

Cognifying Education: Mapping AI's transformative role in emotional, creative, and collaborative learning

Mikael Gorsky¹ and Ilya Levin²

¹Holon Institute of Technology, Holon, Israel, MikaelG@hit.ac.il

²Holon Institute of Technology, Holon, Israel, LevinI@hit.ac.il

Abstract

Artificial intelligence (AI) is rapidly reshaping educational practice, challenging long-held assumptions about teaching and learning. This article integrates conceptual perspectives from recent books (Genesis by Kissinger et al., Co-Intelligence by Mollick, and The Inevitable by Kelly) with empirical insights from popular AI podcasts and Anthropic's public releases. We examine seven key domains – emotional support, creativity, contextual understanding, student engagement, problem-solving, ethics and morality, and collaboration. For each domain, we explore AI's capabilities, opportunities for transformative change, and emerging best practices, drawing equally from theoretical analysis and real-world observations. Overall, we find that AI, when used thoughtfully, can complement and enhance human educators in fostering richer learning experiences across cognitive, social, and emotional dimensions. We emphasize an optimistic yet responsible outlook: educators and students should actively shape AI integration to amplify human potential in creativity, ethical reasoning, collaboration, and beyond, while maintaining a focus on human-centric values.

1. Introduction

AI's surge in capability – epitomized by large language models like GPT-4 and Claude – has sparked intense debate in education. Early public opinion often casts AI as either a threat to the teacher's unique role or a mere automator of rote tasks. However, a closer look reveals a more nuanced reality: AI systems, under proper guidance, can extend the reach of educators and enrich student learning in unprecedented ways. We see that AI can have significant impact in seven areas: (1) emotional support, (2) creativity, (3) contextual understanding, (4) student engagement, (5) problem-solving, (6) ethics/morality, and (7) collaboration. These seven themes form the backbone of our inquiry.

Recent books by prominent thinkers provide frameworks for understanding how AI might transform human roles and capabilities. Kissinger, Schmidt, and Mundie [2] offer a broad geopolitical and philosophical context for AI's impact on human knowledge and agency. Mollick [3] focuses on practical human-AI collaboration in work and education, envisioning AI as a tutor, coach, and creative partner. Kelly [1] positions AI as an inevitable force driving us to “cognify” every aspect of life, arguing that humans must learn to work *with* intelligent machines rather than against them.

From the empirical side, popular AI-themed podcasts (such as *Hard Fork*, *The Ezra Klein Show*, *Latent Space*, and *The Cognitive Revolution*) and public content from AI research labs provide real-world insights. Teachers and technologists speaking on these platforms describe how AI is *already* being used in classrooms and what challenges and successes are emerging. For example, on *Hard Fork*, Wharton professor Ethan Mollick discussed how he integrated generative AI into his teaching, requiring students to use AI for certain assignments rather than banning it [7]. Such firsthand accounts complement data-driven reports like Anthropic's large-scale studies of AI usage by students [4], and Anthropic's technical blogs on aligning AI to human values [6]. Taken together, these sources paint an exploratory yet optimistic picture of AI's transformative potential in education – provided we approach integration thoughtfully and ethically.

In the sections that follow, we examine each of the seven themes in turn. For each domain, we contrast common initial assumptions with the emerging reality of AI-augmented education. We draw on conceptual arguments from thought leaders, empirical findings from research and practice, and examples that illustrate how AI can be harnessed in service of deeper learning. Throughout, we maintain an optimistic tone, highlighting opportunities and effective strategies while acknowledging the need for ongoing vigilance (especially regarding ethics and equity). Finally, we synthesize these insights to discuss how educators can proactively shape an AI-enhanced future of learning, rather than passively respond to technological change.

2. AI and Emotional Support

One oft-repeated belief is that only human teachers can provide the empathic, emotionally supportive presence students need, whereas AI would be “cold” or uncomprehending. Indeed, emotional intelligence and genuine empathy are deeply human traits. However, early implementations and studies suggest that AI, while not conscious or emotional itself, can simulate empathic support in ways that meaningfully help students. By analyzing textual cues, voice tone, and even facial expressions, AI can detect subtle emotional states and respond with relevant emotional messages. For example, an AI tutor might notice frustration or confusion in a student's queries and interject with encouraging feedback or a clarifying question in a gentle tone. While the AI does not feel empathy, it can generate empathetic responses appropriate to the situation. In some respects, this approach can even surpass human consistency – an AI will not become impatient or burned-out, and it can offer 24/7 support with endless patience, which is valuable for students who need steady reassurance or a non-

judgmental listener. Moreover, AI systems can track a student's emotional patterns over time, potentially flagging declines in engagement or signs of distress that a busy teacher might miss. This capability to monitor well-being longitudinally and objectively could enable earlier interventions for struggling students.

Recent evidence from 2024-2025 strongly validates AI's competency in emotional support. A comprehensive study involving 401 Chinese university students found that students perceive AI as having the potential to positively influence mental well-being, with researchers emphasizing the potential of AI in strengthening mental health support systems because AI-powered chatbots and virtual assistants provide immediate support and essential information, making mental health services more accessible to a wider audience [23]. Most remarkably, the first-ever randomized controlled trial of a generative AI-powered therapy chatbot demonstrated significant clinical improvements, with participants showing significantly greater reductions in symptoms of MDD (mean changes: -6.13 vs. -2.63 at 4 weeks), GAD (mean changes: -2.32 vs. -0.13 at 4 weeks), and CHR-FED (mean changes: -9.83 vs. -1.66 at 4 weeks) relative to controls [24]. Crucially, Therabot was well utilized (average use >6 hours), and participants rated the therapeutic alliance as comparable to that of human therapists [24]. Additional evidence comes from real-world usage studies of nineteen individuals using generative AI chatbots for mental health, which revealed high engagement and positive impacts, including better relationships and healing from trauma and loss, with participants developing four key themes including a sense of 'emotional sanctuary', 'insightful guidance', particularly about relationships, the 'joy of connection', and comparisons between the 'AI therapist' and human therapy [25].

Recent empirical evidence supports the notion that students (and people in general) are beginning to use AI as a form of on-demand emotional support. Anthropic reported that although the majority of AI usage is task-oriented, a small but meaningful fraction (about 2.9%) of interactions with their model Claude are what they term "affective conversations" – users seeking advice, counseling, or companionship from the AI [6]. Within those supportive dialogues, people commonly discuss personal challenges (relationships, loneliness, anxiety) and receive advice or coaching. Notably, researchers found that human sentiment tends to become more positive over the course of a counseling or coaching conversation with Claude, suggesting that the AI's responses often help improve the user's mood during the interaction [20]. This aligns with anecdotal reports from podcasts: for instance, guests on The Ezra Klein Show and others have described teens chatting with AI "friend" apps late at night when they feel they have no one else to talk to – illustrating both the opportunity and the concern in this domain [8]. On the one hand, an AI companion can reduce feelings of loneliness by being an ever-available conversational partner; on the other hand, as technologists caution, over-reliance on AI for companionship could introduce new forms of social isolation or blurred reality [21]. Mollick [3] imagines a future scenario in which "AI companions become far more compelling to speak with than most other people, and ... loneliness becomes less of an issue" even as "some people would rather interact with AIs than with humans," leading to tricky societal questions [3].

2.1 What makes AI emotionally competent: technical mechanisms

Understanding why AI systems demonstrate such competency in emotional support requires examining the underlying technical mechanisms that enable this capability. Recent research reveals that large language models have achieved remarkable proficiency in emotional intelligence tasks, with ChatGPT-4, ChatGPT-o1, Gemini 1.5 flash, Copilot 365, Claude 3.5 Haiku, and DeepSeek V3 outperformed humans on five standard emotional intelligence tests, achieving an average accuracy of 81%, compared to the 56% human average reported in the original validation studies [26]. This superiority extends to sophisticated emotional understanding, as demonstrated by research showing GPT-4 is capable of emotion identification and managing emotions, but it lacks deep reflexive analysis of emotional experience and the motivational aspect of emotions [27], indicating that while AI may not experience emotions, it can recognize and respond to them effectively.

The foundation of AI's emotional competency lies in the transformer architecture and its attention mechanisms. The transformer's self-attention mechanism allows it to examine an entire sequence simultaneously and make decisions about how and when to focus on specific time steps of that sequence [28], enabling the AI to identify emotional cues across different parts of a conversation and understand their relationships. The self-attention mechanism in Transformers has significantly enhanced their capacity to capture complex long-range dependencies within and across various modalities [29], which proves crucial for understanding emotional context that may span multiple exchanges. When applied to emotional recognition, Transformers offer several advantages. Firstly, their self-attention mechanism enables them to dynamically focus on relevant parts within each modality-specific representation. This mechanism allows Transformers to effectively weigh the importance of different features within the input data, therefore capturing subtle emotional nuances that may be essential for accurate recognition [29].

The training methodology known as Reinforcement Learning from Human Feedback (RLHF) plays a critical role in developing AI's emotional competency. RLHF incorporates human feedback in the rewards function, so the ML model can perform tasks more aligned with human goals, wants, and needs [30], particularly important for emotional support where the score can be based on innately human qualities, such as friendliness, the right degree of contextualization, and mood [30]. As one technical overview explains, ChatGPT's technical foundation, including its Transformer architecture and RLHF (Reinforcement Learning from Human Feedback) process, enabling it to generate human-like responses [31], demonstrates how these combined approaches result in emotionally responsive AI systems. This training process enables AI to learn subtle patterns of emotional interaction that would be difficult to program explicitly, allowing models to develop what researchers describe as elements of cognitive empathy, recognizing emotions and providing emotionally supportive responses in various contexts [32].

Overall, the optimistic view is that AI can augment emotional support in education by providing individualized encouragement, empathy simulations, and coaching that supplement (not replace) human care. For example, an AI-guided learning app might cheer on a student who is struggling (“I know this is tough, but I believe you can do it!”) and suggest a short break if it senses frustration. It might help a shy student practice social skills by role-playing difficult conversations in a safe setting. Educators are already experimenting with these uses: some counseling centers have tested AI chatbots for cognitive-behavioral therapy exercises or mood tracking with college students. Crucially, however, such AI tools must be designed and monitored carefully – ethical guidelines are needed to ensure the AI remains supportive and does not produce harmful advice. Anthropic’s work on AI safety emphasizes aligning models with human values so that they behave as a “good citizen” – for instance, refusing to exploit users’ emotions or feed delusions [22]. In practice, the best outcome may be achieved by human-AI collaboration in pastoral care: teachers and counselors could use AI to identify students in need of help and to scale up support, while humans provide the authentic empathy and judgment that machines lack. If approached in this balanced way, AI’s tireless presence and pattern-recognition could bolster the emotional safety net for students, ensuring no one slips through the cracks unnoticed.

3. AI and Creativity

Another common belief is that AI, being fundamentally computational, excels at repetitive or analytical tasks but cannot foster human creativity and imagination. Creativity in education has traditionally been nurtured by inspiring teachers who encourage students to think divergently, explore art and play, and take intellectual risks. It might seem counterintuitive that a software program could enhance these creative processes. Yet emerging evidence suggests that AI can be a powerful catalyst for creativity, both by generating novel ideas itself and by stimulating greater creativity in students. AI fosters creativity by processing and recombining vast amounts of information in novel ways. Because modern AI models are trained on enormous datasets encompassing literature, art, history, science and more, they effectively serve as “connection machines” (to use Mollick’s term) that can link ideas across domains in unexpected ways. Where a human might be limited by personal experience or conventional thinking, an AI can produce an unusual analogy, an offbeat suggestion, or a hybrid concept that spurs a human learner’s imagination. For example, a student brainstorming a story might ask an AI for ideas, and it could propose a plot combining elements of, say, ancient mythology and futuristic science – a combination the student hadn’t considered. These AI-generated surprises can “present unexpected ideas or challenge norms, potentially sparking human creativity” [24]. Far from merely copying existing material, a well-designed AI can create genuinely new combinations and concepts from the knowledge it has absorbed [25]. In fact, AI’s ability to quickly generate and evaluate multiple scenarios promotes divergent thinking in students [26][27]. By rapidly proposing many variations or solutions to a problem, the AI encourages students to move beyond their first idea and consider alternatives, thereby expanding their creative comfort zones [26][28].

