Pakistan Journal of Social Sciences Review (PJSSR)
ISSN (e) 2959-8052 (p) 2959-8044
Volume 5, Issue 4, 2026

The Impact of Artificial Intelligence on Natural Human Abilities: A Critical Review with Recommendations for Children, Youth, and Communities

Muhammad Idrees[1], Dr. Muhammad Nawaz Khan[2], Dr. Sameer Ul Khaliq Jan[3*]

[1]ISCTE – Instituto Universitário de Lisboa, PhD researcher in Social Work, Lisbon, Portugal.

[2]Assistant Professor Department of Sociology, Shaheed Benazir Bhutto University Sheringal Dir Upper, Khyber Pakhtunkhwa, Pakistan.

[3]Lecturer, Department of Social Work, Shaheed Benazir Bhutto University, Sheringal, Dir Upper, Khyber Pakhtunkhwa, Pakistan.
sameer@sbbu.edu.pk[3*]

DOI: https://doi.org/

Keywords
Artificial intelligence, human cognition, critical thinking, cognitive offloading, children and youth, emotional intelligence, social capital, digital literacy, bioecological systems theory, social work

Article History
Received on 7 Mar 2026
Accepted on 10 Apr 2026
Published on 18 Apr 2026

Copyright @Author
Corresponding Author: *

Dr. Sameer Ul Khaliq Jan
Lecturer, Department of Social Work, Shaheed Benazir Bhutto University, Sheringal, Dir Upper, Khyber Pakhtunkhwa, Pakistan.
sameer@sbbu.edu.pk[3*]

Abstract

The rapid, largely unregulated integration of artificial intelligence (AI) into everyday life raises pressing questions about long-term effects on distinctively human cognitive, social, emotional, and creative capacities. This narrative review critically examines peer-reviewed research, theoretical scholarship, and interdisciplinary commentary published between 2000 and 2025, with emphasis on studies from 2018 onwards. Drawing on Bronfenbrenner's bioecological systems theory as an organising framework, the review analyses how AI operates across microsystem, mesosystem, exosystem, and macrosystem levels to reshape the conditions under which human abilities develop. Evidence reviewed suggests that habitual, unreflective AI use is associated with reduced engagement in critical reasoning, memory consolidation, empathetic attunement, creative effort, and authentic communication. At the same time, AI, when used intentionally and equitably, can meaningfully augment human capacity and support vulnerable populations. The most significant risks fall disproportionately on children and youth, whose neurological and social development remains formative, and on marginalised communities whose structural disadvantages are likely to be amplified by poorly governed AI systems. The article concludes with differentiated, evidence-grounded recommendations for families, educational institutions, social work practitioners, community organisations, and policymakers, and calls for a reconceptualization of digital literacy as a civic and developmental necessity.

1 Introduction

Few technological shifts in recent history have had as profound an impact on daily life as the rapid expansion of artificial intelligence. Within a decade, AI has moved from a specialised research domain to an ambient infrastructure: embedded in search engines, healthcare diagnostics, creative tools, social media platforms, and educational software. The release of large language model applications accessible to the public from 2022 onwards has extended this reach still further, placing conversational AI systems in the hands of children, students, professionals, and older adults alike. According to the OECD (2023), AI-related patent filings doubled between 2010 and 2020, and by 2024 an estimated 77 per cent of devices in use globally incorporated some form of AI functionality (McKinsey Global Institute, 2023). These figures gesture toward a transformation that is not merely technological but anthropological.

Scholarly and policy conversations about AI have, with some notable exceptions, remained centred on economic productivity, labour displacement, and algorithmic bias. These are important questions. However, they leave a more fundamental question underexplored: what happens to distinctively human abilities when machines routinely take over the cognitive, emotional, and social work that once required human effort? When a seven-year-old asks a voice assistant rather than working through a problem, when an adolescent turns to an AI chatbot in moments of distress, when a student submits AI-generated writing as their own, something about human development and human connection is being altered. Whether it is altered for better or worse, and under what conditions, is precisely what requires careful examination.

This review does not adopt the pessimistic position that AI is inherently damaging, nor the optimistic one that its benefits will naturally distribute fairly and that human capacities will adapt without cost. Both positions are empirically weak and analytically lazy. The defensible position, and the one argued here, is that the effects of AI on human ability are real, context-dependent, structurally uneven, and urgently in need of informed, policy-relevant scholarship. The review is grounded in Bronfenbrenner's bioecological systems theory, which provides a principled framework for understanding AI not as an external force acting on passive individuals, but as a set of technological conditions that alter the nested environments in which human development takes place. This review contributes by integrating AI impacts across ecological systems.

