




Stanford University
Human-Centered
Artificial Intelligence



FEBRUARY 2022
INDUSTRY BRIEF

An initiative of the HAI Corporate Members Program

Education and AI



Stanford University
Human-Centered
Artificial Intelligence

INDUSTRY BRIEF

Introduction

Effective learning solutions are fundamental to creating a higher quality of life that is accessible to everyone. At HAI, we support researchers who are empowering our understanding of learning across cultures, ages, identities, and abilities, which as a result will improve the experience of learners, teachers, and institutions. The community at large is leveraging technologies and data to revolutionize and democratize educational opportunities by making how we learn more scientific, learner-centered, and collaborative.

In this Industry Brief, we present a cross-section of the key research at HAI and more broadly at Stanford, including within Stanford's new Transforming Learning Accelerator, that can help shape the future of education. You will find researchers

developing AI that facilitates high-quality personalized learning at scale, creates grading systems that augment the teaching and assessment experience, and automatically designs intentional learning environments and infrastructures that help all individuals thrive.

We believe motivating the development of technology that uplevels all learners requires efforts and discussions at the intersection of research, industry, and government. Through highlighting how we could creatively disrupt and augment existing methods in education by employing AI, we hope thought leaders, investors, operators, teachers, learners, policy and decision-makers will find the brief useful as a point of further exploration and begin meaningful conversations.

HAI's mission is to advance AI research, education, policy and practice to improve the human condition. To learn more about HAI, visit hai.stanford.edu



Stanford University
Human-Centered
Artificial Intelligence

INDUSTRY BRIEF



“Education is one domain where AI research has great potential to develop solutions that improve the daily experiences of students and lifelong learners. We’ve shown that machine learning systems can automate part of the role of a skilled human tutor, allowing the human tutor to concentrate on what they do best. This is one demonstration of how technology can help broaden access to high quality and personalized education. AI technology does this, not by removing humans from the equation, but rather by putting them at the center and giving them tools to be more effective.”

—Peter Norvig, *HAI Distinguished Education Fellow*



Stanford University
Human-Centered
Artificial Intelligence

INDUSTRY BRIEF

Education

1 Personalized and Adaptive Learning

- Natural Language Interactions
- Augmented Teaching

2 Learning Sciences

3 Learning Infrastructure

1 EDUCATION | PERSONALIZED AND ADAPTED LEARNING | Natural Language Interactions

★ WHAT'S NEW?

Because of recent advancements in large language models, complex NLP tasks are more feasible by applying out-of-the-box algorithms accompanied by domain adaptation. Researchers on campus are leveraging these models to pioneer natural language applications such as chatbot engines to serve demands in teacher training and fuel improvements in student proficiencies, especially in language learning or other communication intensive settings. Furthermore, AI algorithms can quantify and evaluate responses from learners that were never previously encountered, then proceed to predict student performance, deploy better-suited exercises, and grant personalized feedback.

📌 WHY DOES THIS MATTER NOW?

According to a GlobalNewsWire report, amidst the pandemic, the global market for private tutoring was estimated at \$123.8B in 2020 and is projected to reach \$201.8B by 2026.¹ The demand for private and personalized education--driven by increased competition in admissions and the rising prevalence of lifelong and remote learning--points to an opportunity for education firms and startups to supply comprehensive, scalable, and interactive technologies powered by natural language processing and understanding. Natural language technologies flexibly enhance student education on digital platforms in contexts such as online courses, self-guided explorations, and supplemental classes. Besides being able to readily streamline current methods of instruction, they have proven benefits of increasing student learning rates as well as evoking higher learning times.

👁️ EYE ON CAMPUS

Researchers on campus are...

- Building in-house algorithms to rate and classify open-ended text. The models are embedded in chatbots that can identify English writing proficiency and in digital kiosks to assist the community in better navigating buildings.
- Exploring targeted question generation by fine-tuning pre-trained language models for deep knowledge tracing (LM-KT) to predict the probability of correct student answers and generalize to unseen questions. The LM-KT is also used to generate new questions conditioned on a given student and target difficulty. Their findings suggest that the model succeeds in generating novel, well-calibrated language translation questions for second language learners from a real online education platform. This step in personalization allows systems to adapt from an infinite space of questions and enables the fine-grain ability to choose questions, rather than selecting from heuristics or template-based methods driven by a small pool of predefined questions.



