by Pradeep Varakantham
and Sidney Tio

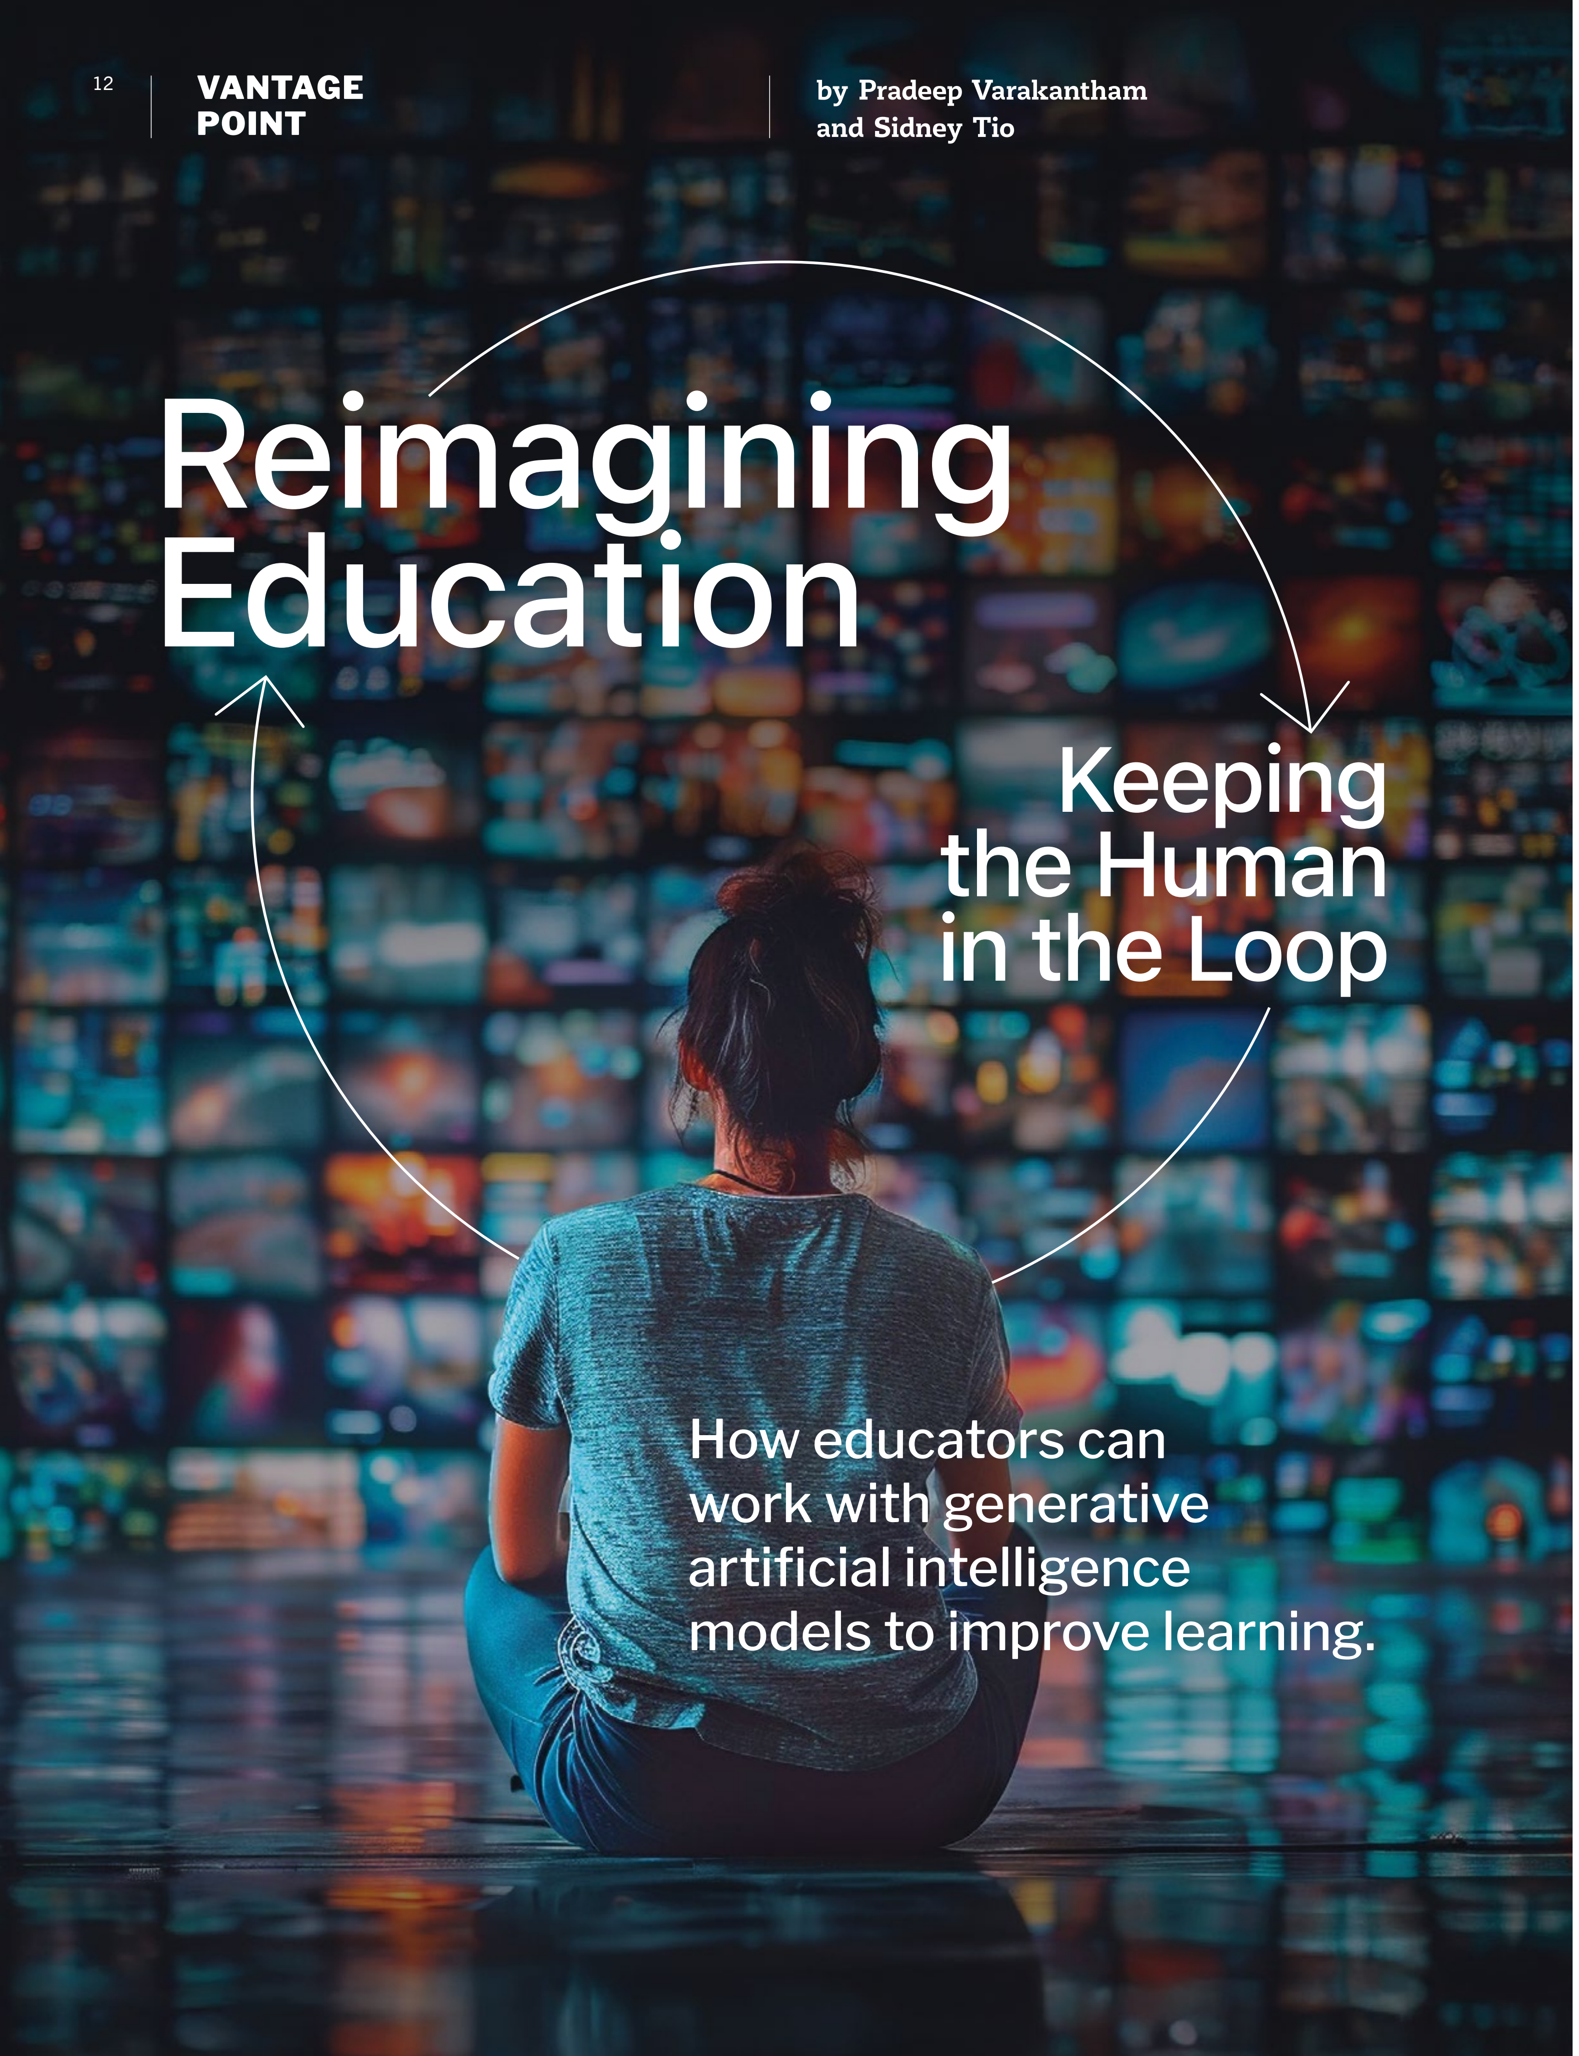

Reimagining Education

Keeping the Human in the Loop

How educators can work with generative artificial intelligence models to improve learning.

Artificial intelligence (AI) helps break the mould of one-size-fits-all education by creating personalised learning paths that adapt to each student's pace and style.

By combining human expertise with AI capabilities, educators can create learning experiences that are both structured and flexible, thus getting the best of both worlds.

While promising, AI used in educational contexts must carefully navigate privacy concerns, ensure fairness across all student groups, and support appropriate learning progression.

Education has traditionally operated under a model where a teacher follows a fixed curriculum, either in a classroom or through standardised video lessons on a learning management platform. While this approach effectively standardises education for large groups, it faces significant limitations. Central to this model is the concept of a 'one-size-fits-all' curriculum, which assumes that all students will progress at the same pace and follow a uniform learning path. While this structure may work for some, it often neglects the diverse needs, prior knowledge, and learning speeds of individual learners.

In this article, I explore how AI is reshaping education by combining human expertise with adaptive learning systems. I begin by examining the limitations of traditional classroom models, where fixed timelines and standardised curricula often fail to meet individual student needs. Next, I delve into how different AI approaches—from supervised learning to reinforcement learning with generative models—can create more personalised and effective learning experiences. Finally, I address the practical challenges of implementing AI in education, including data privacy, fairness concerns, and the crucial role of human oversight in maintaining educational quality. Throughout, I emphasise that the goal is not to replace teachers but to enhance their capabilities and create more responsive learning environments that work for every student.

