Artificial Intelligence and Machine Learning: A Systematic Review of Trends, Techniques, and Applications in Online Tutoring

Drashti Mehta[1] , Prof. Jeel Visani

[1]Student, Gyanmanjari Innovative University, Bhavnagar, Gujarat, India.

Abstract - *The last decade has seen Artificial Intelligence (AI) and Machine Learning (ML) migrate from experimental labs into the heart of daily digital tools [1]. Education, specifically online tutoring, stands out as a primary beneficiary of this transition. This paper evaluates research published between 2020 and 2025 to pinpoint how these technologies actually function in modern learning environments. Our review shows that deep learning, Natural Language Processing (NLP), and supervised learning aren't just secondary features anymore; they are actively redesigning the student experience through automated feedback and custom learning tracks [2], [4]. However, the data also highlights persistent friction points—mainly algorithmic bias, opaque "black box" decision-making, and privacy risks [1], [3]. By weighing these advancements against their inherent risks, this review provides a grounded perspective on the current state of AI-driven education.*

Key Words: **Artificial Intelligence, Machine Learning, Online Tutoring, Intelligent Tutoring Systems (ITS), Natural Language Processing, Deep Learning, Educational Technology, Algorithmic Bias**

1.INTRODUCTION

It is hard to ignore how quickly AI and Machine Learning have moved from "tech buzzwords" to essential infrastructure. While often lumped together, the distinction matters: AI acts as the broader umbrella for "intelligent" task execution, while ML is the specific engine that helps systems improve through data exposure [1], [6]. From healthcare diagnostics to high-frequency trading, these tools are everywhere. Education was perhaps slower to adapt, but the recent global shift to remote learning acted as a massive catalyst.

During the pandemic-era lockdowns, the cracks in traditional digital education namely static PDFs and a total lack of engagement became impossible toIgnore. Simply dumping content onto a screen failed to facilitate real learning. This gap sparked a rush toward adaptive systems. Today, MOOCs and Intelligent Tutoring Systems (ITS) use AI to track student behavior in real-time [5]. Unlike a standard classroom where one teacher manages thirty different learning speeds, these systems can pivot instantly—adjusting difficulty levels or offering a specific practice problem the second a student hits a wall [2].Yet, this rapid adoption has outpaced our regulatory frameworks. Researchers are increasingly vocal about the "hidden" costs of AI in the classroom [3]. Can we ensure a model isn't biased against certain demographics? How do we protect the massive amounts of data these students generate? This paper explores these questions by reviewing the key technological trends of the 2020–2025 period [3], [4]. We aim to move past the hype and examine the practical techniques and ethical hurdles that will define the next decade of online tutoring.

2.METHODOLOGY

Rather than chasing narrow case studies or isolated experiments, this review utilizes a systematic approach centered on existing survey papers and high-level reviews [5]. We chose this macro-level focus to capture a broader snapshot of how AI and Machine Learning are maturing in the educational sector. Individual experiments can often be skewed by specific local variables; by synthesizing insights from multiple large-scale reviews, we can present a more stable and realistic view of where the technology actually stands today [1], [2].

Data Sources and Search StrategyThe weight of a systematic review rests on the quality of its sources. To ensure a high standard of credibility, we pulled literature from major academic powerhouses: IEEE Xplore and SpringerLink for their technical depth in computer science, and ScienceDirect for its focus on the intersection of technology and pedagogy. To fill any gaps and catch high-impact, open-access research, we also utilized Google Scholar and ResearchGate.

Our search strategy relied on specific keyword clusters, including "Artificial Intelligence," "Machine Learning," "Online Tutoring," and "AI Tutors." We used

Boolean operators to narrow the results, which allowed us to filter out irrelevant hits while still capturing papers that used varied terminology for the same underlying concepts [4].