- Developing a novel natural language processing targeting mechanism, which relative to more traditional multiple-choice targeting, is able to provide optimal and adaptive feedback from fewer student interactions and generalize to previously unseen prompts.



- Developing a language learning chatbot that converses with students interactively on college-related topics and provides adaptive feedback. Evaluation shows increase in student fluency, higher engagement, and users voluntarily spent 2.1x more time interacting with EnglishBot compared to listening or repeating pre-existing materials.



“The best educators understand and help their students through an ongoing dialogue surrounding ideas, approaches, and reasons. As AI tools become more facile at using natural language they will join this dialogue to supplement and support teachers, to scaffold the kind of joyful, individualized learning experience that is currently inaccessible to most students. At the same time, as classroom education becomes increasingly intertwined with online experience it will provide new data and new opportunities to develop AI tools.”

–Noah Goodman,
Associate Professor of
Psychology, Computer
Science, and Linguistics;
HAI Faculty Affiliate

[1] Global Private Tutoring Market to Reach \$201.8 Billion by 2026. (2021, July). GlobeNewswire.

1 EDUCATION
PERSONALIZED AND ADAPTED LEARNING | **Augmented Teaching** 

★ WHAT'S NEW?

AI-augmented teaching demonstrates the potential for educators to cater to specific learning needs by tracing student skill progression and to alleviate teachers from repetitive efforts of offering identical feedback. Where institutions were previously limited to stale or outdated curriculums, teachers can use such systems to adaptively sequence practice activities as well as generate real-time customized narration to boost student engagement and performance. As digital testing data increases, scalable assessments also allow instructors to focus on communicating crucial concepts and minimize tedious grading without compromising on the quality of feedback since new algorithms are able to associate mistakes to specific parts of an individual solution and explain common misconceptions.

📌 WHY DOES THIS MATTER NOW?

In an Economic Policy Institute Report, projections reveal a national teacher shortage of sizable magnitude, growing, and worse than expected when indicators of teacher quality are taken into account—with high-poverty schools suffering most acutely.¹ And this massive gap in the teaching workforce is exactly at odds with the desire and need to deliver personalized learning experiences. However, because of the advent of scalable machine learning infrastructure that can interpret student data, design roadmaps, and support each learner on an individual basis, the impact of educators may now be amplified within their respective classrooms.

👁️ EYE ON CAMPUS

Researchers on campus are...

- Proposing Variational Deep Knowledge Tracing (VDKT) that not only traces a student's knowledge change using deep recurrent neural networks but also incorporates new latent variables representing stochasticity in the learning process such as partial understanding, making slips, and guessing answers. The VDKT outperforms baseline and previous state of the art methods by evaluation on two Duolingo language learning datasets and provides interpretable analyses and offers insights into students' stochastic behavior in language learning.
- Introducing a human-in-the-loop "rubric sampling" approach to tackle the "zero-shot" (no training examples) feedback challenge for code education. They provided autonomous feedback for first students working on an introductory programming assignment with accuracy outperforming data-hungry algorithms and results approaching human-level fidelity. This rubric-sample requires minimal teacher effort, associates feedback with specific parts of the student's solution, and articulates student's misconceptions in the language of the instructor.

- Proposing an approach for automatically adaptively sequencing practice activities for an individual student. This new approach builds on the construction of curriculum graphs and advancing students through a graph using multi-armed bandits, and is effective because it requires fewer hyperparameters and a smaller dataset. They found that for a Korean language learning game, the adaptive algorithm had a statistically significant positive effect on learning efficiency metric defined using in-game performance.
- Creating an interactive, tablet-based learning platform with a multi-step math task designed using Common Core State Standards. They showed that embedding learning activities into narratives boosted children's engagement, and integration of tutoring chatbot improved learning outcomes on the assessment.

