Review Article | ISSN: 3108-3021

CORTEXPLORE | Journal of CORTEXPLORE

The Offloading Window: A Conceptual Framework for AI-Assisted Cognition in Early Childhood: A Conceptual and Theoretical Contribution

Divyaansh Kumar Gupta[1] | Komal Jain[2]

Abstract

Children now have daily use of artificial intelligence (AI) like voice response systems, conversational agents (e.g., chatbots), and Digital Learning Systems that provide an AI assist to support their learning. The use of such technology is expected to offer children personalized chances for educational enhancement and cognitive assistance when used appropriately. However, there is limited research on the effect of technology on the developmental/maturational process of children through cognitive offloading; thus, there is insufficient evidence of one or more elements leading to cognitive development through cognitive offloading by adults (i.e., complete adults) who have already developed their cognitive capacity. This paper integrates the current cognitive psychology, developmental neuroscience, and generative-AI-in-education literature to address the above-stated gap in knowledge. The proposed constructs are hierarchically related: The Developmental Offloading Hypothesis (DOH), a general principle confirming that cognitive offloading results in different kinds of effects when components of competence form as opposed to when they are being applied; The Offloading Window (OW), the contextual condition that defines whether the DOH is developmentally active for any particular child at any point in time based on their individual capacity and pattern of use; and the Scaffolded Cognitive AI Integration Framework (SCAIF), the operational framework for translating the principles of DOH and OW into practical guidelines that differentiate scaffolds from substitutions. Each construct is defined behaviourally, evaluated against alternative theoretical accounts (including the Extended Mind, Distributed Cognition, Vygotskian scaffolding, and

Submitted : **June** 02, 2026

Published : **June** 30, 2026

DOI: doi.org/10.65320/jce.vol.1.issue3.23

1* Mukesh Patel Institute of Technology Management and Engineering, Mumbai, India. (divyaanshkumargupta24@gmail.com)

2 Independent Researcher & Psychologist, Mumbai, India. (withkomal7@gmail.com)

Corresponding Author

This paper is blockchain-verified. Scan the QR code to verify its authenticity.

Cognitive Load Theory), and expressed as a set of falsifiable propositions. The paper concludes with a roadmap for the empirical work required to test these propositions, including a planned mixed-methods study of parents and early-childhood educators.

Keywords: Artificial Intelligence | Cognitive Offloading | Early Childhood Development | Executive Functions | Developmental Framework

I. Introduction

1.1 AI and the Changing Landscape of Early Childhood

Artificial intelligence has moved out of the laboratories of specialists and become part of the daily fabric of family life. Before they read, children are now regularly listening to voice assistants, stories told by AI, adaptive learning apps and conversational chatbots. Unlike the earlier educational technologies that primarily stored or presented information, modern generative systems can answer, solve problems, compose and converse in human-like ways of helping. Systematic reviews of generative AI in education reveal a rapid acceleration in adoption since 2023, alongside a persistent evidence gap in downstream cognitive effects (Alfarwan, 2025; Bouguettaya et al., 2025; Garzón et al., 2025). This acceleration poses a question prior to any specific technology: how does the increasing presence of artificial intelligence in the child's environment interact with the development of the child's own intelligence?

1.2 Early Childhood as a Critical Period

Ages three to six are some of the most important years for human cognitive development. During this window, children develop working memory, attention control, language, self-regulation, and problem-solving ability through active engagement with their environment (Diamond, 2013; Piaget, 1952; Vygotsky, 1978). The increased neuroplasticity means that the repeated use of cognitive functions during these years has a disproportionate impact on the neural structures that support later learning (Center on the Developing Child, 2016). During this time, kids aren't just learning; they're creating the tools for all future learning. Any technology that changes what a child does when confronted with a difficult or unfamiliar task therefore directly affects this construction process.

1.3 Cognitive Offloading and AI-Assisted Cognition

What's new is that generative AI systems can be participants in cognition, not just repositories of it: they generate explanations, do things, and talk in ways that previous tools could not. Relying on external, searchable memory alters the content of what adults remember internally, a phenomenon called the "Google effect" (Sparrow et al., 2011). Even the mere presence of a personal smartphone can measurably reduce available cognitive capacity (Ward et al., 2017). Both findings concern adults offloading from an already developed cognitive base.

1.4 Scaffolding or Substitution?

Adults who use a calculator or a navigation app are just offloading tasks they already know how to do. Many of the capacities that an educational AI might take over are ones the child has not yet

CORTEXPLORE | Journal of CORTEXPLORE

built. Young children are in a different position: The distinction developed in this paper is between AI that scaffolds, supporting a process that the child still completes, and AI that substitutes for that process altogether. This paper is not an argument to avoid AI altogether. Instead, it suggests, the question should not be whether children use AI but what function it serves in a particular activity and what the child's own cognitive labour must be

1.5 Positioning, Gap, and Contribution

This paper is a conceptual and theoretical contribution. It does not report primary data. Instead, it integrates two literatures that have developed largely independently: cognitive offloading research, concentrated on adults, and critical-period developmental neuroscience, largely silent on AI, to build new theoretical constructs and propose testable hypotheses for future empirical work. Jensen et al. (2025) note that early claims about generative AI's educational impact substantially outpaced the available evidence in the months following its mainstream release; the present paper responds to that caution by building theory deliberately ahead of, rather than in place of, data. The Developmental Offloading Hypothesis (DOH) presents a theory that explains how offloading produces various effects when people develop and apply their existing competencies. The Offloading Window establishes a specific framework which details when this difference emerges as a developmental factor. The Scaffolded Cognitive AI Integration Framework (SCAIF) transforms both systems into operational guidelines that are ready for use. The paper presents its closing section through a research plan, which uses mixed methods to study parents and early-childhood educators while keeping the project in its initial research phase.