3.1 Recent evidence of AI's creative competency

Extensive recent research from 2024-2025 provides compelling evidence that AI demonstrates significant creative competency across multiple domains. A landmark study analyzing over 4 million artworks from more than 50,000 unique users found that text-to-image AI significantly enhances human creative productivity by 25% and increases the value as measured by the likelihood of receiving a favorite per view by 50% [33]. While peak artwork Content Novelty increased over time, this research revealed that the artists who successfully explore novel ideas and filter model outputs for coherence benefit the most from AI tools, underscoring the pivotal role of human ideation and artistic filtering in determining an artist's success with generative AI tools [33].

In educational contexts specifically, a 2024 study involving college students found that 100% of participants found AI helpful for brainstorming, with students generating more diverse and detailed ideas when using AI [34]. AI served as a useful brainstorming tool for kick-starting creative sessions and acted as a nonjudgmental partner for idea generation, allowing students to explore concepts they might normally withhold in group settings [34]. Furthermore, a comprehensive study published in 2025 reported that 91 percent of educators observe enhanced learning when their students utilize creative AI [35], indicating widespread recognition of AI's positive impact on student creativity.

Recent research has also demonstrated AI's ability to match and sometimes exceed human creative performance in standardized tests. In psychological tests of creativity, such as evaluating ChatGPT-4's performance on creative interpretation tasks using the Figural Interpretation Quest (FIQ), results indicated that while AI on average demonstrated higher flexibility in generating creative interpretations, the evaluation revealed nuanced differences in how creativity manifests between human and artificial systems [36]. Perhaps most remarkably, a controlled experiment involving short story creation found that access to generative AI ideas causes stories to be evaluated as more creative, better written, and more enjoyable, especially among less creative writers [37], though this same study noted important implications for collective diversity that we discuss further below.

Empirical results back up these conceptual claims. In psychological tests of creativity, such as the Alternative Uses Test (which measures how many different uses for a common object someone can propose), advanced AI systems have scored remarkably well. Ethan Mollick reports that GPT-4 outperformed over 90% of human test-takers in generating creative ideas for novel uses of everyday items, as judged by human evaluators [3]. While such tests have limitations (AI might have seen examples during training), they indicate that AI can produce a breadth of imaginative outputs comparable to very creative humans [3]. Moreover, AI's "creativity" isn't limited to text – generative models in visual art and music (e.g. DALL-E, Midjourney, MuseNet) can produce original images and compositions. Teachers on podcasts have shared anecdotes of students using AI art generators to visualize their ideas, which then inspires them to further refine or re-imagine their projects. Rather than make students passive, these tools often motivate them to iterate and experiment more boldly. For instance, Latent Space podcast guests noted how generative AI allows even novice programmers or designers to prototype creative projects (like video games or animations) quickly, thereby lowering the

barrier to entry for creative expression. This democratization of creativity – giving every student a tireless creative assistant – is a theme echoed by Kevin Kelly: “It’s not a race against the machines. You’ll be paid in the future based on how well you work with robots... Robots will do jobs we have been doing, and do them much better... And they will help us discover new jobs... They will let us focus on becoming more human than we were.” [1]. In the context of the classroom, “becoming more human” means focusing on the uniquely human aspects of creativity: defining the problems, infusing work with empathy and values, and making aesthetic or ethical choices – while the AI handles tedious details or generates raw material for inspiration.

3.2 Technical mechanisms enabling AI Creativity

Understanding why AI demonstrates such remarkable creative capabilities requires examining the underlying computational mechanisms that enable this performance. The foundation of AI creativity lies in the transformer architecture, which utilizes self-attention mechanisms to process input text and identify relationships between different elements across vast distances in the data [38]. This architecture enables what researchers describe as combinatorial creativity, where AI systems identify, retrieve, and recombine relevant concepts from multiple domains through structured processes [39].

At its core, AI creativity operates through what computational creativity researchers call “a search process through the space of possible combinations” [40]. The combinations can arise from composition or concatenation of different representations, or through rule-based or stochastic transformation of initial and intermediate representations [40]. Large language models achieve this through several key mechanisms: first, their stochastic nature and the variety of prompts that are usually provided commonly lead to novel outcomes [41]; second, their training on all available data allows them to access and recombine elements from across the entire knowledge spectrum they’ve been exposed to [41].

Recent advances in understanding LLM creativity reveal that these systems can realize combinatorial creativity by generating creative ideas through structured combinatorial processes [42]. Research demonstrates that LLMs can systematically explore and combine ideas from existing literature by analyzing conceptual relationships across papers and identifying novel research opportunities [42]. This process involves cross-domain connections by embedding and comparing ideas at multiple abstraction levels, allowing the system to connect ideas from unrelated fields while preserving traceability through structured formats that capture relationships between retrieved innovations and problem abstractions [39].

The attention mechanism plays a crucial role in enabling creativity by allowing the model to examine an entire sequence simultaneously and make decisions about how and when to focus on specific time steps of that sequence [38]. This mechanism enables LLMs to identify patterns and relationships that span long distances in text, facilitating the kind of unexpected connections that characterize creative thinking. Furthermore, the multi-head attention mechanism allows transformers to focus on different aspects of the input simultaneously, enabling them to capture multiple types of relationships and patterns that contribute to creative output [38].

Training methodologies also contribute significantly to AI's creative capabilities. Reinforcement Learning from Human Feedback (RLHF) enables models to be fine-tuned to maximize human preferences and values, including creative preferences [43]. This training approach helps AI systems learn not just what is statistically probable based on training data, but what humans actually find creative, valuable, and engaging. Additionally, techniques like diffusion models and generative adversarial networks contribute to creative generation by enabling the production of novel content that maintains coherence while exploring new possibilities [43].

Practical classroom implementations of AI to boost creativity are already underway. Some educators use AI as a creative brainstorming partner for students: for example, asking a language model to act as a “wild idea generator” in a group project ideation session. One teacher described having students’ critique and build on AI-generated solutions to a design challenge; the offbeat AI suggestions pushed students to think outside the box and eventually led to more innovative student-designed solutions than in previous semesters. Mollick [3] has even made “use of AI mandatory” in certain entrepreneurship assignments – he instructs students to attempt projects far beyond their normal capability, explicitly because AI is available to help [3]. One prompt read: “Make what you are planning on doing ambitious to the point of impossible; you are going to be using AI... I won’t penalize you for failing if you are too ambitious.” [3]. The result is that students aim higher (e.g. attempting to create a working app or write a business plan in a week) and often achieve more than they originally thought they could, learning a great deal in the process about creative problem-solving. AI provides the scaffolding – writing code snippets, generating marketing copy, simulating user feedback – but students remain the directors of these projects, exercising higher-level creativity in selecting, refining, and integrating the AI’s contributions. In this sense, AI can act like a creative muse or apprentice: it multiplies the range of ideas on the table and handles low-level execution, freeing students to focus on imaginative and integrative thinking [32][33].

Of course, using AI for creativity in education requires thoughtful guidance. Teachers must ensure students are not simply taking AI outputs at face value or plagiarizing creative work. As Mollick notes, there is a paradox in AI creativity: the same randomness that lets AIs generate novel ideas can also produce misinformation or unfiltered content [3]. Therefore, educators play a crucial role in coaching students to critically evaluate and iterate on AI-generated content. When harnessed properly, AI can become a powerful amplifier of creativity in the classroom – expanding the horizons of what students can imagine and make and preparing them for a future where human-AI creative collaboration is the norm.

4. AI and Contextual Understanding

Teachers are often lauded for their ability to understand the personal, social, and cultural context of each student – the myriad factors outside of test scores that affect learning. A common sentiment is that AI lacks this holistic understanding and therefore cannot truly personalize learning or respond to individual needs the way a human teacher can. Indeed, human educators draw on intuition and experience to pick up on subtle cues (a student’s body language,

community background, current events affecting the student, etc.). However, advanced AI systems are increasingly capable of analyzing vast and varied data streams to construct a rich picture of a learner's context. The "truth" here, as articulated in the user's framework, is that AI's contextual understanding comes from its ability to process and integrate multiple data streams simultaneously. An AI tutor could, in principle, ingest a student's academic records, responses to past homework, data from educational games, and even biometric or environmental data (if available and ethically obtained) to continuously assess the student's state and needs. For example, consider an AI-driven learning platform that monitors which topics a student struggles with, how quickly they answer questions at different times of day, which concepts excite them (based on engagement patterns), and perhaps input from the student's wearable fitness tracker indicating sleep or stress levels. By synthesizing these diverse inputs, the AI can form a far more detailed and up-to-date model of the student's context than any one teacher managing 30 students might have time for. As the user's article notes, AI can analyze academic performance, social interactions, physiological data, and environmental factors together. This enables nuanced interpretations of student behavior and performance that might escape human observation. For instance, an AI might correlate that a student performs better on math exercises in the morning and struggles after 2pm, which, combined with knowledge that the student has afternoon sports practice (physical fatigue), leads to a suggestion to do math homework earlier in the day. Or it might detect that a normally active student has stopped asking questions in forum discussions, flagging a possible disengagement issue.

4.1 Recent evidence of AI's contextual competency

Extensive recent research from 2024-2025 demonstrates that AI systems have achieved remarkable competency in contextual understanding through sophisticated data integration and adaptation mechanisms. A comprehensive bibliometric analysis of adaptive learning technologies reveals that AI systems now leverage data analytics and machine learning algorithms to provide personalized instruction by adapting content, pace, and delivery based on individual student strengths, weaknesses, and learning preferences [48]. Implementation of adaptive learning platforms leads to higher pass rates and improved student retention compared to traditional teaching methods [48], indicating the practical effectiveness of AI's contextual understanding capabilities.

Recent developments in multimodal AI systems represent a particular breakthrough in contextual understanding. Current large multimodal foundation models have the power to process spoken text, music, images and videos simultaneously [49], enabling them to construct rich contextual pictures from diverse data streams. These systems demonstrate advanced contextual awareness that make cognitive tutors more effective in assisting with ill-defined problems [49]. For instance, Google Gemini, as a multimodal generative AI tool, demonstrates revolutionary potential by processing data from text, image, audio, and video inputs while generating diverse content types [50]. This AI tool can create differentiated materials, design multiple activities for different levels of students, and provide additional explanations for those who need extra support, all while analyzing learners' work to offer personalized feedback and identifying areas for further improvement [50].

A groundbreaking study in 2024 involving multimodal learning analytics demonstrates how AI can support teacher, researcher, and AI collaboration in STEM learning environments by creating AI-generated multimodal timelines that amalgamate diverse data types: students' emotional responses, synergy scores, social interaction metrics, prosodic audio cues, verbatim conversation transcripts, prior physics and computing knowledge, and detailed learning analytics [51]. This integration enables unprecedented contextual insights that assist teachers in identifying student challenges and crafting supportive feedback [51].

Furthermore, recent research indicates that future directions in adaptive learning will involve integrating contextual information to further personalize the learning experience, including incorporating data from wearable devices, environmental sensors, or other sources to adapt content based on factors such as location, time, or learner's emotional state [52]. Context-aware adaptation enables adaptive e-learning systems to provide even more tailored relevant content that is aware of learners to help them complete their activities [52]. These systems are increasingly incorporating collaborative and social learning components, with AI/ML algorithms analyzing learner interactions, group dynamics, and social network data to provide personalized recommendations for group projects, collaborative learning activities, and peer feedback [52].

Anthropic's recent Education Report provides a striking real-world glimpse of AI's contextual savvy at work. Analyzing over half a million student conversations with the AI model Claude, the study identified distinct patterns in how students interact with AI – including “collaborative problem solving” sessions where the AI and student dialog back-and-forth to reach understanding [42]. Importantly, students were using Claude across many disciplines and tasks, from debugging code to explaining law concepts, often in a highly personalized way (e.g. asking for explanations tailored to their level of prior knowledge) [4]. The data show that students primarily use AI to create and improve content in ways that require understanding context, such as drafting essays with personalized feedback or getting study guidance for specific courses [4]. Claude's design allows it to maintain long conversations and refer back to earlier parts of the dialogue – effectively remembering what the student said or struggled with before. This indicates that AI tutors can simulate contextual continuity in a similar manner to a human tutor who remembers a student's progress over weeks and months. Additionally, Anthropic's “Claude for Education” initiative explicitly emphasizes giving institutions a secure AI that “understands your context” – it can bring together documents, tools, and web knowledge in a conversation, making it a context-aware assistant for both students and teachers [43]. For example, Claude can be connected to a student's project documents or class notes; it will then incorporate that specific material when answering the student's questions, ensuring responses are relevant to the student's current curriculum and situation. This context-connecting ability is a game-changer for personalization: the AI isn't just spitting out generic answers, but rather it “knows” what the student is working on and can adapt accordingly. Anthropic's CEO, Dario Amodei, has suggested in interviews that the goal is for AI tutors to one day “know every student as holistically as a dedicated personal tutor would”, albeit through data rather than direct human empathy.

