

The Eurasia Proceedings of Educational and Social Sciences (EPESS), 2026

Volume 53, Pages 97-110

ICETI 2026: International Conference on Education in Technology and Innovation

AI Behind the Adult, Not in Front of the Child: A Design Framework for Adult-Mediated AI in Early Childhood

Elif Ozturk

Giresun University

Abstract: Artificial intelligence (AI) in education has largely been positioned as a child-facing tutor, conversational agent, or content-delivery system. In early childhood, this orientation warrants particular caution because executive functions, self-regulation, and social-emotional competencies develop through responsive relationships, play, embodied experience, and co-regulation with adults. This theory-informed conceptual design study addressed three questions: which developmental, pedagogical, and ethical principles should structure AI use in early childhood; how these principles can be translated into a system architecture; and what validation pathway is required before such a system can make defensible claims. Drawing on four evidence streams (executive-function and social-emotional development, ecological and embodied learning, human-centered AI ethics, and measurement and implementation science), the study proposes the Behind-the-Adult AI framework, an adult-mediated design framework initially specified for children aged 4–6 years. Six design principles were derived and operationalized in a layered architecture: brief, bounded, non-conversational digital developmental tasks generate behavioral indicators under transparent adaptation rules; a provisional, non-diagnostic profile organizes longitudinal information; and AI operates only behind the adult, synthesizing observations from center, home, and school into hypotheses for human review. A five-stage validation pathway with explicit decision gates separates design claims from evidence claims. The paper argues that developmentally responsible AI in early childhood should be evaluated not by how convincingly it replaces human interaction but by how effectively it strengthens the adults and environments surrounding the child, and it offers a transferable design heuristic for researchers, developers, and policymakers.

Keywords: Artificial intelligence, Early childhood education, Executive functions, Co-regulation, Human-centered design

Introduction

Artificial intelligence (AI) is entering early childhood education (ECE) before the field has reached consensus on what an appropriate role for AI should be. Current applications include adaptive platforms, conversational agents, social robots, automated assessment, generative content, and teacher-facing analytics. Reviews spanning the past five years describe meaningful possibilities alongside unresolved developmental, ethical, and governance concerns, and they consistently note that the evidence base is young, methodologically heterogeneous, and short on longitudinal studies (Berson et al., 2025; Chen, 2025; Ljungcrantz, 2026; Su & Yang, 2022). The central design question is therefore not simply whether AI can be used with young children, but which functions should be delegated to AI, which should remain human, and how the boundary between the two should be engineered.

This distinction matters because early learning is not a scaled-down version of school-age instruction. Young children are still building regulatory and representational capacities on which later learning depends. Executive functions (EFs)—particularly inhibitory control, working memory, and cognitive flexibility—enable children to hold goals in mind, resist prepotent responses, shift between rules, and coordinate behavior in context (Diamond, 2013; Miyake et al., 2000). Together with emotional and attentional regulation, these capacities underpin school readiness (Blair & Raver, 2015) and predict concurrent and later social and behavioral outcomes (Stucke &

- This is an Open Access article distributed under the terms of the Creative Commons Attribution-Noncommercial 4.0 Unported License, permitting all non-commercial use, distribution, and reproduction in any medium, provided the original work is properly cited.

- Selection and peer-review under responsibility of the Organizing Committee of the Conference

© 2026 Published by ISRES Publishing: www.isres.org

Doebel, 2024). Crucially, they develop through repeated use in socially and emotionally meaningful activity: adults scaffold children's problem solving (Vygotsky, 1978; Wood et al., 1976), and warm, responsive co-regulation precedes and supports self-regulation (Rosanbalm & Murray, 2017). High-quality preschool curricula accordingly depend on responsive pedagogy, active engagement, and adaptation to children's individual, cultural, and linguistic contexts (National Academies of Sciences, Engineering, and Medicine, 2024).

The rapid growth of child-facing generative systems sharpens this tension. A University of Cambridge study of generative-AI toys used by children under five documented difficulties with pretend play, multi-party social reciprocity, and emotionally contingent responses, indicating that fluent language generation should not be equated with developmentally adequate participation in play or relationships (Goodacre & Gibson, 2026). The picture is not uniformly negative: in a randomized experiment with children aged 3–6 years, dialogic questioning delivered by a conversational agent improved story comprehension much as a human partner did, although children's communication patterns with the two partners differed (Xu et al., 2021). At the same time, analyses of commercially available children's apps have shown that educational claims frequently outpace learning science (Hirsh-Pasek et al., 2015) and that manipulative design features intended to prolong engagement are widespread (Radesky et al., 2022). The issue is therefore less whether AI can produce benefits than which interfaces and forms of authority are developmentally defensible.

Policy is moving in the same direction. UNICEF's updated guidance calls for AI that protects children's privacy, safety, fairness, transparency, development, and well-being (UNICEF, 2025). England's Department for Education now expects educational AI products to mitigate cognitive deskilling, avoid giving complete solutions by default, refrain from anthropomorphic or engagement-maximizing design, involve educators and child-development experts, and support every claim with transparent evidence (Department for Education [DfE], 2026). The European Union's Artificial Intelligence Act goes further by prohibiting AI systems that infer emotions in educational institutions, except for medical or safety purposes (Regulation (EU) 2024/1689, 2024). Taken together, these developments point toward a design agenda in which developmental appropriateness is treated as an architectural constraint rather than a post hoc safety feature.

General frameworks for human-centered AI already emphasize human control, reliability, transparency, and accountability (Amershi et al., 2019; Shneiderman, 2020), and community-wide work on the ethics of AI in education has articulated principles for responsible use (Fu & Weng, 2024; Holmes et al., 2022). However, these frameworks were largely developed for adult users or older learners. ECE reviews, in turn, identify risks but rarely specify a concrete architecture that states where AI should sit relative to the child, what it may and may not decide, and what evidence must precede any claim of benefit (Berson et al., 2025; Ljungcrantz, 2026). Earlier work by the author argued that AI in early childhood STEM education should be human-centered, ethically constrained, ecologically situated, and linked to teacher and parent capacity rather than treated as an autonomous substitute for pedagogy (Ozturk, 2025). The present study extends that position into a testable system architecture.