The three constructs form a hierarchical structure that should be studied as a single system instead of as independent elements. DOH states that the timing of developmental processes determines how offloading functions. The offloading window is the contextual condition: it specifies when, for a particular child, capacity, and pattern of use, the hypothesis is actually in play. The operational framework of SCAIF becomes the decision structure that is activated after the hypothesis, together with its contextual condition, determines which activities hold developmental value. The research process moves from domain-general through child- and capacity-specific windows to reach activity-specific SCAIF applications. The first three sections of the document present the information in sequential order through Sections 4–6. Figure 2 in Section 6 presents a single integrated diagram that shows the entire process, from mechanism development to outcome measurement.

II. Cognitive Offloading: Evidence from Adult Populations

2.1 Cognitive Offloading and the Extended Mind

Throughout history, people have used outside resources to handle their mental tasks, such as written notes, calendars, maps, and calculators. The process of cognitive offloading describes how people use outside resources and their surroundings to perform tasks that would otherwise require more mental effort (Risko & Gilbert, 2016). The process functions as an adaptive system instead of a deficit because working memory possesses limited storage space, which enables people to handle basic tasks so they can focus on other responsibilities (Baddeley, 2000). The extended-mind thesis by Clark and Chalmers (1998) advances this concept through their argument, which states that cognitive operations extend past the brain and include tools that consistently function as elements of mental processes. A notebook, together with a calculator and modern AI assistants, functions as an element which extends the cognitive system instead of working as a standalone tool. The perspective shows that children who use AI with ease and regularity develop AI systems, which become an essential part of their mental structure. The new perspective on this topic creates a different question about offloaded mental activities which children need to practice internally versus those they can perform without practice.

2.2 Benefits, Risks, and Contested Findings

Research shows that offloading results vary between positive and negative outcomes, so a conceptual paper about this topic needs to present both perspectives. Sparrow et al. (2011) found that people remember less information they expect to find again externally and instead remember where to find it. Risko and Gilbert (2016) established a universal framework which combined memory with attention and planning to demonstrate that people achieve better immediate task results through offloading. Ward et al. (2017) showed that having a smartphone around reduces cognitive capacity for attention-demanding tasks even when the device remains unused. Research shows that AI-assisted education leads to better student learning results when instructors use guided inquiry instead of direct answer delivery according to separate studies. Hu et al. (2025) describe generative AI systems which function as a "Socratic playground" by using questions to maintain student reasoning instead of providing direct answers. Roe and Perkins (2024) state that the current focus on generative AI dependency fails to detect situations where students use AI to increase their learning independence. The research shows that offloading does not create negative effects, but its results depend on how users operate and design systems which matches the scaffolding/substitution model from Section 4.

2.3 Cognitive Offloading in the Age of Generative AI

Modern generative systems operate differently than previous tools because they perform cognitive operations, which include explanation creation and solution generation and dialogue maintenance. Firth et al. (2019) review evidence that habitual internet and device use is already associated with changes in sustained attention and memory strategy among adult users, and automation bias research shows that people increasingly defer judgment to automated systems even where independent reasoning would serve better (Mosier et al., 2001). Zhai's (2024) systematic review shows that students who depend too much on AI dialogue systems develop lower student independence when solving problems and their critical thinking abilities deteriorate. Older learners primarily provide this evidence, but it directly pertains to the mechanism this paper focuses on. Gu and Ericson (2025) show through their AI literacy study that users develop AI system knowledge through educational programs about system operations instead of reducing their exposure to AI systems. The combined research on adults and adolescents shows different dangers together with possible solutions, but it fails to reveal what happens to children who have not achieved their intended developmental potential. That is the gap the next section turns to.

III. Critical Periods, Neuroplasticity, and Executive Function Development

3.1 Critical Periods and Neuroplasticity

According to Knudsen (2004), developmental neuroscience defines critical periods as sensitive times when the brain has the highest level of neural reaction to environmental stimuli. The first critical period for development occurs between three and six years old with the brain growing rapidly and increasing synaptic connections as well as neuroplasticity during this time period, In this

CORTEXPLORE | Journal of CORTEXPLORE

process, the repeating activation of neural circuits strengthens the circuits, whereas connections that are not used are pruned (Johnson, 2001; Kolb & Gibb, 2011). This principle is often referred to as "neurones that fire together, wire together." Neural connections are made through repeated cognitive efforts rather than passive exposure. Piaget (1952) and Vygotsky (1978) made predictions about the impact of behavioural cognitive processes rather than neural processes, both suggesting that children develop their knowledge by interacting with their environment than simply receiving information. For AI exposure, the implications are direct; both the context of exposure and the child's cognitive functions during the exposure are the most important element for development.

3.2 Executive Functions and Effortful Learning

Executive functions (EF)—working memory, inhibitory control, and cognitive flexibility—allow for a person to maintain information in their mind, avoid distractions while completing one task (or maybe several), and change between tasks as situations and demands change (Diamond, 2013). These EF capabilities positively predict academic and social accomplishments later in life; they are also not a given or inherent to individuals, but rather the product of extended experiences that require extended focus or attention (eg; reading a chapter book), planning for future actions (eg; activities of daily living), and engaging in self-monitoring or checking (eg; keeping track of homework), and are directly related to the growth of the pre-frontal cortex in a child (Casey et al., 2005).

