INVITED COMMENTARY

The interplay of learning, analytics and artificial intelligence in education: A vision for hybrid intelligence

Mutlu Cukurova [iD]

UCL Knowledge Lab, University College London, London, UK

Correspondence
Mutlu Cukurova, UCL Knowledge Lab, University College London, 23–29 Emerald St, London WC1N 3QS, UK.
Email: m.cukurova@ucl.ac.uk

Abstract

This paper presents a multidimensional view of AI's role in education, emphasising the intricate interplay among AI, analytics and human learning processes. Here, I challenge the prevalent narrow conceptualisation of AI as tools in Education, exemplified in generative AI tools, and argue for the importance of alternative conceptualisations of AI for achieving human–AI hybrid intelligence. I highlight the differences between human intelligence and artificial information processing, the importance of hybrid human–AI systems to extend human cognition and posit that AI can also serve as an instrument for understanding human learning. Early learning sciences and AI in Education Research (AIED), which saw AI as an analogy for human intelligence, have diverged from this perspective, prompting a need to rekindle this connection. The paper presents three unique conceptualisations of AI: the externalisation of human cognition, the internalisation of AI models to influence human mental models and the extension of human cognition via tightly coupled human–AI hybrid intelligence systems. Examples from current research and practice are examined as instances of the three conceptualisations in education, highlighting the potential value and limitations of each conceptualisation for human competence development, as well as the perils of overemphasis on approaches that replace human learning opportunities with AI tools. The paper concludes with advocacy for a broader approach to AIED that goes beyond considerations on the design

and development of AI and includes educating people about AI and innovating educational systems to remain relevant in an AI ubiquitous world.

KEYWORDS
artificial intelligence, educational technology, future of education, generative AI, human cognition, hybrid intelligence, learning analytics

HUMAN INTELLIGENCE AND ARTIFICIAL INFORMATION PROCESSING

Artificial Intelligence (AI) is often defined as the simulation of intelligence in machines. Intelligence is a complex and multifaceted concept that encompasses several abilities. It does indeed include the capacity to learn, understand, reason, make decisions and adapt to new situations (Russell & Norvig, 2010). It expands beyond commonly considered cognitive abilities to include emotional and social components (Martınez-Miranda & Aldea, 2005), acknowledging that intelligence is not just about how well one reasons but also about how well one interacts with the world and others in it.

Intelligence is not only about what is certain, decontextualised, disembodied, tokenised and reduced to its parts, so that it is predictable and controlled. It is also about understanding things that are fluid. It is about the ability to live with, and survive despite uncertainty, and that what we are seeing as parts may indeed be wholes at another level. Today, most AI we see in educational research and practice considers AI as tools that have been developed to replace decision-making processes through analysis of big data, and prediction of the best value for a designated outcome variable, which is conveyed through a user interface. Intelligence, in my opinion, is more than that.

What do I mean by AI as tools? Let us have a look at it through the most prominent AI tool of today, OpenAI's ChatGPT.[1] Most state-of-the-art language models today including ChatGPT are based on a transformer architecture. During pre-training, a large-scale dataset of sentences is used as input to the transformer architecture. The inputs, for example, masked-out words or paired sentences, are processed automatically and the neural network model is optimised to reconstruct the original text (Demszky et al., 2023). First, an input is fed into the neural network, and it passes through the network's layers to produce an output. This process is known as the forward pass, where each layer's output is the input for the next layer, culminating in a final output from the network and it provides some predictions on masked words, next sentences and so on. Since the actual masked words and next sentences in the original text are known, based on the differences between actual and model-predicted labels, a loss function is calculated. This function measures how far the network's prediction is from the actual result and backpropagation is used to minimise the loss by adjusting the weights of the network. This process of forward pass, loss calculation, backpropagation and weight update is repeated over many iterations (or epochs) across the entire training dataset leading to the final pre-trained large language model. After the pre-training stage, large language models (LLMs) are commonly fine-tuned to improve their performance. This is the subsequent process of refining the model on a smaller, more specific dataset to adapt it to a particular domain or task (eg, through reinforcement learning with human feedback). They are also further prompt tuned which involves optimising the input prompts to guide the pre-trained model's behaviour on specific tasks without actually changing the model's parameters. After these large language models are trained, they are

used to help or replace decision-making processes for the particular task of text generation and their predictions are conveyed through a user interface.

To a certain extent, this is a simplified description, but it broadly covers the essence of LLMs' training process without getting into the details (eg, attention mechanisms (Vaswani et al., 2017)). This approach leads to the current best-performing generative AI models and tools of the day. However, due to their non-transparent nature (ie, they are black-box systems trained on vast amounts of data and their internal workings are complex), lack of understanding of real-world contexts and tendency to generate incorrect information, their value for certain learning tasks might be limited. They also lack clear, reliable and valid measures of success in educational contexts which makes it difficult for performance evaluations in real-world implementations (Chang et al., 2024). This is not to say they do not have any value for educational purposes since they can still be very valuable for certain productivity gain tasks (eg, generating initial draft content to be reviewed) and diagnosis purposes (eg, predicting particular language issues with relatively high accuracy to support teachers in their prioritisation of interventions for their learners).

Many researchers and practitioners of AIED already know these in detail, so the readers might be wondering why we need to keep reminding ourselves about the gist of how these models work, and why this is important in this paper on hybrid intelligence. It matters because when we think about AI in these terms it does not cover all the terms we would use when we are describing human intelligence, it sounds more like artificial information processing, rather than AI. Information processing is an important aspect of intelligence, but diminishing the whole concept to it would be a mistake. This is one of the reasons LLMs are sometimes referred to as 'stochastic parrots' (Bender et al., 2021). Stochastic in that they generate content based on probability analysis, and parrot because they do not appear to have an understanding of the meaning of anything they generate. LLMs are good language models, but they are not models of the real world in which the meaning is situated. Yet, this is not a competition between human intelligence and artificial information processing. We do not necessarily need more replications of human intelligence in machines for hybrid intelligence. Humans are very good at many things that today's AI is still pretty poor at, and AI is good at some others. Machines are much ahead of humans on some variables like computing floating point arithmetic, yet way behind on others like cognitive flexibility and long-term planning in unusual situations. This is not to say that humans are more intelligent than machines, or vice versa, we are differently abled.

Recognition of these differences in abilities provides a strong argument for the value of hybrid intelligence systems which are tightly coupled human–AI systems where both entities interact smoothly and dynamically, leveraging their respective strengths. In educational contexts, hybrid intelligence systems can significantly enhance human competence development by combining human cognitive flexibility, reflective long-term planning and real-world contextual understanding with AI's data processing capabilities and analytics. If the relationship holds (which also requires humans to develop certain competencies as briefly covered in Section "Educating people about AI"), hybrid intelligence systems would be invincible. This relationship, as will be detailed in later sections, not only holds the potential to achieve productivity gains in human tasks but also helps in avoiding human cognitive atrophy, and its potential amplification, by ensuring that humans remain engaged in critical thinking and decision-making processes.

AIED AND THE DIRECTION OF RESEARCH TOWARDS AI AS AN APPLIED TOOL

When we start thinking about intelligence in these broad terms, it becomes clear that considering AI as an applied tool is only one part of a much bigger picture. AI can also be considered

as a method for understanding human intelligence and learning, an opportunity to understand the differences and similarities between humans and artificial information processing. In this sense, studying human learning and studying AI are naturally intertwined. Interestingly, this way of thinking was indeed prevalent in the early roots of the learning sciences and AIED communities. Yet, it seems that we are losing this connection. In a review piece mapping publications from past and recent AIED conference proceedings and International Journal of AI in Education manuscripts on the coordinates of using AI as an applied tool versus AI as an analogy to human intelligence to study learning processes; Rismanchina and Doroudi (2023) showed that although early publications had well-distributed contributions on these coordinates, there was only one single publication in 2021 focusing on the use of AI as an analogy to human intelligence for studying learning. Why is that?