Theoretical and Empirical Background

Executive Functions and Self-Regulation as Developmental Capacities

The unity-and-diversity account distinguishes separable but correlated EF processes such as inhibition, updating, and shifting (Miyake et al., 2000), while developmental syntheses identify inhibitory control, working memory, and cognitive flexibility as core EFs that support reasoning, planning, and goal-directed behavior (Diamond, 2013). A complementary distinction between relatively "cool" EFs, deployed in affectively neutral problems, and "hot" EFs, deployed in motivationally and emotionally significant situations, is particularly relevant in early childhood, when both are changing rapidly and are highly sensitive to context (Zelazo & Carlson, 2012). Emotional competence in the preschool years, in turn, contributes to later social competence (Denham et al., 2003), and well-implemented social and emotional learning programs improve skills, behavior, and achievement in school-age samples (Durlak et al., 2011).

Two measurement implications follow. First, early EF performance is better represented as a pattern over time than as a fixed trait: response accuracy alone can conceal meaningful differences between a child who responds slowly and carefully, a child who responds impulsively, and a child whose performance becomes unstable across a session. Second, there is a principled argument that EF batteries in early childhood may function better as formative composites of task performance than as indicators of a single reflective latent construct (Willoughby et al., 2016). Both points caution against reducing young children's performance on digital tasks to a single global score.

Transfer, Embodiment, and the Limits of Task Training

The case for an ecological rather than task-centric model is strengthened by intervention evidence. Meta-analyses indicate that working-memory training produces gains on trained or closely related tasks but little reliable far transfer (Melby-Lervåg et al., 2016; Sala & Gobet, 2017). In the preschool period specifically, Trujillo-Trujillo et al. (2026) reviewed 35 studies of EF interventions in typically developing children aged 3–6 years; across the 17 randomized trials included in their meta-analysis, effects were small for verbal working memory ($d = 0.13$) and inhibitory control ($d = 0.10$) and almost null for nonverbal working memory ($d = 0.02$). These findings do not imply that structured EF practice is ineffective; they caution against assuming that improvement on trained tasks will yield broad developmental transfer.

By contrast, a meta-analysis of fundamental motor-skill interventions in 10 studies involving 2,039 preschool children reported a moderate effect on EFs (SMD = 0.40) and a smaller effect on social-emotional competence (SMD = 0.16; Li et al., 2026). Diamond and Lee (2011) likewise concluded that EF gains depend on repeated practice with progressively increasing challenge and that approaches which also address children's emotional, social, and physical development may be especially effective. Although these literatures are not directly comparable, the contrast is consistent with an embodied and relational view of early learning in which movement, regulation, social participation, and cognitive control operate together.

Relational and Ecological Foundations

Bronfenbrenner's (1979) ecological perspective directs attention to interactions among the settings in which a child lives, rather than to the child in isolation, and the Nurturing Care Framework treats responsive caregiving, safety, and opportunities for early learning as interdependent conditions for development (World Health Organization [WHO] et al., 2018). Neurobiological evidence offers a converging, if preliminary, perspective: in parent-child dyads, prefrontal brain-to-brain synchrony emerged during cooperation but not competition, predicted subsequent cooperative performance, and mediated the association between parent and child emotion regulation (Reindl et al., 2018). Such findings should not be overinterpreted as causal prescriptions, but they reinforce a broader principle: the relational context is not noise around learning; it is part of the developmental system.

Human-Centered AI and the Early Childhood Boundary Problem

Human-centered AI frameworks argue that high levels of automation and high levels of human control can be combined, and that systems should make clear what they can do, how well, and when humans should intervene (Amershi et al., 2019; Shneiderman, 2020). For young children, these principles require a developmental interpretation. Human oversight is not only an accountability mechanism; it is part of the learning mechanism, because adult timing, contingent language, emotional availability, and scaffolding organize children's attention and regulation. Moreover, a common ingredient of child-facing "affective" AI—inferring emotion from facial or vocal signals—rests on weak scientific ground, since facial movements do not map reliably onto specific emotional states across people and contexts (Barrett et al., 2019). The boundary problem in early childhood AI is therefore twofold: deciding which interface the child should encounter and deciding how much interpretive authority a computational system should hold over a developing child.

Purpose and Research Questions

The purpose of this study was to develop and justify the Behind-the-Adult AI framework, a theory-informed design framework in which AI supports the adults and environments around young children rather than acting as the child's interlocutor. The framework separates two functions that are often collapsed in adaptive education: (a) bounded, non-conversational, child-facing digital developmental tasks that generate behavioral indicators under transparent adaptation rules, and (b) adult-facing AI that integrates longitudinal and cross-context information and generates hypotheses for human review. Three research questions guided the study:

- RQ1: What developmental, pedagogical, and ethical principles should structure AI use in early childhood?
- RQ2: How can these principles be translated into an architecture that links bounded digital behavioral indicators with adult-mediated support across center, home, and school?

- RQ3: What validation pathway is required before such a system can make defensible claims about measurement, personalization, or developmental effectiveness?

The framework is currently pre-efficacy, and no claim is made that systems built on it improve child outcomes. Its contribution at this stage is conceptual and architectural. The initial specification targets children aged 4–6 years, although the principles are intended to be extensible to the wider early childhood and early primary period.

Method

Research Design

This study is a theory-informed conceptual design study rather than an efficacy trial or systematic review. Within the typology of conceptual articles proposed by Jaakkola (2020), it combines a theory synthesis, which integrates previously separate literatures, with a model approach, which specifies relationships among constructs in a form that can be tested. Conceptual framework analysis is appropriate when a problem spans multiple bodies of knowledge and the aim is to articulate linked concepts into a coherent, revisable structure (Jabareen, 2009). Methodologically, the work corresponds to the analysis-and-exploration and design phases of educational design research (McKenney & Reeves, 2019; Wang & Hannafin, 2005) and to the development phase of complex-intervention research, in which program theory is articulated before feasibility and evaluation (Skivington et al., 2021).

Evidence Streams and Selection

The synthesis drew purposively on four evidence streams: (a) EF, self-regulation, and social-emotional development; (b) ecological, relational, embodied, and play-based early learning; (c) AI in ECE and responsible human-centered AI; and (d) measurement, psychometric validation, and implementation science. Within each stream, peer-reviewed systematic reviews and meta-analyses were prioritized where available, particularly those published since 2020; foundational theoretical and developmental works were retained when they defined constructs or paradigms used in the architecture. Policy and rights-based documents from UNICEF, WHO, the DfE, and the European Union were used to define safety and governance constraints, not to establish intervention efficacy. Because the aim was conceptual integration rather than exhaustive estimation, selection was theory-driven rather than systematic, a choice that is acknowledged as a limitation.

