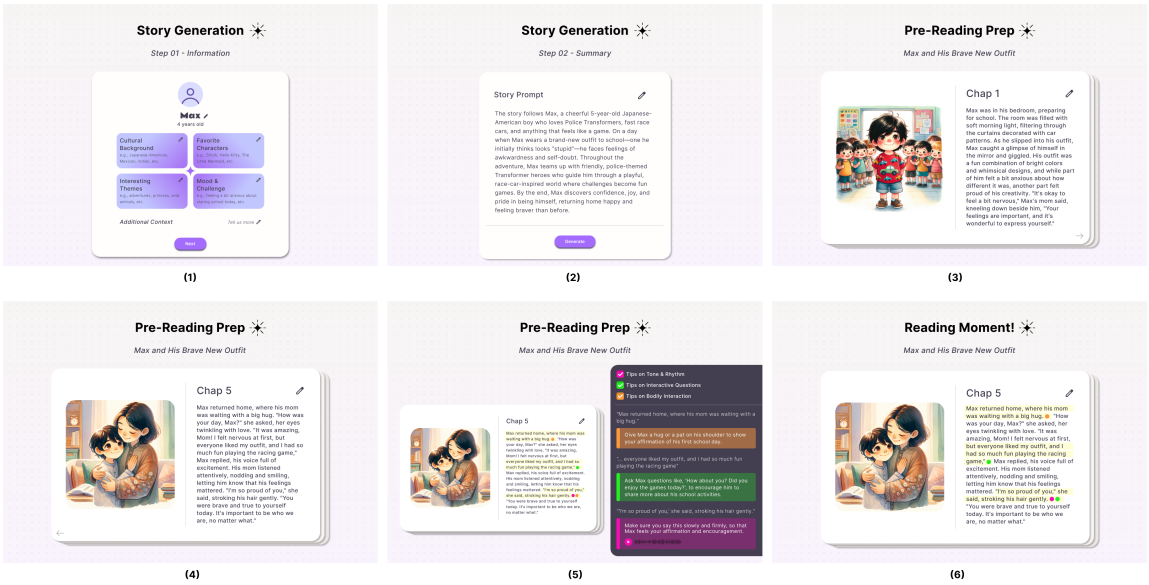

Fig. 2. Pre-Story Workflow for Personalized Story Generation and Pre-Reading Preparation. (1) Parents first provide structured contextual inputs about the child, including background, interests, and current emotional challenges. (2) The system synthesizes these inputs into an editable story prompt, allowing parents to review and refine the narrative intent. (3–4) A personalized illustrated storybook is then generated, grounding the narrative in the child's lived experience. (5) Based on the story content, DreamTales produces optional, theme-specific pre-reading tips (e.g., tone and rhythm, interactive questions, bodily interaction), which parents can selectively enable. (6) During reading, these tips are visually anchored to specific sentences or paragraphs and highlighted through color cues, supporting real-time awareness without interrupting the storytelling flow.

3.2.2 ***During-Session: Passive Listening (The Invisible AI).*** Crucially, during the storytelling session, the system provides no visible or audible feedback. A concealed microphone passively records the interaction. The passive nature of the recording ensures the natural flow of conversation, preserving the intimacy and integrity of the parent-child dyadic space. The system's purpose here is purely observational, gathering the raw data needed for post-session analysis.

3.2.3 ***Post-Session: Reflective Dashboard and Emotional Synchrony.*** Within 2 minutes of the session ending, the app processes the recorded audio via a dual-stream multimodal analysis pipeline to generate the Reflective Dashboard. First, Pyannote performs speaker diarization to separate parent and child speech turns with 92% accuracy (pilot-tested), while Whisper transcribes the audio with word-level timing (±0.2s latency) to align linguistic and acoustic data.

Feature extraction combines acoustic and linguistic analysis: a pre-trained Speech Emotion Recognition (SER) model extracts prosodic features (pitch, intensity, tempo) and interaction markers (turn-taking duration, vocal overlap), while a fine-tuned RoBERTa model analyzes transcribed text for moment-to-moment sentiment and emotional valence (positive/neutral/negative). These features are fused into a normalized valence score (0–100) using a weighted algorithm:

$$\text{Valence}(t) = (P(\text{Happy}) \times 100) + (P(\text{Neutral}) \times 50) - (P(\text{Negative}) \times 50) \tag{1}$$

generating parallel time-series streams for parent and child. Emotional synchrony is computed as

$$\text{Synchrony}(t) = 1 - |\text{Valence}_{\text{Parent}}(t) - \text{Valence}_{\text{Child}}(t)| \tag{2}$$