[1] García, E. & Weiss, E. (2019, March 26). The teacher shortage is real, large and growing, and worse than we thought. Economic Policy Institute.



“The pandemic has highlighted the potential impact of new AI systems that can adapt to not only a student’s cognitive state, but also to the full complex context of a human learner. The last 1.5 years further stressed the demands for instructors, yielding a new need and opportunity for AI tutors that amplify the capacity of human teachers. Such AI teaching assistants can support students with many tasks, as well as automatically intelligently loop in human instructors at appropriate times.”

–*Emma Brunskill, Assistant Professor of Computer Science; HAI Faculty Affiliate*

2 EDUCATION LEARNING SCIENCES

★ WHAT'S NEW?

Researchers are using modeling techniques to help instructors understand both distinct individual learning processes and abilities across diverse student groups in an effort to inform the organization of future curriculums and assessments that are useful in real-time scenarios. Their contributions enable fine-grained analysis of student behaviour, which opens up new avenues in predicting learning outcomes. AI is able to interpret data in various realistic learning environments to classify student learning processes and strategies and subsequently give human instructors the resources needed to administer real-time feedback and interventions.

↑ WHY DOES THIS MATTER NOW?

Previously, instructors relied heavily on standardized and periodic testing as a proxy for student's grasp of the course material that may overlook nuances particular to an individual learner. Now equipped with scalable AI algorithms, teachers are empowered to monitor student progress throughout their learning experience by building a real-time and multi-dimensional picture of the student based on behavioral data. Given this information, models can determine features and strategies emblematic of successful learners, which in turn may aid teachers in designing more productive materials and help lagging learners course-correct immediately, rather than after cumulative assessments that occur less frequently.

👁️ EYE ON CAMPUS

Researchers on campus are...

- Developing variation Bayesian inference algorithms for Item Response Theory (model for understanding humans based on question responses) that is scalable, faster than heuristic algorithms, and more accurate than gold-standard methods. The open-source implementation is easily usable in tasks like grading exams such as the SAT, PSAT, or in learning multi-dimensional student interactions.
- Modelling and analyzing student inquiry strategies in open-ended computer-based learning environments. They demonstrate inquiry strategies influence learning outcome, and also are predictive of academic achievement. Furthermore, a clustering algorithm is developed to segment students into different groups with similar learning strategies, which along with the inquiry-based environment, have the potential to serve as an assessment tool for teachers and tutors.
- Extracting semantically meaningful features from student logged interaction data in a simulation-based task to predict their overall problem-solving ability and understand their specific problem-solving practices.
- Presenting a new assessment framework (Posterlet) that collects student strategies and processes when they independently learn. Student choices to receive critical feedback and revision predict and measure learning outcomes, which can be used to develop as well as evaluate models of instruction and assessments.
- Suggesting actionable and individualized interventions to help students stuck in wheel-spinning (repeatedly failing an education task for learning a skill). The trained model requires minimal expert input and uses interpretable machine learning to identify correct interventions matching expert-level prescription.
- Exploring online recurrent neural network models for classifying students in open-ended and interactive environments that can also provide real-time guidance.



“The application of Artificial Intelligence to education is in its infancy, but the prospects are excellent. On one side of the coin, researchers are automating difficult aspects of traditional instruction such as providing refined, individualized feedback and the automated generation of timely examples. On the other side of the coin, AI is enabling fundamentally new models of learning such as students teaching a computer agent and using augmented reality to help children with autism learn to recognize emotions in their family members. AI may make great in-roads in the challenge of making superior learning accessible to all.”

–Dan Schwartz, I. James Quillen Dean and Nomellini & Olivier Professor of Educational Technology at Stanford Graduate School of Education; HAI Faculty Affiliate

3 EDUCATION
**LEARNING
INFRASTRUCTURE**

★ WHAT'S NEW?