The learning-science literature supports this position from a different perspective. Bjork and Bjork's (2014) recent review of the desirable-difficulties framework, and Kapur's (2008) work on productive failure, illustrates that moderate levels of difficulty and calling on the user to use their own judgment about how to respond to a problem will lead to the development of a robust memory of that event even though both will feel easy or difficult at the time of engagement. Developmental neuroscience and learning sciences come together and support the same conclusion, that the process of creating learning and developing the supporting neural mechanisms associated on with this learning requires ongoing, effortful engagement; thus, effortful engagement is not an inconvenience of the learning process; but rather, it is what creates learning and develops the supporting neural mechanisms associated with the learning process.

3.3 Linking Executive Function to AI-Supported Learning

At this juncture, we must juxtapose developmental neuroscience with the AI-in-education literature, as neither currently does so directly. If executive-function capacities are built through repeated, effortful activation of the same cognitive circuits, then an AI system's design choices are not developmentally neutral: a system that completes a working-memory-demanding step on a child's behalf removes one repetition of the activation that would have strengthened that circuit. A system that instead prompts, questions, or delays its own answer preserves that repetition. The difference between answer-generating AI and Socratic-style generative AI design exists as a real distinction, as Hu et al. (2025) have established. The distinction which Hu et al. (2025) made between answer-generating AI and Socratic-style generative AI systems now exists as a developmental-neuroscience mechanism instead of a general teaching method. The research shows that single AI interactions do not create significant brain development changes, but scientists track brain changes through AI exposure patterns which create probabilistic impacts on brain development. This point is developed formally as

CORTEXPLORE | Journal of CORTEXPLORE

a dimension of the Offloading Window in Section 5.

IV. The Development Offloading Hypothesis (DOH)

4.1 From Cognitive Offloading to Developmental Offloading

Externalisation as a consequence of existing competence An adult who has learned to add may offload arithmetic tasks to a calculator. Cognitive offloading research has approached externalisation because of existing competence. The paper argues that this assumption is a scope limitation, not a general truth. But the offloading literature does not say what happens if an external system becomes available before the internal competence has been built. The Developmental Offloading Hypothesis (DOH) addresses that specification gap. Its central claim is that the developmental impact of offloading depends on the maturational status of the capacity being offloaded: offloading an established skill is developmentally neutral or beneficial; offloading a skill that is still forming runs the risk of foreclosing the practice through which it would have formed.

4.2 Core Proposition and Boundary Conditions

Proposition (DOH): Cognitive offloading is developmentally neutral or beneficial when it supports an already-established competence and potentially disruptive when it substitutes for a competence that is still forming.

Stated this way, the hypothesis has clear boundary conditions, which earlier statements of this argument left implicit. First, it applies specifically to skills within a child's current zone of active construction, excluding all other cognitive activities and skills already stable for that child. Second, it concerns patterns of repeated use rather than isolated instances; a single occasion of AI assistance is not claimed to be developmentally consequential on its own. Third, the hypothesis is agnostic about content domain: it makes the same prediction for language, numeracy, or social problem-solving, provided the capacity in question is still under construction for the child concerned. Outside these conditions (mature skills, one-off use, or capacities the child is not yet developmentally positioned to construct regardless of AI exposure), the hypothesis makes no claim.

4.3 Competence Formation versus Competence Application

The distinction underlying DOH is between competence formation, the process by which a new capacity is built through effortful practice, and competence application, the use of an already-built capacity to complete a task efficiently. During formation, effort, error, and uncertainty are not obstacles but mechanisms (Bjork & Bjork, 2014; Kapur, 2008). During application, the same effort is often simply a cost, and offloading it is a reasonable efficiency gain. The output alone will make children who finish tasks with AI support appear successful, but their true ability to develop skills remains unknown. The evaluation of AI-assisted learning needs to focus on the learning process instead of concentrating only on the final results. SCAIF operationalises this distinction directly in Section 6.

4.4 Operational Definitions

The argument contains three terms, which I will define right here because each term should be observable through direct evidence instead of depending on unverified claims. The child engages in the core cognitive operation that the activity targets during the AI scaffolding process, while the system offers prompts, structure, and feedback that the child must respond to. Substitution is AI involvement in which the system performs the core cognitive operation itself, and the child's role is reduced to initiating the request and receiving the output. Dependency is a behavioural pattern, observable over repeated occasions rather than any single one, in which a child defaults to requesting AI assistance for a task type before attempting it independently, at a frequency that measurably exceeds the child's independent-attempt rate for comparable tasks. These definitions are deliberately behavioural rather than purely conceptual, which is a requirement for the empirical agenda in Section 8.

4.5 Competing Theoretical Explanations

DOH is not the only lens available on this evidence. A conceptual paper should state plainly why a new construct is warranted rather than an existing one stretched to fit. Three cognitive-mechanism accounts are the most direct competitors, because each, like DOH, makes a claim about what offloading does to the cognitive system rather than about how instruction should be designed.

The extended-mind view (Clark & Chalmers, 1998), introduced in Section 2.1, would predict that a child's fluent, well-integrated use of an AI tool is itself a form of cognitive development, an expanded cognitive system, rather than a foreclosure of one. Distributed cognition, the broader research program from which the extended-mind thesis emerged (Solberg, M. & Hutchins, E., 2023), makes a related but weaker claim: cognitive processes are properly analysed at the level of a system spanning person and tool without necessarily implying that the tool becomes part of the person's own cognition, as Clark and Chalmers (1998) argue. Both views are compatible with a great deal of AI-assisted activity that looks developmentally unremarkable, because both treat the person-plus-tool unit, not the person's internal capacity alone, as the relevant object of analysis. However, neither specifies what happens when the child has yet to build the capacity being distributed. Both theories were developed to describe how already-competent adults extend established cognition into tools, not how a capacity forms in a system where the tool is present from the outset. That silence is precisely where DOH's developmental-timing variable does its work: DOH does not dispute that a stable child-AI system may function well as a unit; it claims that whether the child's own contribution to that unit is exercised or bypassed matters differently depending on whether the underlying capacity is still forming.

