Historical threads, missing links, and future directions in AI in education

Ben Williamson & Rebecca Eynon (2020) Historical threads, missing links, and future directions in AI in education, Learning, Media and Technology, 45:3, 223-235,

The final copy of this editorial is open access and available here https://www.tandfonline.com/doi/full/10.1080/17439884.2020.1798995

Artificial intelligence has become a routine presence in everyday life. Accessing information over the web, consuming news and entertainment, the performance of financial markets, the ways surveillance systems identify individuals, how drivers and pedestrians navigate, and how citizens receive welfare payments are among myriad examples of how AI has penetrated into human lives, social institutions, cultural practices, and political and economic processes. The effects of the algorithmic techniques employed to enable AI are far-reaching and have inspired considerable epochal hype and hope, as well as dystopian dread, although they remain largely opaque and weakly understood outside of the social networks of technical experts (Rieder 2020). The profound social and ethical implications of AI, however, are becoming increasingly apparent and the objects of significant critical attention. AI is at the centre of controversies concerning, for example, automation in workplaces and public services; algorithmic forms of bias and discrimination; automated reproduction of inequalities and disadvantage; regimes of data-centred surveillance and algorithmic profiling; disregard of data protections and privacy; political and commercial micro targeting; and the power of technology corporations to control and shape all sectors and spaces they penetrate, from whole cities and citizen populations to specific collectives, individuals or even human bodies (Whittaker et al 2018). Numerous ethical frameworks and professional codes of conduct have been developed to attempt to mitigate the potential dangers and risks of AI in society, though important debates persist about their concrete effects on companies or the way such frameworks and codes may serve to protect commercial interests (Greene, Hoffman & Stark 2019).

The current instantiation of AI on the web, on smartphones, in social media, and in spaces via interconnected objects and sensor networks has a much longer history than some recent epochal claims would suggest. Histories of AI stretch back at least as far as the birth of computer science and cybernetics in the 1940s. The term 'artificial intelligence' itself was coined as part of a project and workshop at Dartmouth College in the mid-1950s. From the 1960s to the 90s, punctuated by periods of 'AI winter', AI research and development focused first on encoding principles of human reasoning to simulate human intelligence, and then on 'expert systems' that emulated the procedural decision-making processes of experts based on defined knowledge bases. After 2010 AI gradually returned under a new paradigm, not as simulated human intelligences or programmable expert systems but as data-processing systems that can learn and make predictions from classifying and correlating huge quantities of 'big data'. Computational processes including data analytics, machine learning, neural networks, deep learning and reinforcement learning underpin most contemporary forms of AI. AI is, perhaps, just a new catch-all name for a range of statistical, mathematical, computational and data scientific practices and developments that each have their own complex and intertwined genealogies, but it also signifies a particular unique nexus of these historical strands (Schmidhuber 2019, 2020). Modern AI is not focused on creating computational 'superintelligences' ('strong AI') but ideally on developing machines that can learn from their own experience, adapt to their contexts and uses, improve their own functioning, craft their own rules, construct new algorithms, make predictions, and carry out automated tasks without requiring control or oversight by human operatives (Alpaydin 2016; Mackenzie 2017).

Interests in the application of AI to education (AIed) have a long history too - with a range of social and ethical implications that this special issue is intended to identify and examine. Our aim in this editorial introduction to the special issue 'Critical perspectives on AI in education' is to provide some historical perspective to the collection of papers and to current hyperbole about AIed and associated ideas about adaptive systems, pedagogic agents, personalized learning, intelligent tutors, and automated governance. As with the AI field more generally, we do not see contemporary AIed as the result of a simple linear history; instead, we understand the contemporary AIed moment as the result of a set of convergences that, among many genealogical threads, include (1) several decades of AIed research and development in academic centres and labs, (2) the growth of the commercial education technology (edtech) industry, (3) the influence of global technology corporations on education, and (4) the emergence of data-driven policy and governance. In this editorial, then, we want to historicize AI in education not by reproducing a linear timeline but by tracing out some of the genealogical threads and convergences that have led to the contemporary fascination with applying AI to education. As a method,