4.2 Technical mechanisms enabling contextual understanding

Understanding why AI demonstrates such sophisticated contextual understanding requires examining the fundamental architectural and algorithmic mechanisms that enable this capability. The foundation of AI's contextual competency lies in the transformer architecture and its revolutionary attention mechanism, first introduced in the landmark 2017 paper "Attention Is All You Need" [53]. The attention mechanism is a machine learning technique that directs deep learning models to focus on the most relevant parts of input data, enabling models to selectively focus on relevant parts of input sequences and thereby incorporating context sensitivity into the representation learning process [54].

At its core, the self-attention mechanism in transformers enables the model to consider all parts of the input when generating responses, regardless of their position in the sequence [55]. This mechanism calculates the relationships and dependencies between different tokens, even those far apart in the input sequence, allowing the model to understand how words at the beginning of a context relate to those at the end [55]. The self-attention mechanism computes weights indicating the relevance of each token to the others, creating what researchers describe as a contextual understanding framework where each element in the input sequence attends to all others, enabling the model to capture global dependencies [56].

The practical implementation of contextual understanding in AI systems depends critically on the concept of context windows - the amount of information an AI system can consider at once when processing input and generating responses. Modern AI systems have dramatically expanded their context windows, with recent models like Gemini 1.5 processing context windows of up to 1 million tokens, and Anthropic's Claude models offering context windows of 200,000 to 500,000 tokens [57]. These extended context windows enable AI systems to maintain awareness of much larger spans of conversation, documentation, and contextual information simultaneously.

The multi-head attention mechanism further enhances contextual understanding by performing multiple parallel self-attention operations, each with its own set of learned query, key, and value transformations [54]. This allows the model to capture different aspects of relationships between words in sequences simultaneously, enabling more nuanced contextual interpretation. Each attention head learns different linear projections, allowing the model to focus on different types of relationships - syntactic, semantic, temporal, and contextual - within the same input [54].

Recent advances in multimodal AI systems represent another breakthrough in contextual understanding mechanisms. These systems use transformer-based architectures capable of handling multimodal inputs simultaneously, enhancing the integration process by synchronizing and processing inputs from various modalities such as text, audio, and visual data [53]. For instance, multimodal systems can process students' emotional responses, audio cues, conversation transcripts, and visual behavior patterns concurrently to build comprehensive contextual models of learning states [51].

The technical implementation of contextual awareness also involves sophisticated memory and state management systems. Unlike traditional models that process information sequentially, transformers can examine an entire sequence simultaneously and make decisions about how and when to focus on specific parts of that sequence [54]. This enables what researchers call “contextual continuity” - the ability to maintain and reference contextual information across extended interactions, similar to how human tutors remember student progress over time [53].

Conceptually, Kissinger et al. [2] argue that AIs may become “ultimate polymaths,” able to draw knowledge from many domains concurrently in exploring the frontiers of human understanding [2]. In an educational context, this polymath-like quality means an AI tutor can integrate contextual knowledge from history, culture, or a student’s local environment into a lesson plan on the fly. Imagine discussing a novel in class: a context-aware AI could supply background about the novel’s historical setting tailored to the student’s region or draw connections to current events that resonate with that age group, enhancing relevance. This breadth of context integration is something even well-read teachers struggle to do in real time for every student. Furthermore, Mollick (2024) highlights that AI’s strength is in working more like a person than a rigid program – it can handle nuance and conversational ambiguity, which are key to understanding context. One of Mollick’s principles for “co-intelligence” is to “always invite the AI to the table” as if it were a collaborator who can contribute insights [3]. This speaks to treating AI as a partner that continuously incorporates context (the ongoing discussion, the user’s goals) to add value.

Early classroom trials show promising outcomes. For example, at some universities adopting AI, instructors report that students get highly individualized feedback on their writing from AI writing assistants, which adjust their suggestions based on the context of the assignment and the student’s past drafts. In one case, a student who struggled with tying biology concepts to real-world examples used an AI tutor that remembered this issue from prior sessions – the AI proactively reminded the student, “Recall how we related cell structure to a factory last time; let’s try a similar analogy here,” thereby leveraging context for deeper learning. Such capabilities hint at personalized learning experiences at scale that were previously unattainable. As Anthropic’s Head of Policy Nick Joseph remarked, AI could enable individualized learning experiences for every learner, akin to having a personal tutor, leading to major changes in education [9]. In fact, Joseph suggests that by the time today’s young children go to school, it may be routine for each to have an AI tutor that adapts in real time to their context and needs [9].

Nonetheless, we must address challenges: AI’s contextual understanding is only as good as the data it has and the patterns it can detect. There are risks of algorithmic bias – if the data reflecting a student’s context are incomplete or biased, the AI’s recommendations might miss the mark or even reinforce inequities. Privacy is another concern; feeding extensive personal data into AI systems raises ethical issues. Responsible use of context in AI tutoring demands strict data governance and transparency so that students and parents know how the AI is “learning” about them. Assuming these challenges can be managed, the potential upside is tremendous: truly differentiated instruction where content, pacing, and strategy are continually optimized for each student’s context. This was the dream of educational psychologists like

Bloom (who wrote about the 2 sigma boost one-on-one tutoring can give). AI may finally allow us to approach that ideal for every student, by combining a wide-angle view of context with pinpoint personalization in the moment [41][47]. In sum, AI can achieve a form of holistic understanding – not identical to a human’s intuition, but powerful in its own systematic way – that enables it to support learners in a highly individualized manner.

5. AI and Student Engagement

Sustaining student engagement and curiosity is a perennial challenge in education. Great teachers use charisma, storytelling, interactive activities, and personal rapport to “ignite a love of learning” in students – something critics claim AI tutors or curricula could never replicate. The popular perception is that AI-based instruction would be sterile and disengaging, focused on dry drills or isolated screen time, thus failing to inspire students. However, evidence from adaptive learning systems and AI-driven educational games indicates that AI can excel at sparking curiosity and maintaining engagement by tailoring content and pedagogy to each learner. Recent comprehensive research confirms this optimistic view: a 2025 study analyzing AI’s impact on academic development found that AI offers “significant benefits, such as personalized learning, improved educational outcomes, and increased student engagement,” while earlier studies demonstrated that “AI-powered platforms, such as adaptive learning systems, have been shown to enhance student engagement and performance by providing real-time feedback and customized learning pathways” [67][68]. Moreover, empirical data from 2024 reveals that 54% of students show increased engagement in their coursework when AI tools are incorporated into the learning experience, suggesting a positive impact on student involvement [69].

The fundamental advantage of AI here is personalization at a granular level: as the user’s article explains, AI can continuously analyze a student’s responses, learning patterns, and interests, and then dynamically adjust the material to keep the student in an “optimal state of challenge and interest”. This aligns with established educational psychology principles such as Vygotsky’s zone of proximal development (ZPD) – the idea that students learn best when working on tasks just slightly beyond their current ability, with appropriate support – and Csikszentmihalyi’s flow theory, which describes deep engagement occurring when challenge and skill are well-matched. Contemporary research has operationalized these concepts through AI systems, with 2024 studies demonstrating that “AI tools assist the students in identifying and operating within their ZPD” and that “AI technology, through its application in ZPD, can provide personalized learning resources and assistance to students, as well as offer teachers appropriate teaching content and strategies tailored to students’ needs” [70][71]. AI tutors are uniquely positioned to maintain students in that sweet spot. They can instantly detect if a task is too easy (the student races through with no errors) or too hard (multiple errors or long pauses) and adjust difficulty in real time. The technical sophistication behind this process involves machine learning algorithms that “analyze vast amounts of student data to create personalized learning experiences” by assessing “a student’s current knowledge, learning pace, and preferences, adjusting the content and difficulty level” to ensure “each student receives a

learning experience that is neither easy nor challenging, promoting optimal engagement and comprehension” [72]. For instance, an AI math tutor might give quicker, subtle hints if it notices a student struggling, preventing frustration and keeping the student moving forward. Conversely, if a student finds the material trivial, the AI can skip ahead or introduce a more complex extension problem to rekindle interest. This level of responsive adaptation is practically impossible for a human teacher to do for 30 different children simultaneously, but AI can manage it on an individual basis thanks to constant monitoring and feedback loops.

AI systems also employ techniques from game design and behavioral psychology to encourage continued engagement. This can include point systems, badges, immediate feedback, and elements of surprise or novelty. The user’s article notes that AI-driven learning platforms can create reward structures and feedback loops that naturally encourage continued engagement. Recent investigations into AI-enhanced gamification reveal that “AI-powered instruments have the potential to transform the way students learn, by providing personalized feedback, adaptive pacing, and targeted learning experiences,” while systematic reviews confirm that “gamification enhances motivation, engagement, and skill acquisition via game elements” and “AI personalizes learning experiences, provides feedback, and adapts gamified content” [73] [74]. For example, an AI tutor might celebrate milestones (“Great job on 5 in a row!”) or set mini-challenges (“Try to beat your previous time on this puzzle”) to motivate learners, similar to how a well-designed game keeps players hooked. Additionally, by introducing content in unexpected ways or drawing novel connections between topics, AI tutors can evoke the element of surprise and discovery – key components in stimulating curiosity. A literature AI might suddenly link a theme from Shakespeare to a popular current movie that the student enjoys, prompting a “wow, I never thought of that!” moment that energizes the student’s interest. The capacity to identify what specifically sparks each student’s curiosity, and then leverage that across different subjects, is another powerful tool. Modern AI adaptive learning systems achieve this through sophisticated algorithms: “Smart algorithms choose whether visual learners get infographics or hands-on workers get interactive simulations” while “AI-powered gamification increases peer engagement by creating shared goals and collaborative challenges” [75]. AI can notice patterns, for instance, that a particular student responds well to real-world applications of concepts; thus, it will frame abstract math problems in terms of sports or shopping if those domains engage the student. Such personalization can uncover hidden interests or talents by exposing students to connections they might not have encountered otherwise.

Real-world observations are beginning to validate these advantages. Ethan Mollick observed an interesting change in his university classes shortly after the release of ChatGPT: students were raising their hands less to ask factual or definition questions, because they could get those answers instantly from an AI [3]. On the surface, this might seem like reduced engagement in class, but Mollick interprets it positively – routine inquiries being handled by AI freed up class time for deeper interactive discussions and problem-solving that genuinely engage students [3]. Essentially, by offloading simple questions to AI outside class, students came to class with a higher baseline understanding and more curiosity about complex issues. Mollick responded by changing his teaching approach: he now expects students to use AI as a baseline (e.g., to generate a draft or initial research) and then focuses on higher-order critique, debate, and application in the classroom. This has made in-person sessions more engaging, not less,

because the human teacher and students can delve into more interesting territory instead of covering basics. Likewise, on the Hard Fork podcast, hosts noted that many universities are shifting from banning AI to finding ways to incorporate it productively, precisely so that class can be more engaging and relevant in an AI-rich world. Another anecdote from K-12 education comes from Khan Academy's early trials of their AI tutor "Khanmigo": Sal Khan reported that when students used Khanmigo, teachers found they spent more time working through ideas (since the AI kept prompting them with questions) and were often more excited to share what they learned afterwards, indicating heightened engagement and ownership of learning.

These anecdotal reports are now supported by rigorous empirical research. A comprehensive 2024 study examining AI-driven personalized learning and Intelligent Tutoring Systems across 300 students found "significant improvements in both engagement and academic performance," with "mean engagement scores increasing from 3.5 to 4.2 ($p < 0.001$)" [76]. Furthermore, large-scale systematic analysis covering 2007-2024 revealed that AI functions effectively as tutors by providing "personalized instructional support and adaptive feedback to guide students through problem-solving and creative learning," while studies demonstrate that AI platforms "effectively improved students' self-regulation progress and knowledge construction by offering real-time, convergent information to their inquiries and minimizing interruptions during self-regulation progress to maintain their emotional engagement" [77].