1.1 Geographic and Socio-Economic Context: The Case of Pakistan

The global spread of AI takes on a specific meaning when we examine the socio-economic framework of Pakistan. Global data shows that AI-related patent filings doubled in a single decade (Khan et al., 2024). However, the impact of this technology in Pakistan is shaped by a massive "youth bulge," as over 60 percent of the country's population is currently under the age of 30. Within the macrosystem of Bronfenbrenner's theory, Pakistan faces a unique and difficult challenge because rapid digital adoption is happening at the same time as deep gaps in digital literacy and infrastructure (Khan et al., 2024). This evidence suggests that the effect of AI on human abilities in this region is not a simple or neutral technological change. Instead, it is a shift that may widen the existing gap between wealthy urban centers and disadvantaged rural communities. By analyzing these statistics, we can see that without national policies to ensure fair access, AI may reinforce social inequality rather than reduce it.

The entry of AI into the Pakistani microsystem changes the immediate settings where children develop basic skills. In rural areas with fewer resources, such as parts of Khyber Pakhtunkhwa, there is often a lack of teacher

supervision and limited infrastructure (Hassan et al., 2026). This may lead students to rely too much on AI tutoring applications for their daily work. This behavior triggers "cognitive offloading," where students use AI to avoid the difficult work of solving math problems or writing essays (Khalid et al., 2025). This research indicates that while AI can act as a "mind partner," an over-reliance on these tools for quick solutions leads to shallow engagement with learning materials. When we examine this trend through the lens of human development, it becomes clear that using AI as a constant "shortcut" during formative years may weaken the mental effort required for independent critical thinking. At the mesosystem level, the lack of coordination between Pakistani families and schools creates a significant risk. Research suggests that high levels of parental engagement in digital spaces, if not handled mindfully, can predict behavioral difficulties in children (Li, 2025). Most parents in Pakistan must manage these new technologies without any professional or institutional guidance, often making choices based on convenience rather than developmental evidence. This is particularly concerning for adolescents who are forming their identities. Furthermore, many AI tools lack the cultural and linguistic nuances needed for the Pakistani context (Integrating AI in Pakistani ESL classrooms, 2025). This lack of localization means that AI-generated content may replace authentic local expression with a uniform, Westernized style. Studying this gap reveals that when the home and school environments do not work together to guide AI use, the technology may overshadow the unique cultural and linguistic voice of the Pakistani experience.

However, a balanced view shows that AI can also be a tool for social improvement if it is governed with clear goals. In Pakistan, AI tools can offer better educational access to students in remote areas and help rural communities join the global digital economy (Hassan et al., 2026).

Additionally, studies have shown that when AI tools are used intentionally, they can improve critical thinking scores by providing structured support (Ishaque et al., 2026). This data implies that the problem is not the technology itself, but how it is integrated into the student's life. The challenge for Pakistani policymakers is to move away from simply accepting technology and instead create a coordinated ecological response. By setting rules that encourage "reflective reliance" rather than "blind obedience," Pakistan can ensure that AI supports, rather than replaces, the human relationships and mental effort that human growth requires.

2. Methodology

2.1 Study Design

This study adopts a critical narrative review design, which is well-suited for synthesizing and interpreting interdisciplinary literature characterized by conceptual diversity and methodological heterogeneity. Unlike systematic reviews that emphasize exhaustive inclusion and quantitative aggregation, critical narrative reviews facilitate theoretically informed analysis, enabling the identification of patterns, contradictions, and knowledge gaps while generating integrative insights across diverse bodies of scholarship.

2.2 Search Strategy and Data Sources

A structured literature search was conducted between January and March 2025 across four major academic databases: PsycINFO, Web of Science, Scopus, and Google Scholar. These databases were selected to ensure comprehensive coverage of psychology, social sciences, education, and interdisciplinary research on artificial intelligence.

Search terms were developed iteratively and applied both individually and in Boolean combinations. The primary search strings included:

- "Artificial intelligence" AND "cognitive ability."
- "AI" AND "critical thinking."
- "AI" AND "memory" OR "cognitive offloading."
- "AI" AND "emotional intelligence."
- "AI" AND "children" OR "youth."
- "Digital technology" AND "human development."
- "AI" AND "social capital."