The AI-literacy view Gu and Ericson (2025) would say the way the child is taught about the tool is what makes the difference in the outcome, not the age at which the child first sees it; thus, if an appropriate AI is used with a well-established child relationship, it could be a benefit rather than a detriment to the child. The agency-centred view (Roe & Perkins, 2024) warns that dependency-based frameworks greatly underestimate instances where AI enhances a child's ability to reach to something compared to exerting effort. The views represented by DOH are incomplete without considering the developmentally-timed nature of the tool; DOH could also be viewed as not explaining the difference in operation of the specific child and adult for which they create their binary of completeness. DOH posits that resolving the differing opinions will rely on gathering the empirical evidence requested in Section 8 rather than on just saying one way or another.

4.6 Relationship to Developmental and Educational Theories

A second set of theories has relevance to the claims made in this paper for a different reason than that they compete to explain the same mechanism — they serve as established frameworks for developing instructional support, and therefore a new developmental construct has

CORTEXPLORE | Journal of CORTEXPLORE

to clearly articulate what it contributes to the existing frameworks.

Vygotsky's Zone of Proximal Development and scaffolding (Vygotsky, 1978) are the closest relatives of the three. Scaffolding theory specifies that support should be calibrated to extend a learner just past their current independent ability and should be withdrawn as competence grows. SCAIF borrows this principle directly (Section 6). What scaffolding theory does not specify is the substrate: it was developed for human-to-human instructional support and is silent on whether a non-human, always-available, infinitely patient assistant changes the calculus of when and how quickly support should fade. An adult tutor fades support partly because sustained tutoring is costly and partly because a human tutor reads a child's frustration cues; neither constraint applies to an AI system in the same way, so the fading schedule that made scaffolding developmentally safe in Vygotsky's (1978) account cannot simply be assumed to transfer. The Offloading Window's frequency dimension (Section 5.3) is, in effect, an attempt to specify a fading condition for a support source that has no natural incentive to fade on its own.

Cognitive Load Theory (Sweller, 1988) provides an additional perspective which operates independently from other theories by focusing on how to handle the working memory requirements of learning tasks to prevent students from exceeding their limited cognitive resources while learning new information for developing their mental frameworks. The paper presents an argument which Cognitive Load Theory uses to classify AI systems which eliminate task steps as load reduction systems although this approach is considered beneficial. DOH's disagreement is not with the mechanism but with the default assumption: Cognitive load theory does not distinguish between loads that are extraneous (irrelevant to what the task is meant to build) and loads that are precisely the productive difficulties the task exists to provide (Bjork & Bjork, 2014; Kapur, 2008). An AI system reducing load on a working-memory-formation task is, on DOH's account, not necessarily reducing extraneous load; it may instead be removing the germane load the activity was designed to exercise. Cognitive Load Theory provides no principled way to tell these apart because it was not built with a formation-versus-application distinction in mind; that distinction (Section 4.3) is what DOH contributes to this specific gap.

The two families of comparison in this section make the same point from different directions. The cognitive-mechanism theories (4.5) do not specify a developmental-timing variable; the instructional-design theories (4.6) specify how support should be delivered but not whether a still-forming capacity changes what “appropriate support” means when the support source is an AI system rather than a human instructor. DOH, the Offloading Window, and SCAIF are proposed as the minimal addition needed to close both gaps at once, rather than as a wholesale replacement for any of the four theories discussed here.

V. The Offloading Window

5.1 Defining the Offloading Window: Period, Mechanism, or Condition?

The offloading window has sometimes been described loosely as a “developmental period". That framing invites confusion with critical or sensitive periods in neuroscience, which are properties of the brain's general responsiveness (Knudsen, 2004), not of any specific technology or behaviour. To be precise, the Offloading Window is not itself a biological period. It is a contextual condition that holds whenever a specific cognitive capacity is (a) still under active construction in a specific child, and (b) exposed to a specific pattern of external assistance capable of substituting for it. The condition does not depend on a specific age range, although children between three and six years old develop their essential skills during this period, which makes the condition common among children of this age group. Nor is it a mechanism in the causal sense: the mechanism is the reduced-repetition pathway described in Section 3.3; the offloading window names the condition under which that mechanism becomes active for a given child and skill. This distinction matters practically: the window is assessed per capacity and per child, not applied as a blanket label to an age group.

5.2 Boundary Conditions for Window Activation

Because the window is a condition rather than a period, it needs a stated activation rule rather than an age cutoff. Four conditions are proposed as jointly necessary: the window is considered active for a given child only when all four conditions—capacity and activity—are met simultaneously, rather than when any one condition is present in isolation.

Developmental stage of the specific capacity, not the child's age in general. The relevant test is whether this particular capacity, working memory for this task type, for instance, is still under active construction for this child, since children of the same age vary in which capacities are stabilised.

Cognitive domain. The condition applies to domains in which offloading is behaviourally possible at all, such as memory, planning, problem-solving, and language production, and does not apply uniformly across all domains a child engages in, since some activities have no meaningful external-assistance analogue.

Task complexity. The activity must be complex enough that AI assistance could plausibly perform the core operation rather than a trivial or peripheral part of it; a task with no meaningful cognitive core to replace does not activate the condition regardless of AI presence.

Frequency and pattern of AI use. Consistent with Section 4.2's boundary condition on repeated use, the pattern of exposure must reach habitual or repeated use for the developmental capacity in question, rather than an isolated or occasional instance.