> Genealogical analysis traces how contemporary practices and institutions emerged out of specific struggles, conflicts, alliances, and exercises of power…. [I]ts intent is to problematize the present by revealing the power relations upon which it depends and the contingent processes that have brought it into being. (Garland, 2014, 372)

Genealogies seek to trace out the historical processes, conditions, conflicts, bifurcations and con- and disjunctures out of which contemporary practices emerged. Applied to a field such as AIed, these contingent genealogical threads and twists include disciplinary conflicts and encounters, technological developments, funding schemes, methodological advances and sectoral encounters between academic research and commercial imperatives. We cannot in this brief editorial produce a full genealogy of AIed - though we think that is a necessary future task - but instead highlight some particular historical trajectories that help situate the present collection. Following that, we also identify some key missing elements in existing research on AIed, and some future directions for further work.

AIed research

'AI in education' has existed as a coherent academic research field since at least the 1980s, as marked by the first publication of the *International Journal of Artificial Intelligence in Education* in 1989 and the formation of the International AI in Education Society (IAIED) in 1993, although it was preceded by development of Intelligent Tutoring Systems Computer Assisted Instruction systems in the 1960s and 70s (Alkhatlan & Kalita 2018; Selwyn 2019). As a field AIed has developed along two complementary strands of activity: the development of AI-based tools for classrooms, and the use of AI to understand, measure and improve learning (Holmes, Bialik & Fadel 2019). Particularly as a field of research on learning, AIed is closely related to the learning sciences and cognitive science - the 'cognition, technology and education nexus' as Pea (2016, 51) has termed this disciplinary confluence - as well as to developments in learning analytics and educational data mining that have unfolded over several decades (Buckingham Shum & Luckin 2019).

In the second issue of the *International Journal of Artificial Intelligence in Education*, published in 1989, Schank and Edelson (1989: 3) claimed that AI is 'intimately bound up with education':

> AI concerns itself with getting machines to read, to reason, to express themselves, to make generalizations, and to learn. These issues are the stuff of education after all. … AI people are in a unique position to improve education…. [W]e can couple our expertise in computer technology with our theories of learning and understanding to build computer-based instructional systems that will have a positive impact on education. (Schank & Edelson, 1989, 3-4)

For Schank and Edelson, AI—and the 'AI people' who create it—would not only have a positive impact on education as it was formally organized in terms of curriculum and instructional design, but would demand far-reaching changes to the ways that teaching and learning were understood and practised. In fact, the authors were writing from the newly established Institute for Learning Sciences at Northwestern University, a centre set up outside of the discipline of educational research with funding from a Chicago-based corporate consulting firm and the Advanced Research Projects Agency (ARPA) of the US Department of Defense.

A few years later, in a review of applications of AI in education completed for RAND, McArthur, Lewis and Bishay (1995, 42) argued that AI applications such as Intelligent Tutoring Systems (ITS) 'can significantly improve the speed and quality of students' learning'. Like Schank and Edelson before them, however, they also saw such successes merely as indicators of potentially profound transformations:

> The technologies that make it possible to automate traditional methods of teaching and learning are also helping to create new methods … and to redefine valued educational goals and learning outcomes. (McArthur et al, 1995, 42)

Realizing this transformative ambition, they suggested, would require bringing 'high-tech companies into better cooperation with educational technology research and classroom practice' (43). The future prospects for AI in education, they further argued, would be to 'challenge and even threaten' existing teaching and learning practices, replacing them with methods such as computer-based 'individualized tutoring' and 'inquiry- or project-based learning' in ways that might '*transform* schools and classrooms, not improve them in any simple sense', and 'offer new goals and practices for teaching and learning' (72, original italic).

Twenty-five years later, many of the animating aims and transformative commitments of the 'AI people' working in the field of education remain distinctly similar, even if the underlying computational logics of AI have shifted from programmable expert systems to big data analytics, machine learning, neural networks and deep learning (Knox, Williamson & Bayne 2020). Intelligent tutoring systems have continued to evolve in AIed research labs and commercial settings for 40 years (du Boulay 2019). From around 2005, new research fields of educational data mining and learning analytics began to emerge, focused on the analysis of 'big data' in education, and on the development of new professional positions for 'education data scientists' and other analytics experts (Fischer et al 2020). Although learning analytics and mining big educational data - whether at the institutional or individual level - are not entirely synonymous with AI, they are increasingly genealogically intertwined and related, as signified by the publication of a double special issue of the *British Journal of Educational Technology* dedicated to learning analytics and AIed as a single research topic in 2019. The field is also mutating and evolving, with new positions opening up for so-called 'learning engineers' who possess hybrid forms of expertise crisscrossing the computing, data and learning sciences (Williamson 2020).