AI in education can also be conceptualised to externalise, to be internalised or extend human cognition (Cukurova, 2019). As the first conceptualisation, in the externalisation of cognition, certain human tasks are defined, modelled and replaced by AI as a tool. In the second conceptualisation, AI models can be used to help humans change their representations of thought, through the internalisation of these models. At last, AI models can be used to extend human cognition as part of tightly coupled human and AI hybrid intelligence systems. It is important to note that in such systems, changes in both agents are expected to be observed through their interactions and the whole emergent intelligence is synergistic, that is, it is expected to be more than the sum of each agent's intelligence, both human and artificial (Cukurova, 2019).

If we try to map these conceptualisations on the coordinates of Schneiderman's human control and agency versus automation through AI (Shneiderman, 2020), perhaps most traditional educational technology could be considered to have a low allowance for human agency and low automation built into them (see the AIED-HCD conceptual framework in Figure 1). With the initial proliferation of AIED research field, many researchers had the ambitious goal of creating systems that are as perceptive as human educators through the automation of certain tutoring behaviours, which led to significant developments in intelligent tutoring systems (ITSs). A similar trend in increased popularity of the externalisation conception can also be observed in most of today's LLMs and generative AI tools.

Mainly driven by the cognitivist approaches that consider learning only as an intracranial activity of information processing, ITSs have the main goal of adapting to the individual levels of mastery and needs of each student, tailoring the content, pacing and feedback accordingly. The field of AIED has multiple examples of early successful ITSs including Carnegie
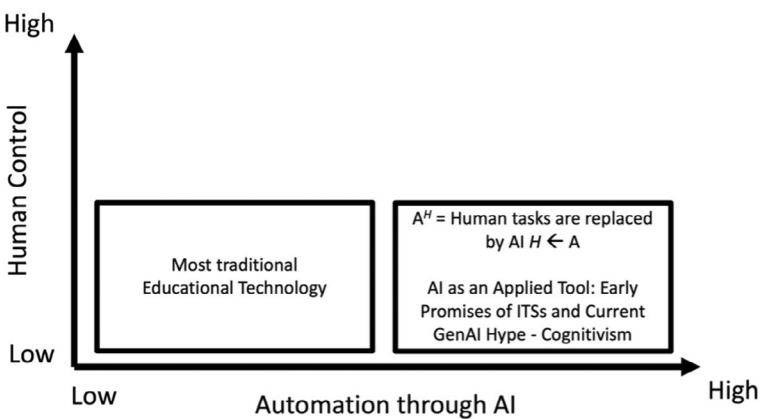

FIGURE 1   The AIED-HCD conceptual framework for human–AI interaction in education for human competence development—Human tasks are replaced by AI.

Learning (Ritter et al., 2015), Duolingo (Von Ahn, 2013) and Assistments (Heffernan & Heffernan, 2014) and similar tutors are also emerging with the use of LLMs in educational chat interfaces (eg, Khanmigo). It is important to highlight that perhaps one of the reasons for the success of these systems compared to relatively less successful ones in the field is that they have engaged with end users (eg, teachers, students and school admins) and taken the dynamics of education systems, school environments and classrooms very seriously early on. For Carnegie learning, for instance, there are physical course books and other resources in addition to the tutoring system itself. There is significant guidance, advice and support for how teachers and students should interact ahead of any use of the tutoring systems, as well as while they are being used. So, there is significant human agency support, both for the teachers and the students, at the deployment phase of these ITS examples even the tools themselves have pedagogical action externalisation and automation at the system levels.

ITSs have traditionally focused on pedagogical task automation in digital environments. However, these systems can also work with data from physical spaces and can focus on a range of affective, metacognitive and engagement tasks. For instance, an example from Kawamura et al. (2022) is using multimodal data to detect students' engagement states while working with an ITS. The system processes data from students' heart rate, seat pressure and facial recognition to model students' level of awakeness and aims to provide suggestions for rest or adjust the content or feedback accordingly.

These systems are also not limited to traditional human-generated educational content but can be delivered with AI-generated synthetic media. Recent experimental research investigating the potential of using AI-generated synthetic video to create viable educational content for ITSs, no significant differences were observed in learning gains and learner experience between the two conditions of students learning from a recorded human lecture versus from an AI-generated synthetic media that deliver the same content (Leiker et al., 2023). Admittedly, this is from a relatively small sample of 83 adult learners in an explorative study, but these models are getting better by the day, and the promising results justify further large-scale explorations of their potential.

Such tutoring systems that externalise particular pedagogical tasks to model and automate their support have been a significant area of research in AIED and they work very well for various domains and knowledge acquisition tasks. More long-term evaluation studies and RCTs with larger sample sizes and better-controlled conditions in the real world would always be welcome but there is indeed good evidence both at the individual studies level including but not limited to SQL-Tutor (Mitrovic & Ohlsson, 1999), ALEKS (Craig et al., 2013), Cognitive Tutor (Pane et al., 2014), ASSISTments (Koedinger et al., 2010) and also at the meta-reviews level. For instance, Van Lehn (2011) found that the effectiveness of intelligent tutoring systems was nearly as effective as average human tutors; Ma et al. (2014) found similar results both when compared to no tutoring or large group human tutor instruction; Pane et al. (2014) found evidence of the relative effectiveness of online tutors over conventional teaching; in Kulik and Fletcher (2016)'s work, the median effect was observed raising test scores of 0.66 standard deviations over conventional levels, or from the 50th to the 75th percentiles, and du Boulay (2016) summarised some metareviews in his work and showed that these systems achieve positive results in the delivery of knowledge acquisition, particularly for the subjects of maths, language learning and algebra.

RESEARCH EVIDENCE ON EFFECTIVENESS VERSUS REAL-WORLD IMPACT OF AI IN EDUCATION

Considering these systems are not necessarily new, and the evidence about their effectiveness is not necessarily new, one question to ponder upon as a community is why they are still not prevalent in mainstream education.

There are various reasons why this is the case and each one of these reasons would probably require a paper on its own to be discussed in detail. But, let me attempt to mention a few reasons, from my point of view, briefly here. First, there are numerous factors influencing the adoption and use of AI in education that are broader than the effectiveness of the specific AI technology. These include but are not limited to policy landscape, institutional governance, pedagogical culture, technological infrastructure and social support mechanisms provided to teachers. For instance, in our recent work looking at the factors influencing teachers' adoption of AI in schools with about 800 schoolteachers (Cukurova et al., 2023), we observed that although AI-tool-related factors were indeed important, they were not necessarily the most important factors influencing the teachers' engagement with AI in schools. Not generating any additional workload, teachers' knowledge of and confidence in using AI, increasing teacher ownership, generating support mechanisms for help when needed and assuring that ethical issues are minimised, were also essential for the adoption of AI in schools. So let us never forget that the tools we are working on are not only closed engineering systems but are part of a large socio-technical ecosystem, and many factors will influence their adoption and effectiveness.

Second, education with fully automated systems that externalise human cognition to deliver educational practice can be argued to dehumanise learning. When AI in education is considered in a narrow sense, as lonely individual learners working on their own with an AI system, this might indeed lead learners to prioritise information gathering and declarative knowledge acquisition over tacit knowledge and wisdom which comes through rich experiences in the real world. Particularly, if these systems are considered as a replacement for human interactions with each other, and with the real world, then the knowledge that comes through experience and practical acquisition of an embodied skill can be replaced with tokens of representations that are far off the actual construct. Learning is not only about absorbing information and education is not only about learning. These are also about developing social competence, emotional intelligence and various metacognitive abilities in real-world interactions as well as serving other societal needs. Fully automated systems are unlikely to deliver such experiences as a whole in the near future, even if it is ever possible. However, these technologies can also be used as opportunities for increasing students' interactions with adults, and with each other, in affluent, private schools (see, for instance, Alpha School, Austin). So, the impact of AI in education is not only dependent upon the AI tool itself but the particular learning design and the instructional design in which they are embedded.