When any one of these four conditions is absent (the capacity is already stable, the domain does not admit offloading, the task is trivial, or the exposure is a one-off), the Window is not active, and DOH's core proposition (Section 4.2) does not apply to that instance. This activation rule is what allows the four dimensions introduced next (Section 5.3) to function as continuous moderators of severity once the window is already active, rather than as the activation test itself.

5.3 Distinguishing the Offloading Window from Related Constructs

Three adjacent constructs need to be held apart. 'Critical' or 'sensitive' periods describe general neural responsiveness to environmental input (Knudsen, 2004); the 'offloading window' describes a specific condition of exposure during such a period, not the period itself. 'Cognitive offloading' (Risko & Gilbert, 2016) describes the behaviour of delegating a task; the 'offloading window' describes the developmental context in which that delegation becomes consequential rather than neutral. Educational scaffolding (Vygotsky, 1978) is a design principle: support calibrated to extend a learner just past current ability while preserving active participation. The offloading window does not prescribe an instructional method; it identifies when scaffolding is developmentally preferable to substitution. None of the three existing constructs, individually, specifies the interaction between developmental timing and the delegation of a specific, unstabilised capacity to an

CORTEXPLORE Journal of CORTEXPLORE

external system. That interaction is the Offloading Window's distinct claim, and it is deliberately narrower than any construct it borrows from.

5.4 Dimensions of the Offloading Window

Once the boundary conditions in Section 5.2 establish that the window is active, four dimensions determine its severity for a given child and task: the developmental maturity of the targeted capacity (an emerging skill is more exposed than a stabilised one); cognitive dependency, the degree to which the child defaults to external assistance rather than attempting the task as defined behaviourally in Section 4.4; effort preservation, the amount of the child's own cognitive work the interaction still requires; and frequency of exposure, since the concern is a cumulative pattern rather than a single instance. These four dimensions are proposed as the measurable indicators the framework in Section 6 requires: each is observable in principle without direct neural measurement, which is what makes empirical testing (Section 8) feasible in the near term.

5.5 Propositions

The following are three propositions that will be empirically tested in the future rather than have already been established through research.

P1: The developmental benefit of using AI-assisted offloading will be greater when the child is forming their competency than when they are demonstrating their competency, provided the tasks are matched.

P2: The benefit received during formation will depend on the amount of effort preserved; if greater effort is preserved during the formation, then the benefit will be smaller.

P3: The benefits will accumulate through many exposures rather than from one-time exposures.

Thus, all three of these propositions will have an empirical test related to them and will be included in the empirical research agenda outlined in Section 8.

VI. Scaffolded Cognitive AI Integration Framework (SCAIF)

6.1 Rationale and Theoretical Foundations

SCAIF exists to answer the question of what parents, educators, and designers should do today in an activity that is appropriate for the given circumstance (Risko & Gilbert, 2016; Diamond, 2013; Knudsen, 2004; Vygotsky, 1978) and to provide a structure for determining the appropriateness of the use of AI for a given activity. The answer to the question of what to do right now is provided by means of the integration of off-loading theory, the impact of developmentally-related neural changes and executive function (particularly with respect to AI) into a single structure based upon the bandwidth of AI, the degree of potential for collaboration through the use of AI, and the likelihood of successful collaborative use of AI over the course of time. Sections 1.2 through 1.5 of this report establish that SCAIF provides the operational layer of the hierarchy; meaning that it provides a connection between an activity that is within the limits of an identified Offloading Window and the appropriate mode of collaboration/partnership to be utilized for the implementation of AI for the activity.

6.2 Core Design Principle and the Four Guiding Principles

SCAIF's operating question is, for a child at a specific developmental stage engaged in a particular activity, what role should AI play, and what cognitive effort must the child retain on their own? Four principles operationalise this question:

Developmental calibration: assistance should track which capacities are actively forming for this child, not technological convenience.

Review Article | ISSN: 3108-3021

CORTEXPLORE Journal of CORTEXPLORE

Scaffolding before substitution: AI should guide, prompt, and support before it completes work outright.

Preservation of cognitive effort: activities should retain productive struggle even where AI support is present.

Developmental transparency: the adult overseeing the interaction should be able to state which developmental objective the AI's involvement serves, since a principle that cannot be stated cannot be checked.

6.3 Modes of AI Involvement and the Developmental Application Matrix

SCAIF classifies AI participation into four modes of increasing intensity: observer (present, non-intervening; the child completes the task independently); prompter (offers questions, hints, or cues without supplying answers); collaborator (shares the cognitive load without assuming it); and substitute (performs the task on the child's behalf). *Figure 1* explains the four graduated modes of AI involvement within the Scaffolded Cognitive AI Integration Framework (SCAIF). As AI involvement increases from left to right, the child's independent cognitive effort correspondingly decreases.

Figure 1. The four graduated modes of AI involvement within the SCAIF

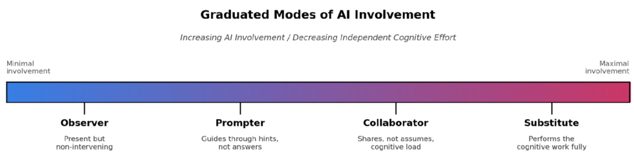

Source: Authors

Table 1 maps these four modes against seven common developmental goals, following the logic set out in Section 4.4: goals that draw on capacities typically still under construction at ages three to six, such as working memory and attention regulation, tolerate observer and prompter modes but not substitution; goals closer to fact retrieval or routine administration, where the targeted capacity is less central to the activity's developmental purpose, tolerate a fuller range of modes.