Learning infrastructures that support environments for learning are leveraging data and AI to not only provide better customized classroom and remote experiences, but also to suggest policies that align with intentional objectives (e.g., school assignment choices that maximize factors like diversity). By empowering and training local institutions and teams with tooling and resources powered by machine learning, high quality, adaptable, and human-centric learning can be deployed to a breadth of geographic locations despite language, cultural, and teaching differences.

↑ WHY DOES THIS MATTER NOW?

According to a 2021 McKinsey survey on the lingering effects of the pandemic on student learning, 24% of parents are not convinced to choose in-person instruction for fall, and will select homeschool or other alternative options.¹ The reluctance to return to physical buildings warrants innovation for virtual learning models for the foreseeable future, serving demands of both homeschooling families and of school districts. Separate from safety concerns, some parents feel remote learning and hybrid approaches have positively impacted their children's social-emotional and mental health.

As we transition into schools of the future, firms innovating in edtech and specialists can collaborate to define new types of effective learning enabled by AI and create infrastructures to distribute high-quality content to a global student audience. In this future, each student will be given more flexibility and choice in defining their individual learning journey.

[1] Dorn, E., Hancock, B., Sarakatsannis, J., & Viruleg, E. (2021, July 27). COVID-19 and education: The lingering effects of unfinished learning. McKinsey & Company.



“One of the most helpful ways to augment education is to assist with the difficult task of examining student work and recognizing student mistakes. Computer systems that can do this would make feedback, grading, and class preparation easier and more effective. Ten years ago, algorithms could perform at about one tenth of the accuracy of humans, in the context of understanding student pseudo-code, a bellwether task. Just this year, we built a deployable system to make real classrooms smarter. This transformative technology is one ingredient that we will use in our broader effort to make education's digital future more inspiring, more joyful, and less lonely.”

–Chris Piech, Assistant Professor of Computer Science Education; HAI Faculty Affiliate

3 EDUCATION
**LEARNING
INFRASTRUCTURE**

CONTINUED

EYE ON CAMPUS

Researchers on campus are...

HAI
RESEARCH

- Designing high quality open-source course-in-a-box CS curriculum easily adaptable to a variety of local teaching practices, languages, and cultures already delivered to over 1,000 highschool students around the globe. Joint teaching methods consisting of localized instructors and undergraduate student-teachers in conjunction with automated feedback are scalable and provide high-impact professional development experiences.
- Collaborating with The San Francisco Unified School District (SFUSD) to develop an optimization and simulation engine to suggest and evaluate potential student school assignment policies to promote predictability, proximity, and diversity. Using this tool, they recommended a policy passed by the San Francisco Board of Education that will take effect for 2023-24 school year.

HAI
RESEARCH

- Deploying a meta-learning model that adapts to give feedback to student code for a new programming question with only a few annotated solutions. This is the first successful deployment of machine learning based feedback to open-ended code. The approach generated feedback for 16,000 student exam-solutions in an online programming course, where it wasn't previously possible for instructors to provide feedback.
- Presenting a workflow for knowledge externalization in science text comprehension and prototyping tools to assist learners in organizing knowledge by effectively creating active diagrams without disrupting reading. This is powered by natural language processing on key terms to reduce cost of producing diagrams while maintaining cognitive effort required for comprehension.





The Transforming Learning Accelerator (TLA), part of Stanford's Vision, aims to combine science and design in order to advance novel approaches to understanding learning, and create a sustainable model to turn discovery into scalable and equitable solutions.

Among other initiatives, researchers are...

- Strengthening child literacy by introducing a literacy support network amplifying educator-parent-student connections through messages, AI, and tutoring
- Using AI analysis of mobile games to track and treat a continuum of early childhood learning
- Automating online assessment of reading ability to enable personalized education

Let's work together.

Are you planning to deploy AI initiatives for education within your organization or company?

We're interested in your ideas for improving learning solutions as a partner, researcher, supporter, or co-creator.

Reach out to learn about our projects and collaboration opportunities.




Stanford University
Human-Centered
Artificial Intelligence

INDUSTRY BRIEF

Industry Take

Industry will play a critical role in scaling the applications of AI research. For this reason, it is a goal and privilege of HAI to convene stakeholders from industry in addition to those in academia, government, and civil society to address the technical and societal challenges posed by AI.