As Luckin et al (2016, 12) envisaged in an influential report on the prospects for AIed, a 'marketplace' of thousands of AI components will eventually combine to 'enable system-level data collation and analysis that help us learn much more about learning itself and how to improve it'. With this marketplace of AI components in place, Luckin and coauthors envisioned both new AIed tools for classrooms and enhanced analytical capacity for the measurement of learning and teaching at multiple levels:

> Once we put the tools of AIEd in place ... we will have new and powerful ways to measure system level achievement. … AIEd will be able to provide analysis about teaching and learning at every level, whether that is a particular subject, class, college, district, or country. This will mean that evidence about country performance will be available from AIEd analysis, calling into question the need for international testing. (Luckin et al 2016, 48)

As the president of the International AI in Education Society, Luckin has rapidly translated AIed R&D and advocacy into mainstream, media-friendly predictions which posit not only the transformative effects of AI on education but on human intelligence and learning with a focus on augmentation:

> The real power that AI brings to education is connecting our learning intelligently to make us smarter in the way we understand ourselves, the world and how we teach and learn. For the first time we will be able to extend, develop and measure the complexity of human intelligence — an intellect that is more sophisticated than any AI. This will revolutionise the way we think about human intelligence. (Luckin 2020)

From this brief overview, we see the contested nature of AIed, with a complex history of methodologies and disciplines, and to varied claims of how AI will enable learning scientists to develop 'transformative' or 'revolutionary' understandings of human cognition, learning and intelligence.

In this issue, Perrotta and Selwyn (2020) highlight how learning sciences research performed with machine learning, an applied AI method, superimposes algorithmic complexity on reductionist and contested understandings of human learning. They note that applied AI techniques of 'automated knowledge discovery' such as pattern recognition and correlational analysis are based on a mechanical, inductivist epistemology that assumes all patterns are interpretable in the same standardized ways across all cultures and contexts. They highlight instead how AI-generated patterns reflect the specific situations from which they are gathered and are imprinted with professional, disciplinary and economic contingencies, as they exemplify in their case study of a controversy over the use of predictive machine learning in personalized learning and adaptive learning development. The article is a salient reminder that, despite its veneer of objectivity, AIed is infused with politics, embodies particular sets of values and entails new distributions of power in educational research to data science experts with particular ontological and epistemological commitments. AIed is also specifically located in a longer history of knowledge production, theory generation, and the development of epistemic expertise in the learning sciences and educational data science that is related to, but not reducible to, commercial and economic interests.

AIed as commercial edtech

AI in education is not just a pursuit of educational data scientists and learning science specialists. It is also a major commercial concern of educational technology (edtech) companies, which have sought to bring multiple forms of AI-based products to market, and of powerful philanthropic and investment actors that support AIed startups as part of the development of adaptive personalized learning software enacted by machine learning (Selwyn 2019). Rising interests in 'big data' in education, educational data mining and learning analytics from the mid-2000s were accompanied by powerful framing discourses of adaptive systems and personalized learning

(Williamson 2017). These same discourses of personalized learning are now used to justify, promote and market AI-based solutions in education (Bulger 2016; Boninger, Molnar & Saldana 2020).