The temporal scope covered publications from 2000 to 2025, with particular emphasis on studies published from 2018 onwards, reflecting the rapid advancement of contemporary AI technologies.

2.3 Inclusion and Exclusion Criteria

Studies were included based on the following criteria:

- Peer-reviewed journal articles, scholarly book chapters, and major institutional reports (e.g., OECD, UNESCO, WHO)
- Publications in the English language
- Direct relevance to artificial intelligence and its relationship with human cognitive, social, emotional, or developmental outcomes
- Empirical, conceptual, or theoretical contributions to the topic

Studies were excluded if they:

- Focused exclusively on technical or engineering aspects of AI without relevance to human capabilities
- Lacked clear scholarly grounding (e.g., non-academic opinion pieces)
- Were duplicate records across databases

2.4 Study Selection Process

The initial search yielded approximately 340 records. Following title and abstract screening, 186 articles were retained for full-text review. Of these, 74 sources were ultimately included based on their methodological rigor, conceptual relevance, and contribution to the research objectives.

Study selection was guided by a critical interpretive approach, prioritizing works that demonstrated theoretical depth, empirical robustness, and substantive insight into the relationship between AI and human development. Importantly, divergent and contradictory findings were intentionally retained to avoid artificial consensus and to capture the complexity of the field.

2.5 Analytical Approach

The analysis was informed by Bronfenbrenner's bioecological systems theory, which provided a multi-level framework for examining the influence of AI across individual, interpersonal, institutional, and societal contexts. Additionally, cognitive offloading theory guided the exploration of mechanisms through which AI technologies may shape cognitive processes.

Rather than employing quantitative aggregation, the study utilized a thematic and interpretive synthesis approach, identifying recurring patterns, theoretical tensions, and emerging risks and opportunities associated with AI integration in human development.

2.6 Limitations and Potential Bias

Several limitations inherent to the study design should be acknowledged. First, as a narrative review, the study relies on interpretive judgment in the selection and analysis of literature, which may introduce selection bias. Second, the restriction to English-language publications may exclude relevant studies from non-English-speaking contexts, thereby limiting global representativeness. Third, reliance on published and peer-reviewed sources may contribute to publication bias, as studies reporting non-significant findings are less likely to be included in the literature.

Furthermore, given the rapidly evolving nature of AI technologies, the available evidence

remains temporally constrained, and the long-term developmental implications are not yet fully understood. These limitations highlight the need for cautious interpretation and underscore the importance of future longitudinal and cross-cultural research.

3. Theoretical Framework

3.1 Bronfenbrenner's Bioecological Systems Theory

The organising framework of this review is Bronfenbrenner's bioecological model, originally articulated in The Ecology of Human Development (1979) and refined in collaboration with Morris (2005). In its mature formulation, the model holds that development results from proximal processes, sustained reciprocal interactions between an active human being and the persons, objects, and symbols in their immediate environment, occurring within a layered system of nested ecological contexts. The microsystem comprises the immediate settings of development: the family, school, and peer group, where children first learn to reason, communicate, form attachments, and regulate emotion. The mesosystem refers to linkages among microsystems. The exosystem encompasses settings that influence development indirectly. The macrosystem constitutes the overarching cultural ideologies, economic structures, and policy frameworks that shape available resources. The chronosystem captures the dimension of historical change, particularly relevant given the speed of the AI transformation.

Bronfenbrenner's framework is appropriate here because it refuses reductionist accounts that locate the causes of developmental outcomes purely in individuals or purely in structures and insists on the bidirectionality of development. AI, from this perspective, is not simply something that happens to people. It is a set of tools and relational interfaces that enter the microsystem, reconfigure the mesosystem, operate through exosystemic institutions, and are shaped by macrosystemic values. Understanding AI's effects requires analysis at all these levels simultaneously.

3.2 Cognitive Offloading Theory

Complementing the ecological framework is cognitive offloading theory, developed most systematically by Risko and Gilbert (2016). Cognitive offloading describes the use of external devices to reduce the internal cognitive demands of a task. Risko and Gilbert distinguish between immediate offloading, which is situational and reversible, and habitual offloading, which becomes a stable feature of how an individual approaches cognitive challenges. The neurological basis for developmental concern lies in what is well established in neuroplasticity research: cognitive capacities are maintained, strengthened, or weakened in response to patterns of use (Draganski et al., 2006; Kolb & Whishaw, 2015). A skill that is not practised loses its neural substrate. What the offloading framework contributes analytically, read alongside Bronfenbrenner, is an account of the mechanism through which AI alters development at the microsystemic level: not by preventing children from developing abilities, but by systematically reducing the occasions on which they are required to exercise them.