2.1 Selection Process and Workflow

Refining the raw search results into a usable dataset involved a four-stage "funnel" process, similar to established automation workflows in systematic reviews [4]:

1. **Identification:** Initial broad-spectrum search across all databases using the primary keyword strings.
2. **Screening:** A quick manual review of titles and abstracts to discard papers that didn't explicitly bridge the gap between AI techniques and practical education.
3. **Full-Text Eligibility:** Remaining papers were read in full. At this point, we enforced a strict cutoff: only peer-reviewed surveys or review papers published between 2018 and 2025 remained in the running.
4. **Final Extraction:** The final "shortlist" was analyzed to pull out specific data points—primarily common algorithms, real-world tutoring applications, and documented ethical or technical hurdles [3], [6].

2.2 Inclusion and Exclusion Criteria

We set strict boundaries to keep the review from becoming too diluted. To make the cut, a study had to directly connect AI/ML with a practical educational

use case, specifically within digital or online settings. The 2018–2025 window was chosen because it covers the specific "boom" in deep learning and LLM adoption [3].

Conversely, we actively excluded "theory-only" papers that focused on math or abstract models without any application to learning [1]. We also ignored non-peer-reviewed content, such as industry blogs or marketing whitepapers, to ensure our conclusions remain grounded in verified academic research rather than corporate hype.

3. TRENDS IN AI AND ML (2020–2025)

The half-decade between 2020 and 2025 stands as a watershed moment for AI. We've moved past the era where these tools were just "cool tech demos" in lab settings; they are now the invisible backbone of everyday digital life [1]. As these models became more accessible and computationally cheaper, they naturally bled into the education sector [3]. The result was a fundamental pivot: online learning shifted from being a static repository of videos and PDFs to a dynamic, responsive environment that actually "listens" to the student.

3.1 Recent Technological Developments in Educational AI

The most disruptive force in this window has undeniably been the explosion of Generative AI and Large Language Models (LLMs). Early tutoring software was notoriously rigid, relying on "if-then" logic that felt more like a glorified FAQ page than a teacher. Generative models changed that by allowing for real-time, context-aware dialogue [4]. Students can now dig deeper, asking "why" or requesting a simpler analogy, and the system can pivot its explanation on the fly.

Deep learning has also matured significantly [2]. By utilizing multi-layered neural networks, platforms can now handle messy, unstructured data [6]. This has been a game-changer for language learning (via sophisticated speech recognition) and for grading open-ended essays—tasks that previously required ahuman eye to interpret intent. This "intelligence" is largely fueled by the democratization of cloud computing. Platforms like AWS and Google Cloud have lowered the barrier to entry, allowing even small-scale EdTech startups to rent the massive computing power needed to run these models without needing a multimillion-dollar server room.We are also seeing the rise of "multimodal" AI [3]. Instead of just processing text, these systems can "see" a photo of a handwritten math equation or "hear" a student's pronunciation, synthesizing different types of input to provide a more holistic feedback loop [5].

3.2 Commercial Implementations: AI in the Wild

Several major players have already set the blueprint for how this looks in practice:

- **Duolingo:** Rather than a fixed curriculum, it uses ML to crunch data on millions of users. If the system notices you're consistently tripping over French verb conjugations, it doesn't just move on—it algorithmically reshuffles your future lessons to reinforce those specific weak points [2].
- **Khanmigo (Khan Academy):** This acts less like an answer key and more like a Socratic tutor. Instead of "giving away" the solution, it uses AI to nudge students with hints, forcing them to do the heavy lifting themselves.
- **Coursera:** It uses recommendation engines similar to Netflix or Amazon. By mapping your professional goals against your past course performance, it builds a custom "learning path" that feels tailored to your career trajectory [5].

3.3 The Shift in E-Learning Frameworks

We are witnessing a total breakdown of the "one-size-fits-all" model. Traditional digital education was a linear track everyone started at point A and ended at point B at the same speed. Modern frameworks are far more "liquid." They use real-time engagement metrics to decide if a student should skip ahead or loop back for more practice [3].