TRADITIONAL CLASSROOMS: STATIC IN A FLUID AGE

In the traditional classroom model, time is a fixed variable, while learning is flexible—students must master a set of skills within a predetermined time frame, regardless of their starting point or learning speed. This structure leaves faster learners unchallenged and slower learners unsupported. Teachers, despite their best efforts, have limited capacity to tailor their instruction to the unique needs of every student. Consequently, students who require additional support may fall behind, while those ready for more advanced material remain unengaged.

The rigidity of traditional curricula further compounds this issue. Curricula designed by experts, though informed by educational research, are often slow to adapt to new developments in knowledge and skills. This problem is particularly acute in fast-evolving fields like technology, where the knowledge taught may become obsolete within a few years. Revising curricula requires substantial time and effort from educators, administrators, and policymakers. As a result, students may be learning outdated material that fails to meet the demands of modern industries.

Another challenge of traditional education models is their reliance on standardised assessments. Examinations and tests provide a snapshot of student performance at a particular moment, but they rarely capture a student's complete learning trajectory. Assessment systems prioritise 'what' students know at a specific time over 'how' they learn, making it difficult to identify learning gaps or recognise alternative approaches students might use to solve problems. Designing diverse and non-repetitive assessments also places a heavy burden on human instructors and trainers.

ENTER AI IN EDUCATION

There are two major components of AI, one which is focused on how humans learn (machine learning) and the other that is focused on how humans decide (also known as reinforcement learning or RL). To reimagine education through the lens of AI, it is essential to understand how different AI methods can help transform traditional learning paradigms for humans. AI-powered educational systems leverage various learning approaches–supervised learning, unsupervised learning, and RL–each offering unique capabilities to enhance the human learning experience. Alongside these, advancements in generative models and natural language processing (NLP) have opened up new frontiers for personalised, and interactive learning experiences.

Supervised learning: Structured adaptation and personalisation

Supervised learning, a foundational AI technique, trains models using labelled datasets, where input-output pairs are explicitly defined. This approach is widely applied in educational settings for tasks that require structured decision-making and outcome predictions.

> Different AI approaches–from supervised learning to reinforcement learning with generative models–can create more personalised and effective learning experiences.

Some key applications include automated grading and feedback, where AI models trained on past student submissions can evaluate assignments, essays, and quizzes with high accuracy, offering detailed feedback in areas such as grammar, structure, and concept clarity. AI models can also perform student performance prediction by analysing historical data such as attendance, past assessments, and engagement patterns. Through supervised learning, at-risk students can be identified, and early intervention recommendations can be made to educators. Additionally, adaptive learning platforms can leverage supervised learning to propose tailored content based on a student's previous performance and reinforce concepts where improvement is needed.

Despite its strengths, supervised learning requires large, well-annotated datasets and is limited to recognising patterns from past data, making it less effective for evolving, exploratory learning experiences.

Unsupervised learning: Discovering hidden patterns and structures

Unsupervised learning, which works with unlabelled data, focuses on identifying patterns, clusters, and structures within the data without explicit human intervention.

In educational applications, this capability is particularly valuable for student clustering and segmentation, where AI models can group students based on their learning behaviours, preferences, and challenges, allowing for customised teaching strategies that cater to distinct learning styles. Unsupervised learning can also create curriculum optimisation by analysing student interactions with various learning materials and identifying sequences of content, leading to optimal knowledge retention and mastery. Unusual learning patterns can be uncovered that may indicate students struggling with specific topics or concepts that require deeper intervention; this would be a form of anomaly detection.

Unsupervised learning provides powerful insights into large-scale educational data but requires careful interpretation and validation by human instructors to ensure meaningful educational impact.

Reinforcement learning: A more active mode

While supervised and unsupervised learning with NLP (see following section) play critical roles in educational settings, RL introduces a more dynamic and interactive dimension. Unlike the 'passive' nature of supervised learning where models make predictions from fixed data, RL models interact with their environment and learn from trial-and-error experiences. This interactivity is especially relevant in personalised and adaptive learning, where AI must continuously adapt to the needs of instructors and individual students.

To that end, we propose to develop human-in-the-loop systems to assist both instructors and students in educational settings, by utilising RL and generative models to interact in a natural way with students and instructors.