Leading venture investors, positioned at the frontlines of startup innovation, can provide a unique perspective on the impact and role of AI technologies in education.





Industry Take



“From the K – 12 and higher education segments to the corporate and lifelong learning categories, AI is providing key benefits across the multi-trillion dollar education sector. In the formal schooling environments, AI is delivering key insights and data to educators so they can create more personalized experiences for learners and focus on individual or small group instruction. When it comes to adult learning and the future of work, AI is allowing companies to better understand how they can continue to remain competitive by offering their employees critical opportunities to reskill and upskill. In these ways, AI is reshaping what education will look like over a lifetime for the next generation of learners.”

–*Amit Patel, Managing Director, Owl Ventures*



“All people deserve equal access to the future--technological innovations in learning and workforce skilling at scale are key to realizing this vision. There is no more important development than the mainstream movement of AI to support the personalization and adaptivity of learning, skilling and assessment platforms. We’ve already seen success with automated grading in higher education and subject-specific coaches for grammar, math, and writing. The future of education will be characterized by this kind of “invisible learning and assessment” – highly efficacious education that occurs when an AI-driven platform augments and reduces the burden on educators and accelerates the learning of students.”

–*Deborah Quazzo, Managing Director, GSV Ventures*



“The successful application of AI to edtech is best evidenced by tools that boost productivity, for example with personalized coaching or targeted content consumption. One of our companies develops NLP algorithms to provide just-in-time coaching for litigators on appropriate case precedents. The algorithms can also identify case chronology, significant actors in disputes, and other meaningful data that help accelerate the learning curve for legal teams. You can see how this same formula, applied in schools or in the workforce, can help learners focus on the most high-impact tasks and topics for their education and work.”

–*Alex Spiro Latsis, Partner, Brighteye*



Stanford University
Human-Centered
Artificial Intelligence

INDUSTRY BRIEF



**CORPORATE
ENGAGEMENT**

Corporate Engagement

Stanford HAI seeks corporate members who will enable it to lead with the unparalleled interdisciplinary breadth and excellence of Stanford University, and a laser focus on human-centered development and deployment of AI technology. We invite engagement from companies that share our mission to advance AI research, education, policy, and practice to improve the human condition.



Are you prepared for the next wave of change?

THINGS YOU SHOULD KNOW.

52%

of Fortune 500 companies were extinguished by digital disruption between 2000 and 2014¹

\$15.7T

in value will be added by AI to the global economy by 2030³

50%

of the S&P 500 in 2018 was forecasted to be replaced in just ten years²

\$1.4B

in annualized value can be gained by AI-led transformation of a Fortune 500 company⁴

References:

[1] Wang, R. (2014, February 18) "Research Summary: Sneak Peeks From Constellation's Futurist Framework And 2014 Outlook On Digital Disruption," Constellation Research.

[2] Anthony, S. D., Viguier, S. P., Schwartz, E. I., & Landeghem, J. V. (2018a, February). 2018 Corporate Longevity Forecast: Creative Destruction is Accelerating. Innosight.

[3] "PwC's Global Artificial Intelligence Study: Sizing the prize," PricewaterhouseCoopers, Retrieved March 1, 2021.

[4] "Incorporate enterprise A.I. now or risk getting disrupted," Fortune + C3. Ai, Retrieved March 1, 2021.



Education Today

A RIPE OPPORTUNITY FOR INNOVATION

414M

higher ed students are projected in 2030, double from 207M today and with all of the growth happening online¹

\$400B

are deployed globally by enterprises training their workforces²

500M

of learners around the world visit MOOCs in a given month³

\$1T

of market value is forecasted to be occupied by Digital Learning by 2026, close to doubling the growth rate projected before COVID-19⁴

References:

[1][4] Dawn of the Age of Digital Learning. (2020, January 14). GSV Ventures.
[2][3] 2020 Education Outcomes Report. (2020, December 7). Owl Ventures.