4. AI and Natural Human Abilities: Evidence and Critical Analysis

4.1 Critical Thinking and Analytical Reasoning

Critical thinking, the reflective, reasoned process of evaluating claims, identifying assumptions, and constructing well-supported arguments, is widely regarded as a foundational competency for education, professional life, and democratic citizenship (Facione, 1990; Paul & Elder, 2006). It is also the ability most directly threatened by habitual reliance on AI systems designed to generate confident, fluent, apparently authoritative

output without accompanying evidence of reasoning.

Pennycook and Rand (2019) demonstrated that individuals who engage in more analytical thinking are significantly less likely to endorse misinformation, including AI-generated misinformation. Their findings imply, inversely, that anything systematically reducing such engagement raises the risk of epistemic vulnerability. This concern gained urgency with the emergence of large language models. Bender et al. (2021) argued in their influential 'stochastic parrots' paper that the danger is not merely that language models make errors, but that the fluency of their output creates a persuasive veneer that discourages critical verification. This was corroborated by Alkaissi and McFarlane (2023), who found that a substantial proportion of AI-generated academic references were entirely fabricated yet presented in the formal conventions of scholarly citation with sufficient plausibility to pass initial scrutiny.

Perkins (2023) found that the dominant pattern of AI use among university students was not critical engagement but uncritical substitution, students submitting AI-generated text with minimal modification, not because they lacked integrity but because institutional conditions provided little incentive to do otherwise. This points to a systemic rather than individual problem. Gerlich (2025), in a study involving over 600 university students across multiple countries, found a statistically significant negative association between habitual AI tool use and performance on standardised measures of critical thinking, with modest but consistent effect sizes. The theoretical logic of cognitive offloading, combined with neuroplasticity evidence, suggests that AI use and declining critical thinking may operate in a reciprocal feedback loop, warranting urgent longitudinal investigation.

4.2 Memory, Information Retention, and the Ecology of Knowing

Human memory is not a passive archive. As Schacter (2001) described, it is a dynamic, reconstructive, emotionally shaped process integral to identity, reasoning, and the flexible application of knowledge. What we know, in the deep sense, is information actively integrated into our cognitive architecture, available for rapid deployment and novel combination, qualitatively different from the ability to retrieve a correct answer from an external source.

The foundational empirical work on how digital technology alters this form of knowing was conducted by Sparrow, Liu, and Wegner (2011), who demonstrated that when people believe information is easily retrievable online, they are less likely to encode it in memory and more likely to remember how to find it rather than what it is. Roediger and Karpicke's (2006) work on retrieval practice offers a compelling reason to be concerned: the act of retrieving information from memory, rather than accessing it externally, is among the most powerful known mechanisms for consolidating long-term retention and building flexible, transferable knowledge. When AI systems make retrieval unnecessary, they remove one of the most important conditions for deep learning. The practical consequences are not trivial: physicians who cannot reason through a differential diagnosis without immediate AI support, engineers who cannot identify structural problems without algorithmic prompting, represent a real and growing form of professional de-skilling with implications for the quality and safety of essential services. Ward et al. (2017) extended these findings to the smartphone itself, demonstrating that the mere presence of a device on a desk, even switched off, significantly reduced available cognitive capacity, with sobering implications for AI systems that proactively offer assistance.

4.3 Emotional Intelligence, Empathy, and Social Cognition

The development of emotional intelligence, understood following Mayer, Salovey, and Caruso (2004) as involving the perception, facilitation, understanding, and management of emotion, is a relational and experiential achievement, not primarily a cognitive one. Emotional attunement is built through repeated, often effortful engagement with the emotional lives of others: through misunderstanding and repair, the discomfort of emotional disagreement, and the gradual acquisition of a nuanced emotional vocabulary. These are not skills that can be learned from instruction alone; they require social experience of a particular kind.