Natural Language Processing (NLP): Bridging communication gaps

A subfield of AI that combines supervised and unsupervised learning (including generative models), NLP plays a critical role in enabling intelligent interactions between students and educational platforms.

Chatbots and virtual tutors are AI-driven, and they provide instant feedback and explanations, helping students clarify doubts without waiting for instructor intervention. NLP-based tools also offer automated writing assistance, which assesses student essays for grammatical accuracy, coherence, and originality, as well as provide suggestions for improvement in writing style and structure. AI-powered translation tools also enable language translation and accessibility, facilitating access to educational resources in multiple languages, while ensuring inclusivity and personalised learning experiences.

By leveraging NLP, educational platforms can create more interactive and inclusive learning environments, offering students continuous support across different subjects and linguistic backgrounds.

BUILDING THE HUMAN EXPERIENCE INTO AI

In education, the mastering of complex subjects such as calculus requires a structured and well-conceptualised progression of knowledge. For example, before learners can understand calculus, they must have a solid foundation in algebra, which in turn relies on a firm grasp of arithmetic. These learning pathways, often codified through years of pedagogical research and curriculum design, rely heavily on human expertise to structure the optimal learning experience. However, designing such curricula is not only resource-intensive but also requires a deep understanding of both the subject matter and the cognitive processes of learners.

On the other hand, modern video games present a stark contrast, where progression is often achieved through trial and error. Players engage in repeated attempts, gradually building skills through experience, exploration, and experimentation, without a predefined learning path. In such cases, the absence of a structured curriculum means that players must rely on intuition, adaptive learning, and sometimes even external resources like online guides or communities in order to advance. Despite the lack of formal scaffolding, games are successful at engaging learners and facilitating skill acquisition, albeit in an unstructured manner.

This highlights a key challenge in learning and development across different domains: how can we balance structured learning pathways with adaptive, exploratory approaches that cater to individual learning styles? The promise of human-in-the-loop RL with generative models offers an exciting potential solution by combining the expertise of human instructors with the computational power of AI-driven adaptive learning systems.

Instructor perspective: Enhancing curriculum design

For instructors, the integration of generative models with RL in a human-in-the-loop paradigm provides a powerful tool to streamline and optimise the curriculum design process. This approach allows for iterative interactions and the following benefits.

- Human instructors set the learning objectives and granularity: Instructors can define the overarching learning goals, desired proficiency levels, and key learning milestones while providing input on the granularity of details needed in the curriculum.

- Generative models and RL algorithms automate content generation and adaptation: AI models can dynamically generate instructional materials, exercises, and assessments based on the instructor's input. These models can also adapt to feedback, identifying potential gaps in the content and refining the curriculum iteratively.
- RL models adjust curriculum dynamically: By continuously analysing student performance data, RL models can suggest modifications to the curriculum structure, ensuring alignment with diverse learning paces and preferences.

This synergy reduces the cognitive load on educators while ensuring that the curriculum remains responsive to evolving student needs and pedagogical trends.

Student perspective: Personalised and adaptive learning

From the student's viewpoint, human-in-the-loop planning presents a transformative shift towards personalised and adaptive learning experiences. The benefits of this framework include the following:

- Iterative learning with formative assessments: Students engage with learning materials through ongoing check-ins and quick evaluations (known as "formative assessments" in education). Unlike traditional standardised tests that come at the end of a course, these AI-powered assessments can adapt to each student's progress. For example, if a student struggles with algebra, the system might generate more practice problems focused on their specific pain points, while another student might receive more advanced challenges. Students receive instant feedback on what they understand well and where they need help, allowing them to adjust their learning in real time. These personalised check-ins provide valuable data about what each student knows and how they are progressing.
- Adaptive content selection by RL models: Based on student interactions, RL models can recommend the subsequent concepts to explore, ensuring a tailored learning path that addresses specific weaknesses and builds on existing strengths.
- Continuous feedback and nudging: Generative models can provide targeted nudges and scaffolding to guide students through challenging concepts, fostering deeper understanding and engagement.

Through this iterative cycle, students benefit from a more personalised, data-driven approach to learning that evolves in response to their progress, helping them achieve mastery more efficiently.