Design Synthesis Procedure

External evidence was translated iteratively into an internal design corpus comprising a theoretical framework, age-specific developmental goals, a 12-week curriculum architecture, specifications for EF digital tasks, a social-emotional learning extension, and an observation and assessment system. These documents were developed for reference implementation within NEUROSEEDU, an early childhood development program founded by the author (see Conflict of Interest). They functioned as design artifacts, not as empirical data, and program-specific details are reported here only insofar as they illustrate how the framework's principles can be operationalized. The synthesis proceeded in four steps: (1) defining the developmental and ethical design problem; (2) mapping evidence-based propositions onto design principles and requirements (RQ1); (3) translating requirements into system components, information flows, and boundaries (RQ2); and (4) specifying falsifiable validation questions and decision gates that must be passed before efficacy claims are made (RQ3).

Quality and Transparency Criteria

Following recommendations for conceptual research (Jaakkola, 2020; Jabareen, 2009), four criteria guided the work. *Traceability* required that every design decision be linked to an explicit evidence-informed proposition (Table 1). *Internal coherence* required that components, information flows, and ethical boundaries be mutually consistent (Figures 1 and 2; Table 3). *Falsifiability* required that each architectural claim be paired with a testable validation question (Table 4). *Reflexivity* required that the author's dual role as researcher and founder of the program be disclosed and that design claims be explicitly separated from evidence claims throughout the paper.

Results and Discussion

The results are organized by research questions. The first section presents the design principles derived from the synthesis (RQ1); the second describes how these principles were operationalized in the framework's architecture, including its ethical boundaries (RQ2); and the third sets out the staged validation pathway (RQ3). A discussion of the framework's contributions in relation to literature and its limitations follows.

Design Principles for AI in Early Childhood (RQ1)

Six design principles emerged from the synthesis. *Principle 1: Developmental work stays with the child.* Digital components should preserve the cognitive effort through which EFs develop rather than offloading it, consistent with evidence that EF development depends on repeated, progressively challenging practice (Diamond & Lee, 2011) and with policy concerns about cognitive deskilling (DfE, 2026). *Principle 2: Relational work stays with adults.* Functions that depend on contingent, emotionally attuned interaction—co-regulation, scaffolding, pretend play—should be supported rather than simulated (Goodacre & Gibson, 2026; Rosanbalm & Murray, 2017; Wood et al., 1976). *Principle 3: Measure patterns, not traits.* Indicators should be interpreted longitudinally and across tasks and informants, because early EFs are context-sensitive and may be better represented formatively than as a single latent trait (Willoughby et al., 2016; Zelazo & Carlson, 2012).

Principle 4: Design for transfer, not task gains. Because trained-task improvements generalize weakly (Melby-Lervåg et al., 2016; Trujillo-Trujillo et al., 2026), digital activity should be embedded in movement, play, inquiry, and family routines, and transfer should be measured rather than assumed (Li et al., 2026). *Principle 5: Bound the interface and the authority.* Children should not engage in open-ended conversation with generative AI; systems should not infer emotion from biometric signals, diagnose, label, or prescribe autonomously; and design should not optimize for time-on-device (Barrett et al., 2019; Radesky et al., 2022; Regulation (EU) 2024/1689, 2024; UNICEF, 2025). *Principle 6: Separate design claims from evidence claims.* Claims about measurement, personalization, or developmental benefit should be made only after the corresponding validation stage has been completed (DfE, 2026; Hirsh-Pasek et al., 2015; Skivington et al., 2021). Table 1 summarizes how each principle was translated into a design response and a validation implication.

Table 1. Evidence-informed design principles and design responses

Principle	Evidence-informed proposition	Design response	Validation implication
1. Developmental work stays with the child	EFs develop through repeated, progressively challenging use (Diamond, 2013; Diamond & Lee, 2011)	Tasks preserve cognitive demand; no hints or solutions that remove the child's effort	Test whether adaptation keeps children within an interpretable challenge range
2. Relational work stays with adults	Co-regulation and scaffolding drive early regulation; GenAI toys struggle with play and reciprocity (Goodacre & Gibson, 2026; Rosanbalm & Murray, 2017)	AI is adult-facing only; no child–AI conversation; adults remain the relational partner	Audit for displacement of adult–child interaction; collect practitioner and caregiver accounts
3. Measure patterns, not traits	EFs are separable, context-sensitive, and possibly formative (Miyake et al., 2000; Willoughby et al., 2016)	Multiple short tasks; provisional, non-diagnostic profile instead of a global score	Examine reliability, age sensitivity, and covariance structure before naming dimensions
4. Design for transfer, not task gains	Preschool EF training effects are small; motor and multimodal approaches show larger effects (Li et al., 2026; Trujillo-Trujillo et al., 2026)	Digital tasks embedded in center, home, and school activity with shared language	Evaluate cross-context convergence and far transfer, not task improvement alone
5. Bound the interface and the authority	Emotion inference is unreliable; manipulative design is common; regulation restricts educational AI (Barrett et al., 2019; Radesky et al., 2022; Regulation (EU) 2024/1689, 2024)	No biometric emotion recognition, diagnosis, labeling, autonomous prescription, or engagement maximization	Audit privacy, fairness, explainability, and human oversight with stakeholders
6. Separate design claims from evidence claims	Educational technology claims often outpace evidence (DfE, 2026; Hirsh-Pasek et al., 2015)	Staged validation with explicit decision gates	No efficacy claim before controlled evaluation (Table 4)

Note. EF = executive function; GenAI = generative artificial intelligence.

Architecture of the Behind-the-Adult AI Framework (RQ2)

Embedded Co-Regulation Ecology

The framework uses the term *embedded co-regulation ecology* to describe its central design proposition: developmental support should be embedded in the environments where children regulate, play, communicate, and learn, and those environments should share a coherent support language. Center, home, and school are therefore treated as connected developmental contexts rather than independent service channels. The center provides structured, adult-guided sessions; the home receives brief micro-routines that fit existing family life; and, with family consent, the school contributes limited observations and receives concise bridge messages. The purpose is not to standardize the child across settings but to reduce fragmentation and make transfer observable, in line with ecological and nurturing-care perspectives (Bronfenbrenner, 1979; WHO et al., 2018).