Third, a significant amount of work is still needed to address socio-psychological barriers to the use of AI in educational contexts. Students need to be motivated enough to engage with AI tools in the first place, yet only about 5% of them manage to engage with educational resources long enough on their own to get statistically significant benefits (Weatherholtz et al., 2022). In addition, teachers and learners tend to have confirmation biases and unrealistic expectations from AI (Nazaretsky et al., 2021). Previous research showed that when people are presented with content framed as coming from AI, they tend to judge it as less credible and trustworthy compared to the same content framed as products of educational psychology or neuroscience (Cukurova, Luckin, & Kent, 2020). Similar results are now emerging, when AI-generated content or feedback is presented to teachers and students they tend to judge its quality lower and trust it less if they know that it is AI generated (Nazaretsky et al., 2024). There is emerging work focusing on how we measure and gain the trust of teachers and learners in AI, but considerably more research is needed in this space to optimise their trust (eg, Nazaretsky et al., 2022).

Fourth, particularly with the goal of keeping the talk focused on AI design and development rather than opening it up to broader ecosystem-level issues, one of the most significant limitations of these AI tools that externalise tasks to automate them is that they are

commonly built on computational cognitive models that overlook the intricacies of socio-cultural learning occurring beyond an individual's mind (Doroudi, 2023). For many learning scientists who are aligned with constructivist learning theories, this approach is considered too simple to represent the complexities of the kind of learning they are interested in.

ALTERNATIVE CONCEPTUALISATIONS OF AI

In this regard, another conceptualisation of AI can provide significant opportunities. AI can also be conceptualised as computational models of complex learning phenomena for humans to internalise and change their representation of thought (Kent et al., 2021). These lead to relatively low automation systems that allow high human agency and control (see the AIED-HCD conceptual framework in Figure 2).

In learning sciences literature, it is well established that learning may be facilitated with mental models (Johnson-Laird, 1983). Mental models help us explain and predict how people interact with the world, and how they explain, understand, solve anticipated events and communicate. While mental models are internal structures, they can be exteriorised when triggered by interaction (Bransford et al., 2000). The assumption in this conceptualisation of AI is that computational models can offer an externalisation trigger for mental models. By that, they can serve as a learning affordance, and the learning outcomes can be observed through changes in learners' language, and their ability to explain, predict or diagnose phenomena (Kent et al., 2021). So, although everyone comes to a learning situation with their own mental model of what success looks like, as well as the models that learners hold about themselves (Byrne, 1992), we can use data to model what success looks like with computational approaches and present these models back to people and trigger them to refine their mental models.

This conceptualisation allows more opportunities for researchers interested in understanding and designing learning environments through the lens of constructivist and socio-cultural learning perspectives. It allows opportunities to focus on AI models to help develop learner competence through rich learning experiences and reflecting on these experiences with the help of computational models.

In my research, we have been mainly focusing on such open-ended learning environments and trying to design analytics and AI models that would support teachers and learners in these constructivist contexts. For instance, we were investigating students who were
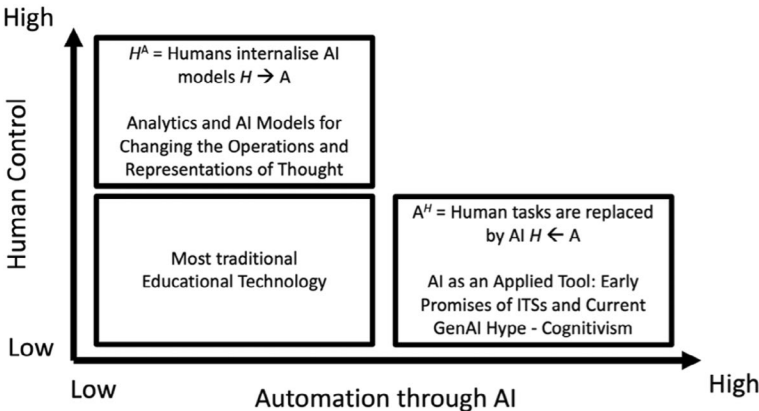

FIGURE 2 The AIED-HCD conceptual framework for human–AI interaction in education for human competence development—Humans internalise AI models.

engaged in solving open-ended design problems (eg, Cukurova et al., 2016) and have been collecting data through 2D and 3D video for face and gesture analysis, tracking the speed and distance between their hands, logging their physical computing kit interactions, their mobile tool reflections and notes, as well as their self-declared emotions (Spikol et al., 2018).

If we take the first conceptualisation of AI, we can use these kinds of data to build machine learning classifications of success in these environments. First, using these different modalities of data, one can calculate various input measures as independent variables. For instance, about the group itself such as the number of faces looking at the screen, their mean distances, their hand movement in space and speed, their audio features, gaze directions, emotional states, etc. but also from the context such as the amount of time each group spends on different phases of the collaborative problem-solving process such as ideation (problem scoping), activities to solve the problem and reflection on potential solutions. Then, you can ask researchers/teachers/experts to use existing learning sciences frameworks and evaluation metrics to label groups' competence while watching these interaction videos and meticulously labelling them (eg, Cukurova et al., 2018).

This leads to having multimodal measures of potentially relevant variables as input and labelled competencies of groups as the ground truth output. Then it becomes a machine learning problem of building reliable and accurate classification and regression models for the tasks of predicting competence classes or scores (eg, Spikol et al., 2018). In order to see the value of different modalities of data, we also train models with various features removed to often find out that the best results are achieved when multimodal data are used compared to unimodal predictions (Cukurova, Giannakos, & Martinez-Maldonado, 2020).

These models at this stage still have significant technical problems, and they are usually prototypical, rather than reliable tools for immediate real-world use. Regardless, often time the ultimate goal of such prediction models is to generate some kind of a dashboard for teachers and learners to directly intervene in the practice. For instance, Aslan et al. (2019) used a multimodal learning analytics (MMLA) tool with a dashboard to provide teachers with help to prioritise and structure their interactions with students. Their results indicate that when teachers are using the dashboard, they spend statistically significantly less amount of time in close monitoring actions and more amount of time in scaffolding activities. Students also appear to spend less time in bored states, and there can be a positive impact on their learning gains.

THE CHALLENGE OF USING AI AS A TOOL TO DIRECTLY INTERVENE IN TEACHING AND LEARNING

However, the use of AI that directly intervenes in the practice of teaching and learning in such constructivist learning environments has significant challenges. These issues broadly relate to threatened human agency, the challenges of predictions in social contexts and the normativity issue of not being able to decide what is actually, or eventually, good or bad in a complex social learning situation (summarised in Figure 3). Addressing these issues would require a stronger alignment between human values and AI goals which is a significant challenge highlighted by various other scholars (eg, Christian, 2021; Russell, 2019).

Of course, prediction issues also include well-documented challenges of algorithmic bias, transparency and accountability of AI tools (eg, Baker & Hawn, 2022), but I think the concern is even bigger than these. The issue is that in the design and development of these tools, we are used to doing engineering work where we can in effect see how the gears mesh to understand how things work. Whereas I am not sure if it is always possible to make a mathematical narrative or a model to explain, or predict, how a complex social system behaves. Is it always possible to explain or predict all aspects of human learning and human
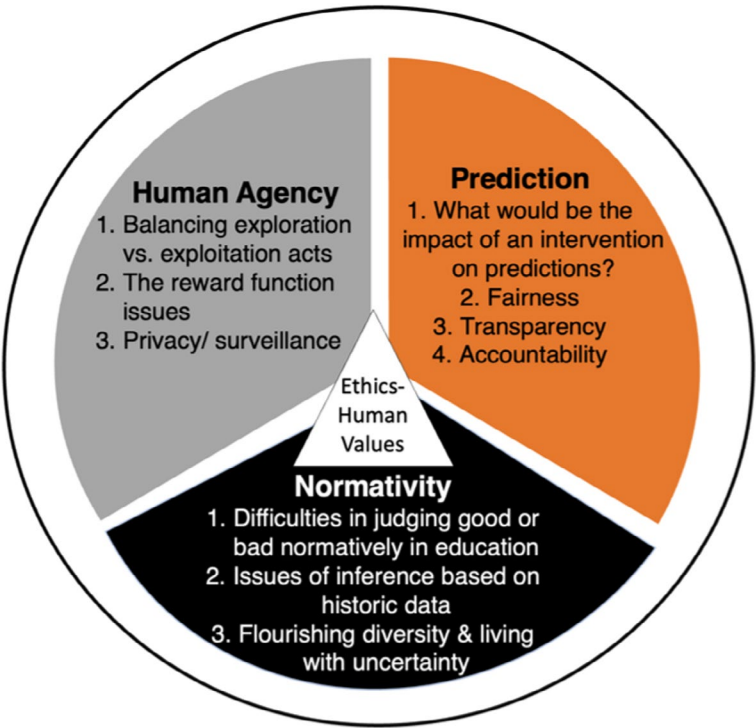
FIGURE 3   Main categories of issues related to using AI that directly intervenes in the practice of teaching and learning.