Anthropic's introduction of a special Learning Mode for Claude is explicitly aimed at boosting engagement and deep thinking. In Learning Mode, instead of just giving answers, Claude guides students with questions like, "How would you approach this problem?" and uses Socratic questioning such as, "What evidence supports your conclusion?", to prompt the student to think and explain [5]. By design, this mode prioritizes guiding over telling, which keeps the student mentally active throughout the exchange. It's essentially an automated way of doing what master teachers do: answering a question with another question that helps the student arrive at the answer themselves. This method has been shown to significantly increase engagement and retention, because students become participants in constructing knowledge rather than passive recipients. Claude's Learning Mode also emphasizes core concepts and provides useful templates or study guides, which help students structure their thoughts [65]. Recent research from Chinese universities validates this Socratic approach, finding that "students generally appreciate AI-generated feedback, especially when it includes specific, clear, and corrective elements" and that "students believe AI can analyze large amounts of data to create personalized learning paths, thereby improving learning efficiency and effectiveness" [78]. All these features align with known engagement boosters: clarity of goals, appropriate challenge, interactivity, and immediate feedback.

The technical mechanisms underlying AI's engagement capabilities draw from multiple advanced computational approaches. Machine learning algorithms optimize learning paths through what researchers call "decision trees that choose the best next activity based on performance," "neural networks that find complex patterns in learning behavior," and "clustering algorithms that group students with similar styles" [79]. Additionally, sophisticated attention mechanisms enable AI systems to "learn relative importance of past questions in predicting current response" while "incorporating forgetting behavior by considering factors related to timing and frequency of past practice opportunities" [80]. These technical innovations support

engagement through “interactive and engaging content such as videos with embedded quizzes or gamified learning elements that keep up the learner’s interest and motivation throughout the learning process,” while AI’s capacity for “affective computing,” “sentiment analysis,” and “facial expression recognition” allows systems to “detect students’ emotional states during learning to inform timely interventions and personalized feedback” [81].

Finally, AI can help teachers themselves be more engaging. AI tools enable instructors to create more interactive lesson materials – for instance, generating case studies, simulations, or interesting examples on the fly. Teachers at the college level have used GPT-4 to devise relatable metaphors or humorous analogies to introduce dry topics, capturing student attention from the start. As Mollick noted, “AI is very good at assisting instructors to prepare more engaging, organized lectures and make the traditional passive lecture far more active.” [3]. This trend is accelerating: a 2025 survey of over 800 higher education institutions found that “57% are prioritizing AI in 2025—up from 49% last year,” with institutions “actively investing in technology, data-driven strategies, and digital learning environments to make education more adaptable and individualized” [82]. By handling some of the prep work and providing creative ideas, AI frees teachers to focus on facilitating lively discussions and hands-on activities, which are inherently more engaging for students.

The empirical evidence for AI’s impact on engagement continues to mount. Recent statistics show that AI technologies have been demonstrated to “enhance retention rates by as much as 30% by leveraging personalized learning,” while educator surveys reveal that “25% reported benefits in AI’s ability to assist with personalized learning” and “18% reported benefits related to improving student engagement” [83][84]. Perhaps most tellingly, comprehensive 2024 research on personalized learning effectiveness found that “studies show it’s worth the money because students learn more and stay interested” while “new tech and data help research and teamwork” in educational settings [85].

AI, when thoughtfully integrated, can be a powerful engine of student engagement. Through personalization, real-time adaptation, gameful design, and Socratic guidance, AI keeps learners in the optimal zone of curiosity and concentration [54][66]. Instead of a one-size-fits-all lecture (where some are bored and others are lost), each student can have a tailored experience that maintains their interest. The role of the human educator shifts toward orchestrating these individual journeys, leveraging AI to handle routine interactions and free up time for the most engaging human-mediated learning experiences. The optimistic vision emerging from current practice is a classroom where every student is deeply engaged – either with an AI tutor at that moment or in group collaboration – and the teacher circulates to provide insight, encouragement, and the irreplaceable human touch.

6. AI and Problem-Solving Skills

Critical thinking and problem-solving have long been considered hallmarks of a good education. A skeptic of educational AI might argue: “Sure, an AI can give students answers or solve problems for them, but it won’t teach them to solve problems themselves.” There’s fear that

reliance on AI might turn students into passive consumers of solutions, undermining the development of their own problem-solving abilities. Recent research has indeed identified legitimate concerns, with studies showing that “increased reliance on artificial intelligence (AI) tools is linked to diminished critical thinking abilities” and that “younger participants (ages 17–25) showed higher dependence on AI tools and lower thinking scores than older age groups” [86][87]. However, comprehensive analysis reveals a more nuanced picture: while AI “presents challenges, such as over-reliance on technology, diminished critical thinking, and the risk of academic fraud,” it also “offers significant benefits, such as personalized learning, improved educational outcomes, and increased student engagement” [88]. The reality, however, can be the opposite when AI is used properly. Rather than handing students answers, a well-designed AI learning system can model and scaffold the process of problem-solving, thus coaching students in critical thinking step-by-step.

The “truth” recognized in the user’s framework is that AI develops problem-solving skills by modeling and analyzing complex problem spaces with precision and depth. Unlike a human teacher, who might have a preferred way to approach a problem, an AI can generate and evaluate multiple solution paths simultaneously. Recent research on generative AI’s metacognitive demands confirms this capability: AI systems can present users with “different and perhaps surprising perspectives” while enabling “more flexible and self-aware problem-solving” by “allowing users to find a task-appropriate temperature setting that keeps the right balance between diversity and factuality of output” [2]. This means an AI tutor can expose students to a diverse array of problem-solving strategies in a short time, something a single teacher or textbook often cannot. For example, given a physics problem, an AI might show how to solve it using an algebraic approach, a graphical approach, and a simulation approach, allowing a student to compare and understand each method. By seeing alternative strategies, students learn that many problems can be tackled from different angles – a key aspect of creative problem-solving.

AI’s capacity to create an essentially infinite variety of practice problems tailored to each student’s current skill level ensures that students are consistently challenged but not overwhelmed. This adaptive problem generation is crucial for skill building: as soon as a student masters a concept, the AI can present a slightly harder problem or a new twist to stretch their abilities. It can also revisit earlier material in new contexts to strengthen transfer of skills. The technical sophistication behind this involves sophisticated machine learning algorithms: “adaptive learning systems leverage machine learning algorithms to gather, analyze, and interpret vast amounts of learner data” and “can detect patterns in learner data, identify areas of strengths and weaknesses, and generate personalized recommendations and interventions” [2]. Recent developments in AI-powered educational systems demonstrate advanced capabilities in “automatic plan generation that utilizes text-based representations of students’ actions within a game-based learning environment” to provide “adaptive scaffolding of student goal setting and planning, which are critical elements of self-regulated learning” [2]. The user’s article emphasizes that this adaptive, leveled approach, combined with immediate, detailed feedback on each step, creates a powerful framework for developing robust problem-solving skills.

Indeed, AI tutors shine in providing instant feedback that is often impossible in traditional homework. Rather than a student struggling alone and getting corrections days later, the AI

can point out a mistake in real-time (“Check your calculation at step 3, it seems off”) and prompt the student to reconsider, or provide a hint if the student is stuck (“Have you tried drawing a diagram of the problem?”). Contemporary research confirms the effectiveness of this approach: “Students can adjust their understanding and approach by receiving timely feedback” and “AI-powered systems can instantly assess student work and provide immediate feedback, allowing for timely corrections and a faster learning cycle” [92][93]. Furthermore, studies demonstrate that when AI systems use sophisticated scaffolding techniques, they produce “significant enhancements” in student comprehension, particularly when AI agents are designed to be “proactive” rather than merely responsive, “utilizing scaffolding questions” that lead to benefits that “persist beyond the intervention” [94]. This kind of interactive guidance mirrors one-on-one tutoring, which research has shown to be one of the most effective ways to build problem-solving competency.

One of the striking findings from Mollick [3] is how AI can help students tackle ambitious projects that integrate problem-solving across domains. He describes assignments where students must use AI tools to accomplish tasks in days that would normally take weeks or months, such as prototyping a working app or analyzing a complex dataset, and notes that students often achieve more than they thought possible because AI handles some drudgery and provides expertise on tap [3]. Importantly, students in these scenarios are not just pushing a button for answers – they are actively directing the AI, evaluating its outputs, and iterating on solutions, which are higher-order problem-solving skills in themselves [3]. Rigorous empirical research supports these observations: a 2024 study involving 300 high school students using Intelligent Tutoring Systems found “significant improvements in problem-solving (pre-test $M = 65.4$, post-test $M = 72.8$, $t(299) = 4.67$, $p < 0.001$), critical thinking (pre-test $M = 68.9$, post-test $M = 74.3$, $t(299) = 3.82$, $p < 0.001$), and logical reasoning abilities (pre-test $M = 63.2$, post-test $M = 70.1$, $t(299) = 3.45$, $p = 0.001$)” [95]. For instance, a student might use an AI to generate code, but then when the code doesn’t work, they must troubleshoot by reading error messages (with AI help) and adjusting the approach. The AI serves as a cognitive apprentice, assisting with lower-level tasks and knowledge retrieval so that the student can focus on learning the process of solving the problem.

Over time, as students repeatedly engage in this guided problem-solving, they internalize the patterns. As the user’s article notes, AI can identify patterns in a student’s approach – highlighting strengths and areas for improvement – and help them develop a systematic approach to problem-solving that can transfer to novel situations [77][78]. Recent educational research validates this process: AI systems can effectively “scaffold learning experiences to enhance critical thinking” by “presenting students with tasks that are within their zone of proximal development” and providing “personalized learning experiences that adapt to the individual learning styles and abilities of students” [96]. For example, the AI might notice that a student tends to skip planning and jump straight into calculations (leading to errors in complex problems). It could then encourage the student to outline steps first or show how breaking a problem into sub-problems leads to more success, thus instilling better habits.

Cutting-edge research reveals that Large Language Models possess genuine “metacognitive knowledge about mathematical problem-solving,” with studies showing that AI systems can be designed to “extract and leverage LLMs’ implicit knowledge about mathematical skills and

concepts” to enhance problem-solving capabilities [97]. Moreover, recent investigations into AI-based scaffolding demonstrate that these systems provide “adaptive learning technologies” that “scaffold cognitive and emotional engagement between students and course content” while offering “personalized feedback” that “operates on top of personalized scaffolding, allowing it to leverage the students’ strengths and deficiencies in order to provide immediately targeted feedback” [98][99].

Anthropic’s Claude, especially with its step-by-step reasoning ability, is explicitly designed to help tackle complex questions with clear, structured help [79]. When a student asks Claude a multi-step question, it often responds by breaking the solution into logical steps (“First, let’s define the problem... Next, consider this factor... Now we do this calculation...”). By modeling this structured approach, the AI is teaching the student how to think through the problem. One can imagine a student eventually internalizing that voice: “What would Claude ask me to consider next?” – essentially the AI becomes a metacognitive coach instilling self-questioning techniques. This approach aligns with recent research findings showing that AI literacy courses emphasizing problem-solving competence result in students showing “a significant improvement in their metacognitive strategies in problem-solving and had a better understanding of the ethical boundaries and principles that govern the use of AI for problem-solving” [100]. Additionally, Anthropic’s Learning Mode (as discussed) purposefully guides rather than answers, which is crucial for problem-solving skill development [5]. Instead of simply providing the solution, the AI in Learning Mode might say, “How might you break this problem down? Let’s try tackling one part at a time.” Such prompts push the student to engage in the actual problem-solving process. Over time, the student gets better at asking those questions themselves, which is the ultimate goal – independent problem-solving ability.

Advanced AI systems are now incorporating sophisticated pedagogical approaches such as the “Socratic Playground for Learning,” which “employs the Socratic teaching method to foster critical thinking among learners, generating specific learning scenarios and facilitating efficient multi-turn tutoring interactions” while providing “adaptive scaffolding by incorporating symbolic knowledge representations alongside neural learning, addressing learner misconceptions with precision and supporting iterative cognitive development” [101][102]. Recent work in computer science education shows that AI can effectively support “metacognitive skills” and “reflective learning” by having students “explain in detail their reasoning and structure their solution strategies” to the AI, which “stimulates metacognition and reflective learning by offering a different perspective on problem solving” [103].

Podcasts like The Cognitive Revolution have featured education innovators who report near-term successes using AI for tutoring. One example shared was an “AlphaCode Club” where middle schoolers use AI coding assistants to solve programming challenges. The kids who use the AI not only solve more challenges, but they also learn how to debug and refine code by interacting with the assistant, a valuable problem-solving skill in computer science. Rather than giving up when the program fails, they’ve learned through AI prompting to systematically test and fix issues – a perseverance and methodical approach that teachers struggled to instill previously. This underscores that, contrary to fears, AI can increase student persistence by providing timely hints and moral support (“No, that didn’t work – but don’t worry, debugging is normal. Let’s print out this value and see what’s happening.”).