The wider system comprises six mutually supporting components: Center, Family, School, Data, Science, and Education. The first three are developmental contexts; Data supports multi-source observation; Science governs validation and revision; and Education refers to practitioner preparation and supervision. As Figure 1 shows, this arrangement deliberately locates AI within an adult-facing support layer rather than within the child-adult relational channel.




Figure 1. Embedded co-regulation ecology of the framework

Note. The center–home–school triangle forms the developmental surface of the framework. The support layer informs adults only; arrows indicate reciprocal information and transfer, not automated control.

Bounded Digital Developmental Tasks

The child-facing component of the framework consists of *digital developmental tasks*: brief, game-like activities whose scientific function is measurement and practice rather than open-ended entertainment. In the reference implementation for children aged 4–6 years, the proposed task set includes 14 EF-oriented tasks and five social-emotional tasks. The EF set samples inhibition, working memory, cognitive flexibility, sustained and selective attention, planning, and dual-task coordination. Each task is built around a recognizable experimental paradigm—for example, go/no-go and stop-signal logic for inhibition, a dimensional-change rule-switching structure for flexibility (Zelazo, 2006), and a child-adapted flanker structure for selective attention (Rueda et al., 2004)—with a small set of observable indicators, bounded session duration, and explicit difficulty parameters (Table 2).

The aim is not to reproduce laboratory tests on a tablet but to preserve their cognitive demand in age-appropriate, low-language, culturally reviewable interactions. Prior work has shown that brief, tablet-based EF and self-regulation tasks can be administered to preschool children with acceptable validity and reliability (Howard & Melhuish, 2017), which supports the feasibility of this measurement layer while leaving its validity in the present context to be established.

The social-emotional extension adds five domains: emotion knowledge, emotional self-awareness, theory of mind and perspective taking, emotion regulation, and empathy and prosocial decision-making. These domains are sampled through repeated scenario-based choices rather than through generative conversation, child self-disclosure, or biometric emotion detection. Because social-emotional behavior is strongly contextual and culturally shaped, these indicators are treated more cautiously than EF metrics and require convergent validation against external measures and human observations.

Table 2. Proposed architecture of digital developmental tasks

Domain	Representative paradigm or task logic	Illustrative behavioral indicators	Adaptation principle
Inhibition	Go/no-go; stop-signal; interference control	Accuracy; commission and omission errors; response time; stop success; interference cost	Adjust pacing, cue salience, and distractor or stop demands within predefined limits
Working memory	Visual-spatial sequence or span; ordered recall	Span length; serial-position errors; recall latency; sequence accuracy	Adjust sequence length, presentation support, and response demand
Cognitive flexibility	Rule switching; dimensional change	Switch accuracy; switch cost; perseverative errors	Increase rule complexity only after stable performance; provide non-punitive scaffolds
Attention	Sustained attention; flanker-type selective attention	Hit and false-alarm rates; omissions; RT variability; vigilance decrement; congruency cost	Adjust session length, pace, and distractor load without optimizing for prolonged engagement
Planning and dual-task	Multi-step planning; coordinated inhibition and working memory	Planning latency; excess moves; completion efficiency; dual-task cost	Increase steps or concurrent load only after prerequisite performance is stable
Social-emotional learning	Emotion in context; perspective taking; regulation choices; prosocial scenarios	Choice appropriateness; response latency; confusion patterns; strategy diversity	Adjust scenario complexity and support; never frame emotion or developmental variation as pathological

Note. RT = response time. Indicators are design specifications awaiting psychometric evaluation, not validated scores.

A Provisional Dynamic Developmental Profile

Across repeated sessions, task indicators are summarized into a provisional dynamic profile rather than a fixed score. The reference specification contains 11 dimensions: inhibition, working memory, cognitive flexibility, attention, learning rate, self-regulation, emotion knowledge, emotional self-awareness, theory of mind, emotion regulation, and empathy and prosocial decision-making. The term *profile* is intentionally operational: at this stage, these dimensions are design variables, not validated latent traits, clinical constructs, or norms. A low value in one domain is not interpreted as a deficit; it is a prompt to examine whether the pattern is stable over time, appears across tasks, and converges with observations from adults who know the child in natural contexts. Likewise, strong digital performance does not establish generalization. The profile is designed to organize longitudinal evidence and support questions, not to close them.

AI Behind the Adult: Three Constrained Functions

The framework reserves AI for three adult-facing functions. First, AI can support caregivers with brief, contextually timed prompts derived from current support goals and previously recorded family responses—for example, a reminder of a transition phrase or co-regulation routine that has previously been useful. These prompts are not intended to replace parental judgment or to provide continuous behavioral coaching. Second, AI can synthesize multi-week practitioner notes and structured indicators so that longitudinal patterns become easier to notice; an output might flag that rule-switching difficulty has recurred across sessions while reflective language has increased. Such outputs are framed explicitly as hypotheses, and the practitioner decides whether a pattern is meaningful, whether contextual explanations are more plausible, and whether support should change. Third, AI

can help translate multi-source information into a family-readable developmental summary at the end of a support cycle, under data-minimization, consent, and interpretability requirements. The system is not authorized to produce diagnostic labels, clinical risk scores, or autonomous intervention prescriptions. Figure 2 shows how these functions are sequenced so that human judgment remains the final authority.




Figure 2. Human-in-the-loop personalization cycle

Note. Transparent adaptation rules operate within the bounded digital tasks; AI is used only for longitudinal and cross-context synthesis for adult review, which prevents an algorithmic output from becoming a final developmental judgment.

From Digital Indicators to Real-World Transfer

In the reference implementation, the digital layer is embedded within a 12-week program. For the 4–6-year group, a center session lasts approximately 55 minutes and is followed by a brief caregiver bridge. Sessions follow a predictable sequence—welcome and safety, body and rhythm, core module activity, reflection and verbalization, and closure with a family bridge—and the curriculum moves through phases of foundation, skill activation, reinforcement and transfer, and synthesis. EF goals are therefore practiced in movement, rule-based play, inquiry, social interaction, and reflective language, not only in digital tasks. This creates a deliberately demanding criterion for personalization: a digital pattern is valuable only if it helps adults ask better questions and adjust support in ways that are developmentally plausible and observable beyond the device. The unit of personalization is thus not the next item on a screen but the ecology of support surrounding the child.