competence development? AI was coined as a term in the Dartmouth College summer school proposal in 1956 based on the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it (McCarthy et al., 2006). Since then, we somehow live under the impression that if only we could find them, there would be formulas and models to somehow predict all aspects of human learning. Yet, maybe to find out that to develop what such learning is, we just have to go through the same irreducible steps as the system itself. Maybe some aspects of learning just come through the slow experience of living those learning experiences. This itself makes the time spent on them more meaningful in the sense that, we just cannot jump ahead to get the answer with a prediction telling us what would be the most productive next step to take in these complex socio-cultural learning environments.

On the other hand, if we take the second conceptualisation of AI, as computational models for humans to internalise; these can be considered as opportunities to describe the learning processes in more precision, rather than aiming for 'the potentially impossible task' of prescribing acts for the future based on predictions in a complex social learning process. In this sense, AI models become opportunities for thinking about learning. As Seymour Papert (2005) famously noted 'You can't think about thinking without thinking about thinking about something'. This suggests that to engage in the process of thinking, one must have something specific to think about. It underscores the importance of the objects of our thoughts in shaping the processes and pathways through which our thinking unfolds. In this sense, AI models can be objects to think about human learning. In learning analytics research, we often utilise this idea of 'from clicks to constructs' (eg, Knight & Buckingham Shum, 2017) aiming to describe precisely the relationship between the digital traces of data collected and the educationally meaningful constructs we are interested in supporting.
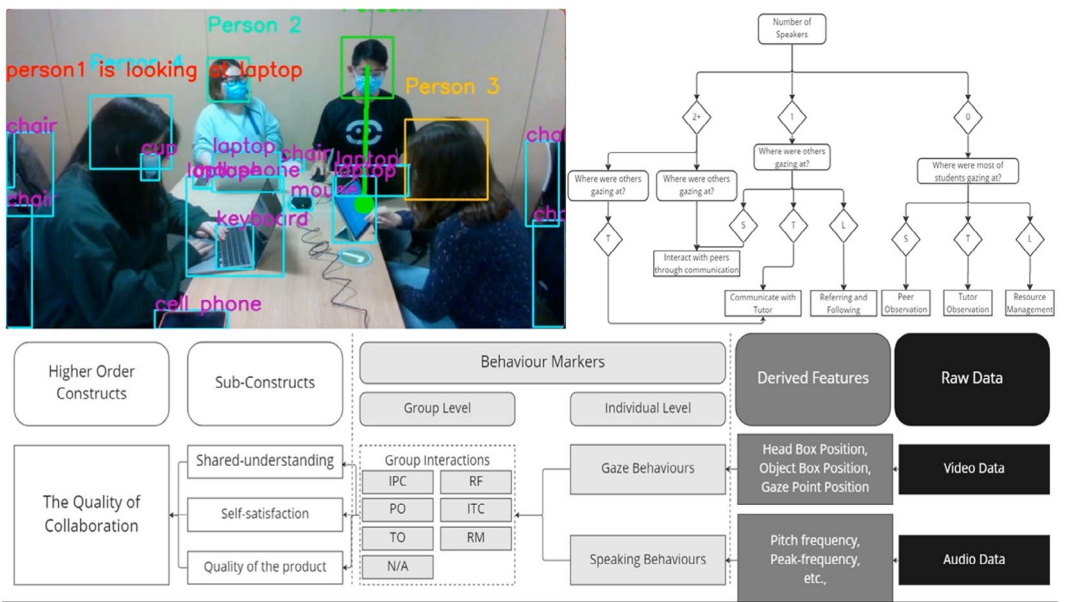
FIGURE 4   An example of the transitioning 'From digital traces to constructs of collaboration'.

Below is an example that focuses on the concept of collaboration, moving from digital traces of audio data and video data processed with computer vision to model group interaction behaviours and their connections to outcome measures of shared understanding, satisfaction and product quality (Zhou et al., 2024; Figure 4).

Sometimes, such models can give the impression that our attempts are for, or only, to build prediction models for prescriptive suggestions accordingly. If such an impression is taken forward, as the constructs are broken into small measurable components, more mechanistic measures of *how much, how fast* or *how precisely* a component is completed may be prioritised. In the collaboration process model, for example, if one keeps eye contact for one more second, starts speaking with a peer one second earlier, clicks the resource button faster or attempts to answer a question one more time, they would develop the expected competence in collaboration. For certain learning experiences, more and faster completion of a component, or even the more precise yet not flexible completion of it, may not lead to improvements in the quality of learning. Such measures are great for measuring the capabilities of machines, but not for all aspects of human learning. Taking a more modest goal of using AI models for describing these learning processes in more detail and precision can still allow significant opportunities for researchers, teachers and learners.

VALUE OF MAKING LIVED EXPERIENCES VISIBLE TO END USERS

For end users, the models provide specific and precise feedback opportunities to improve their awareness of the lived experiences and keep them motivated to engage with relevant activities in the future. Building on the same example, for instance, visualisations of speech time percentages and the types of group interactions in each group can be used as feedback after group interactions. Similarly, different types of group interactions students engage with during the timeline of their activity and relevant feedback based on these interactions for self-awareness, reflection and motivation for future interactions can be created and sent

to students (eg, Zhou et al., 2021) and teachers (Pozdniakov et al., 2022). In my view, these models are tools for making sense of the phenomena being modelled and sharing this sense with the end users, but they are not necessarily the mirrors of the reality representing the complex process of collaborative learning itself.

This recognition of the limitations of the AI models also provides opportunities for more realistic interactions between the model outputs and end users. For instance, we have been working on models that detect speech in groups, transcribe it from speech to text and label the text with groups' challenge moments in their collaborative discourse (Suraworachet et al., 2024). Then, these detected challenge moments are also sent back to students as feedback with visualisations and further explanations based on the threshold values on certain aspects and diagnosed challenges, building upon the human interpretation of these values and providing suggestions on how to address these challenges in their next group activities. Rather than assuming what is being modelled and detected as the reality, we also provide students with their own transcriptions, as well as the details on how these models are built, and their episode-level and sentence-level predictions. So, they see the sense-making process we use in detail and interpret it themselves when such feedback is valuable, or when it is safe to be ignored.

When evaluating the value of AI models with this conceptualisation for feedback to end users, it is always interesting to observe what meaning learners and teachers generate from the visible information from these models, what actions they take based on such an understanding, what is the accountability of this understanding and what are their concerns related to their values and moral considerations (Erickson & Kellogg, 2000). End users are often aware of the limitations of these models and agree about the incomplete nature of them. However, they still find them valuable to increase their awareness of their own learning activities, and also awareness of others' behaviours in these socio-constructivist learning environments (Zhou et al., 2021). The accountability of this awareness tends to influence both their motivation and engagement with the learning experience, as well as their regulation (including co-regulation and socially shared regulation (Hadwin et al., 2011)). Therefore, making the lived experiences of learning more visible and explicit with computational models still has significant value for teachers and learners regardless of their potential for accurate future predictions with prescriptive suggestions.