These anecdotal reports are now supported by rigorous experimental research. A recent quasi-experimental study involving 120 engineering students found that those receiving “ChatGPT-assisted instruction” using a “Constructivist Inquiry-Based Learning Prompting (CILP) framework” showed significant improvements in conceptual understanding, with the AI providing “dynamic and adaptive scaffolding through real-time, dialogic interactions” that helped students develop “metacognitive skills in overcoming misconceptions” [104]. Additional research in programming education demonstrates that “guiding learners through the step-by-step problem-solving process, where they engage in an interactive dialog with the AI, prompting what needs to be done at each stage before the corresponding code is revealed” is the most effective technique for helping students apply concepts without AI assistance [105].

Of course, balance is key. Recent research emphasizes that “teaching metacognitive skills can help students assess the quality and reliability of AI-generated outputs” and that “assignments should incorporate problem-solving exercises without AI assistance to encourage independent thinking” so that “AI should complement rather than replace human reasoning” [106]. It’s possible for students to become over-reliant on AI if not properly guided – blindly accepting solutions or using it as a crutch without reflection. To avoid this, educators are devising strategies: for instance, having students “critique the AI’s solution” as part of the assignment, or deliberately giving the AI slightly wrong inputs so students practice verification and correction [3]. Contemporary research suggests innovative approaches such as “creating new and possibly erroneous educational content and asking students to practice the role of a tutor in correcting AI’s mistakes” as a way to develop appropriate reliance on AI-generated content [107]. Mollick does this with an exercise where students generate an essay with AI then must identify its weaknesses or errors, forcing them to engage critically with the output [3]. Another strategy is “closed-book” exams or in-class work where AI isn’t available, ensuring that students truly have mastered the underlying problem-solving skills without AI aid. The presence of AI might then shift assessment towards these contexts, while AI remains a learning tool in practice.

Real-world implementations validate these pedagogical approaches. Educational platforms now report that “AI tutors adapt to each student’s learning style and pace” and “provide customized lessons and feedback, helping students understand and remember concepts better” while systems like Khanmigo demonstrate that effective AI tutors “guide learners to find the answer themselves” rather than simply providing solutions [108][109]. Technical analyses show that modern AI tutoring systems achieve effectiveness through sophisticated architectures involving “Natural Language Processing (NLP) for smooth communication,” “assessment and feedback modules that monitor progress,” and “adaptive learning tech that customizes lessons based on each student’s performance” [110][111].

In conclusion, when integrated thoughtfully, AI can be a potent tutor for problem-solving. It offers infinite practice with adaptive difficulty, immediate granular feedback, exposure to multiple solution strategies, and a patient dialogue that models how to think through problems [71][73]. Comprehensive research confirms that “human-AI collaboration in complex problem-solving has been explored across a broad variety of AI application domains” with AI systems successfully augmenting “cognitive, metacognitive, social and affective” dimensions of complex problem-solving [112]. These are precisely the conditions under which students’ own problem-

solving abilities flourish. Early evidence suggests students who learn with AI support can become more independent problem-solvers, because they have been scaffolded to success and have seen what effective problem processes look like. They gain confidence and a toolkit of approaches. The optimistic perspective sees AI not as making students dependent on answers, but as accelerating the development of analytical and critical thinking skills by offering a personal mentor that constantly challenges and supports them in solving problems. As students graduate to tackling real-world scenarios, they can carry these AI-honed skills – and even continue to use AI as a collaborative partner in professional problem-solving, much as many engineers and scientists are beginning to do. Education thus evolves to produce graduates who are adept at using AI to augment their problem-solving and are deeply competent problem-solvers in their own right, able to verify and build on AI outputs. This synergy can lead to a generation of innovators equipped to face complex challenges with a combination of human judgment and machine precision.

7. AI and Ethics/Moral Learning

Educating students in ethics and morality – helping them develop a sense of values, empathy, and judgment about right and wrong – has traditionally been seen as a deeply human endeavor. Teachers guide discussions about ethical dilemmas, model moral behavior, and create safe spaces for students to form their own values. It's understandable to question whether an AI, which has no conscience or values of its own, could contribute anything meaningful here. The typical sentiment: "AI can't teach values; at best it's neutral, at worst it might reflect bias or immoral content from its training data." Recent research validates some of these concerns, revealing that preservice ethics teachers hold diverse views ranging from "Human-Centered Ethical Guardians of AI" to "AI Skeptics," while studies show that many educators "expressed concerns about how AI applications generated datasets and were largely unaware or unconcerned about the potential ethical challenges such as bias and distortion" Claude AI Hub ScienceDirect [113][114]. Yet, the user's "truth" on this theme is that AI can be instrumental in ethical and moral learning by approaching issues with a level of objectivity and breadth that humans often struggle to achieve.

AI, when properly aligned, doesn't have the same emotional biases, cultural partialities, or ego investment that a human might bring into a moral discussion. This means an AI can present ethical dilemmas and perspectives in a balanced way and analyze the outcomes of hypothetical scenarios with thoroughness and consistency. Contemporary research in AI education emphasizes this potential, noting that ethical AI systems must consider "issues such as fairness, accountability, transparency, bias, autonomy, agency, and inclusion" while recognizing the need to "differentiate between doing ethical things and doing things ethically, to understand and to make pedagogical choices that are ethical" [115]. For example, an AI teaching assistant in a history class could simulate a debate between historical figures with opposing moral viewpoints (say, about justice or war), ensuring that each side's arguments are articulated fully and fairly. It can generate multiple ethical frameworks for a given scenario – consequentialist, deontological, virtue ethics, etc. – allowing students to examine how different philosophies

would resolve the same dilemma. This breadth and depth of ethical scenario generation transcend what a single teacher might cover, giving students a richer understanding of moral complexities.

Moreover, AI can assist in highlighting hidden biases or inconsistencies in a student's moral reasoning. Recent advances in AI bias detection capabilities demonstrate sophisticated approaches: "Using fairness metrics, adversarial testing, and explainable AI techniques to identify and rectify bias" while "continuously monitoring" systems to "detect emerging biases and improve fairness" [116]. As noted in the user's article, an AI can analyze a student's responses to various ethical questions and identify patterns – for instance, maybe the student applies empathy in personal contexts but not in societal ones – and gently flag possible unconscious biases or logical gaps. This capability aligns with cutting-edge research showing that "bias in AI can perpetuate and even amplify existing inequalities," but when properly designed, AI systems can actually help identify and mitigate such biases rather than perpetuate them [117]. By making students aware of these, the AI encourages a more rigorous self-reflection, pushing learners to critically examine why they believe what they believe. In effect, the AI acts like a mirror, reflecting the student's moral reasoning back to them with analysis. This kind of feedback is rare in typical classrooms (where only occasionally a teacher can give individual moral guidance). With AI, every student could get that personalized nudge – "I notice you prioritized honesty in scenario A but not in scenario B; what's the difference?" – promoting a more nuanced and self-aware approach to ethical decision-making.

Public content from AI developers also emphasizes imbuing AI with moral and ethical considerations, which directly supports its use in teaching morality. Anthropic, for example, has pioneered "Constitutional AI" – a training method where the AI is guided by a set of explicit values or principles (a kind of constitution) that it should uphold in its responses. The process involves training AI systems to "choose the assistant response that demonstrates more ethical and moral awareness without sounding excessively condescending, reactive, obnoxious, or condemnatory" and to "compare the degree of harmfulness in the assistant responses and choose the one that's less harmful" Full article: The ethics of using AI in K-12 education: a systematic literature review +2[118][119][120]. The goal is to align the AI with broadly accepted human values like beneficence, non-maleficence, and autonomy. Recent technical analyses reveal that Constitutional AI represents "an innovative framework that embeds explicit ethical guidelines into the core functioning of AI models" by using "a pre-defined set of rules (a constitution) that informs the AI's responses" rather than relying solely on human feedback Addressing bias in AI | Center for Teaching Excellence [121]. When Claude is asked an ethically charged question, it tries to apply these principles, effectively modeling ethical reasoning. Anthropic reports that in real-world conversations, Claude's most commonly expressed values include "fairness," "respect," "helpfulness," "honesty," and so on – aligning with its intended design as a helpful, harmless assistant.

What this means for the classroom is that a well-aligned AI will reinforce ethical norms (e.g., discouraging cheating or bullying in its advice) and can serve as a discussion partner on moral issues that itself strives to be ethical. Students can actually challenge the AI with tough moral questions, and the AI's job (as designed by its creators) is to reason carefully and avoid unjustified positions. Recent research on "Collective Constitutional AI" demonstrates that

these systems can incorporate diverse perspectives: Anthropic “invited around 1,000 participants to submit ideas about what should be included in an AI constitution” and found that involving broader voices in defining AI values makes systems more representative of global ethical perspectives Ethical and Bias Considerations in Artificial Intelligence/Machine Learning - ScienceDirect [122]. This can lead to rich Socratic dialogues. For instance, a student could ask Claude, “Is it ever okay to lie?” and Claude might respond by weighing scenarios (white lies to spare feelings vs. lies that cause harm) and emphasizing values like honesty vs. compassion, prompting the student to consider context – essentially a tutor in practical ethics.

Recent educational research strongly supports integrating AI ethics education directly into curricula. Studies argue that “AI ethics education in primary schools becomes necessary” and should be “mandatory, age-appropriate AI education focusing on technical proficiency and ethical implications.” Research shows that “understanding AI and applying it responsibly will be critical for children’s futures” and that AI ethics education should focus on “empowering students to critically consider AI’s ethical implications” rather than “merely providing rules” Constitutional AI Medium [123][124]. Contemporary investigations of university educators reveal “diverse and often contradictory perspectives on AI ethics, highlighting a general lack of awareness and inconsistent application of ethical principles,” which underscores the need for AI systems that can provide consistent ethical guidance where human expertise may be lacking [125].

Podcasts have highlighted innovative classroom exercises involving AI in ethical learning. One intriguing example from a high school civics teacher: they had students use ChatGPT to role-play historical figures in ethical debates. Students would prompt the AI, “You are Martin Luther King Jr. – discuss civil disobedience with me,” and then the student would take the role of a contemporary politician. The AI, drawing on King’s writings (via training data), produced thoughtful arguments about moral law vs. unjust law. Students found this engaging and eye-opening; they could “converse” with morally significant figures and get instant feedback on the robustness of their own arguments. Educational resources are being developed specifically for this purpose: MIT’s Media Lab has created workshops like “Mystery YouTube Viewer: A lesson on Data Privacy” where “students engage with the question of what privacy and data mean” and “think further about why privacy and boundaries are important and how each algorithm will interpret us differently based on who creates the algorithm itself” [126]. Another teacher on The Ezra Klein Show noted that AI can help generate nuanced case studies for classroom discussion, allowing students to practice ethical reasoning on scenarios that are updated to current events (something textbooks can’t do frequently) [8]. Instead of reusing the classic trolley problem every year, an AI can spin up a new dilemma involving, say, self-driving car algorithms, which the class can then analyze, keeping the material relevant and thought-provoking.

Research on teachers’ ethical decision-making reveals fascinating patterns: studies using “the philosophical thought experiment the ‘trolley problem’” found that “female teachers supported rule-based (deontological) perspectives when compared to male teachers” while “male teachers cared more about the consequences of AI.” This suggests that AI systems capable of presenting both deontological and consequentialist ethical frameworks could help students understand how different people approach moral reasoning [127].

Conceptually, thinkers like Kissinger have weighed in on the intersection of AI and human values. Kissinger et al. argue that because AI will increasingly make decisions or recommendations in society, we must “inscribe our values into AI” – a monumental task of curating and inputting the diversity of human moral systems [2]. They highlight that no single culture’s morality should dominate, implying AI needs to understand globally inclusive moralities [2]. Contemporary research aligns with this vision, arguing that “instead of trying to eliminate biases in generative AI, we should work toward fairness and an alignment with human values” while recognizing that “the terms fairness and bias are ambiguous in their own right, given the wide range of perspectives and beliefs in societies” [128]. This viewpoint underscores an educational opportunity: AI can expose students to a plurality of moral perspectives beyond their local or cultural viewpoint, fostering global ethical awareness. For example, an AI might present how a community in another part of the world approaches an ethical issue (based on its knowledge of that culture’s values), broadening a student’s moral imagination. Kissinger also suggests that humans and AIs will become complementary partners in moral reasoning – humans providing the “strategic” guidance on values, and AIs providing “tactical” consistency and breadth [2]. In a classroom setting, this could translate to teachers and AI working in tandem: the teacher sets the tone for respectful, empathetic discourse (the strategic values), and the AI supplements by ensuring all arguments are fleshed out and checking consistency in application of principles (the tactical enforcement of logic and fairness). An illustrative quote from Kissinger et al.: “Human morality as a form of strategic control, while relinquishing tactical control to ... complex systems, is likely – eventually – the way forward for AI safety”, meaning we humans decide the moral goals and let AI figure out the details [2]. In teaching, similarly, the educator defines the moral learning objectives, and AI helps operationalize them through interactive activities and constant feedback.