Table 1. Developmentally Appropriate Modes of AI Involvement

Developmental Goal	Observer	Prompter	Collaborator	Substitute
Working Memory Development	✓	✓	Limited	✗
Attention Regulation	✓	✓	Limited	✗
Problem Solving	Limited	✓	✓	✗
Creativity and Storytelling	Limited	✓	✓	Limited
Language Development	Limited	✓	✓	Limited
Fact Retrieval	Limited	✓	✓	✓
Routine Administrative Tasks	Limited	Limited	✓	✓

Source: Authors

Review Article | ISSN: 3108-3021

CORTEXPLORE Journal of CORTEXPLORE

6.4 Causal Pathways and Measurable Indicators

A recurring criticism of frameworks like SCAIF is that they name constructs without specifying how those constructs connect. The proposed pathway is as follows: the Developmental Offloading Hypothesis specifies the general mechanism: offloading affects formation and application differently; the Offloading Window specifies the contextual condition under which that mechanism is active for a given child and skill (Section 5.1); SCAIF's mode of AI involvement is the manipulable variable determining how much effort preservation occurs within that condition; effort preservation and frequency of exposure (Section 5.3) are the measurable mediators; and the developmental outcome, strengthened or under-practised executive-function capacity, is the endpoint.

Figure 2. Proposed Causal Pathway

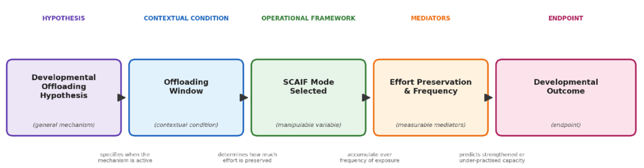

Source: Authors

Figure 2 illustrates the proposed causal pathway linking the Developmental Offloading Hypothesis (DOH), the Offloading Window (OW), and the Scaffolded Cognitive AI Integration Framework (SCAIF). The framework illustrates the hypothesised progression from the developmental mechanism through contextual conditions and AI involvement to measurable mediators (effort preservation and frequency of exposure), culminating in developmental outcomes. Each arrow represents a testable theoretical relationship rather than an assumed causal link. Each arrow in that chain corresponds to a testable relationship rather than an assumed one, which is what separates this proposal from a purely classificatory framework: mode selection is hypothesised to predict effort preservation, and effort preservation, accumulated over frequency, is hypothesised to predict outcome; neither relationship is asserted outright.

6.5 Operationalising the Offloading Window: Decision Logic

In practice, three questions operationalise the framework for a specific activity: is the targeted capacity still developing for this child; does the AI interaction support or replace the process that would develop it; and how much of the child's own cognitive effort remains in the interaction. An activity where the answers indicate a developing capacity, a replaced process, and low retained effort sits in what this paper terms the developmental substitution zone; an activity where the process is supported and effort is retained sits in the developmental scaffolding zone.

Figure 3. Operationalising the Offloading Window: A Decision Sequence for Zone Classification

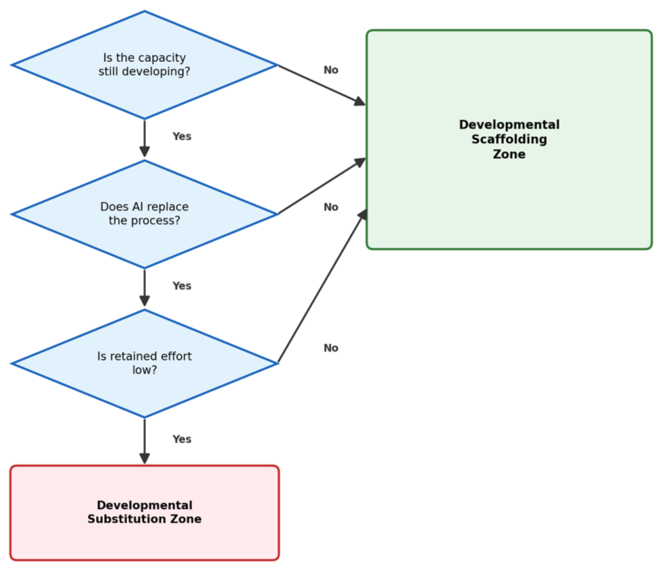

Source: Authors

Figure 3 illustrates the sequential logic for classifying AI-assisted learning activities into either the Developmental Scaffolding Zone or the

Developmental Substitution Zone. The offloading window is activated only when the targeted capacity is still developing, AI replaces the developmental process, and independent cognitive effort is substantially reduced. This gives adults a repeatable check rather than a one-time classification, since a single AI tool can sit in either zone depending on how it is used on a given occasion.

VII. IMPLICATIONS FOR PRACTICE

7.1 Parents and Educators

For parents, the practical shift recommended by SCAIF is to move from monitoring the amount of AI a child uses to monitoring how it is used; specifically, whether a child's question receives a direct answer or a prompt that encourages further reasoning is more important than the total exposure time. Moments of difficulty are not necessarily problems requiring resolution; per the desirable-difficulties literature (Bjork & Bjork, 2014), some are the mechanism of learning itself, and removing them removes the mechanism along with the discomfort. For educators, the same logic applies at the level of activity design: an AI tool that offers hints and follow-up questions is doing pedagogically different work from one that generates finished output, even when both are labelled "educational AI", and even though Hu et al. (2025) Socratic-style systems demonstrate that the former is achievable at scale rather than only in principle. Educators are also well placed to align AI involvement with a lesson's actual developmental target, reserving fuller AI involvement for activities where the point is exposure or creative collaboration rather than skill construction, per the logic of *Table 1*.