Ethical Design Boundaries

The framework's ethical position is not that AI should be absent from ECE, but that its interface and authority should be developmentally bounded. The distinction between child-facing digital tasks and child-facing generative AI is central: children may complete brief digital tasks, but they do not converse with a generative system. No design goal seeks to maximize time-on-device; sessions are intentionally finite, and adaptation maintains an interpretable challenge range rather than prolonging engagement. These boundaries are consistent with child-centered guidance on safety, privacy, non-discrimination, transparency, and well-being (UNICEF, 2025), with expectations to mitigate cognitive deskilling, avoid manipulative engagement, and preserve pathways to human help (DfE, 2026), and with ECE-specific concerns about privacy, developmental impact, bias, and regulatory gaps (Berson et al., 2025). Table 3 operationalizes the boundaries as a matrix of permitted and prohibited functions.

Table 3. Human-centered AI boundary matrix

Permitted or intended AI function	Prohibited or out-of-scope function	Rationale
Summarize longitudinal, structured observations for adults	Diagnose a developmental, psychological, or medical condition	Pattern detection is not clinical inference; diagnostic validity is not established
Generate hypotheses for practitioner review	Prescribe or change an intervention without human approval	Preserves professional judgment and keeps responsibility traceable
Personalize brief caregiver prompts within approved content boundaries	Replace caregiver or teacher interaction, or present itself as the child's relational partner	Protects co-regulation and relational learning
Support cross-context synthesis and plain-language reporting	Assign fixed labels, intelligence categories, or deficit identities	Development is dynamic and context-sensitive; labels risk stigma and construct overreach
Use bounded task data such as choices, accuracy, timing, and error patterns	Use open-ended child-generative-AI conversation or infer emotion from facial or vocal data	Reduces exposure to unpredictable content, relationship simulation, and scientifically weak inference
Adjust transparent task parameters within predefined developmental limits	Optimize for session length, compulsive engagement, advertising, or dark patterns	The objective is developmental information and practice, not engagement maximization

Validation Pathway and Decision Gates (RQ3)

Developmentally responsible architecture is not validated by theoretical coherence alone. Consistent with an argument-based approach to validity, in which each intended interpretation and use of a score must be supported by evidence (Kane, 2013), and with professional testing standards (American Educational Research Association et al., 2014), the framework separates design claims from evidence claims and proposes a five-stage validation program with explicit decision gates (Table 4).

Table 4. Staged validation pathway and decision gates

Stage	Primary question	Illustrative evidence	Decision gate before progression
1. Content, ethics, and usability	Are tasks, scenarios, interfaces, and AI outputs developmentally appropriate, interpretable, culturally reviewable, and safe?	Expert review; cognitive interviews; child and adult usability testing; accessibility and risk assessment	Revise any construct–task mismatch, foreseeable harm, or comprehension problem
2. Measurement validation	Do indicators show acceptable reliability, developmental sensitivity, and convergence without overclaiming latent constructs?	Test–retest and interrater reliability; external criteria; EFA/CFA where justified; invariance and fairness checks	No diagnostic or efficacy interpretation until psychometric evidence supports the intended use
3. Feasibility and implementation	Can the 12-week ecology be delivered with fidelity and acceptable burden across center, home, and school?	Recruitment, retention, fidelity, dosage, usability, qualitative interviews; pre–mid–post and follow-up data	Proceed only if implementation and data quality are adequate
4. Controlled effectiveness	Does the integrated model improve outcomes beyond comparison conditions, and do gains transfer?	Waitlist or active comparison; later cluster RCT; effect sizes; transfer; follow-up	Claims limited to demonstrated outcomes, populations, and contexts
5. Scale and continuous audit	Does performance remain safe, fair, and effective across settings and populations?	Post-deployment monitoring; fairness audits; privacy review; output-drift review; replication	Pause or revise functions when harms, drift, or inequities emerge

Note. EFA = exploratory factor analysis; CFA = confirmatory factor analysis; RCT = randomized controlled trial.

The first stage is content, ethics, and usability validation. Early childhood educators, developmental psychologists, measurement specialists, child-safety and ethics experts, and AI specialists evaluate construct–task alignment, developmental appropriateness, cultural fit, wording, visual demands, and foreseeable harms, while usability testing with children and adults examines comprehension, burden, accessibility, and unintended strategies rather than satisfaction alone. The second stage is measurement validation. Analyses should examine test–retest stability where theoretically appropriate, internal consistency only where the item structure justifies it, measurement error, age sensitivity, and convergence with established performance-based and rating measures. The 11-dimensional profile should not be submitted directly to confirmatory analysis as if its specified dimensions were established factors; instead, the empirical covariance structure—and the possibility of formative rather than reflective measurement (Willoughby et al., 2016)—should determine whether the proposed dimensions are distinguishable, redundant, or in need of revision.

The third stage is feasibility and implementation research across the full ecology, addressing acceptability, demand, implementation, practicality, and adaptation (Bowen et al., 2009). The 12-week program uses pre-, mid-, and post-program checkpoints, with follow-up where feasible; practitioner observations, caregiver reports, child-appropriate reflection, and optional teacher observations are triangulated rather than treated as interchangeable. Fidelity data are essential because a null finding is uninterpretable if the center, family, or digital components were not delivered as intended. Only after feasibility and measurement validity are established should controlled effectiveness studies test developmental outcomes, progressing from a waitlist-controlled pilot to adequately powered cluster-randomized designs. Analyses should report effect sizes with confidence intervals, examine differential effects and attrition, and test transfer beyond trained tasks. Finally, algorithmic fairness analyses should assess whether task performance or AI-generated flags differ systematically by language background, gender, socioeconomic context, disability status, or device characteristics, and post-deployment monitoring should continue at scale.

Discussion

With respect to RQ1, the synthesis indicates that the principles most relevant to early childhood AI are not primarily technical but developmental: preserving the child's cognitive effort, preserving adults' relational role, interpreting performance as provisional patterns, designing for transfer, bounding both interface and authority, and separating design from evidence. These principles extend general human-centered AI frameworks (Amershi et al., 2019; Holmes et al., 2022; Shneiderman, 2020) by specifying that, for young children, human oversight is part of the learning mechanism rather than an external safeguard. They also answer a recurring call in ECE reviews to move from listing risks to specifying concrete design constraints (Berson et al., 2025; Ljungcrantz, 2026; Su & Yang, 2022).