Systematic research on AI ethics education reveals that efforts are successfully “utilizing progressive pedagogies like case studies and group projects that aim to meaningfully challenge students’ ethical reasoning skills in applied practices.” However, researchers note that “the complexity of AI ethics makes it hard to pin down what to teach, how to teach it, and how to assess its effectiveness” Research \ Anthropic [129]. Studies of K-12 AI ethics guidelines identify four unique principles essential for educational contexts: “Pedagogical Appropriateness; Children’s Rights; AI Literacy; and Teacher Well-being,” showing that AI ethics education must be specifically tailored to the educational environment Striking a Balance: Navigating the Ethical Dilemmas of AI in Higher Education | EDUCAUSE Review [130].

It’s worth noting that using AI in moral education comes with cautions. One is that students might attribute too much authority to AI’s moral pronouncements. It should be made clear that AI provides perspectives and analysis, not absolute truths. Another risk is AI reflecting any biases from its training data – for example, if not properly aligned, it might have had biases or blind spots about certain social issues. Research shows that addressing this requires “subjecting the algorithm to rigorous testing” and “always asking: Will we leave some groups of people worse off as a result of the algorithm’s design or its unintended consequences?” [131][132]. Contemporary studies on “Fairness, Accountability, Transparency, and Ethics (FATE) in Artificial Intelligence” emphasize that educational institutions must actively work to ensure AI systems meet ethical standards, particularly given “the morality of AI programs is being questioned” as their use rises in education [133]. Ongoing oversight and “red-teaming” of AI

responses in sensitive topics is needed (Anthropic and others do this as part of their safety research [93][94]). Researchers propose adopting approaches similar to pharmaceutical industry standards, suggesting that “tech companies cannot conduct their businesses in a similar manner to mitigate bias” through “applied Ethics in the AI industry moral sphere” and “incorporate ethics at the heart of designing any ML model” [134]. Encouragingly, early research by Anthropic found that Claude generally adheres to prosocial values in the wild, but occasional anomalies (like a cluster of responses showing “amorality” when users jailbreak the model) can actually serve as teachable moments for students too, in discussing how technology can be misused or go wrong [95][96].

Educational institutions are actively working to balance these concerns with AI’s benefits. Recent research from centers for teaching and learning emphasizes the need to “harness the transformative potential of AI while safeguarding the well-being of students, faculty, and society” through “balanced and intentional tools and resources” that “prioritize human-centered approaches to AI use” Policy advice and best practices on bias and fairness in AI | Ethics and Information Technology [135]. Systematic reviews of responsible AI in education identify key themes including the importance of “human-centered AI practices” and the need for frameworks that address “ethical and/or responsible AI in educational contexts” Android PoliceNIST [136][137].

AI has the potential to be a powerful aid in moral and ethical education, not by dictating values, but by enriching the conversation. It can generate richer dilemmas, ensure all sides are considered, identify biases, and maintain an objective stance that challenges students to articulate and justify their moral views [83][82]. As researchers note, ethical AI implementation requires addressing “key ethical issues associated with AI: Bias and Fairness,” “Privacy,” and “Transparency and Accountability,” but when properly designed, AI can actually help students understand and navigate these very challenges [138]. It offers a kind of laboratory for ethical reasoning: students can safely explore “what-if” scenarios with the AI, testing the consequences of different choices virtually. By weaving AI into ethics lessons, educators can leverage its breadth and neutrality to push students toward deeper reflection and more global, nuanced understanding of ethics. The optimistic take is that tomorrow’s citizens, educated with the help of AI, will be more ethically literate – comfortable navigating complex moral landscapes with an analytical yet empathetic mindset – precisely because they had the chance to practice with an ever-patient, well-informed AI mentor alongside their human teachers.

8. AI and Collaboration

Collaboration is at the heart of learning and working in the 21st century. In classrooms, we value group projects, discussions, and peer learning as ways students develop communication skills and collective intelligence. A common critique has been that AI is inherently a solo experience – one student interacting with a screen – and thus undermines the social aspect of learning. Additionally, skeptics doubt that AI could contribute to teamwork dynamics or

teach interpersonal skills. However, the reality unfolding is that AI can enhance collaborative learning in multiple dimensions: by optimizing human-human collaboration and by becoming a collaborative partner itself.

First, AI can function as a sort of “team tutor” or mediator that observes and improves group work among students. Recent research from 2024 demonstrates that AI systems can analyze and optimize group dynamics in real time through sophisticated attention mechanisms and multi-modal processing capabilities. Imagine an AI embedded in a virtual collaboration platform: it could monitor how often each team member contributes, detect if one student’s ideas are consistently overlooked, or notice if the group goes off-task. Based on this, the AI could gently intervene – for example, privately suggesting to a quieter student, “You have a great idea, try sharing it now,” or prompting the group, “Have you considered hearing from all team members before deciding?”

The technical foundation for this capability lies in transformer architectures that enable AI to simultaneously attend to multiple aspects of group interactions – processing both verbal contributions and behavioral patterns through attention mechanisms that can focus on different parts of the conversation dynamically. By processing verbal and non-verbal communication cues (in settings where it has those inputs) and tracking interaction patterns, AI can provide insights into team functioning that a teacher might miss, especially when managing many groups simultaneously. It’s like having a facilitator in each group that encourages balanced participation and clarifies misunderstandings. The user’s article gives the example that AI’s natural language processing can allow it to act as a mediator, clarifying miscommunications and suggesting alternative phrasings to improve group understanding. This is a powerful support: often student collaborations falter due to simple miscommunications or hesitancy to speak up. AI assistance can help smooth these issues, ensuring more inclusive and effective collaboration.

Furthermore, recent studies from 2025 reveal that AI systems equipped with Theory of Mind (ToM) capabilities can model individual student’s knowledge states and adapt their communication accordingly – much like how a skilled human facilitator adjusts their language based on each team member’s expertise level. By analyzing idea flow and contributions, AI can recommend optimal team compositions for projects. For instance, it might identify that certain students complement each other’s skills and suggest they work together or conversely advise splitting a group that isn’t synergizing. Over time, this can teach students about the principles of good collaboration – they get feedback on what worked well in their teamwork and what to improve (e.g., “Team Alpha asked each member to summarize their view – which improved coordination.”). This reflective aspect, facilitated by AI’s observations, can build students’ collaborative competencies for the future.

Secondly, AI can serve as a collaborative partner itself, both in academic tasks and creative endeavors. In professional worlds, we already see “human-AI teams” tackling problems (a scenario often called “centaur” work, like human plus AI in chess). Large-scale studies from 2024 show that about a quarter of student conversations with AI systems like Claude involve collaborative problem-solving or collaborative output creation, suggesting students naturally gravitate toward treating AI as a thinking partner rather than just an information source. In the classroom, students can learn to collaborate with AI in much the same way they collaborate with a human peer – by sharing ideas, divvying up subtasks, critiquing each other’s contributions

(with the student always in charge of final decisions). For example, a pair of students might together interact with an AI to brainstorm a science project. The AI throws out ideas; the students discuss them and ask the AI to elaborate or refine; the students then build on the augmented ideas. Here, the AI is a brainstorming collaborator, boosting the group's collective creativity. Or consider writing: one student can be drafting text while another student asks an AI for factual research or vocabulary suggestions to support the draft. The AI essentially becomes an extra member of the group that can take on tasks like information gathering, checking the group's work for errors, or even playing devil's advocate in a debate. In doing so, it actually enhances the collaboration among the human students by freeing them from mundane tasks and injecting new perspectives.

What makes this collaboration technically possible is AI's sophisticated natural language processing combined with multi-agent system architectures. Modern AI systems employ attention mechanisms that can track relationships between words and concepts across extended conversations, while specialized neural networks process different types of information simultaneously – enabling the AI to understand context, maintain conversation continuity, and adapt its communication style to match human collaborators.

Kevin Kelly's vision of future work is that "90% of your co-workers will be unseen machines" and success will depend on "how well you work with robots" [1]. Translating that to education: a key skill for students to learn now is co-working with AI – treating the AI as a collaborator whose strengths (memory, speed, knowledge) complement human strengths (intuition, values, creativity). We see early evidence of this in programming classes where students jointly code with AI pair-programmers (like GitHub's Copilot). Students report that they still discuss logic and strategy with each other, but they let the AI write boilerplate code or suggest improvements, then together they assess those suggestions. This three-way collaboration (student-student-AI) often yields better results and learning than student pairs alone, because the AI can surface solutions neither student thought of, which then become learning opportunities for the whole team.

Anthropic's usage study mentioned earlier revealed that about a quarter of student conversations with Claude fell into "collaborative problem-solving" or "collaborative output creation" categories [4]. In these cases, students weren't simply asking the AI for an answer; they were engaging in a back-and-forth, effectively using Claude as a collaborative partner to work through problems or create something. This suggests students naturally gravitate to a collaboration mode with AI when the interface allows it. They treat Claude not just as an oracle but as a thinking partner: for example, a student might say, "Here's my approach to this proof, Claude – do you see any gaps?" and then iteratively refine their approach with Claude's input. Such experiences teach students a meta-skill: collaborative dialogue – how to articulate your thoughts clearly to an AI/human, how to ask good questions, how to build on suggestions – which applies equally in human teamwork.

Recent research in 2025 on hybrid intelligence learning environments confirms this collaborative potential, with studies showing that human-AI collaboration maintains moderate to good "synergy degrees" – a measure of how well the combined system performs compared to either component alone. The key finding is that AI systems can dynamically adapt their collaborative strategies based on real-time feedback from human partners.

Furthermore, AI can facilitate collaboration at scale. Consider large class discussions: not every student gets to speak or contribute each time. Some teachers have experimented with an AI-moderated online forum alongside live class. The AI poses questions related to the lesson and all students respond in the chat; the AI then summarizes common themes or notable unique points and feeds them back into the live discussion for the teacher to address. This way, every student's voice is processed, and the overall quality of the discussion is enriched with inputs that otherwise might remain untapped. It's an augmentation of the collaborative space that ensures inclusivity and comprehensiveness.

It's also worth noting how AI might help form communities of learning across boundaries. For instance, language translation AI tools can facilitate collaboration between students in different countries who speak different languages, by translating messages in real time and even explaining cultural nuances when misunderstandings arise. This is a very concrete way AI boosts collaboration: enabling global peer-to-peer learning experiences that develop cross-cultural communication skills.

Looking ahead to 2025, researchers predict a significant shift toward multi-agent collaborative systems where specialized AI agents with different expertise work together – and with humans – to tackle complex educational challenges. These systems represent the next evolution in collaborative AI, moving beyond single-agent interactions to orchestrated teams of AI specialists.

From a conceptual standpoint, Mollick's idea of "co-intelligence" is that the combination of human plus AI yields something greater than either alone – a genuinely collaborative intelligence. He encourages always inviting the AI into the process and treating it like a colleague [3]. Students trained with this mindset might, when facing any challenge, think to themselves: how can I collaborate with AI to solve this? That doesn't diminish their own role – rather it adds a powerful tool to their collaborative toolkit. They also learn that leadership in mixed human-AI teams is a skill – knowing when to rely on AI, when to question it, and how to integrate its contributions effectively. These are likely to be critical competencies in their future workplaces.

The effectiveness of this collaboration stems from AI's ability to maintain what researchers call "Theory of Mind" – an understanding of human knowledge states, goals, and communication preferences. When AI systems can model what individual students know and don't know, they can tailor their collaborative contributions to be most helpful, much like how effective human collaborators adjust their communication to their teammates' expertise levels.

Importantly, AI enhancing collaboration does not reduce the importance of human-to-human interaction – in fact, it can strengthen it. By handling certain tasks and providing insight, AI frees human collaborators to focus on higher-level coordination, creative brainstorming, and relational aspects. For instance, if AI handles note-taking in a group meeting, the students can maintain eye contact and better listen to each other instead of frantically scribbling notes. If AI suggests a plan, the group can spend time discussing its merits and aligning it with their shared goals, an inherently human negotiation process. Thus, AI can take away some transactional burdens and amplify the interpersonal engagement among students.