This has also popularized "microlearning." Recognizing that digital attention spans are fragmented, platforms are breaking hour-long lectures into five-minute, high-impact modules. Perhaps most importantly, this tech is redefining not replacing the teacher's role. By offloading the "drudge work" of grading and data entry to an AI, educators are finally being freed up to focus on high-level mentoring and emotional support, which are the parts of teaching a machine still can't replicate [1], [6].

4. MACHINE LEARNING: THE ENGINES OF MODERN TUTORING

The "magic" behind platforms like Coursera or Khan Academy isn't the result of rigid, old-school programming. You can't code a million "if-then" rules to account for every student's behavior. Instead, these systems rely on data-driven models that "learn" by crunching massive amounts of user interaction data [1]. Our review of current research highlights three dominant architectures: supervised learning, unsupervised learning, and the increasingly dominant field of deep learning [6].

4.1 Supervised Learning: Predicting the Student Curve

Supervised learning is currently the workhorse of the education sector because it excels at one thing: prediction [2]. Think of it as a model trained with an "answer key." By feeding the system years of historical data where the outcomes (passing, failing, or dropping out) are already known the model learns to spot the early warning signs of a struggling student.

Algorithms like Random Forests and Support Vector Machines (SVM) act as a high-tech early warning system [6]. They look at variables like how many times a student re-watched a video or their average "time-to-complete" for a quiz. If a current student starts mimicking the behavior of past students who eventually dropped out, the system flags them. This allows for proactive intervention suggesting a simpler module or alerting a human mentor before a student actually fails [2].

4.2 Unsupervised Learning: Finding the "Hidden" Student Types

What happens when you don't have an answer key? This is where unsupervised learning comes in.
Rather than looking for a specific outcome, these models look for "shapes" and "clusters" in raw data [1]. In a massive online course with 10,000 students, you can't manually categorize everyone. Using techniques like K-means clustering, the AI can automatically group students based on their actual habits [6].

For example, the system might discover a "Late-Night Binger" group that watches four hours of video on Sundays, or a "Steady Reviser" group that spends 20 minutes a day. By identifying these clusters, platforms can adapt. They might send a "Late-Night Binger" a different type of notification than a "Steady Reviser." It's also used for content organization automatically grouping similar study materials together so that the platform's library stays organized without a human librarian having to tag every single file [5].

4.3 Deep Learning and the Quest for Conversational AI

Deep learning is the most "heavy-duty" tool in the shed. Built on multi-layered neural networks, it's designed to handle the "messy" data that humans excel at but computers traditionally hate: speech, handwriting, and natural language [1], [3].

The biggest breakthrough here is in Natural Language Processing (NLP). Traditional software struggles with the nuance and sarcasm of human speech. Deep learning models, however, can parse the intent behind a student's question [4]. This is what makes modern AI tutors feel conversational rather than robotic; they can actually "understand" a Follow-up question and provide a relevant, natural-sounding explanation.

However, there is a catch. Deep learning is a "data-hungry" and "resource-heavy" beast [1]. It requires massive datasets and high-end server clusters to function. This has created a digital divide: while global platforms can afford to deploy sophisticated LLMs, smaller institutions or local non-profits are often forced to stick with simpler, less "intelligent" models because the cost of entry is simply too high [3].

5. PRACTICAL APPLICATIONS IN EDUCATIONAL ECOSYSTEMS

While the math behind machine learning is fascinating, its true value is measured by how it changes the student's daily experience. In the real world, supervised and deep learning models aren't just abstract concepts; they are the gears inside a larger machine [1]. Based on our analysis of the current literature, these applications generally fall into three pillars: Intelligent Tutoring Systems (ITS), high-speed automated feedback, and the shift toward hyper-personalized learning "journeys" [3].

5.1 Intelligent Tutoring Systems (ITS): The Digital Mentor

The "holy grail" of EdTech has always been to replicate the experience of a one-on-one human tutor. This is where Intelligent Tutoring Systems (ITS) come in. Unlike a standard website that just hosts videos, an ITS actually "converses" with the learner [2].