With respect to RQ2, the framework makes three architectural contributions. First, it reframes the role of AI. Most debates begin with the child–AI interface—how well an agent can tutor, converse, recognize emotion, or sustain engagement. The Behind-the-Adult AI framework begins instead with the adult–child developmental ecology and asks where computational support can add value without occupying a relational role. Generative and analytic intelligence is therefore shifted away from the child-facing surface toward adult-facing synthesis, while short, bounded digital tasks are retained where precise behavioral indicators are useful. Evidence that conversational agents can support specific outcomes (Xu et al., 2021) is not inconsistent with this choice; rather, the framework sets a higher burden of proof for functions that may displace co-regulation, social reciprocity, pretend play, or adult observation, whose developmental consequences of substitution remain uncertain (Goodacre & Gibson, 2026). Second, the framework separates adaptive task rules from AI decision support. Not every adaptive feature requires generative AI, and not every AI output should directly change a child's experience. In the framework, trial- and level-based adjustments are transparent and constrained, whereas AI operates at a different layer to synthesize patterns for human interpretation. This layering improves inspectability and makes it easier to identify which part of a system is responsible for a given change or error, a property emphasized in guidelines for human–AI interaction (Amershi et al., 2019). Third, the framework treats transfer as a design target rather than an assumed consequence. The modest effects reported for preschool cognitive training (Trujillo-Trujillo et al., 2026) and the weak far transfer of working-memory training (Melby-Lervåg et al., 2016; Sala & Gobet, 2017) challenge the common inference that better task performance equals broader developmental improvement; linking digital evidence to adult-guided activity, family routines, and school observations makes that inference testable.

With respect to RQ3, the framework treats restraint as a scientific feature. The dynamic profile is explicitly provisional, AI produces hypotheses rather than decisions, social-emotional tasks require cultural validation, and no efficacy claim is made before controlled research. This stance responds to a recurring problem in educational technology, where product claims often mature faster than the evidence base (Hirsh-Pasek et al., 2015), and it aligns with policy that asks developers to state intended use, support claims with evidence, monitor developmental impacts, and avoid manipulative design (DfE, 2026). A research-first architecture can make these obligations part of the design process rather than an external compliance layer. More broadly, the framework suggests a transferable design heuristic: when a developmental function depends heavily on reciprocal human interaction, AI should preferentially support the human system around that interaction rather than simulate the relationship itself. This heuristic does not prohibit all child-facing AI; it allocates the burden of proof to those who would place AI in front of the child.

Limitations

Some limitations should be noted for this study. First, this is a conceptual design paper; it contains no child outcome data and cannot establish effectiveness. Second, the evidence synthesis was purposive and theory-driven rather than systematic, so relevant studies may have been omitted, and the 11-dimensional profile remains a design specification pending psychometric validation. Third, the framework was developed in Türkiye and incorporates culturally reviewable scenarios, but its cross-cultural portability is unknown; adaptation should be treated as revalidation, not translation alone. Fourth, digital behavioral indicators may be influenced by motor speed, device familiarity, language comprehension, fatigue, sensory preferences, or motivation, and multi-source triangulation reduces but does not eliminate such construct-irrelevant variance. Fifth, AI-assisted synthesis can introduce its own biases—over-weighting frequently recorded behaviors, reproducing linguistic bias in practitioner notes, or creating false coherence from sparse observations—so human review is necessary but not sufficient, and systematic audits are also required. Finally, the restriction on open-ended child–AI interaction is a normative design choice grounded in developmental caution, not evidence that all conversational AI is harmful; comparative research is needed to identify which child-facing AI functions, if any, can meet a sufficiently high developmental and ethical standard.

Conclusion

This study asked which principles should structure AI use in early childhood, how they can be translated into a system architecture, and what evidence must precede claims about such a system. In response to the first question, six principles were derived that keep developmental work with the child and relational work with adults, treat performance as provisional patterns, target transfer, bound the interface and authority of AI, and separate design from evidence. In response to the second, the Behind-the-Adult AI framework operationalizes these principles as a layered architecture: brief, bounded digital developmental tasks generate interpretable behavioral indicators; transparent rules adapt task difficulty; a provisional profile organizes longitudinal information; AI synthesizes cross-context patterns for adult review; and center, home, and school remain the primary sites of developmental transfer. In response to the third, a five-stage pathway with explicit decision gates specifies that content validity, measurement validity, and feasibility must be established before effectiveness is tested and before any claim of benefit is made.

The framework deliberately separates innovation from proof. Its next scientific task is not to claim that it works, but to test whether its measures are valid, its boundaries are safe, its implementation is feasible, and its ecological design produces meaningful transfer. Developmentally responsible AI in early childhood should ultimately be evaluated not by how convincingly it replaces human interaction, but by how effectively it strengthens the adults and environments surrounding the child.

Recommendations

On the basis of the framework, some recommendations are offered. Firstly, *researchers* should test the proposed validation pathway empirically, beginning with expert and usability studies and psychometric evaluation of the digital indicators, and should compare adult-mediated and child-facing AI configurations with outcomes that include transfer, adult–child interaction quality, and potential harms.

Secondly, *developers* of early childhood technologies should treat developmental appropriateness as an architectural constraint: separating transparent adaptation rules from AI decision support, avoiding biometric emotion inference and engagement-maximizing design, and publishing the evidence behind each product claim.

Thirdly, *practitioners and teacher educators* should be prepared to interpret AI-generated outputs as hypotheses rather than conclusions; pre-service and in-service programs should therefore include data literacy, ethical reasoning, and co-regulation practice.

Lastly, *policymakers* should require developers of AI for young children to specify intended use, the locus of human authority, and the stage of validation reached, and should support independent evaluation and fairness auditing before large-scale adoption.

Scientific Ethics Declaration

* The author declares that the scientific ethical and legal responsibility of this article published in EPESS Journal belongs to the author.

* This is a conceptual design study that did not involve human participants, identifiable personal data, or any experimental intervention; therefore, ethics committee approval was not required. Future empirical studies within the proposed validation pathway will be conducted with institutional ethics approval, informed parental consent, age-appropriate child assent, and data-protection procedures consistent with applicable legislation.

Conflict of Interest

* The author declares a potential conflict of interest: the author is the founder of NEUROSEEDU, an early childhood development program within which a reference implementation of the framework described in this article is being developed and is the originator of the framework. The article reports on a conceptual framework and does not report efficacy data for any program or product. The author commits to reporting future validation and effectiveness findings irrespective of whether they support the framework.