Fig. 3. Representative Emotional Synchrony Plot captured during the pilot study (Family B). The blue line tracks the parent's consistent engagement (valence), while the dark line captures the child's fluctuations. The system correctly identified a "Desynchrony Event" at T8 (distraction) and a "Fatigue Drop" at T14, validating the pipeline's ability to map qualitative behavioral states to quantitative metrics.

visualized on the dashboard via a dual-line plot (parent: blue, child: dark gray) with time-stamped annotations. The dashboard prioritizes interpretability over complexity: it highlights key moments (e.g., "High Synchrony: Shared excitement at 1:45") and "missed mirror moments" (e.g., "Child's valence dropped at 3:10—did you notice their tone shift?"), avoiding numerical scores or evaluative labels. Additional non-judgmental cues include open-ended prompts ("What do you think led to the shared laughter here?") and actionable suggestions ("Next time, try pausing when your child's valence rises to deepen engagement"). The dashboard is accessible for 72 hours post-session, with an option to save insights to a private journal for long-term reflection.

3.3 Pilot User Study

We deployed the DreamTales prototype with the same two parent-child dyads to explore experiential qualities of the reflective workflow. The pilot aimed to surface value signals, breakdowns, and how personalized storytelling, pre-session priming, and post-session reflection shaped interaction dynamics—providing grounding for key findings rather than generalizable conclusions.

4 Key Findings and Design Insights

Synthesizing insights from formative exploration, prototype use, and pilot deployment, we identify four core themes that clarify how reflective, non-intrusive AI aligns with parents' values in intimate storytelling contexts.

Preference for Reflection over Real-Time Mediation. Parents consistently expressed discomfort with real-time AI intervention, describing prompts or guidance as disruptive, cognitively taxing, and distracting from the child. In contrast, the reflective companion model was viewed as more supportive and less evaluative, as it preserved uninterrupted, parent-led emotional flow during storytelling. This preference underscores the importance of respecting parental agency and authority by shifting AI support outside live interactions, aligning with caregivers' desire for authentic intimacy.

Personalized Story Foundations as a Synchronization Catalyst. Personalized storytelling emerged as a key driver of parent-child emotional alignment, with the pre-session process of articulating the child's interests, experiences, and emotional challenges already initiating pre-attunement. Anchored in the child's lived context, these stories created a shared emotional reference point that fostered heightened engagement and conversational reciprocity during reading. Beyond narrative relevance, personalization established the affective conditions necessary for synchrony and co-regulation, functioning as a relational catalyst rather than mere content.

Value of Non-Judgmental Reflective Cues. Parents responded positively to reflective feedback focused on awareness and interpretation, as visualizations of emotional synchrony and qualitative prompts helped them notice otherwise overlooked shifts in the child's engagement. Critically, they resisted numerical scores or prescriptive judgments, which risked introducing performance pressure or anxiety. Non-judgmental, open-ended cues aligned more closely with caregiving values, emphasizing that reflective AI in relational contexts should prioritize interpretability over measurement.

Practical Constraints and Technical Challenges. Two key limitations surfaced during pilot deployment: story length and linguistic complexity, which led to fatigue and disengagement—particularly for younger children—and technical challenges in child speech diarization and emotion recognition. Generic models struggled to adapt to the noisy, high-pitched, and dynamic nature of caregiving environments, highlighting the need for tailored multimodal analysis pipelines and age-appropriate content optimization in future iterations.

5 Conclusion and Future Work

This work explores an alternative role for AI in intimate caregiving contexts: positioning AI not as an in-situ mediator, but as a Reflective Companion. Through the design and pilot deployment of DreamTales, we demonstrate how shifting AI support from real-time intervention to pre-session preparation and post-session reflection better aligns with parental values of agency, intimacy, and emotional presence. A key insight is that personalized storytelling is not merely a content delivery mechanism, but a central relational practice in itself—with the process of co-constructing story inputs fostering anticipatory reflection, priming emotional attunement, and creating shared affective ground for parent-child interaction. This reframes AI-supported storytelling systems from tools that optimize momentary interaction to infrastructures that sustain meaning-making across time.