The global education business Pearson in particular has supported AI in education enthusiastically for more than a decade, first through support for big data and learning analytics, and later explicitly by advocating and producing AI-based intelligent tutors and other 'adaptive and personalized learning' applications. Pearson's proposed vision of AIEd includes the development of Intelligent Tutoring Systems (ITS) which 'use AI techniques to simulate one-to-one human tutoring, delivering learning activities best matched to a learner's cognitive needs and providing targeted and timely feedback, all without an individual teacher having to be present.' It also promises intelligent support for collaborative working—such as AI agents that can integrate into teamwork—and intelligent virtual reality environments that simulate authentic contexts for learning tasks. Its vision is of teachers supported by their own AIEd teaching assistants and AIEd-led professional development. As its webpage on 'The future of education' states:

> By combining AI with the learning sciences – psychology, neuroscience, linguistics, sociology and anthropology – we gain an understanding of what and how people learn. With AI, how people learn will start to become very different. … AI can adapt to a person's learning patterns. This intelligent and personalized experience can actually help people become better at learning, the most important skill for the new economy. (https://www.pearson.com/news-and-research/the-future-of-education/artificial-intelligence.html)

As the realization of its approach to AIed, Pearson created and launched AIDA, a smartphone based AIed app, as an adaptive tutor that offers personalized, adaptive responses to students (https://www.pearson.com/en-us/learner/products-and-services/learning-apps-development/aida.html). Pearson's launch of AIDA is significant for several reasons. First, it signifies a certain mainstreaming of AIed, translating the long history of research and development in intelligent tutoring systems and pedagogic agents into a marketable product. Second, AIDA is targeted primarily at students as consumers who will pay a subscription for automated support for their studies, part of an emerging trend in 'consumer edtech'. Third and relatedly, it represents the emergence of a new 'shadow education market' of private supplementary tutoring in which the private tutor is an automated smartphone app or a personal robot assistant. Whether AIDA will prove to be a long-term success for Pearson remains to be seen, but we can see it as a material instantiation of commercial aspirations to embed automation in education as a potentially profitable market niche.

The example of Pearson exemplifies how commercial edu-businesses and edtech companies have been pursuing AIed products for many years, and seeking to grow market interest in their products. However, the expansion of AIed has not taken linear pathways. In this issue, Knox (2020) highlights the significant growth of commercial AI in education products in the specific context of China. Catalyzed by massive state investment in AI, new Chinese AIed companies are attracting enormous venture capital investments too. In spring 2020, as the spread of Covid-19 led to school closures, the Beijing-based Yuanfudao edtech company received the largest venture capital investment ever recorded for a startup in the edtech sector for its AI-based homework and tutoring platform, with its US$1bn investment taking its total value to an estimated $7.8bn (Dai 2020). China's adoption of AI in education is the result of large-scale venture capital investment and increasing parental willingness among more wealthy families to pay for private tutoring and supplemental education services and products, but is also driven by public-private partnership arrangements and the strong support of the state for private sector technologies (Knox 2020). It demonstrates the inseparability of the state, the private sector, and consumer markets in AIed development, roll-out and uptake.

However, the huge expansion of AIed in China has also raised significant controversy and concern:

> If schools are capable of tracking every keystroke, knowledge point and facial twitch, they are effectively furnishing either a technology company or the Chinese state with an eternal ledger of every step of a child's development. This is potentially problematic because, whereas the human teacher assumes change, AI assumes continuation. … [A]n intelligent tutoring system could not only store that information and tailor a personalised pathway for the student in the first grade, it may extrapolate that information many years later, when the student is in high school. (Liu 2020, n.p.)

Although the adoption of AIed in China has generated significant interest and concern, its approach to investment in AI in education is also mirrored by the ambitions of the edtech sector and by venture philanthropic supporters in other countries, especially in the US and India (Chamuah & Ghildiyal 2020). In the US context, major supporters and funders of AI-based educational technologies include the Bill and Melinda Gates Foundation, Schmidt Futures, and the Chan Zuckerberg Initiative, the funding and investment vehicles, respectively, of Microsoft founder Bill Gates, ex-Google chair Eric Schmidt, and Facebook founder Mark Zuckerberg. These organizations make significant claims for the effectiveness of AI and personalized learning in raising student achievement, enabling students to develop 'mastery' over knowledge domains, and in reducing inequalities by distributing support to underserved students. They inject millions of dollars into selected organizations as a way of driving up adoption of personalized learning platforms, although their claims that AI improves outcomes or reduces inequalities remain highly contested (Boninger, Molnar & Saldana 2020).