Stanford University
Human-Centered
Artificial Intelligence

INDUSTRY BRIEF



CORPORATE
ENGAGEMENT

Up to **\$11.5T** could be added to **global GDP by 2028** if countries succeed in better preparing learners for the needs of the future economy.¹

COVID-19 sent 1.4B students home, and even before the pandemic, low-income countries were struggling to get students the basic primary level skills.²

Do you want to know how AI can help educate your future workforce?

[1] Schools of the Future: Defining New Models of Education for the Fourth Industrial Revolution. (2020, January 14). World Economic Forum.

[2] United Nations Educational, Scientific and Cultural Organization. (2020, March 24). 1.37 billion students now home as COVID-19 school closures expand, ministers scale up multimedia approaches to ensure learning continuity. UNESCO.



Stanford University
Human-Centered
Artificial Intelligence

INDUSTRY BRIEF

Become a corporate member today.

50% of companies are likely to
miss the window of opportunity.¹

Let's talk.

Learn more about the
Corporate Members Program
and the Stanford advantage.

Panos Madamopoulos,
Managing Director for Industry
Programs and Partnerships

¹ Jacques Bughin, "Wait-and-See Could Be a Costly AI Strategy," *MIT Sloan Management Review*, June 15, 2018.



Stanford University
Human-Centered
Artificial Intelligence

INDUSTRY BRIEF

Appendix



Artificial Intelligence Definitions

Intelligence might be defined as the ability to learn and perform suitable techniques to solve problems and achieve goals, appropriate to the context in an uncertain, ever-varying world. A fully pre-programmed factory robot is flexible, accurate, and consistent but not intelligent.

Artificial Intelligence (AI), a term coined by emeritus Stanford Professor John McCarthy in 1955, was defined by him as “the science and engineering of making intelligent machines”. Much research has humans program software agents to behave in a clever way, like playing chess, but, today, we emphasize agents that can learn, as human beings navigating our changing world do.

Autonomous systems can independently plan and decide sequences of steps to achieve a specified goal without micro-management. A hospital delivery robot must autonomously navigate busy corridors to succeed in its task. In AI, autonomy doesn’t have the sense of being self-governing that is common in politics or biology.

Machine Learning (ML) is the part of AI studying how computer systems can improve their perception, knowledge, decisions, or actions based on experience or data. For this, ML draws from computer science, statistics, psychology, neuroscience, economics, and control theory.

In **supervised learning**, a computer learns to predict human-given labels, such as dog breed based on labeled dog pictures; **unsupervised learning** does not require labels, sometimes making its own prediction tasks such as trying to predict each successive word in a sentence; **reinforcement learning** lets an agent learn action sequences that optimize its total rewards,

such as winning games, without explicit examples of good techniques, enabling autonomy.

Deep Learning is the use of large multi-layer **(artificial) neural networks** that compute with continuous (real number) representations, a little like the hierarchically-organized neurons in human brains. It is currently the most successful ML approach, usable for all types of ML, with better generalization from small data and better scaling to big data and compute budgets.

An **algorithm** lists the precise steps to take, such as a person writes in a computer program. AI systems contain algorithms, but often just for a few parts like a learning or reward calculation method. Much of their behavior emerges via learning from data or experience, which is a sea change in system design that Stanford alumnus Andrej Karpathy dubbed **Software 2.0**.

Narrow AI is intelligent systems for particular tasks, e.g., **speech** or **facial recognition**. **Human-level AI**, or **Artificial General Intelligence (AGI)**, seeks broadly intelligent, context-aware machines. It is needed for effective, adaptable **social chatbots** or **human-robot interaction**.

Human-Centered Artificial Intelligence is AI that seeks to augment the abilities of, address the societal needs of, and draw inspiration from human beings. It researches and builds effective partners and tools for people, such as a robot helper and companion for the elderly.