Building on these findings, our future work will center on two interconnected goals: isolating the relational impact of personalized story creation and validating the Reflective Companion model at scale. First, we will conduct a comparative study contrasting personalized, co-constructed stories with non-personalized baseline stories, treating the story creation process itself as both an intervention and an evaluative lens. This study will examine differential impacts on three core outcomes: parental agency and reflective capacity (measured via post-session interviews and PRF-based qualitative coding), parent-child emotional synchrony (operationalized through vocal synchrony indices and temporal alignment metrics), and subjective experiences of intimacy and naturalness (assessed via parent self-reports and thematic analysis). By reframing the post-story reflective dashboard as an evaluation instrument rather than a core feature, we will isolate how preparatory reflection, narrative relevance, and shared meaning-making shape emotional dynamics.

Second, we will address the practical and technical constraints identified in the pilot—optimizing story length and linguistic complexity for 3–6 year-olds, and refining the multimodal analysis pipeline to better adapt to noisy, high-pitched caregiving environments. Beyond technical improvements, we aim to expand the participant pool to include diverse family structures and cultural backgrounds, ensuring the model's relevance across varied caregiving contexts.

More broadly, this work points toward a design agenda for relational AI systems in caregiving: one that prioritizes when and how AI participates—acting as a catalyst for human reflection, agency, and connection rather than a real-time actor. By centering personalized, reflective support that respects the integrity of parent-child dyads, we open new pathways for AI to enhance, rather than disrupt, the intimate relational practices that underpin early childhood development.

References

[1] Mary S Ainsworth. 1989. Attachments beyond infancy. *American psychologist* 44, 4 (1989), 709.

[2] Sharifa Alghowinem, Huili Chen, Cynthia Breazeal, and Hae Won Park. 2021. Body gesture and head movement analyses in dyadic parent-child interaction as indicators of relationship. In *2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*. IEEE, 01–05.

[3] John Bowlby. 1969. *Attachment and loss*. Number 79. Random House.

[4] Jude Cassidy and Phillip R Shaver. 1999. *Handbook of attachment: Theory, research, and clinical applications*. Vol. 236. Rough Guides.

[5] John Danaher. 2019. *Automation and utopia: human flourishing in a world without work*. Harvard University Press.

[6] Valdemar Danry, Pat Pataranutaporn, Christopher Cui, Jui-Tse Hung, Lancelot Blanchard, Zana Buçinca, Chenhao Tan, Thad Starner, and Pattie Maes. 2025. Human-AI Interaction for Augmented Reasoning: Improving Human Reflective and Critical Thinking with Artificial Intelligence. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–6.

[7] Vedant Das Swain, Qiuyue" Joy" Zhong, Jash Rajesh Parekh, Yechan Jeon, Roy Zimmermann, Mary P Czerwinski, Jina Suh, Varun Mishra, Koustuv Saha, and Javier Hernandez. 2025. Ai on my shoulder: Supporting emotional labor in front-office roles with an llm-based empathetic coworker. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–29.

[8] Griffin Dietz Smith, Siddhartha Prasad, Matt J Davidson, Leah Findlater, and R Benjamin Shapiro. 2024. Contextq: Generated questions to support meaningful parent-child dialogue while co-reading. In *Proceedings of the 23rd Annual ACM Interaction Design and Children Conference*. 408–423.

[9] Peter Fonagy and Mary Target. 1997. Attachment and reflective function: Their role in self-organization. *Development and psychopathology* 9, 4 (1997), 679–700.

[10] Jocelyn Shen, Ying Li, Javaria Hassan, Sharifa Alghowinem, Hae Won Park, Cynthia Breazeal, and Rosalind Picard. 2023. Fostering parent-child interactions through behavioral understanding of synchrony. In *2023 11th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)*. IEEE, 1–4.

[11] Jocelyn J Shen, Jennifer King Chen, Leah Findlater, and Griffin Dietz Smith. 2025. eaSEL: Promoting Social-Emotional Learning and Parent-Child Interaction through AI-Mediated Content Consumption. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–18.

[12] Arietta Slade. 2005. Parental reflective functioning: An introduction. *Attachment & human development* 7, 3 (2005), 269–281.

[13] Feiwen Xiao, Ellen Wenting Zou, Jiaju Lin, Zhaohui Li, and Dandan Yang. 2025. Parent-led vs. AI-guided dialogic reading: Evidence from a randomized controlled trial in children's e-book context. *British Journal of Educational Technology* 56, 5 (2025), 1784–1813.

[14] Zheng Zhang, Ying Xu, Yanhao Wang, Bingsheng Yao, Daniel Ritchie, Tongshuang Wu, Mo Yu, Dakuo Wang, and Toby Jia-Jun Li. 2022. Storybuddy: A human-ai collaborative chatbot for parent-child interactive storytelling with flexible parental involvement. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–21.

Received 22 January 2026