EXPERIMENTAL STUDIES ON HUMAN-AWARE AI TEACHER ALGORITHMS

We evaluate our AI teacher algorithms against baselines using human participants who have undergone training in multiple simulation environments. Here, we present results based on a novel 3D emergency response environment designed to simulate an emergency medical care setting. Players use mouse controls to retrieve and apply medical items in an ambulance to treat patients with various conditions. The game requires correct item selection and application within a time limit, with up to four incorrect attempts permitted per task.

We conducted an experiment with 120 participants, randomly assigned to one of the four groups: (a) Reading Only (control): Learned solely through reading materials, without engaging in gameplay; (b) Random: Performed tasks selected at random from the pool, without replacement; (c) Handcrafted: Followed a predefined task sequence designed by the expert; and (d) SimMAC: Adhered to an adaptively curated task sequence using our AI algorithm.

In our evaluation, we analysed the effectiveness of teacher-guided training in improving post-training performance on the final test. Students trained using our proposed AI-based teacher algorithms significantly outperformed those in the control Random and Handcrafted curricula groups. The results were statistically significant. More importantly, we observed similar results in other environments, one of which emphasises the training of motor skills and another which stresses the cultivation of math skills.

CHALLENGES OF RL WITH GENERATIVE EDUCATION MODELS

While the potential of RL with generative models in education is substantial, several significant challenges must be addressed to ensure their effective implementation for both educators and students.

Data requirements

Unlike traditional AI models that are pre-trained on large datasets, RL requires constant interaction data from students and instructors. For educators, collecting real-time student interaction data at scale is costly and time-consuming, while also raising ethical and privacy concerns. For students, concerns about data privacy and security can affect their willingness to participate fully in digital learning environments. Privacy regulations such as the European Union's General Data Protection Regulation (GDPR) and US federal law Family Educational Rights and Privacy Act (FERPA) impose strict limitations on how student data can be collected, stored, and used. These regulations aim to protect student privacy but simultaneously restrict access to the diverse datasets needed for RL model training.

Privacy-preserving approaches such as federated learning offer a solution by allowing models to learn directly from student devices without transferring data to centralised servers, thereby mitigating privacy concerns. For example, consider an online learning platform that wants to personalise study recommendations for students preparing for university entrance examinations. Due to data privacy regulations, it cannot store student performance data on central servers. By using federated learning, the platform can analyse study patterns locally on students' devices, improving recommendations without compromising privacy.

Personalised learning adaptation

A key challenge in educational settings is designing personalised curricula that adapt to each student's progress without constant expert supervision. For educators, traditional approaches rely heavily on their expertise to structure the learning progression, while for students, personalised pathways can enhance engagement and comprehension. RL offers a breakthrough solution by employing AI agents as 'proxy students' to explore relationships between different learning scenarios.

An example of the above approach is in emergency medical training, where RL agents can analyse patterns across diverse medical cases, such as stabilising a patient with asthma versus handling both asthma and trauma concurrently, to determine shared medical principles and skills. This would ultimately benefit both students and educators in monitoring the learning progress.

Exploration versus exploitation trade-off

In RL, striking a balance between exploration (trying new strategies) and exploitation (using known effective methods) is crucial. For educators, this balance is essential to avoid disrupting established curricula, while for students, too much experimentation may lead to confusion and hinder learning progress. Generative models can help address this challenge by simulating potential student responses, enabling exploration without negatively impacting real students.

For instance, in a math learning application, an RL-based system might explore different teaching approaches, such as interactive quizzes versus video lectures. By simulating student responses, it can determine which method is more effective for different types of learners without disrupting actual learning schedules.

Data sparsity and generalisation

Students typically follow guided pathways, resulting in limited exploration of all possible learning trajectories, thereby hindering RL models from generalising to students who deviate from the norm. For educators, this can pose challenges in designing inclusive learning experiences. Generative models can help mitigate this issue by creating synthetic learning trajectories that fill in data gaps, offering a more comprehensive training dataset. Synthetic learning trajectories refer to artificially created data simulations that imitate real-world scenarios, improving how generative models learn and make decisions. However, ensuring that these synthetic trajectories are both realistic and representative of actual student behaviour remains a significant technical challenge.