To pull this off, the system runs four synchronized models under the hood [5]:

- **The Domain Model:** The "brain" containing the actual subject matter the rules, facts, and logic of the topic.
- **The Student Model:** A dynamic profile that follows the learner like a shadow, recording every click, hesitation, and mistake to map out what they actually know versus what they've memorized.
- **The Tutor Model:** The pedagogical engine. It looks at the student's struggles and decides how to nudge them. It doesn't just hand out answers; it offers hints or "scaffolding" to help the student find the solution themselves.
- **The User Interface:** The front-end layer where all this data-crunching is translated into a natural, accessible conversation.

5.2 Moving Beyond Multiple Choice: Automated Feedback

For years, digital grading was stuck in the "binary" world of multiple-choice questions it was either right or wrong. There was no room for nuance. Thanks to the leap in Natural Language Processing (NLP), that ceiling has finally been broken [4].

Modern AI can now do the heavy lifting of evaluating open-ended essays and complex math proofs. Instead of just slapping a grade on a paper, these systems act as a "first responder." They can flag a logical fallacy in a paragraph or spot a specific calculation error halfway through a long equation [6]. This instant feedback loop is a game-changer for remote students who might otherwise wait days for a human instructor to grade their work. By the time the feedback arrives in a traditional setting, the student has often moved on; with AI, the correction happens while the concept is still fresh in their mind.

5.3 Hyper-Personalized Learning Pathways

The "one-size-fits-all" curriculum is a relic of the industrial age. It assumes every student learns at the same speed, which we know isn't true. AI-driven platforms have replaced this "fixed track" with "liquid" pathways [3], [5].

If a student breezed through a module on basic algebra, the AI recognizes the lack of friction and automatically "levels them up" to more challenging content to prevent boredom. Conversely, if the system detects a student is struggling, it doesn't just let them fail; it slows down the pace and offers alternative resources perhaps a visual simulation instead of a text-heavy chapter [2]. This level of adaptability ensures that the platform wraps itself around the student's specific needs, rather than forcing the student to fit into a rigid, pre-made mold [1].

6. SYSTEMIC HURDLES, ETHICS, AND THE PATH FORWARD

While the potential for AI in the classroom is undeniable, the "smooth integration" narrative is often more marketing than reality. We cannot ignore the significant friction points that come with automating education [1]. Far from being a universal "fix," poorly implemented AI risks entrenching existing social and digital inequalities rather than dismantling them [3].

6.1 The Privacy Trade-off: Data as the New Tuition

To function, AI tutors require a constant stream of high-resolution data. These systems aren't just looking at test scores; they are tracking how long a student pauses on a specific paragraph, how many times they delete a sentence, and even their emotional tone in chat logs [5].

This creates a massive privacy paradox. We are effectively building deep behavioral profiles on students often minors who cannot fully grasp the long-term implications of this data trail. If this "weakness profile" follows a student through their academic career, it could lead to pigeonholing or unfair labeling long before they've had a chance to improve. Protecting this data isn't just a technical task; it's a moral one [1].

6.2 The Bias Loop: When Algorithms Discriminate

Machine learning models are only as "fair" as the data they were fed. If the training sets are skewed toward specific demographics, the AI will naturally mirror those biases [3]. We see this most clearly in automated essay grading, where models might penalize students with non-standard accents or different cultural writing styles.

In education, fairness isn't optional. If an algorithm systematically underestimates students from a specific background, it doesn't just give a bad grade it creates a barrier to future opportunity. We need to move past "trusting the math" and start auditing these models for social equity [1].

6.3 The "Black Box" and the Death of Transparency

One of the most frustrating aspects of advanced deep learning is its lack of "explainability." This is the "Black Box" problem: the system gives a result, but it can't tell you why [1], [6].
In a classroom, "because the computer said so" is an unacceptable answer. If a student is marked down, they have a right to a logical explanation they can learn from. Without transparency, we lose the "teachable moment." There is a desperate need for Explainable AI (XAI) systems designed to show their work so that teachers and students can stay in the loop [3], [4].