7.2 Technology Designers and Policymakers

For designers, the implication is that default behaviour matters: a system whose default is to answer completes the child's task by design; a system whose default is to ask a clarifying question or offer a graduated hint preserves developmental effort by design. Pinho et al.'s (2025) review of generative AI governance in education points out that most current frameworks deal with safety, privacy, and content appropriateness but not with developmental design defaults like this one. SCAIF is in a good position to fill this gap. For policymakers, current guidance, including UNESCO's (2023) generative-AI guidance for education, concentrates similarly on data protection, bias, and appropriate use rather than on age-differentiated design standards. The Offloading Window and SCAIF suggest that policy attention to the developmental quality of engagement, not only duration of exposure, is a distinct and currently under-addressed regulatory question, particularly for the three-to-six age band that existing K–12-focused guidance does not target (Alfarwan, 2025; Nadelson, 2025).

VIII. FUTURE RESEARCH AGENDA

8.1 A Roadmap for Empirical Validation

Section 5.4 describes an outline for how these propositions of Section 5.4 will be tested (not to be assumed), therefore this section functions as a roadmap for this process rather than being a research proposal itself. The theoretical foundation is provided by the contribution made by this manuscript while the empirical program outlined below will represent a future line of inquiry developed off of this empirical foundation. This empirical program will consist of three sequential phases.

Stage 1: Descriptive base. In order to establish a base upon which to test the hypotheses posed in P1-P3, a descriptive exploratory mixed methods survey will be conducted with both parents and early childhood educators (target N ≈ 100). The purpose of this survey is to identify existing patterns of AI exposure for children ages 3-6, identify any behavioural correlates of that AI exposure, and what strategies current practitioners are using as they attempt to balance providing children with assistance versus encouraging children's independent efforts. This phase is descriptive & hypothesis-generating in nature. Due to the constraints of recruitment, a convenience sample will be utilized (Etikan, Alkassim, & Abubakar, 2016), meaning there is no intention of using this as a test for the hypotheses posed in P1-P3, but rather it is the necessary precursor to phases II & III below.

Stage 2: Direct tests of P1-P3. Comparative designs will contrast children's task performance and behaviour when using scaffold-mode AI assistance versus substitute-mode AI assistance on matched formation-stage tasks, which will be sequenced after Stage 1 and will require significantly more resources.

Stage 3: Longitudinal and neuroscientific extension. Cohort designs following children across preschool and primary years, paired where feasible with non-invasive neurocognitive methods (EEG, fNIRS, eye-tracking) to connect the behavioural indicators in Section 5.3 to their neural correlates. Given the ethical and logistical demands of this stage with young children, it is proposed as a later rather than immediate priority.

8.2 Cross-Cultural and Contextual Extensions

AI access, parenting norms, and educational philosophy vary considerably across settings, and the DOH propositions as stated avoid claiming how these contextual factors moderate the core relationships. Comparative work across countries and across socioeconomic and digital-access contexts is identified as a further extension once Stages 1–2 above have established a baseline, to test whether the Offloading Window's dimensions operate uniformly or are moderated by the baseline exposure level and cultural context.

IX. Conclusion

This paper has shown that when using AI to help us think, it can either make us smarter, or it can cause us to lose brainpower. The paper explains that these two effects depend upon whether we are already using our brain's capability (for example, to write on paper) or whether we are developing that capability (for example, learning how to ride a bike). There are several other theories on this, such as the extended mind theory, distributed cognition, artificing (AI literacy), agency-centered frameworks, Vygotskian scaffolding and cognitive load theory that explain how to use AI, scaffold learning (in developing minds), and retain capacity (in those experiencing loss). In Section 4.5 and 4.6, I have detailed all of the and points of view from all of these theories and what they disagree on.

This paper provides testable propositions that will be proven by empirical evidence, not by future argument, as would be the case in Section 8. The practical question here is whether you ask your tool a question or whether the tool provides you with an answer to a question when you need it. When the distinction between those two types of interaction occurs often enough, you will know whether the AI tool is working to scaffold a developing mind or if the tool is working as a substitute for the additional development of the skill.

References

1. Alfarwan, A. (2025). Generative AI use in K–12 education: A systematic review. *Frontiers in Education, 10*, Article 1647573. https://doi.org/10.3389/feduc.2025.1647573
2. Baddeley, A. D. (2000). The episodic buffer: A new component of working memory? *Trends in Cognitive Sciences, 4*(11), 417–423. https://doi.org/10.1016/S1364-6613(00)01538-2
3. Bjork, E. L., & Bjork, R. A. (2014). Making things hard on yourself, but in a good way: Creating desirable

difficulties to enhance learning. In M. A. Gernsbacher and J. Pomerantz (Eds.), *Psychology and the real world: Essays illustrating fundamental contributions to society (2nd edition).* (pp. 59-68). New York: Worth. https://bjorklab.psych.ucla.edu/publication/bjork-e-l-bjork-r-a-2014-making-things-hard-on-yourself-but-in-a-good-way-creating-desirable-difficulties-to-enhance-learning-in-m-a-gernsbacher-and-j-pomerantz-eds-psycholo/