Funding

* The author received no external funding for the preparation of this article.

Acknowledgements or Notes

* This article was presented as an oral presentation at the International Conference on Education in Technology and Innovation (www.iceti.net) held in İstanbul/Türkiye on September 17-18, 2026.

References

- American Educational Research Association, American Psychological Association, & National Council on Measurement in Education. (2014). *Standards for educational and psychological testing*. American Educational Research Association.
- Amershi, S., Weld, D., Vorvoreanu, M., Fournery, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for human-AI interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (pp. 1–13). Association for Computing Machinery. <https://doi.org/10.1145/3290605.3300233>
- Barrett, L. F., Adolphs, R., Marsella, S., Martinez, A. M., & Pollak, S. D. (2019). Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. *Psychological Science in the Public Interest*, 20(1), 1–68. <https://doi.org/10.1177/1529100619832930>
- Berson, I. R., Berson, M. J., & Luo, W. (2025). Innovating responsibly: Ethical considerations for AI in early childhood education. *AI, Brain and Child*, 1(1), Article 2. <https://doi.org/10.1007/s44436-025-00003-5>
- Blair, C., & Raver, C. C. (2015). School readiness and self-regulation: A developmental psychobiological approach. *Annual Review of Psychology*, 66, 711–731. <https://doi.org/10.1146/annurev-psych-010814-015221>
- Bowen, D. J., Kreuter, M., Spring, B., Cofta-Woerpel, L., Linnan, L., Weiner, D., Bakken, S., Kaplan, C. P., Squiers, L., Fabrizio, C., & Fernandez, M. (2009). How we design feasibility studies. *American Journal of Preventive Medicine*, 36(5), 452–457. <https://doi.org/10.1016/j.amepre.2009.02.002>
- Bronfenbrenner, U. (1979). *The ecology of human development: Experiments by nature and design*. Harvard University Press.
- Chen, J. J. (2025). From Turing’s conception of machine intelligence to the evolution of AI in early childhood education: Conceptual, empirical, and practical insights. *AI, Brain and Child*, 1(1), Article 1. <https://doi.org/10.1007/s44436-025-00002-6>
- Denham, S. A., Blair, K. A., DeMulder, E., Levitas, J., Sawyer, K., Auerbach-Major, S., & Queenan, P. (2003). Preschool emotional competence: Pathway to social competence? *Child Development*, 74(1), 238–256. <https://doi.org/10.1111/1467-8624.00533>
- Department for Education. (2026, January 19). *Generative AI: Product safety standards*. GOV.UK. <https://www.gov.uk/government/publications/generative-ai-product-safety-standards>

- Diamond, A. (2013). Executive functions. *Annual Review of Psychology*, 64, 135–168. <https://doi.org/10.1146/annurev-psych-113011-143750>
- Diamond, A., & Lee, K. (2011). Interventions shown to aid executive function development in children 4 to 12 years old. *Science*, 333(6045), 959–964. <https://doi.org/10.1126/science.1204529>
- Durlak, J. A., Weissberg, R. P., Dymnicki, A. B., Taylor, R. D., & Schellinger, K. B. (2011). The impact of enhancing students' social and emotional learning: A meta-analysis of school-based universal interventions. *Child Development*, 82(1), 405–432. <https://doi.org/10.1111/j.1467-8624.2010.01564.x>
- Fu, Y., & Weng, Z. (2024). Navigating the ethical terrain of AI in education: A systematic review on framing responsible human-centered AI practices. *Computers and Education: Artificial Intelligence*, 7, Article 100306. <https://doi.org/10.1016/j.caeai.2024.100306>
- Goodacre, E. J., & Gibson, J. L. (2026). *AI in the early years: Examining the implications of GenAI toys for young children*. PEDAL Centre, University of Cambridge. <https://doi.org/10.17863/CAM.126270>
- Hirsh-Pasek, K., Zosh, J. M., Golinkoff, R. M., Gray, J. H., Robb, M. B., & Kaufman, J. (2015). Putting education in “educational” apps: Lessons from the science of learning. *Psychological Science in the Public Interest*, 16(1), 3–34. <https://doi.org/10.1177/1529100615569721>
- Holmes, W., Porayska-Pomsta, K., Holstein, K., Sutherland, E., Baker, T., Buckingham Shum, S., Santos, O. C., Rodrigo, M. T., Cukurova, M., Bittencourt, I. I., & Koedinger, K. R. (2022). Ethics of AI in education: Towards a community-wide framework. *International Journal of Artificial Intelligence in Education*, 32(3), 504–526. <https://doi.org/10.1007/s40593-021-00239-1>
- Howard, S. J., & Melhuish, E. (2017). An early years toolbox for assessing early executive function, language, self-regulation, and social development: Validity, reliability, and preliminary norms. *Journal of Psychoeducational Assessment*, 35(3), 255–275. <https://doi.org/10.1177/0734282916633009>
- Jaakkola, E. (2020). Designing conceptual articles: Four approaches. *AMS Review*, 10(1–2), 18–26. <https://doi.org/10.1007/s13162-020-00161-0>
- Jabareen, Y. (2009). Building a conceptual framework: Philosophy, definitions, and procedure. *International Journal of Qualitative Methods*, 8(4), 49–62. <https://doi.org/10.1177/160940690900800406>
- Kane, M. T. (2013). Validating the interpretations and uses of test scores. *Journal of Educational Measurement*, 50(1), 1–73. <https://doi.org/10.1111/jedm.12000>
- Li, Y., Ler, H. Y., Zhang, D., & Su, L. (2026). Fundamental motor skill interventions significantly improve executive functions and social–emotional competence in preschoolers: A meta-analysis. *Frontiers in Psychology*, 16, Article 1721589. <https://doi.org/10.3389/fpsyg.2025.1721589>
- Ljungcrantz, L. (2026). The interaction of AI and early childhood education: A state-of-the-art review 2020–2024. *Early Childhood Education Journal*, 54(5), 3565–3581. <https://doi.org/10.1007/s10643-025-02079-3>
- McKenney, S., & Reeves, T. C. (2019). *Conducting educational design research* (2nd ed.). Routledge.
- Melby-Lervåg, M., Redick, T. S., & Hulme, C. (2016). Working memory training does not improve performance on measures of intelligence or other measures of “far transfer”: Evidence from a meta-analytic review. *Perspectives on Psychological Science*, 11(4), 512–534. <https://doi.org/10.1177/1745691616635612>
- Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, A. H., Howerter, A., & Wager, T. D. (2000). The unity and diversity of executive functions and their contributions to complex “frontal lobe” tasks: A latent variable analysis. *Cognitive Psychology*, 41(1), 49–100. <https://doi.org/10.1006/cogp.1999.0734>
- National Academies of Sciences, Engineering, and Medicine. (2024). *A new vision for high-quality preschool curriculum*. The National Academies Press. <https://doi.org/10.17226/27429>
- Ozturk, E. (2025). Artificial intelligence in early childhood STEM education: A review of pedagogical paradigms, ethical issues, and socio-political implications. *Journal of Education in Science, Environment and Health*, 11(2), 108–125. <https://doi.org/10.55549/jeseh.800>
- Radesky, J., Hiniker, A., McLaren, C., Akgun, E., Schaller, A., Weeks, H. M., Campbell, S., & Gearhardt, A. N. (2022). Prevalence and characteristics of manipulative design in mobile applications used by children. *JAMA Network Open*, 5(6), Article e2217641. <https://doi.org/10.1001/jamanetworkopen.2022.17641>
- Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act). (2024). *Official Journal of the European Union*, L, 2024/1689. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- Reindl, V., Gerloff, C., Scharke, W., & Konrad, K. (2018). Brain-to-brain synchrony in parent-child dyads and the relationship with emotion regulation revealed by fNIRS-based hyperscanning. *NeuroImage*, 178, 493–502. <https://doi.org/10.1016/j.neuroimage.2018.05.060>
- Rosanbalm, K. D., & Murray, D. W. (2017). *Caregiver co-regulation across development: A practice brief* (OPRE Brief No. 2017-80). Office of Planning, Research, and Evaluation, Administration for Children and Families, U.S. Department of Health and Human Services. <https://acf.gov/opre/report/co-regulation-birth-through-young-adulthood-practice-brief>