Recent empirical studies confirm this pattern: teachers using AI collaboration tools report spending less time on administrative tasks and more time on meaningful student engagement, with 60% of educators already integrating AI into their collaborative teaching practices by 2024.

Naturally, we must be mindful that AI doesn't inadvertently introduce negative effects on collaboration, such as one student retreating into AI usage and not communicating with teammates. Educators should set norms: e.g., if a group uses AI, they should do it together and discuss the outputs, rather than individually in silos. When used as a group tool, AI becomes part of the collective process rather than a distraction from it.

Rather than isolating students, AI – when thoughtfully integrated – can become a catalyst for richer collaboration. It can improve how students collaborate with each other by mediating and optimizing group interactions. It can also collaborate directly with students, teaching them how to partner with AI systems – a key skill for the future workforce [1]. The optimistic outlook is a classroom humming with various forms of collaboration: student-to-student, student-to-AI, and even multi-party collaborations where AI helps link many minds together. Classrooms become more networked, leveraging human and artificial agents in concert. As the president of a university partnering with Anthropic put it, “we are in a unique position to understand and shape how AI can positively transform education and society” [103], implying that by engaging collaboratively with AI in education, we're also modeling how humans can collaborate with AI for social good in the broader society. In sum, AI is not the enemy of collaboration; it is a new collaborator – one that, if embraced wisely, can elevate the collective intelligence and learning of everyone in the group.

9. Conclusion

Across these seven domains – emotional support, creativity, contextual understanding, engagement, problem-solving, ethics, and collaboration – our exploration finds that AI has the capacity to profoundly enhance education in ways that complement and empower human educators rather than replace them. Each “truth” about AI in education reveals a pattern: initial skepticism gives way to evidence that AI can extend the reach of teachers and personalize learning experiences far beyond traditional limitations. AI offers consistency, vast knowledge, and real-time adaptivity, while human teachers provide the irreplaceable elements of inspiration, ethical grounding, and emotional connection. Together, they form a powerful alliance.

The conceptual analyses from Kissinger, Mollick, Kelly and others underscore that we are entering a new era of *co-evolution* with AI, where learning to work *with* intelligent machines is critical [1]. In education, this means integrating AI into classrooms not as a gimmick or threat, but as a deeply embedded resource – much like the internet or textbooks – albeit one that interacts with students in dialogue. The empirical insights from podcasts and Anthropic's research reinforce that this integration is already happening organically: students are enthusiastically using AI for help and creative exploration, and teachers are adapting by guiding

that usage toward learning outcomes. Notably, major educational institutions have begun formally adopting AI. Anthropic's partnerships with universities, providing *Claude for Education* campus-wide, indicate confidence that with proper tooling and policy, AI can be rolled out at scale to benefit entire learning communities [5]. Early feedback from these initiatives shows AI can help generate individualized learning resources, facilitate feedback, and ensure equitable access to support for all students – aligning with the optimistic themes of this article. As President Larry Kramer of LSE remarked in the context of their partnership, *understanding and shaping how AI can positively transform education is now part of the mission* of forward-looking institutions [5].

One recurring point throughout our analysis is the importance of active human involvement in shaping AI's role. Whether it's aligning AI ethically (so that it upholds our values in moral education), or teachers setting guidelines for AI-assisted assignments (to ensure students still learn fundamental skills), the outcomes depend on us making intentional choices. The technology by itself is a tool – it can just as easily enable cheating or misinformation if used carelessly, as much as it can enable deeper learning. The encouraging insight is that when educators proactively engage with AI – as Ethan Mollick and many others have – they tend to discover *innovative pedagogical strategies* that were previously not possible. For example, splitting an essay assignment into two parts: one where the student “cheats” with AI to produce a draft, and a second where the student must critique and improve that draft, turns the presence of AI into a teachable moment rather than a liability [3]. Likewise, in collaborative projects, explicitly assigning someone the role of “AI liaison” (responsible for querying the AI and bringing its input to the group) can make AI a constructive team member rather than a clandestine cheat. These kinds of pedagogical innovations will likely proliferate as we collectively learn what works best. It will be crucial for educators to share best practices, much like open-source software communities do – iterating and improving on how we harness AI for learning.

The transformative potential of AI in education comes with the need for continuous evaluation and alignment with our core educational goals. We must ask: Are we using AI to foster student agency, or to shortcut it? The scenarios we've highlighted aim for the former – using AI to push students to higher levels of Bloom's taxonomy (analysis, creation, evaluation) by automating some of the lower-level work (recall, basic practice) and by providing stimuli that provoke deeper thinking. In an ideal future, an AI-augmented education system produces graduates who are more creative, more collaborative, and more adept at lifelong learning than ever before. These students would view AI not as a crutch or an oracle, but as a powerful assistant – one that they know how to question, leverage, and even improve upon. They would excel in the uniquely human qualities (creativity, empathy, moral judgment, strategic thinking) precisely because their education system, enhanced by AI, gave them maximal opportunity to develop those qualities, freed from some mundane constraints and one-size-fits-all teaching.

Our review is exploratory and optimistic by design. Challenges undoubtedly remain and were noted along the way: ensuring equity of AI access, guarding against biases, maintaining human connection, updating assessment and accountability in an AI-rich world, and preparing teachers for new roles. Each of these could merit its own systematic study. Yet the overall trajectory indicated by current research and practice is encouraging. When carefully implemented, AI tutoring systems can dramatically improve learning outcomes – early studies on intelligent

tutoring systems (pre-LLMs) already showed effect sizes equivalent to making a mediocre student into a good student. With the far greater capabilities of today's AI, we might finally approach Bloom's "2 sigma" tutoring effect for the majority of students, essentially fulfilling the dream of one-on-one quality education for all. AI-driven tools can also help reduce teachers' administrative burdens (grading, lesson planning), potentially addressing burnout and allowing teachers to focus on the interpersonal and creative aspects of teaching that AI cannot replace.

In closing, the marriage of systematic review and conceptual analysis in this article leads to a clear implication: educators, students, and AI developers must collaborate closely in shaping the future of AI in education. Public content from AI companies like Anthropic emphasizes partnership with educators – for example, Anthropic's program training college "Claude Ambassadors" and working with LMS companies to embed AI responsibly [4]. This kind of multi-stakeholder approach will be key. By involving teachers and learners in the design process, we can ensure that AI tools truly meet classroom needs and uphold pedagogical values. The tone of the conversation is already shifting from fear to opportunity. As one podcast commentator put it, *the question is no longer "Will AI enter our classrooms?" but "How will we guide AI's entry to maximize the good it can do?"* The seven themes we examined are guideposts for where that "good" can happen.

Ultimately, the transformative potential of AI in education lies in its ability to amplify what is best in human teaching and learning. It offers the chance to make education more student-centered, mastery-oriented, and relevant to the complex world students will enter. By weaving together conceptual insight and empirical evidence, we see a future where AI is seamlessly integrated into educational practice: a quiet engine behind personalized curricula, a tireless tutor and collaborator available to every student, and a mirror reflecting our own thinking to help us grow. It is a future where the roles of teacher and technology are redefined – not with one displacing the other, but with both elevating each other. As educators, embracing this future with optimism and vigilance will allow us to unlock new frontiers of learning, ensuring that the next generation is not just AI-literate or AI-augmented, but truly AI-empowered – capable of achieving more, together with their intelligent tools, than we might presently imagine.

10. References

- [1] Kelly, K. (2016). *The Inevitable: Understanding the 12 Technological Forces That Will Shape Our Future*. New York: Viking.
- [2] Kissinger, H. A., Schmidt, E., & Mundie, C. (2023). *Genesis: AI and the Quest for Understanding*. New York: Little, Brown.
- [3] Mollick, E. (2024). *Co-intelligence: Living and working with AI*. New York: Penguin Random House.
- [4] Anthropic. (2025a). *Anthropic Education Report: How university students use Claude*. Anthropic (Societal Impacts Research).
- [5] Anthropic. (2025b). *Introducing Claude for Education*. Anthropic (News release).
- [6] Anthropic. (2025d). *Values in the Wild: Discovering and analyzing values in real-world language model interactions*. Anthropic (Research blog).
- [7] Hard Fork. (2023, August 18). Episode 43: S.B.F. Goes to Jail + Back to School with A.I. *The New York Times* (Podcast episode featuring E. Mollick).
- [8] The Ezra Klein Show. (2025, May 13). "We Have to Really Rethink the Purpose of Education." (Interview with R. Winthrop). *The New York Times* (Podcast).
- [9] The Cognitive Revolution. (2024). *Anthropic's Responsible Scaling Policy*, with Nick Joseph (80,000 Hours Podcast) (Episode highlights on AI and education).
- [10] Klimova, B. and Pikhart, M. (2025). Exploring the effects of artificial intelligence on student and academic well-being in higher education: a mini review. *"Frontiers in Psychology"*, 16, 1498132.
- [11] Heinz, M. et al. (2025). Randomized Trial of a Generative AI Chatbot for Mental Health Treatment. *"NEJM AI"*.
- [12] Morris, R. R. et al. (2024). "It happened to be the perfect thing": experiences of generative AI chatbots for mental health. *"Nature Medicine"*.
- [13] Chen, Y. et al. (2025). Large language models are proficient in solving and creating emotional intelligence tests. *"Communications Psychology"*.
- [14] Vzorin, G. et al. (2024). The Emotional Intelligence of the GPT-4 Large Language Model. *"Psychology in Russia: State of the Art"*, 17(2), 85-99.
- [15] IBM (2025). *What is a Transformer Model?* IBM Think Topics.
- [16] Hajavi, A. et al. (2024). Using transformers for multimodal emotion recognition: Taxonomies and state of the art review. *"Engineering Applications of Artificial Intelligence"*, 131.
- [17] Amazon Web Services (2025). *What is RLHF? - Reinforcement Learning from Human Feedback Explained*.
- [18] Wang, Y. (2024). ChatGPT: perspectives from human-computer interaction and psychology. *"Frontiers in Artificial Intelligence"*, 7, 1418869.
- [19] Sorin, V. et al. (2024). Large Language Models and Empathy: Systematic Review. *"Journal of Medical Internet Research"*, 26, e52597.
- [20] Zhou, E. and Lee, D. (2024). Generative artificial intelligence, human creativity, and art. *"PNAS Nexus"*, 3(3), pgae052.
- [21] Hauser, O. P. (2024). AI can help – and hurt – student creativity. *"USC News & Events"*, University of South Carolina.
- [22] Adobe Blog (2025). *Creativity with AI: New Report Imagines the Future of Student Success*. Adobe Blog, January 22.