VALUE OF CONTRIBUTING TO LEARNING SCIENCES LITERATURE

Another value of this conceptualisation of AI as models for making sense of the phenomena being modelled is that they enable opportunities for clarification and communication of researchers' concepts in a more detailed, precise and formal language; generating potential insights into complex learning processes to advance learning theory (Giannakos & Cukurova, 2023). For instance, in our recent work, we have been using five-channel multimodal data to make sense of collaborative problem-solving processes. Instead of looking at each group's collaborative problem-solving activities separately, we merged all sequences from all groups in our dataset for a given task, then looked at clusters of patterns emerging using optimal matching algorithms and Ward's clustering, identifying three different clusters of CPS patterns in multichannel data streams at a granularity level that is not possible with traditional statistics. We then looked at the transitional and structural differences of each CPS pattern type with hidden Markov models and epistemic network analysis to discover that they are associated with different performance outcomes (Ouyang et al., 2023).

In turn, such insights can provide opportunities for learning theory to be further improved (Doroudi, 2023). A recently published literature review investigating the relationship between

learning theories and models in MMLA research indicated that such models have the potential to contribute to learning theory (Giannakos & Cukurova, 2023). For instance, in their work on embodied learning and maths education using insights from computational models and eye-tracking data, Abrahamson et al. (2015) revisit, refine and elaborate further on some of the seminal claims from Piaget's theory of genetic epistemology (Piaget, 1970) (eg, his insistence on the role of situated motor–action coordination in the process of reflective abstraction).

HUMAN–AI HYBRID INTELLIGENCE SYSTEMS

Going back to three conceptualisations of AI, this leaves us with the last corner of high automation and high human agency, the corner of human cognition being extended with AI in tightly coupled human–AI hybrid intelligence systems (see the AI-HCD framework in Figure 5). The top right corner of the AI-HCD framework is not covered as much in this paper, mainly because we have yet to see substantial work on this front in AIED. At best, the current complementarity paradigm is to make a better match of what humans can do and what AI can do with the problems to be tackled to achieve productivity gains at tasks rather than making humans more intelligent. More commonly, for any given job, we tend to give up our agency to AI to complete a task for us, which in turn is expected to improve the performance of task completion. This inclination to employ AI for task performance is not only appealing but also reflects a natural human propensity for automating processes, which is a trend evident throughout human history (Lubars & Tan, 2019). However, we must be judicious in selecting the tasks we delegate to AI as the over-reliance on AI could lead to the atrophy of critical competencies in the long term. For instance, there are ongoing attempts to automate qualitative coding processes with the use of large language models, as this can increase the productivity gains in generating labelled qualitative data (Barany et al., 2024). However, on many occasions, qualitative coding is not only done to generate final labelled datasets but
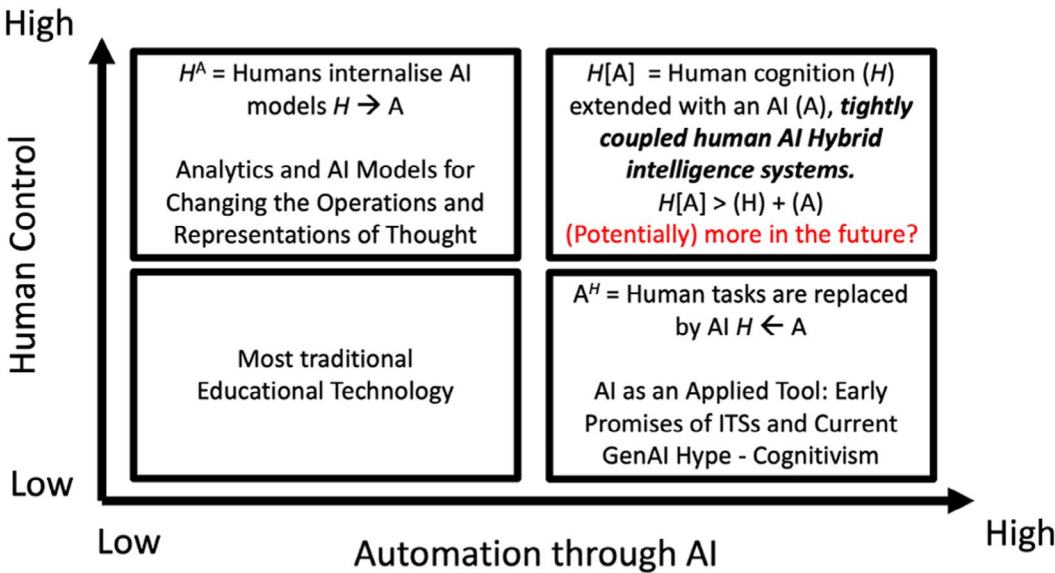

FIGURE 5 The AIED-HCD conceptual framework for human–AI interaction in education for human competence development—Human cognition extended with AI in tightly coupled hybrid intelligence systems.

also to improve researchers' and practitioners' meaning-making competence through their engagement with the reflective coding process. Similarly, for the literature reviews with LLMs, the goal is not only to identify gaps in the literature but to improve one's understanding of the research undertaken on a particular topic. Full automation is unlikely to be valuable with these tasks in the long-term for human competence development unless their use cases are well structured with learning design principles that prioritise human agency.

There is also significant concern regarding the human-in-the-loop correction of AI-generated content, as this process might lead to convergence towards the AI-generated content and labels rather than critically evaluating them to ground them in our own understanding. Similar to today's society, in which most people converge towards the first suggestion of a Google search engine rather than critically engaging with suggestions to choose the most appropriate one for our goals (Burguet et al., 2015). We might end up in a future where the first generated automated content is taken as 'the truth'. This is a significant concern for the future of education, therefore, just because we can technically automate a task in education does not always mean that we should. Such significant decisions about automation should be evidence informed by adopting an approach that involves testing any automation ideas and observing the outcomes in the long term. Otherwise, the rush towards automation may result in humans forfeiting their fundamental cognitive competencies and evolutionary superiority as a species in the future. A potentially wiser perspective would be to highlight the critical importance of maintaining and enhancing our intrinsic intellectual abilities, which have historically conferred upon us distinct advantages for survival and adaptation. This requires alternative conceptualisations of AI.

On the other hand, for AI models that are internalised by humans, the goal of the model would be to 'fade away' as humans' competence at the task that is modelled develops through their interaction with the model. At last, extending human cognition for intelligence augmentation in tightly coupled human–AI hybrid intelligence systems would require AI to be a synergistic superstructure built on top of the human intelligence structure, in a way that as the interaction with the tool increases, our competence at the task being modelled would also increase.