- Rueda, M. R., Fan, J., McCandliss, B. D., Halparin, J. D., Gruber, D. B., Lercari, L. P., & Posner, M. I. (2004). Development of attentional networks in childhood. *Neuropsychologia*, 42(8), 1029–1040. <https://doi.org/10.1016/j.neuropsychologia.2003.12.012>
- Sala, G., & Gobet, F. (2017). Working memory training in typically developing children: A meta-analysis of the available evidence. *Developmental Psychology*, 53(4), 671–685. <https://doi.org/10.1037/dev0000265>
- Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human-Computer Interaction*, 36(6), 495–504. <https://doi.org/10.1080/10447318.2020.1741118>
- Skivington, K., Matthews, L., Simpson, S. A., Craig, P., Baird, J., Blazeby, J. M., Boyd, K. A., Craig, N., French, D. P., McIntosh, E., Petticrew, M., Rycroft-Malone, J., White, M., & Moore, L. (2021). A new framework for developing and evaluating complex interventions: Update of Medical Research Council guidance. *BMJ*, 374, Article n2061. <https://doi.org/10.1136/bmj.n2061>
- Stucke, N. J., & Doebel, S. (2024). Early childhood executive function predicts concurrent and later social and behavioral outcomes: A review and meta-analysis. *Psychological Bulletin*, 150(10), 1178–1206. <https://doi.org/10.1037/bul0000445>
- Su, J., & Yang, W. (2022). Artificial intelligence in early childhood education: A scoping review. *Computers and Education: Artificial Intelligence*, 3, Article 100049. <https://doi.org/10.1016/j.caeai.2022.100049>
- Trujillo-Trujillo, C. C., Cadavid Ruiz, N., & Sanabria-Mazo, J. P. (2026). Effectiveness of executive function interventions in typically developing preschool children: A systematic review and meta-analysis. *Child Neuropsychology*, 32(4), 598–632. <https://doi.org/10.1080/09297049.2025.2595075>
- UNICEF. (2025). *Guidance on AI and children* (Version 3.0). UNICEF Office of Strategy and Evidence – Innocenti. <https://www.unicef.org/innocenti/reports/policy-guidance-ai-children>
- Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes* (M. Cole, V. John-Steiner, S. Scribner, & E. Souberman, Eds.). Harvard University Press.
- Wang, F., & Hannafin, M. J. (2005). Design-based research and technology-enhanced learning environments. *Educational Technology Research and Development*, 53(4), 5–23. <https://doi.org/10.1007/BF02504682>
- Willoughby, M. T., The Family Life Project Investigators, & Blair, C. B. (2016). Measuring executive function in early childhood: A case for formative measurement. *Psychological Assessment*, 28(3), 319–330. <https://doi.org/10.1037/pas0000152>
- 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. <https://doi.org/10.1111/j.1469-7610.1976.tb00381.x>
- World Health Organization, United Nations Children’s Fund, & World Bank Group. (2018). *Nurturing care for early childhood development: A framework for helping children survive and thrive to transform health and human potential*. World Health Organization. <https://www.who.int/publications/i/item/9789241514064>
- Xu, Y., Wang, D., Collins, P., Lee, H., & Warschauer, M. (2021). Same benefits, different communication patterns: Comparing children’s reading with a conversational agent vs. a human partner. *Computers & Education*, 161, Article 104059. <https://doi.org/10.1016/j.compedu.2020.104059>
- Zelazo, P. D. (2006). The Dimensional Change Card Sort (DCCS): A method of assessing executive function in children. *Nature Protocols*, 1(1), 297–301. <https://doi.org/10.1038/nprot.2006.46>
- Zelazo, P. D., & Carlson, S. M. (2012). Hot and cool executive function in childhood and adolescence: Development and plasticity. *Child Development Perspectives*, 6(4), 354–360. <https://doi.org/10.1111/j.1750-8606.2012.00246.x>

Author(s) Information

Elif Ozturk

Giresun University
Faculty of Education, Department of Early Childhood Education
Giresun, Türkiye

ORCID ID: <https://orcid.org/0000-0002-3764-4526>

*Contact e-mail: elif.ozturk@giresun.edu.tr

* Corresponding author(s) contact email

To cite this article:

Ozturk, E. (2026). AI behind the adult, not in front of the child: A design framework for adult-mediated AI in early childhood. *The Eurasia Proceedings of Educational and Social Sciences (EPESS)*, 53, 97-110