Text by Professor Christopher Manning, v 1.1, November 2020

[Learn more at hai.stanford.edu](https://hai.stanford.edu)

References

EDUCATION & AI

Natural Language Interactions

Conversational AI to Accelerate Skill Integration | Office of Innovation & Technology. (n.d.). Retrieved July 26, 2021, from <https://gse-it.stanford.edu/conversational-ai-accelerate-skill-integration>

Kolchinski, Y. A., Ruan, S., Schwartz, D., & Brunskill, E. (2018). Adaptive natural-language targeting for student feedback. *Proceedings of the Fifth Annual ACM Conference on Learning at Scale*, 1–4. <https://doi.org/10.1145/3231644.3231684>

NLP | Office of Innovation & Technology. (n.d.). Retrieved July 26, 2021, from <https://gse-it.stanford.edu/nlp>

ReportLinker. (2021, July 15). *Global Private Tutoring Market to Reach \$201.8 Billion by 2026*. GlobeNewswire News Room. <https://www.globenewswire.com/news-release/2021/07/15/2263310/0/en/Global-Private-Tutoring-Market-to-Reach-201-8-Billion-by-206.html>

Ruan, S., Jiang, L., Xu, Q., Liu, Z., Davis, G. M., Brunskill, E., & Landay, J. A. (2021). EnglishBot: An AI-Powered Conversational System for Second Language Learning. *26th International Conference on Intelligent User Interfaces*, 434–444. <https://doi.org/10.1145/3397481.3450648>

Srivastava, M., & Goodman, N. (2021). Question Generation for Adaptive Education. *ArXiv:2106.04262* [Cs]. <http://arxiv.org/abs/2106.04262>

Augmented Teaching

Metz, C. (2021, July 20). Can A.I. Grade Your Next Test? *The New York Times*. <https://www.nytimes.com/2021/07/20/technology/ai-education-neural-networks.html>

Mu, T., Wang, S., Andersen, E., & Brunskill, E. (2021). Automatic Adaptive Sequencing in a Webgame. In A. I. Cristea & C. Troussas (Eds.), *Intelligent Tutoring Systems* (Vol. 12677, pp. 430–438). Springer International Publishing. https://doi.org/10.1007/978-3-030-80421-3_47

Ruan, S., He, J., Ying, R., Burkle, J., Hakim, D., Wang, A., Yin, Y., Zhou, L., Xu, Q., AbuHashem, A., Dietz, G., Murnane, E. L., Brunskill, E., & Landay, J. A. (2020). Supporting children's math learning with feedback-augmented narrative technology. *Proceedings of the Interaction Design and Children Conference*, 567–580. <https://dl.acm.org/doi/10.1145/3392063.3394400>

Ruan, S., Wei, W., & Landay, J. (2021). Variational Deep Knowledge Tracing for Language Learning. *LAK21: 11th International Learning Analytics and Knowledge Conference*, 323–332. <https://doi.org/10.1145/3448139.3448170>

The teacher shortage is real, large and growing, and worse than we thought: The first report in the Perfect Storm in the Teacher Labor Market series. (n.d.). *Economic Policy Institute*. Retrieved September 5, 2021, from <https://www.epi.org/publication/the-teacher-shortage-is-real-large-and-growing-and-worse-than-we-thought-the-first-report-in-the-perfect-storm-in-the-teacher-labor-market-series/>

Wu, M., Mosse, M., Goodman, N., & Piech, C. (2018). Zero Shot Learning for Code Education: Rubric Sampling with Deep Learning Inference. *ArXiv:1809.01357* [Cs, Stat]. <http://arxiv.org/abs/1809.01357>

Learning Sciences

Cutumisu, M., Chin, D. B., & Schwartz, D. L. (2019). A digital game-based assessment of middle-school and college students' choices to seek critical feedback and to revise. *British Journal of Educational Technology*, 50(6), 2977–3003. <https://doi.org/10.1111/bjet.12796>

Käser, T., & Schwartz, D. L. (2019). *Exploring Neural Network Models for the Classification of Students in Highly Interactive Environments*. 10. <https://eric.ed.gov/?id=ED599211>