Turkle's (2015) ethnographic research documented a striking pattern: young people who had grown up with significant mediated communication were exhibiting reduced capacity and reduced desire for deep, unscripted, reciprocal conversation. Her central argument, that always-on, always-responsive digital interlocutors alter young people's tolerance for the messiness and unpredictability of real human relationships, is supported by a broader pattern of evidence. Xie et al. (2023) found that adolescents regularly using AI companion apps, while reporting short-term reductions in loneliness, showed reduced motivation to engage in peer relationships and increased anxiety in unstructured social situations. AI companions, by providing consistently agreeable, conflict-free interaction, fail to provide the relational experiences that are developmentally formative: the experience of waiting, of misunderstanding, of repair, of another person's needs mattering as much as one's own. These experiences are not incidental to development; they are its substance.

4.4 Creativity, Intrinsic Motivation, and the Risk of Substitution

The relationship between AI and creativity is the most actively contested domain in this literature. From a self-determination theory perspective, Deci and Ryan (2000) argued over decades of empirical research that intrinsic motivation, the disposition to engage in activities for their inherent satisfaction, is the psychological substrate from which genuine creativity grows. The theoretical implication is clear: when AI provides a high-quality creative output at near-zero effort, it fundamentally alters the cost-benefit calculation that shapes whether individuals invest in developing their own creative capacities.

Grassini (2023) conducted a systematic review identifying both augmentation effects, where AI tools enhanced creative output by reducing technical barriers, and substitution effects, where AI availability reduced the motivation and practice of independent creative generation. Critically, augmentation effects were more pronounced among individuals with existing creative skills and intrinsic motivation, while substitution effects were more pronounced among novices and those with lower baseline creative confidence. For children and students still developing creative capacities, AI availability may displace the very practice through which those capacities are built. Cropley (2023) raises a further concern: creativity is not only an individual cognitive capacity but a social practice through which communities produce meaning, identity, and cultural goods. When creative work is increasingly AI-mediated, the communicative and relational functions of creativity may be eroded in ways individual-level research cannot capture.

4.5 Language, Communication, and Authentic Expression

Language is the medium through which human beings think, form relationships, construct identity, and participate in civic life. Vygotsky's (1978) foundational argument, that language is not merely an instrument of communication but the primary medium of cognitive development, establishes the highest

possible stakes for any technology that alters the conditions under which language is produced. If AI systems routinely take over the work of language production, the consequences extend well beyond writing quality or communication style. Weidinger et al. (2022) identified the risk of skill degradation through automation in their comprehensive risk taxonomy of large language models, drawing on well-established evidence from the automation of other skilled activities. Beyond individual skill degradation, there is a qualitative concern about what AI-mediated communication does to the texture of human expression. AI-generated language, trained to optimise for fluency and broad acceptability, tends toward a polished uniformity that flattens the idiosyncrasy, risk-taking, and relational specificity that make human communication meaningful. A communication culture in which a distinctive human voice is progressively replaced by algorithmically optimised text is one in which the conditions for authentic relationship and cultural diversity of expression are simultaneously diminished.

5. Multi-Level Implications: A Bioecological Analysis

5.1 The Microsystem: Children and Developmental Vulnerability

Children represent the population for whom the developmental implications of AI use are most consequential, precisely because the cognitive and social capacities at risk are not yet fully formed. From infancy through adolescence, the brain undergoes periods of heightened plasticity during which environmental inputs have outsized and sometimes irreversible effects on its structural and functional organisation (Blakemore & Choudhury, 2006; Greenough et al., 1987). Executive functions, including working memory, cognitive flexibility, and inhibitory control, continue to develop into early adulthood and are particularly sensitive to environmental
conditions during middle childhood and adolescence.

Radesky et al. (2023) argued that the introduction of AI into children's learning environments requires careful developmental calibration: AI tools should not be introduced before children have had adequate opportunity to develop the underlying competencies those tools replicate. A child taught to read using AI-assisted tools before they have developed phonemic awareness and decoding skills is not being supported in learning to read; they are being given a prosthetic that bypasses the developmental process through which reading is built. Similar arguments apply to numeracy, writing, and social problem-solving. The issue is not the tool but the developmental sequencing of its introduction.

5.2 The Mesosystem: Youth, Identity Formation, and Institutional Contexts

Adolescence presents a distinct developmental profile. In Erikson's (1968) framework, the central developmental task is identity formation: the construction of a coherent sense of self through exploration, experimentation, and commitment. Healthy identity development requires active engagement with challenges, genuine decision-making, the experience of failure and recovery, and exposure to diverse perspectives. AI systems present a structural challenge to this process in several respects. First, AI that pre-empts or resolves challenges before adolescents have engaged with them reduces the experiential raw material from which identity is constructed. Second, AI-curated social environments, designed to maximise engagement through personalisation, tend to reinforce existing preferences rather than challenge them, creating what Pariser (2011) called the filter bubble, depriving adolescents of the encounters with difference that are essential for identity exploration.