Our interactions with AI systems are influencing us and we are currently lacking long-term impact studies of these interactions as a research community. In order to be able to achieve human–AI hybrid intelligence systems, we need AI models that are able to interact fluidly with us to shape intents and meanings dynamically. The current AI systems lack the ability to update their models based on interactions with users real time. Human–AI hybrid intelligence systems would require interactions with AI models in which AI encourages people to shape their own meanings by leveraging the strength of language as a tool for mutual understanding. People very rarely come to an interaction situation with a specific, or fully formed, meaning in mind. Rather such meaning is shaped dynamically through the interaction itself. This requires AI to help humans shape their intents and knowledge in interactions rather than pushing their predictions using language as a query response approach. Current AI systems' requirement that the meaning should be clearly transferred to an AI for it to generate a relevant output is a significant limitation against the creation of human–AI hybrid intelligence systems. In such a paradigm, the responsibility of meaning making would also be shared between the two agents, human and AI, rather than humans assuming the entire responsibility for the success of the interaction. In such human–AI hybrid intelligence systems, as humans' interaction with the system increases their competence at the task at hand would also increase. There is an urgent need for research in human–AI interaction in education and I hope that there
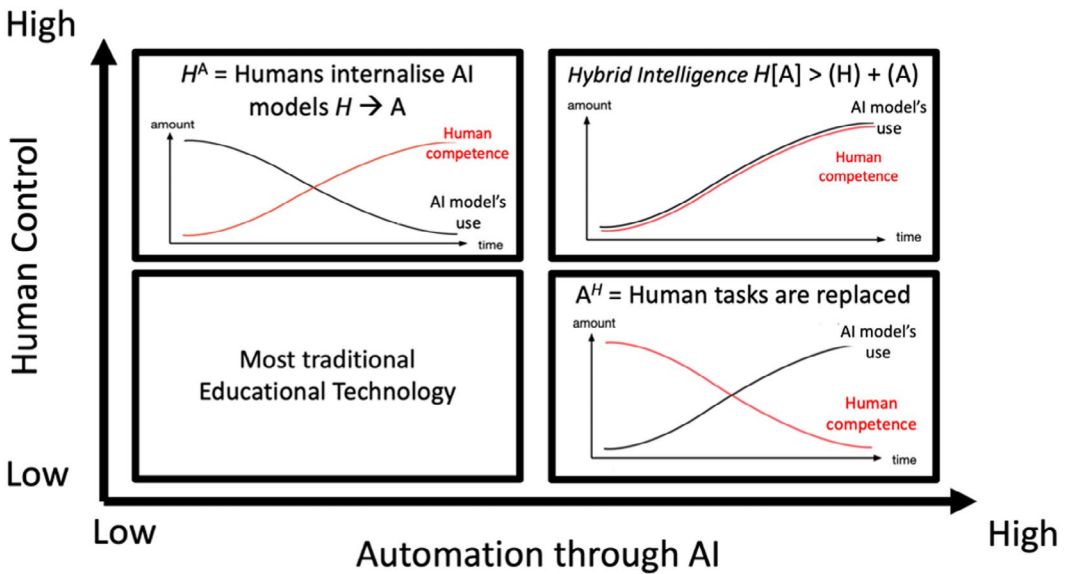
FIGURE 6   The impact of three AI conceptualisations on human competence development in the long term—The AIED-HCD conceptual framework.

will be more examples of developing and implementing human–AI hybrid systems in the future (Figure 6).

AIED IS BROADER THAN THE DESIGN, DEVELOPMENT AND USE OF AI

Most of the paper has been focusing on the design and development of AI in education and its different conceptualisations including human–AI hybrid intelligence systems. However, the conceptualised human–AI hybrid intelligence system here would also require humans to develop a set of competencies to operationalise such systems and would also require significant innovations in our education systems for them to be integrated. Broadly speaking, in addition to the design, development and use of AI systems in education, there are two other AI implications for education that are equally important for such a conceptualisation to make real-world impact in education. One is about educating people about AI and data so that they can learn how to use it safely, effectively and ethically (see Figure 7). Another one is work that focuses on how we should think about innovating our education systems so that they are more compatible with and still relevant to a world heavily influenced by AI (Luckin & Cukurova, 2019).

Educating people about AI

Educating people about AI is indeed about teaching AI to certain groups in population to create AI experts, including AI in our school curricula as well as investing further in our tertiary education institutes to improve capacity to sustain expertise in AI, yet not limited to this. There are many specific research and policymaking initiatives providing guidance in this space (eg, AI4K12 (Touretzky et al., 2019), UNESCO K-12 AI curricula).

FIGURE 7 AI's three main implication areas for education.

Aspects	Progression		
	Acquire	**Deepen**	**Create**
1. **Human-centred Mindset**	Human agency	Human accountability	AI social responsibility
2. **Ethics of AI**	Ethical principles	Safe and responsible use	Co-creating AI ethical rules
3. **AI Foundations and Applications**	Basic AI techniques and applications	Application skills	Creating with AI
4. **AI Pedagogy**	AI-assisted teaching	AI-pedagogy integration	AI-enhanced pedagogical transformation
5. **AI for Professional Development**	AI enabling lifelong professional learning	AI to enhance organizational learning	AI to support professional transformation

FIGURE 8 The draft UNESCO AI competency framework for teachers.

Perhaps more importantly, it is about helping an overwhelming majority of society, if not everyone, to develop relevant AI competencies so that they can use AI effectively and ethically as part of tightly coupled human–AI hybrid intelligence systems. A recent example in this space is the draft UNESCO AI competency framework[2] for teachers (see Figure 8). The key distinctions between AI and previous iterations of ICT tools necessitate the definition of a specific set of competencies for people. For instance, human–AI hybrid intelligence systems would demand a stronger emphasis on competencies related to human agency, ethical considerations, critical thinking and human-centred design in human–AI interactions. Training people to develop their competencies in these aspects has utmost importance for the future of human–AI hybrid intelligence systems.

AI competency training is not limited to gaining fundamental AI knowledge, techniques and skills to apply AI. It is much broader than this. In this version of the UNESCO AI competency framework for teachers, AI techniques and applications is only one of the five main aspects of AI competence. Similar set of competencies would be required from people who are not AI experts but are expected to interact with AI as part of their daily and professional lives. Often time both learning analytics and AIED communities consider the creation of a particular AI tool, model or dashboard as the end goal of research. However, such an approach is limited to making progress towards a future of human–AI hybrid intelligence. It

might be more productive to start considering the tools and models we build as the start of another line of research and practice journey in which end users' competencies are further developed to achieve the expected impact of these solutions. For this, teacher and learner interactions with tools and models should be appropriately scaffolded and supported, rather than assuming that if we left the AI tools in the hands of end users, hybrid intelligence would emerge automatically.

Innovating education systems for an AI-driven world

Finally, AI does not only have implications for our education systems that are direct but there are also multiple indirect implications which necessitate innovations in our existing education systems. The future of tightly coupled human–AI hybrid intelligence systems would require significant advancement in our traditional educational systems, particularly with regard to assessment structures.

For instance, although most skills and competencies we are interested in and expect people to develop through education in modern societies are process driven; often times in traditional education systems, these are only evaluated through the outcomes of this competence rather than the process that leads to it. We need to move towards innovations in our assessment structures that encourage process evaluations rather than outcome evaluations only. Recent advancements in LLMs created significant concerns that most student will use them to submit their assignments rather than write essays themselves. This is a valid concern, but it misses the point in the big picture. Essay writing is most often assigned not because the resulting essay has much value to us, but because the process of writing an essay teaches crucial skills to people: regulation of one's own behaviours to engage in a topic, researching a topic, judging claims' accuracy, synthesising knowledge and expressing it in a clear, coherent and persuasive manner. These skills should be the focus of assessment not *only the final product of what is produced*. This is an example of the kind of innovation we need in our education systems now that AI is ubiquitous.

As an example, in addition to traditional content feedback on students' writings, Suraworachet et al. (2023) have been providing students with personalised behavioural feedback based on their writing engagement analytics using their data from online word-processing platforms like Word or Google Docs as part of their assessment. When students are engaged with the platform data on time, how much the content is edited and what content exactly is added are all logged to be able to model their essay writing engagement trends. Based on the analytics of their engagement, students are sent feedback with key suggestions on how they can improve their engagement (Suraworachet et al., 2021). The feedback has formative suggestions on how they can space their writing practice rather than cramming to complete it all the night before, as well as how previous years' high-performing students engaged with their writing assignments.

When the impact of such analytics feedback interventions is evaluated, it was observed that the intervention had limited impact on students who were already doing well in the course, but significantly boosted the engagement and performance of those students who were struggling and were initially predicted to fail (Suraworachet et al., 2023). This kind of innovation prioritising process evaluations with analytics is an example of the kind of innovation needed in education systems for assessments to be meaningful if we are to move towards the use of human–AI hybrid intelligent systems in the future.

CONCLUDING REMARKS

Although most of the current rhetoric is limited to this, AI is more than a set of applied tools we use in education. As presented in the proposed AIED-HCD conceptual framework here (see Figures 5 and 6), AI can be conceptualised to externalise, internalise or extend human cognition as part of tightly coupled human–AI hybrid intelligence systems. AI can also be a methodology to create computational models of learning as objects to think about learning. While we are doing so, we might notice that some aspects of learning just come through the slow experience of living those learning moments and cannot be fully explained with AI models being hacked with predictions. Still though, if we take the more modest goal of using AI models for describing these learning processes in more detail and precision, they can provide valuable opportunities for feedback, motivation, awareness and contributions to theoretical understanding.