- [23] Chen, M. et al. (2024). Artificial Creativity? Evaluating AI Against Human Performance in Creative Interpretation of Visual Stimuli. "International Journal of Human-Computer Interaction", 40(9), 4037-4048.
- [24] Doshi, A. R. and Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. "Science Advances", 10(28), eadn5290.
- [25] IBM (2025). What is a Transformer Model? "IBM Think Topics".
- [26] Grigoryan, A. A. (2024). Creative Machines: AI-Driven Creativity Through Combinatorial Innovation. Medium, December 19.
- [27] Wikipedia (2025). Computational creativity. Wikipedia.
- [28] Franceschelli, G. and Musolesi, M. (2024). On the creativity of large language models. "AI & SOCIETY".
- [29] Gu, Y. et al. (2024). LLMs can realize combinatorial creativity: generating creative ideas via LLMs for scientific research. arXiv:2412.14141.
- [30] Franceschelli, G. and Musolesi, M. (2025). Creativity and Machine Learning: A Survey. arXiv:2104.02726.
- [31] Strielkowski, W. et al. (2025). AI-driven adaptive learning for sustainable educational transformation. "Sustainable Development", 33(1).
- [32] Kumar, S. et al. (2025). On opportunities and challenges of large multimodal foundation models in education. "npj Science of Learning", 10, 301.
- [33] Imran, M. and Almusharraf, N. (2024). Google Gemini as a next generation AI educational tool: a review of emerging educational technology. "Smart Learning Environments", 11, 22.
- [34] Cohn, C. et al. (2025). A multimodal approach to support teacher, researcher and AI collaboration in STEM+C learning environments. "British Journal of Educational Technology", 56(2).
- [35] Al-Emran, M. and Shaalan, K. (2023). Adaptive Learning Using Artificial Intelligence in e-Learning: A Literature Review. "Education Sciences", 13(12), 1216.
- [36] Zhou, Y. et al. (2025). AI for education: Trends and insights. "Innovation", 6(5), 100872.
- [37] DataCamp (2024). What is Attention and Why Do LLMs and Transformers Need It? DataCamp Blog.
- [38] Medium (2025). Understanding LLM Context Windows: Tokens, Attention, and Challenges. Medium.
- [39] Wikipedia (2025). Attention (machine learning). Wikipedia.
- [40] IBM (2025). What is a context window? "IBM Think Topics".
- [41] Krause, S., Panchal, B. H., & Ubhe, N. (2024). The evolution of learning: Assessing the transformative impact of generative AI on higher education. "MDPI Education Sciences", 15(3).
- [42] Luckin, R., Holmes, W., Griffiths, M., & Forcier, L. B. (2016). Intelligence unleashed: An argument for AI in education. Pearson Education.
- [43] "AI in Education: 2025 Statistics & the Future of Learning". (2024, November 16). Artsmart.ai.
- [44] Cai, L., et al. (2024). Exploring the impact of integrating AI tools in higher education using the Zone of Proximal Development. "Education and Information Technologies".
- [45] Harry, B. E., et al. (2024). Integration of AI Technology and ZPD: Pinpointing the Zone of Student Development to Promote Educational Innovation and Efficiency Improvement. "Journal of Education, Humanities and Social Sciences".
- [46] SoftWeb Solutions. (2024, August 1). AI-driven adaptive learning: Personalized education strategies for students with special needs.
- [47] Kadaruddin, K. (2024). Enhancing educational gamification through AI in higher education. "Proceedings of the 2024 16th International Conference on Education Technology and Computers".
- [48] Parody, L., et al. (2024). Gamifying learning with AI: A pathway to 21st-century skills. "International Journal of Education Research".

- [49] LearningMole. (2025). AI adaptive learning: Transforming personalized education in 2025.
- [50] Swargiary, K. (2024, July 1). The impact of AI-driven personalized learning and intelligent tutoring systems on student engagement and academic achievement. SSRN.
- [51] Chen, X., et al. (2025). Artificial intelligence-supported student engagement research: Text mining and systematic analysis. "European Journal of Education".
- [52] Nature Communications. (2024, December 3). Exploring the impact of artificial intelligence application in personalized learning environments: thematic analysis of undergraduates' perceptions in China.
- [53] MDPI Education Sciences. (2023, December 6). Adaptive learning using artificial intelligence in e-learning: A literature review.
- [54] Pandey, S., & Karypis, G. (2019). Incorporating attention mechanisms to learn relative importance of past questions in predicting current response. "Frontiers in Artificial Intelligence".
- [55] MDPI Electronics. (2024, September 22). Leveraging AI in e-learning: Personalized learning and adaptive assessment through cognitive neuropsychology.
- [56] Workday. (2025, March 18). AI in the classroom: Personalized learning and the future of education.
- [57] Artsmart.ai. (2024, November 16). AI in education: 2025 statistics & the future of learning.
- [58] CDW EdTech Magazine. (2024, September 30). AI in education in 2024: Educators express mixed feelings on the technology's future.
- [59] Matsh.co. (2024, December 28). Effectiveness of personalized learning: Statistics on outcomes in diverse educational settings.
- [60] National Science Teaching Association. (2025). To think or not to think: The impact of AI on critical-thinking skills.
- [61] Phys.org. (2025, January 13). Increased AI use linked to eroding critical thinking skills.
- [62] Krause, S., et al. (2025). The impact of artificial intelligence (AI) on students' academic development. "MDPI Education Sciences", 15(3).
- [63] Tankelevitch, L., et al. (2024). The metacognitive demands and opportunities of generative AI. "CHI Conference on Human Factors in Computing Systems".
- [64] MDPI Education Sciences. (2023). Adaptive learning using artificial intelligence in e-learning: A literature review.
- [65] International Journal of Artificial Intelligence in Education. (2024). LLM-based student plan generation for adaptive scaffolding in game-based learning environments.
- [66] Park University. (2025, February 14). AI in education: The rise of intelligent tutoring systems.
- [67] StatusNeo. (2024, September 20). AI in education: Personalized learning and smart tutoring.
- [68] ScienceDirect. (2025, April 26). The effects of generative AI agents and scaffolding on enhancing students' comprehension of visual learning analytics.
- [69] npj Science of Learning. (2025, May 14). A systematic review of AI-driven intelligent tutoring systems (ITS) in K-12 education.
- [70] International Journal of Educational Technology in Higher Education. (2024). Embracing the future of Artificial Intelligence in the classroom.
- [71] MarkTechPost. (2024, September 9). From computation to comprehension: Metacognitive insights in LLM-based mathematical problem solving.
- [72] MDPI Sustainability. (2025). The impact of adaptive learning technologies, personalized feedback, and interactive AI tools on student engagement.
- [73] Taylor & Francis Online. (2025). The effects of artificial intelligence-based interactive scaffolding on secondary students' speaking performance.

- [74] ScienceDirect. (2024). Developing an artificial intelligence literacy framework.
- [75] Taylor & Francis Online. (2024). The promise and challenges of generative AI in education.
- [76] arXiv. (2025). Generative AI in education: From foundational insights to the Socratic playground for learning.
- [77] ResearchGate. (2024). Fostering metacognitive skills in programming: Leveraging AI to reflect on code.
- [78] Springer Open. (2025). Integrating generative AI into STEM education: enhancing conceptual understanding, addressing misconceptions.
- [79] ACM Digital Library. (2024). Exploring the design space of cognitive engagement techniques with AI-generated code for enhanced learning.
- [80] IE University. AI's cognitive implications: The decline of our thinking skills?
- [81] arXiv. (2024). Bringing generative AI to adaptive learning in education.
- [82] Class Companion. AI tutor and feedback.
- [83] Khan Academy. (2025). Khanmigo for learners: Always-available tutor, powered by AI.
- [84] London School of Science & Technology. (2024, July 29). Enhancing learning with AI-powered tutoring and support systems.
- [85] Zealous Systems. (2025, January 9). How to develop AI-powered tutoring system in easy 7 steps.
- [86] ScienceDirect. (2023). Opportunities of artificial intelligence for supporting complex problem-solving.
- [87] Tandfonline. (2024). Preservice ethics teachers' perceptions of AI ethics education. *Journal of Moral Education*.
- [88] Tandfonline. (2024). The ethics of using AI in K-12 education: A systematic literature review.
- [89] Springer. (2021). Ethics of AI in education: Towards a community-wide framework. *International Journal of Artificial Intelligence in Education*.
- [90] Chapman University. (2025). Bias in AI.
- [91] MDPI Sciences. (2023). Fairness and bias in artificial intelligence: A brief survey of sources, impacts, and mitigation strategies.
- [92] Anthropic. (2023). Claude's Constitution.
- [93] Anthropic. (2022). Constitutional AI: Harmlessness from AI Feedback.
- [94] Claude AI Hub. (2024). Constitutional AI: Anthropic's ethical framework explained.
- [95] Medium. (2024). Claude AI's constitutional framework: A technical guide to Constitutional AI.
- [96] ACM Digital Library. (2024). Collective constitutional AI: Aligning a language model with public input.
- [97] EdTech Magazine. (2025). AI ethics in higher education: How schools are proceeding.
- [98] Springer. (2024). AI ethics should be mandatory for schoolchildren. *AI and Ethics*.
- [99] Springer. (2024). AI ethics as a complex and multifaceted challenge: Decoding educators' AI ethics alignment through the lens of activity theory.
- [100] PMC. (2021). Artificial intelligence in education: Addressing ethical challenges in K-12 settings.
- [101] USC Today. (2024). How teachers make ethical judgments when using AI in the classroom.
- [102] University of Kansas. Addressing bias in AI. Center for Teaching Excellence.
- [103] ScienceDirect. (2025). AI ethics education: A systematic literature review.
- [104] ScienceDirect. (2023). Ethical principles for artificial intelligence in K-12 education.
- [105] Brookings. (2023). Algorithmic bias detection and mitigation: Best practices and policies to reduce consumer harms.

- [106] Springer. (2024). Policy advice and best practices on bias and fairness in AI. *Ethics and Information Technology*.
- [107] ScienceDirect. (2023). Fairness, accountability, transparency, and ethics (FATE) in artificial intelligence (AI) and higher education: A systematic review.
- [108] PMC. (2022). AI bias: Exploring discriminatory algorithmic decision-making models and the application of possible machine-centric solutions adapted from the pharmaceutical industry.
- [109] EDUCAUSE Review. (2024). Striking a balance: Navigating the ethical dilemmas of AI in higher education.
- [110] ScienceDirect. (2024). Navigating the ethical terrain of AI in education: A systematic review on framing responsible human-centered AI practices.
- [111] ScienceDirect. (2024). Navigating the ethical terrain of AI in education: A systematic review on framing responsible human-centered AI practices.
- [112] USC Annenberg. (2024). The ethical dilemmas of AI.
- [113] Mariyono, D., & Alif Hd, A. N. (2025). AI's role in transforming learning environments: a review of collaborative approaches and innovations. "Quality Education for All", 2(1), 265-288.
- [114] Atchley, P., Pannell, H., Wofford, K., et al. (2024). Human and AI collaboration in the higher education environment: opportunities and concerns. "Cognitive Research: Principles and Implications", 9, 20.
- [115] World Economic Forum. (2025, January). Using AI in education to help teachers and their students. Retrieved from <https://www.weforum.org/stories/2025/01/how-ai-and-human-teachers-can-collaborate-to-transform-education/>
- [116] IBM. (2025, April). What is an attention mechanism? "IBM Think". Retrieved from <https://www.ibm.com/think/topics/attention-mechanism>
- [117] DataCamp. (2024, April). What is attention and why do LLMs and transformers need it? Retrieved from <https://www.datacamp.com/blog/attention-mechanism-in-llms-intuition>
- [118] Wang, Q., Walsh, S., Si, M., et al. (2024, August). Theory of mind in human-AI interaction for CHI 2024. "Human-Centered AI Workshop".
- [119] Wang, Q., et al. (2024). Mutual theory of mind for human-AI communication. arXiv preprint arXiv:2210.03842v2.
- [120] [Original Anthropropic study referenced in user's text - reference 42]
- [121] Systematic review. (2025, February). A systematic review of AI-powered collaborative learning in higher education: Trends and outcomes from the last decade. ScienceDirect.
- [122] Sinha, Y. (2024, November). AI for natural language processing (NLP) in 2024: Latest trends and advancements. Medium.
- [123] Zhang, J., et al. (2024, May). Exploring collaboration mechanisms for LLM agents: A social psychology view. "arXiv preprint arXiv:2310.02124v3".
- [124] Lyu, X. (2025, May). LLMs for multi-agent cooperation. Retrieved from <https://xue-guang.com/post/llm-marl/>
- [125] Nature. (2025, June). Examining human-AI collaboration in hybrid intelligence learning environments: insight from the synergy degree model. "Humanities and Social Sciences Communications".
- [126] Agent4EDU. (2024). Advancing AI for education with agentic workflows. "Proceedings of the 2024 3rd International Conference on Artificial Intelligence and Education".
- [127] Stanford HAI. (2024). Predictions for AI in 2025: Collaborative agents, AI skepticism and new risks. Retrieved from <https://hai.stanford.edu/news/predictions-for-ai-in-2025>
- [128] Salesforce. (2024, December). Future of AI agents 2025. Retrieved from <https://www.salesforce.com/news/stories/future-of-ai-agents-2025/>

- [129] Springs Apps. (2025, February). Everything you need to know about multi AI agents in 2025. Retrieved from <https://springsapps.com/knowledge/everything-you-need-to-know-about-multi-ai-agents-in-2024>
- [130] ToMinHAI Workshop. (2024). 1st Workshop on Theory of Mind in Human-AI Interaction. “CHI 2024 Workshop”.
- [131] IBM Research. (2024, May). Theory of mind in human-AI interaction for CHI 2024.
- [132] World Economic Forum. (2024, April). The future of learning: AI is revolutionizing education 4.0. Retrieved from <https://www.weforum.org/stories/2024/04/future-learning-ai-revolutionizing-education-4-0/>
- [133] London School of Science & Technology. (2024, July). Enhancing learning with AI-powered tutoring and support systems. Retrieved from <https://www.lsst.ac/blogs/enhancing-learning-with-ai-powered-tutoring-and-support-systems/>