At the mesosystemic level, the relationship between schools and families in mediating AI use is critically important and critically underexplored. Schools that have implemented AI tools have done so at enormously varying speeds and with varying pedagogical deliberation. Families, for their part, are largely navigating AI without institutional guidance, making decisions shaped more by marketing, peer norms, and convenience than by developmental evidence. The mesosystemic connection between school and family, potentially one of the most important mediators of AI's developmental effects, remains largely unstructured and unsupported.

5.3 The Exosystem and Macrosystem: Communities, Inequality, and Structural Conditions

At the community level, AI's effects must be understood in relation to social capital. Putnam (2000) distinguished between bonding social capital, the ties that bind people within similar groups, and bridging social capital, the ties that connect people across different social groups. Both forms depend on conditions of direct, reciprocal interaction. When AI systems mediate, optimise, or substitute for these interactions, the conditions under which social capital is produced and maintained are altered. At the local level, the digitalisation of community services through AI-driven interfaces may reduce the relational dimensions of service encounters that have historically served important social functions: a claimant navigating a benefits system through an AI chatbot may receive technically adequate service while being deprived of the relational recognition, the experience of being seen and heard by another person, that gives social support much of its protective power.

The distributional dimension of AI's effects is where the analysis becomes most morally urgent. Noble (2018) demonstrated through detailed case analysis that AI search and recommendation systems trained on existing data reproduce and,
in some cases, amplify historical patterns of racial and gender discrimination. Eubanks (2018) extended this to public services, showing that algorithmic decision systems deployed in child welfare, benefits administration, and criminal justice systematically disadvantaged already marginalised populations. Benjamin (2019) coined the term 'the new Jim Code' to describe how race-neutral technical language can encode and perpetuate racial hierarchy. The digital divide literature reinforces this analysis: access to high-quality internet, devices, and the skills needed to use them effectively remains unequal across socioeconomic and geographic lines (van Dijk, 2020). Communities with limited digital literacy are not simply unable to benefit from AI; they are increasingly unable to navigate institutional environments that have been reorganised around AI-mediated processes.

6. Critical Appraisal of the Evidence Base

The scholarship reviewed here, while growing rapidly and in many respects compelling, has significant limitations that must be acknowledged before recommendations can be responsibly drawn. The most persistent methodological problem is the difficulty of establishing causal direction. Most studies are cross-sectional or correlational in design. They can demonstrate that AI or screen use is associated with outcomes but cannot establish whether AI use produces those outcomes or whether individuals with pre-existing characteristics are more likely to use AI in particular ways. The reciprocal causation hypothesis, that weak critical thinking predicts AI dependence and AI dependence further weakens critical thinking, is plausible and suggested by the theoretical frameworks here, but requires longitudinal experimental designs to establish.

A second limitation is the near-exclusive focus on student populations, particularly university

students in high-income English-speaking countries. This sampling bias constrains generalisability considerably. The experience of a university student using an AI writing tool in a well-resourced North American institution is not comparable to that of a child using an AI tutoring application on a shared smartphone in a rural Pakistani village, or an older adult receiving AI-mediated social care in a Portuguese municipality. Third, the field suffers from a severe lag between technological development and research. Large language models of the capability available from 2022 onwards have been in wide use for a short time, and longitudinal effects, arguably the most important for developmental questions, will not be visible in the literature for years or decades. This means that policy and practice decisions about AI in education, social services, and community life are necessarily being made under conditions of significant empirical uncertainty.

Fourth, a critical review must not elide the genuine benefits of AI. AI tools have demonstrated meaningful utility in supporting learners with disabilities, providing educational access in underserved contexts, enabling early identification of child protection concerns, and reducing the administrative burden on social workers (UNESCO, 2023). The question is not whether AI produces only harm or only benefit, but under what conditions, for whom, and with what governance structures its deployment promotes human flourishing rather than diminishing it. Finally, there is a conceptual limitation in the tendency of much of this literature to treat AI's effects as primarily an individual-level phenomenon. An account that focuses only on individual cognitive outcomes, without attending to the family, school, community, institutional, and cultural conditions in which those effects occur, will inevitably be both partial and analytically weak.