Each conceptualisation of AI discussed here might potentially bring some advancements to education, yet they also have significant unintended consequences that should be carefully considered. Our interactions with AI systems are influencing us, and we are currently lacking long-term impact studies of these interactions as a research community. This is a pivotal moment for all of us as scientists and practitioners to envision a future of education that is aligned with our societal values and ensure that AI is used responsibly to achieve this vision. Therefore, this special section on hybrid intelligence has utmost importance to stimulate the line of research on empirical investigations of human–AI interaction in education as well as relevant studies on the implications for improving human competencies in their interactions with AI, and innovating education systems for the future of hybrid intelligence. Based on the emerging evidence in these studies, the future of education with hybrid intelligence systems should be intentional, evidence informed and human centred.

As AIED researchers we finally get the expected attention from society due to recent developments in generative AI, let us not stop questioning who we are as a community, and what we are doing as researchers and practitioners of learning, analytics and AI. Let us not forget that the purpose of scientific research is only realised to the extent that it helps us understand who we are. Research in AI began as an attempt to understand our own intelligence, its atrophy, its augmentation or amplification. Which goal are we striving for? When is it acceptable to entrust core cognitive competencies to an AI tool and when might this pose too great a risk? What are the long-term implications of the tools we are developing on our own competence? We must exercise wisdom in making such decisions and consider alternative conceptualisations of AI in our research and practice. Everything we do as a research community is about humans; it has always been so, and it always should be.

CONFLICT OF INTEREST STATEMENT
No conflict of interest to be declared.

DATA AVAILABILITY STATEMENT
None.

ETHICS STATEMENT
None.

ORCID
Mutlu Cukurova 🅾 https://orcid.org/0000-0001-5843-4854

ENDNOTES
[1] https://chatgpt.com/chat.

[2] https://www.unesco.org/sites/default/files/medias/fichiers/2024/07/UNESCO-Draft-AI-competency-frameworks-for-teachers-and-school-students_0.pdf.

REFERENCES
Abrahamson, D., Shayan, S., Bakker, A., & Van Der Schaaf, M. (2015). Eye-tracking Piaget: Capturing the emergence of attentional anchors in the coordination of proportional motor action. *Human Development*, *58*(4–5), 218–244.

Aslan, S., Alyuz, N., Tanriover, C., Mete, S. E., Okur, E., D'Mello, S. K., & Arslan Esme, A. (2019). Investigating the impact of a real-time, multimodal student engagement analytics technology in authentic classrooms. In *Proceedings of the 2019 Chi Conference on Human Factors in Computing Systems*, Glasgow, UK (pp. 1–12).

Baker, R. S., & Hawn, A. (2022). Algorithmic bias in education. *International Journal of Artificial Intelligence in Education*, *32*, 1–41.

Barany, A., Nasiar, N., Porter, C., Zambrano, A. F., Andres, A. L., Bright, D., Shah, M., Liu, X., Gao, S., Zhang, J., Mehta, S., Choi, J., Giordano, C., & Baker, R. S. (2024). ChatGPT for education research: Exploring the potential of large language models for qualitative codebook development. In *International Conference on Artificial Intelligence in Education* (pp. 134–149). Cham.

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021, March). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, Virtual Event Canada, Canada (pp. 610–623).

Bransford, J. D., Brown, A. L., & Cocking, R. R. (2000). *How people learn: Brain, mind, experience, and school (expanded edition)*. The National Academies Press.

Burguet, R., Caminal, R., & Ellman, M. (2015). In Google we trust? *International Journal of Industrial Organization*, *39*, 44–55.

Byrne, R. (1992). The model theory of deduction. In Y. Rogers, A. Rutherford, & P. Bibby (Eds.), *Models in the mind: Theory, perspective and application* (pp. 11–28). Academic Press.

Chang, Y., Wang, X., Wang, J., Wu, Y., Yang, L., Zhu, K., Chen, H., Yi, X., Wang, C., Wang, Y., Ye, W., Zhang, Y., Chang, Y., Yu, P. S., Yang, Q., & Xie, X. (2024). A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology*, *15*(3), 1–45.

Christian, B. (2021). *The alignment problem: How can machines learn human values?* Atlantic Books.

Craig, S. D., Hu, X., Graesser, A. C., Bargagliotti, A. E., Sterbinsky, A., Cheney, K. R., & Okwumabua, T. (2013). The impact of a technology-based mathematics after-school program using ALEKS on student's knowledge and behaviors. *Computers & Education*, *68*, 495–504.

Cukurova, M. (2019). Learning analytics as AI extenders in education: Multimodal machine learning versus multimodal learning analytics. In *Proceedings of the Artificial Intelligence and Adaptive Education Conference*, Beijing, China (pp. xx1–xx3).

Cukurova, M., Avramides, K., Spikol, D., Luckin, R., & Mavrikis, M. (2016). An analysis framework for collaborative problem solving in practice-based learning activities: A mixed-method approach. In *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge*, Edinburgh, UK (pp. 84–88).

Cukurova, M., Giannakos, M., & Martinez-Maldonado, R. (2020). The promise and challenges of multimodal learning analytics. *British Journal of Educational Technology*, *51*(5), 1441–1449.

Cukurova, M., Luckin, R., & Kent, C. (2020). Impact of an artificial intelligence research frame on the perceived credibility of educational research evidence. *International Journal of Artificial Intelligence in Education*, *30*, 1–31.

Cukurova, M., Luckin, R., Millán, E., & Mavrikis, M. (2018). The NISPI framework: Analysing collaborative problem-solving from students' physical interactions. *Computers & Education*, *116*, 93–109.

Cukurova, M., Miao, X., & Brooker, R. (2023). Adoption of adaptive learning platforms in schools: Unveiling factors influencing teachers engagement. In *Artificial Intelligence in Education*. Springer. https://doi.org/10.1007/978-3-031-36272-9_13

Demszky, D., Yang, D., Yeager, D. S., Bryan, C. J., Clapper, M., Chandhok, S., Eichstaedt, J. C., Hecht, C., Jamieson, J., Johnson, M., Jones, M., Krettek-Cobb, D., Lai, L., JonesMitchell, N., Ong, D. C., Dweck, C. S., Gross, J. J., & Pennebaker, J. W. (2023). Using large language models in psychology. *Nature Reviews Psychology*, *2*(11), 688–701.

Doroudi, S. (2023). What happened to the interdisciplinary study of learning in humans and machines? *Journal of the Learning Sciences*, *32*(4–5), 663–681.

du Boulay, B. (2016). Recent meta-reviews and meta–analyses of AIED systems. *International Journal of Artificial Intelligence in Education*, *26*(1), 536–537.

Erickson, T., & Kellogg, W. A. (2000). Social translucence: An approach to designing systems that support social processes. *ACM Transactions on Computer-Human Interaction (TOCHI)*, *7*(1), 59–83.

Giannakos, M., & Cukurova, M. (2023). The role of learning theory in multimodal learning analytics. *British Journal of Educational Technology*, *54*(5), 1246–1267.

Hadwin, A. F., Järvelä, S., & Miller, M. (2011). Self-regulated, co-regulated, and socially shared regulation of learning. In D. H. Schunk & B. Zimmerman (Eds.), *Handbook of self-regulation of learning and performance* (Vol. *30*, pp. 65–84). Routledge.

Heffernan, N. T., & Heffernan, C. L. (2014). The ASSISTments ecosystem: Building a platform that brings scientists and teachers together for minimally invasive research on human learning and teaching. *International Journal of Artificial Intelligence in Education*, *24*, 470–497.

Johnson-Laird, P. N. (1983). *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness, No. 6*. Harvard University Press.