Käser, T., & Schwartz, D. L. (2020). Modeling and Analyzing Inquiry Strategies in Open-Ended Learning Environments. *International Journal of Artificial Intelligence in Education*, 30(3), 504–535. <https://doi.org/10.1007/s40593-020-00199-y>

Mu, T., Jetten, A., & Brunskill, E. (2020). *Towards Suggesting Actionable Interventions for Wheel-Spinning Students*. 11. <https://eric.ed.gov/?id=ED608061>

Wang, K. D., Salehi, S., Arseneault, M., Nair, K., & Wieman, C. (2021). Automating the Assessment of Problem-solving Practices Using Log Data and Data Mining Techniques. In *Proceedings of the Eighth ACM Conference on Learning @ Scale* (pp. 69–76). Association for Computing Machinery. <https://doi.org/10.1145/3430895.3460127>

Wu, M., Davis, R. L., Domingue, B. W., Piech, C., & Goodman, N. (n.d.). *Variational Item Response Theory: Fast, Accurate, and Expressive*. 12. <https://arxiv.org/abs/2002.00276>

References

EDUCATION & AI

Learning Infrastructure

COVID-19 and education: The lingering effects of unfinished learning | McKinsey. (n.d.). Retrieved September 26, 2021, from <https://www.mckinsey.com/industries/public-and-social-sector/our-insights/covid-19-and-education-the-lingering-effects-of-unfinished-learning>

Improving Diversity and Equity in School Choice | Center for Comparative Studies in Race & Ethnicity. (n.d.). Retrieved July 26, 2021, from <https://ccsr.stanford.edu/projects/improving-diversity-and-equity-school-choice>

Piech, C., Yan, L., Einstein, L., Saavedra, A., Bozkurt, B., Sestakova, E., Guth, O., & McKeown, N. (2020). Co-Teaching Computer Science Across Borders: Human-Centric Learning at Scale. *Proceedings of the Seventh ACM Conference on Learning @ Scale*, 103–113. <https://doi.org/10.1145/3386527.3405915>

Subramonyam, H., Seifert, C., Shah, P., & Adar, E. (2020). texSketch: Active Diagramming through Pen-and-Ink Annotations. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–13. <https://doi.org/10.1145/3313831.3376155>

Wu, M., Goodman, N., Piech, C., & Finn, C. (2021). ProtoTransformer: A Meta-Learning Approach to Providing Student Feedback. *ArXiv:2107.14035 [Cs]*. <http://arxiv.org/abs/2107.14035>

Other

Education 4.0. (n.d.). World Economic Forum. Retrieved September 26, 2021, from <https://www.weforum.org/projects/learning-4-0/>

<https://plus.google.com/+UNESCO>. (2020, March 24). *1.37 billion students now home as COVID-19 school closures expand, ministers scale up multimedia approaches to ensure learning continuity*. UNESCO. <https://en.unesco.org/news/137-billion-students-now-home-covid-19-school-closures-expand-ministers-scale-multimedia>

Dawn of the Age of Digital Learning | by Michael Moe | GSV Ventures | Medium. (n.d.). Retrieved October 4, 2021, from <https://medium.com/gsv-ventures/dawn-of-the-age-of-digital-learning-4c4e38784226>

Outcomes | Owl Ventures. (n.d.). Retrieved October 4, 2021, from <https://owlvc.com/outcomes.php>




Stanford University
Human-Centered
Artificial Intelligence

INDUSTRY BRIEF

Acknowledgments

This brief was produced by the HAI Partnerships team as one of its strategic initiatives, with research and writing led by Alex Ker. We are grateful for the generosity of the following people for providing their time, helpful suggestions, and constructive feedback in the creation of this brief (names listed in no particular order): Alice Hau, Brooke Donald, Carl Wieman, Chris Piech, Dan Schwartz, Denise Pope, Emma Brunskill, Irene Lo, Jiayu He, Joshua Weiss, Noah Goodman, Peter Norvig, Hariharan Subramonyam, Karen Wang, Nick Haber, Paul Kim, William Cai, Peter Norvig, and the HAI Staff.



Stanford University
Human-Centered
Artificial Intelligence