7. Conclusions

This review has argued that the effects of artificial intelligence on natural human abilities are real, structurally significant, and developmentally consequential but neither uniform nor inevitable. The evidence reviewed supports concern about the effects of habitual, unreflective AI use on critical thinking, memory, emotional intelligence, creativity, and communication, particularly for children and youth who are in the formative stages of developing these capacities. At the same time, the evidence does not support technological fatalism or moral panic. AI, used thoughtfully, intentionally, and equitably, can augment human capacity, expand access to learning and support, and strengthen rather than undermine the human abilities that matter most.

The bioecological framework applied throughout this review has illuminated something important: the effects of AI on human development cannot be understood at the individual level alone. They are shaped by the microsystems in which children develop, the mesosystemic linkages between family and school, the exosystemic institutions that deploy AI in ways that shape but are not directly experienced by individuals, and the macrosystemic cultural values and policy frameworks that determine what human capacities are valued, invested in, and protected. Responding to AI's challenges, therefore, requires action at all these levels simultaneously, not a single intervention but a coordinated ecological response.

The analysis has also highlighted distributional dimensions that the field cannot afford to ignore. The costs and benefits of AI do not distribute themselves neutrally. Children in under-resourced environments, communities with limited digital literacy, and populations historically marginalised by algorithmic systems bear a disproportionate share of AI's risks while accessing fewer of its benefits. A scholarship and policy response that attends only to the average effects of AI on the average

user is not only analytically incomplete; it is ethically inadequate. What this review ultimately argues for is developmental intentionality in the governance and use of AI: a commitment, grounded in developmental evidence and ecological thinking, to ensuring that the environments in which children, young people, and communities develop remain rich in the conditions of effortful practice, genuine social experience, exposure to difference, and human relational quality that human flourishing requires. This is not a romantic project. It is a practical and urgent one.

Future research should prioritise longitudinal experimental designs establishing causal direction in the relationships between AI use and specific cognitive and social outcomes. It should attend far more seriously to cultural, economic, and geographic diversity in samples and contexts. It should examine the mesosystemic linkages between home and school as sites of mediated AI use with particular care. And it should engage more directly with the perspectives of children, young people, and community members themselves, the most important and least consulted experts on the experience of growing up and living in a world increasingly shaped by artificial intelligence.

8. Recommendations

8.1 For Families and Caregivers

Caregivers should establish age-appropriate norms around AI use informed by the developmental logic reviewed here: children need to develop the underlying cognitive and social competencies before they routinely delegate them to AI tools. This means developmental sequencing, not prohibition. Families should cultivate deliberate AI-free spaces and practices, shared meals without devices, reading aloud, creative play, outdoor activity, and unstructured peer interaction, not as nostalgic luxuries but as the developmental contexts in which critical thinking, emotional intelligence, and communication capacity are built. Alongside this, families should model reflective AI use: discussing with children and young people what AI does well, what it does poorly, and why questioning its outputs is an act of intelligence rather than technological hostility.

8.2 For Educational Institutions

Schools and universities carry responsibility for the conditions under which cognitive and social capacities develop. Educational institutions should develop explicitly developmental AI policies specifying not only when AI use is permitted or prohibited but why, grounded in a coherent account of what capacities each task is designed to build. AI literacy curricula, covering how AI systems work, their limitations, how to evaluate AI-generated content, and the societal implications of AI, should be integrated into formal education from primary school, with critically deepening treatment through tertiary levels. Assessment practices require urgent rethinking: systems that measure learning through the production of outputs without attending to the process of reasoning are failing to assess what matters most. Process-based assessment, oral examination, and collaborative problem-solving tasks are not merely AI-resistant; they are better aligned with what the evidence says about effective learning. Targeted support is warranted for students from disadvantaged backgrounds who may be less likely to have the metacognitive resources to use AI tools in augmentative rather than substitutive ways.

8.3 For Social Work Practitioners

Social work as a profession is directly implicated in the issues this review has examined. At the direct practice level, social workers should incorporate AI use into their ecological assessment of clients' developmental environments, attending to patterns of use, the ages at which AI tools are introduced, and the quality of adult mediation

around children's AI use. At the community and group practice level, social workers should actively protect and create contexts for human-to-human connection and collective problem-solving that are not AI-mediated. At the policy and advocacy level, social work has both a professional and an ethical obligation to advocate for algorithmic accountability in AI systems used in child welfare, benefits administration, and criminal justice; for the digital inclusion of communities with limited AI access; and for the preservation of human professional judgment in welfare decisions with significant consequences for vulnerable people's lives.