Kawamura, R., Shirai, S., Takemura, N., Alizadeh, M., Cukurova, M., Takemura, H., & Nagahara, H. (2022). Detecting drowsy learners at the wheel of e-learning platforms with multimodal learning analytics. *IEEE Access*, *9*, 115165–115174. https://doi.org/10.1109/ACCESS.2021.3104805

Kent, C., Chaudhry, M., Cukurova, M., Bashir, I., Pickard, H., Jenkins, C., du Boulay, B., & Luckin, R. (2021). Machine learning models and their development process as learning affordances for humans. In *International Conference of Artificial Intelligence in Education*. Springer.

Knight, S., & Buckingham Shum, S. (2017). Theory and learning analytics. In C. Lang, G. Siemens, A. Wise, & D. Gasevic (Eds.), *Handbook of learning analytics* (pp. 17–22). SOLAR.

Koedinger, K. R., McLaughlin, E. A., & Heffernan, N. T. (2010). A quasi-experimental evaluation of an on-line formative assessment and tutoring system. *Journal of Educational Computing Research*, *43*(4), 489–510.

Kulik, J. A., & Fletcher, J. D. (2016). Effectiveness of intelligent tutoring systems: A meta-analytic review. *Review of Educational Research*, *86*(1), 42–78.

Leiker, D., Gyllen, A. R., Eldesouky, I., & Cukurova, M. (2023). Generative AI for learning: investigating the potential of learning videos with synthetic virtual instructors. In *International conference on artificial intelligence in education* (pp. 523–529). Springer Nature Switzerland.

Lubars, B., & Tan, C. (2019). Ask not what AI can do, but what AI should do: Towards a framework of task delegability. *Advances in Neural Information Processing Systems*, *32*, 57–67.

Luckin, R., & Cukurova, M. (2019). Designing educational technologies in the age of AI: A learning sciences-driven approach. *British Journal of Educational Technology*, *50*(6), 2824–2838.

Ma, W., Adesope, O. O., Nesbit, J. C., & Liu, Q. (2014). Intelligent tutoring systems and learning outcomes: A meta-analysis. *Journal of Educational Psychology*, *106*(4), 901–918.

Martınez-Miranda, J., & Aldea, A. (2005). Emotions in human and artificial intelligence. *Computers in Human Behavior*, *21*(2), 323–341.

McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (2006). A proposal for the Dartmouth summer research project on artificial intelligence, August 31, 1955. *AI Magazine*, *27*(4), 12.

Mitrovic, A., & Ohlsson, S. (1999). Evaluation of a constraint-based tutor for a database language. *International Journal of Artificial Intelligence in Education*, *10*, 238–256.

Nazaretsky, T., Ariely, M., Cukurova, M., & Alexandron, G. (2022). Teachers' trust in AI-powered educational technology and a professional development program to improve it. *British Journal of Educational Technology*, *53*, 914–931. https://doi.org/10.1111/bjet.13232

Nazaretsky, T., Cukurova, M., Ariely, M., & Alexandron, G. (2021). Confirmation bias and trust: Human factors that influence teachers' attitudes towards AI-based educational technology. In *CEUR Workshop Proceedings*, Bozen-Bolzano, Italy (Vol. *3042*).

Nazaretsky, T., Mejia-Domenzain, P., Swamy, V., Frej, J., & Käser, K. (2024). AI or human? Evaluating student feedback perceptions in higher education. In *European Conference on Technology Enhanced Learning, EC-TEL 2024*. Springer International Publishing.

Ouyang, F., Xu, W., & Cukurova, M. (2023). An artificial intelligence-driven learning analytics method to examine the collaborative problem-solving process from the complex adaptive systems perspective. *International Journal of Computer-Supported Collaborative Learning*, *18*(1), 39–66.

Pane, J. F., Griffin, B. A., McCaffrey, D. F., & Karam, R. (2014). Effectiveness of cognitive tutor algebra I at scale. *Educational Evaluation and Policy Analysis*, *36*(2), 127–144.

Papert, S. (2005). You can't think about thinking without thinking about thinking about something. *Contemporary Issues in Technology and Teacher Education*, *5*(3), 366–367.

Piaget, J. (1970). *Genetic epistemology*. Columbia University Press.

Pozdniakov, S., Martinez-Maldonado, R., Tsai, Y. S., Cukurova, M., Bartindale, T., Chen, P., Marshall, H., Richardson, D., & Gasevic, D. (2022). The question-driven dashboard: How can we design analytics interfaces aligned to teachers' inquiry? In *LAK22: 12th International Learning Analytics and Knowledge Conference, Frankfurt*, Germany (pp. 175–185).

Rismanchina, S., & Doroudi, S. (2023). *Four interactions between AI and education: Broadening our perspective on what AI can offer education*. https://edarxiv.org/kps79

Ritter, S., Carlson, R., Sandbothe, M., & Fancsali, S. E. (2015). Carnegie learning's adaptive learning products. In *Proceedings of the 8th International Conference on Educational Data Mining*, Madrid, Spain.

Russell, S. (2019). *Human compatible: AI and the problem of control*. Penguin UK.

Russell, S. J., & Norvig, P. (2010). *Artificial intelligence: A modern approach*. Prentice Hall.

Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human–Computer Interaction*, *36*(6), 495–504.

Spikol, D., Ruffaldi, E., Dabisias, G., & Cukurova, M. (2018). Supervised machine learning in multimodal learning analytics for estimating success in project-based learning. *Journal of Computer Assisted Learning*, *34*(4), 366–377.

Suraworachet, W., Seon, J., & Cukurova, M. (2024). Predicting challenge moments from students' discourse: A comparison of large language models to other natural language processing approaches. In *Proceedings of the 14th Learning Analytics and Knowledge Conference*, Tokyo, Japan (pp. 473–485).

Suraworachet, W., Villa-Torrano, C., Zhou, Q., Asensio-Pérez, J. I., Dimitriadis, Y., & Cukurova, M. (2021). Examining the relationship between reflective writing behaviour and self-regulated learning competence: A time-series analysis. In *Technology-Enhanced Learning for a Free, Safe, and Sustainable World: 16th European Conference on Technology Enhanced Learning, EC-TEL 2021, Bolzano, Italy, Proceedings 16* (pp. 163–177). Springer International Publishing.

Suraworachet, W., Zhou, Q., & Cukurova, M. (2023). Impact of combining human and analytics feedback on students' engagement with, and performance in, reflective writing tasks. *International Journal of Educational Technology in Higher Education*, *20*(1), 1–24.

Touretzky, D., Gardner-McCune, C., Martin, F., & Seehorn, D. (2019, July). Envisioning AI for K-12: What should every child know about AI? *Proceedings of the AAAI Conference on Artificial Intelligence*, *33*(1), 9795–9799.

Van Lehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, *46*(4), 197–221.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems*, Long Beach, CA (Vol. 30).

Von Ahn, L. (2013). Duolingo: Learn a language for free while helping to translate the web. In *Proceedings of the 2013 International Conference on Intelligent User Interfaces*, Santa Monica, CA (pp. 1–2).

Weatherholtz, K., Grimaldi, P. J., & Hill, K. M. (2022). *Use of MAP accelerator associated with better-than-projected gains in MAP growth scores*. Khan Academy.

Zhou, Q., Suraworachet, W., & Cukurova, M. (2024). Detecting non-verbal speech and gaze behaviours with multimodal data and computer vision to interpret effective collaborative learning interactions. *Education and Information Technologies*, *29*(1), 1071–1098.

Zhou, Q., Suraworachet, W., Pozdniakov, S., Martinez-Maldonado, R., Bartindale, T., Chen, P., Richardson, D., & Cukurova, M. (2021). Investigating students' experiences with collaboration analytics for remote group meetings. In *Artificial Intelligence in Education: 22nd International Conference, AIED 2021, Utrecht, The Netherlands, June 14–18, 2021, Proceedings, Part I 22* (pp. 472–485). Springer International Publishing.