4. Bouguettaya, S., Pupo, F., Chen, M., & Fortino, G. (2025). A meta-survey of generative AI in education: Trends, challenges, and research directions. *Big Data and Cognitive Computing, 9*(9), Article 237. https://doi.org/10.3390/bdcc9090237
5. Casey, B. J., Tottenham, N., Liston, C., & Durston, S. (2005). Imaging the developing brain: What have we learned about cognitive development? *Trends in Cognitive Sciences, 9*(3), 104–110. https://doi.org/10.1016/j.tics.2005.01.011
6. Center on the Developing Child. (2016). *From best practices to breakthrough impacts: A science-based approach to building a more promising future for young children and families*. Harvard University. https://developingchild.harvard.edu/resources/report/best-practices-breakthrough-impacts/
7. Clark, A., & Chalmers, D. J. (1998). The extended mind. *Analysis, 58*(1), 7–19. https://doi.org/10.1093/analys/58.1.7
8. Diamond, A. (2013). Executive functions. *Annual Review of Psychology, 64*, 135–168. https://doi.org/10.1146/annurev-psych-113011-143750
9. Etikan, I., Musa, S. A., & Alkassim, R. S. (2016). Comparison of convenience sampling and purposive sampling. *American Journal of Theoretical and Applied Statistics, 5*(1), 1–4. https://doi.org/10.11648/j.ajtas.20160501.11
10. Firth, J., Torous, J., Nicholas, J., Carney, R., Rosenbaum, S., & Sarris, J. (2019). The online brain: How the Internet may be changing our cognition. *World Psychiatry, 18*(2), 119–129. https://doi.org/10.1002/wps.20617
11. Garzón, J., Baldiris, S., & Gutiérrez, J. (2025). Systematic review of artificial intelligence in education: Trends, benefits, and challenges. *Multimodal Technologies and Interaction, 9*(8), Article 84. https://doi.org/10.3390/mti9080084
12. Gu, X., & Ericson, B. J. (2025). AI Literacy in K-12 and Higher Education in the Wake of Generative AI: An Integrative Review. In Proceedings of the 2025 ACM Conference on International Computing Education Research V.1 (ICER '25). Association for Computing Machinery, New York, NY, USA, 125–140. https://doi.org/10.1145/3702652.3744217
13. Hu, X., Xu, S., Tong, R., & Graesser, A. (2025). Generative AI in education: From foundational insights to the Socratic playground for learning. *Computers and Education: Artificial Intelligence, 8*, Article 100311. https://doi.org/10.48550/arXiv.2501.06682
14. Jensen, L. X., Buhl, A., Sharma, A., & Bearman, M. (2025). Generative AI and higher education: A review of claims from the first months of ChatGPT. *Higher Education, 89*, 1145–1161. https://doi.org/10.1007/s10734-024-01265-3
15. Johnson, M. H. (2001). Functional brain development in humans. *Nature Reviews Neuroscience, 2*(7), 475–483. https://doi.org/10.1038/35081509
16. Kapur, M. (2008). Productive failure. *Cognition and Instruction, 26*(3), 379–424. https://doi.org/10.1080/07370000802212669
17. Knudsen, E. I. (2004). Sensitive periods in the development of the brain and behavior. *Journal of Cognitive Neuroscience, 16*(8), 1412–1425. https://doi.org/10.1162/0898929042304796
18. Kolb, B., & Gibb, R. (2011). Brain plasticity and behaviour in the developing brain. *Journal of the Canadian Academy of Child and Adolescent Psychiatry, 20*(4), 265–276. https://pmc.ncbi.nlm.nih.gov/articles/PMC3222570/
19. Mosier, K. L., Skitka, L. J., Dunbar, M., & McDonnell, L. (2001). Aircrews and Automation Bias: The Advantages of Teamwork? *The International Journal of Aviation Psychology*, *11*(1), 1–14. https://doi.org/10.1207/S15327108IJAP1101_1
20. Nadelson, L. S. (2025). The use of generative artificial intelligence in K–12 education: An introduction to the special issue of the *Journal of Educational Research*. *The Journal of Educational Research, 118*(6), 531–534. https://doi.org/10.1080/00220671.2025.2514913
21. Piaget, J. (1952). *The origins of intelligence in children*. International Universities Press.
22. Pinho, I., Costa, A.P. & Pinho, C. (2025). Generative AI governance model in educational research: A

scoping review. *Frontiers in Education, 10*, Article 1594343. https://doi.org/10.3389/feduc.2025.1594343

23. Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences, 20*(9), 676–688. https://doi.org/10.1016/j.tics.2016.07.002
24. Roe, J., & Perkins, M. (2024). Generative AI and Agency in Education: A Critical Scoping Review and Thematic Analysis. https://doi.org/10.48550/ARXIV.2411.00631
25. Sparrow, B., Liu, J., & Wegner, D. M. (2011). Google effects on memory: Cognitive consequences of having information at our fingertips. *Science, 333*(6043), 776–778. https://doi.org/10.1126/science.1207745
26. Solberg, M. & Hutchins, E. (2023): *Cognition in the wild*: MIT Press, Cambridge, 1995, pp. 402. ISBN: 9780262581462. *WMU Journal of Maritime Affairs*, 22, 267–272. https://doi.org/10.1007/s13437-023-00305-6
27. Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science, 12*(2), 257–285. https://doi.org/10.1207/s15516709cog1202_4
28. UNESCO. (2023). *Guidance for generative AI in education and research*. UNESCO Publishing. https://unesdoc.unesco.org/ark:/48223/pf0000386693
29. Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. 174 pp., Harvard University Press. https://www.hup.harvard.edu/books/9780674576292
30. Ward, A. F., Duke, K., Gneezy, A., & Bos, M. W. (2017). Brain drain: The mere presence of one's own smartphone reduces available cognitive capacity. *Journal of the Association for Consumer Research, 2*(2), 140–154. https://doi.org/10.1086/691462
31. Zhai, C., Wibowo, S., & Li, L. D. (2024). The effects of over-reliance on AI dialogue systems on students' cognitive abilities: A systematic review. *Smart Learning Environments, 11*, Article 28. https://doi.org/10.1186/s40561-024-00316-7