Imagine learning as navigating a hiking trail. Most students follow the marked paths, taking similar routes to reach their destination. This creates a challenge: we do not have much data about what happens when students take different paths. It is like only knowing about the main trail but not the alternate routes that might work better for some hikers. AI systems need this variety of experiences to learn how to help all types of students effectively. To solve this, we can use AI to imagine and create realistic 'what-if' scenarios like mapping out possible alternate trails that students might take. While this helps fill in our knowledge gaps, making sure these artificial scenarios actually reflect how real students learn remains tricky. It is similar to making sure our imagined hiking paths are actually walkable, not just lines on a map.

An example of where this RL model can be deployed would be a language learning platform which may notice that students primarily practise basic vocabulary, leaving gaps when it comes to acquiring complex grammar topics. By generating synthetic data that provides for diverse learning pathways, this RL model can provide better practice exercises tailored to individual needs.

Bias in training data

Bias in training data is another major concern. RL systems trained on data that primarily represents certain student demographics may fail to generalise to underrepresented groups. For educators, this can lead to unintentional disparities in learning outcomes. For students, this means access to equitable learning experiences may be compromised. Addressing this issue requires thoughtful data collection strategies and fairness-aware algorithms to promote equity in education.

For instance, a career guidance AI system trained predominantly on urban students' data might not provide relevant recommendations to students from rural areas. Incorporating diverse datasets and fairness-aware algorithms can help ensure more equitable outcomes.

Aligning learning with progression

Educational success is rooted in the principle of learning continuity, where new material builds upon existing knowledge. RL-driven educational systems must ensure that learning pathways introduce challenges at an appropriate pace. For educators, this means they can rely on data-driven insights to guide students effectively, while for students, it ensures a smoother learning journey.

For example, in medical training, a student who has mastered how to deal with basic respiratory distress cases should be gradually introduced to more complex scenarios involving multi-system failures, rather than being overwhelmed with advanced cases prematurely. AI-driven systems can analyse learning trajectories to construct personalised pathways that optimise the progression of skills. Or in a coding bootcamp, students might first master basic programming concepts before gradually moving on to handle more complex algorithms. An RL system can track their progress and introduce challenges such as debugging exercises or full-stack development projects at the right time.

By addressing these challenges, RL with generative models has the potential to revolutionise education by providing personalised, adaptive, and scalable learning experiences that cater to the needs of both educators and students while maintaining ethical and privacy standards.

THE PATH FORWARD

The future of RL with generative models in education is promising but requires careful planning, ethical foresight, and continuous technical innovation. To overcome data constraints, researchers are exploring self-supervised learning techniques that enable models to extract meaningful insights from raw, unlabelled data. Additionally, advances in federated learning empower AI models to learn from decentralised student interactions while preserving privacy and complying with data protection regulations.

Generative models will continue to play a critical role in mitigating data sparsity and balancing the exploration versus exploitation trade-off. By generating realistic learning trajectories, these models help bridge gaps in student data and provide richer datasets for RL systems. Simulation-based training environments further enhance these efforts by allowing AI systems to test and refine new teaching strategies without disrupting actual student learning.

The potential impact of RL with generative models on education is profound. Unlike traditional educational methods that follow fixed, linear learning paths, RL-driven education systems can create personalised, adaptive, and equitable learning experiences. By leveraging generative models, these systems can simulate the vast diversity of human learning paths, ensuring that every student, regardless of their background, location, or learning style, has access to high-quality education tailored to their unique needs and pace of learning.

As these technologies evolve, collaboration among educators, policymakers, and technologists will be essential to ensure that RL-based educational solutions are effective, ethical, and inclusive. With the right approach, RL and generative models have the potential to revolutionise education by fostering lifelong learning and empowering students worldwide.

DR PRADEEP VARAKANTHAM
is Professor of Computer Science and Director, CARE.AI Lab, School of Computing and Information Systems at Singapore Management University

SIDNEY TIO
is a computer science doctoral candidate at Singapore Management University