6.4 The Digital Divide: Scaling Inequality

There is a real danger that AI will only benefit the already-privileged. High-end AI tutors require low-latency internet and modern hardware [3]. In rural or economically struggling regions especially in countries like India this creates a massive barrier to entry. If we only deploy these tools in elite schools, we aren't innovating; we're just widening the gap between the "haves" and the "have-nots" of the digital age.

6.5 Future Scope: Beyond the English-Centric Model

Looking ahead, the focus must shift from "more power" to "more inclusion" [2]. This includes:

- **Lighter Models:** We need AI that can run on a $100 smartphone with a spotty 4G connection, not just on high-end desktop rigs.
- **Linguistic Diversity:** Most AI is built on English-centric data. Expanding these systems to support regional languages is essential for making "AI for all" a reality [3].
- **Human-Centric Design:** The goal shouldn't be to replace the teacher, but to build tools that make the human teacher's job more impactful [1].

Ultimately, the next phase of EdTech won't be defined by how fast the AI is, but by how fair and accessible it can become.

7. CONCLUSION

We've officially moved past the point where AI in the classroom is a futuristic "maybe." As this review of the 2020–2025 landscape shows, machine learning is now the quiet engine driving nearly every major digital learning platform [1]. What used to be a clunky, linear experience has been replaced by something far more reactive. The jump in NLP and cloud accessibility has effectively killed off the old "static" model of online schooling, giving us systems that don't just host content, but actually participate in the learning process [4].

The real win here is the end of the "one-size-fits-all" trap. By utilizing student data as a real-time compass, these platforms can now pivot to meet a student where they actually are not where a syllabus says they should be [5]. Whether it's a system offering a Socratic nudge during a math problem or an algorithm instantly grading a complex essay, the old bottlenecks of remote education are finally starting to break [2], [6].

But we can't let the tech hype blind us to the friction. The "black box" nature of these models is a genuine problem; if a student is marked down and the computer can't explain why, the trust is gone [1].We're also facing a massive ethical hurdle with data privacy and demographic bias. If we aren't careful, we're just building a high-tech version of the same old inequalities [3]. In a country as vast and diverse as India, the stakes are even higher. AI has to be an equalizer, not a barrier that only the elite can afford.

At the end of the day, AI is a tool, not a replacement. It can crunch the data, grade the papers, and suggest the paths, but it can't provide the empathy or the cultural context that a human teacher brings to the table [1]. The next decade of EdTech won't be defined by who has the fastest code, but by how well we can use that code to support not replace the human heart of education.

REFERENCES

[1] R. I. Mukhamediev et al., "Review of Artificial Intelligence and Machine Learning Technologies: Classification, Restrictions, Opportunities and Challenges," Mathematics, vol. 10, no. 15, pp. 1–32, 2022.

[2] A. Mishra, "A Comprehensive Review of Artificial Intelligence and Machine Learning: Concepts, Trends, and Applications," International Journal of Scientific Research in Science and Technology, vol. 11, no. 5, pp. 1–14, 2024.

[3] F. Undie, C. O. Edeh, and J. O. Okoye, "Artificial Intelligence and Machine Learning: A Review of State-of-the-Art Trends, Global Developments, and Practical Implications," Journal of Current Trends in Computer Science Research, vol. 4, no. 1, pp. 1–18, 2024.

[4] R. O.-Boateng, E. K. Nyarko, and S. O. Antwi, "Towards the automation of systematic reviews using NLP, ML, and Deep Learning: a comprehensive review," Artificial Intelligence Review, vol. 57, no. 1, pp. 1–45, 2024.

[5] K.-Y. Lin and K.-H. Chang, "Artificial Intelligence and Information Processing: A Systematic Literature Review," Mathematics, vol. 11, no. 11, pp. 1–29, 2023.

[6] Y. K. Rathod, "A Survey of Machine Learning Techniques for Artificial Intelligence," International Journal of Computer Techniques, vol. 11, no. 3, pp. 1–10, 2024.