8.4 For Community Organisations and Policymakers

Civil society organisations, libraries, cultural centres, youth clubs, community health organisations, occupy an irreplaceable role in creating contexts for human development and social connection that are neither family nor school. These organisations should be understood and supported as essential infrastructure for the maintenance of human cognitive and social capacity in an AI-saturated environment. AI literacy in community contexts should be understood not primarily as a technical skill but as a civic capacity: the ability to understand how AI affects one's community, to identify and challenge AI-mediated injustices, and to participate in collective decisions about how AI is used in shared life.

At the governmental level, age-appropriate AI regulation is urgently needed, including minimum age standards for AI systems that simulate social or emotional relationships, mandatory developmental impact assessments for AI tools marketed to children, and transparency requirements for AI systems used in educational contexts. Public investment in digital literacy education, understood broadly to include critical AI literacy, should be treated as a policy priority comparable to basic literacy and numeracy.

Algorithmic accountability legislation, requiring that AI systems used in high-stakes public decisions be subject to independent audit for bias, accuracy, and proportionality, is a basic requirement of social justice. International cooperation on AI governance is essential given the global reach of AI systems and the inadequacy of purely national responses.

References

Habib, A., Harris, L., Pollick, F., & Melville, C. (2019). A meta analysis of working memory in individuals with autism spectrum disorders. *PLoS ONE, 14*(4), e0216198.

An, W., Ding, Z., & Zhu, Z. (2025). Verbal working memory in children and adolescents with autism spectrum disorders: A meta analysis. *Acta Psychologica*.

Han, Y. M. Y., Chan, M. C., & Chan, M. M. Y. (2022). Effects of working memory load on frontal connectivity in children with autism spectrum disorder. *Scientific Reports, 12*, 3456.

Wang, Y., Zhang, Y. B., & Liu, L. L. (2017). Working memory impairments in autism spectrum disorders: Meta analysis review. *Neuropsychological Review, 27*(3), 275–290.

Williams, D., Happé, F., & Jarrold, C. (2014). Working memory and autism: A review of literature. *Research in Autism Spectrum Disorders, 8*(8), 1071–1084.

Rahko, J. S., et al. (2016). Attention and working memory in adolescents with autism spectrum disorder: fMRI study. *Child Psychiatry and Human Development, 47*(2), 245–257.

Baddeley, A. D., & Hitch, G. J. (1974). Working memory. In G. Bower (Ed.), *The Psychology of Learning and*

Motivation (Vol. 8, pp. 47–89). Academic Press.

DSM 5 TR. (2022). *Diagnostic and Statistical Manual of Mental Disorders* (5th ed., Text Revision). American Psychiatric Association.

Landa, R., & Goldberg, M. (2005). Language, social and executive functions in children with autism spectrum disorders. *Developmental Neuropsychology, 27*(3), 311–330.

Lopez, B. R., Lincoln, A. J., Ozonoff, S., & Lai, Z. (2005). Examining the relationship between executive functions and restricted, repetitive symptoms of autistic disorder. *Journal of Autism and Developmental Disorders, 35*(4), 445–460.

Joseph, R. M., & Tager Flusberg, H. (2004). Executive dysfunction and its relation to language ability in children with autism. *Developmental Neuropsychology, 26*(2), 465–486.

Geurts, H. M., Verté, S., Oosterlaan, J., Roeyers, H., & Sergeant, J. A. (2004). How specific are executive functioning deficits in attention deficit hyperactivity disorder and autism? *Journal of Child Psychology and Psychiatry, 45*(4), 836–854.

Kenworthy, L., Yerys, B. E., Anthony, L. G., & Wallace, G. L. (2008). Understanding executive control in autism spectrum disorders in the lab and in the real world. *Neuropsychology Review, 18*(4), 320–338.

Ozonoff, S., Strayer, D. L., McMahon, W. M., & Filloux, F. (1994). Executive function abilities in autism and Tourette syndrome: An information processing approach. *Journal of Child Psychology and Psychiatry, 35*(6), 1015–1032.

Williams, D., Jarrold, C., Grainger, C., & Lind, S. E. (2012). Executive function in autism spectrum disorders. In T. E. Zelazo (Ed.), *The Oxford Handbook of Developmental Psychology, Vol. 1: Body and Mind* (pp. 911–943). Oxford University Press.