Dixon-Roman, Nichols and Nyame-Mensah (2020), in this issue, draw critical attention to the 'racializing forces' of commercial AI in education, through a detailed case study of one commercial AIed application. They argue that AI applications in education may 'inherit sociopolitical relations' and reproduce racializing processes and educational inequities through the 'cloak of algorithmic objectivity'. Their argument reflects the ways that AI and data systems are implicated in race-based profiling and 'discriminatory designs' that reinforce and normalize racial hierarchies and may enforce 'racial correction' (Benjamin 2016, 148). Critical edtech research, for example, has highlighted the role of 'digital redlining' in excluding certain groups from access to knowledge and information based on gender, class and race:

> Digital redlining arises out of policies that regulate and track students' engagement with information technology. … Digital redlining is not a renaming of the digital divide. It is a different thing, a set of education policies, investment decisions, and IT practices that actively create and maintain class boundaries through strictures that discriminate against specific groups. … Digital redlining takes place when policymakers and designers make conscious choices about what kinds of tech and education are 'good enough' for certain populations. (Gilliard & Culik 2016, n.p.)

Such research draws urgent attention to the ways that the architectures of technologies are involved in new forms of exclusion and discrimination. Data-intensive technologies such as AI in education may be additionally discriminatory based on 'machine bias' being encoded in the datasets used to train algorithms. These brief examples highlight how AIed needs to be understood in the historical context of the growth of commercial edtech and the support it has gathered from actors as diverse as state ministries, high-tech venture philanthropies and investment vehicles. AIed also needs to be understood in the history of discriminatory practices, such as digital redlining or the reproduction of racial hierarchies, that may be reinforced rather than ameliorated by the design of commercial edtech.

AI infrastructures

The third historical thread in the development of AIed concerns the ways that transnational technology corporations have sought to embed AI technologies in education. AI has become a major industry concern, with many of the major global tech businesses offering 'AI-as-a-service', such as Google's TensorFlow, Microsoft Azure, IBM Watson, and Amazon's machine learning on AWS. All of these developments have their own long histories of product development which are, in turn, embedded in organizational cultures, hiring practices, and the cultivation of in-house expertise. They reflect the belief, common in the technology and data analytics industries, that big data and AI can solve problems as diverse as business efficiency, scientific discovery, urban management, and educational improvement. As Beer (2019, 30) notes, 'out-of-the-box Artificial Intelligence tools' are presented as 'ready-made and thinking technologies that carry the burden of technique, know-how and method,' and 'locate value in a way that the human alone is unable to do'. Data analytics and AI are understood by this industry as providing the only sensible way of making decisions or solving problems. Part of the history of 'AI-as-a-service' is related to the emergence of 'platform capitalism' as a dominant mode of value creation and capital accumulation in the digital economy (Srnicek 2017). Out-of-the-box AI positions machine intelligence as a platform that can be deployed across a range of sectors, enabling other organizations to deploy ready-made thinking technologies to aid decision-making and improve their service to customers or users. But such arrangements also mean that such organizations are dependent upon global commercial AI infrastructures and on the particular forms of automated intelligence they enable.

Within education, AI-as-a-service applications can be detected in the many ways that global technology firms have begun to provide back-end infrastructure services for other educational institutions or edtech companies. Pearson, for example, previously partnered with IBM to embed Watson APIs in courseware products as part of its early ambitions to create AIed applications. Google's G Suite is familiar as a front-end set of apps for use in classrooms, but it also depends on Google's data-extractive infrastructure and APIs for integrating third-party products. Microsoft has also recently begun to promote its Power Platform for education. The Power Platform is an integrated platform enabling organizations to integrate various applications and data sources into a single data model for analysis, and provides common models and templates for organizations to create their own services and applications, all supported by Microsoft's Azure cloud computing service. It includes Power Apps enabling educators to build their own low-code apps, and Power Virtual Agents to 'enable institutions to easily create and maintain intelligent chatbots', such as 'an intelligent Question Bot that gets smarter and is capable of supplying answers on its own to students, which allows for greater student independence and supports personalized learning' (Microsoft Education Team 2020). Power Platform is in this sense an infrastructural technology to undergird and support educators' and institutions' implementation of AI. Likewise, Amazon has begun promoting a variety of educational services and technologies related to its cloud infrastructure Amazon Web Services (AWS). For example, it claims, 'Using the AWS Cloud, schools and districts can get a comprehensive picture of student performance by connecting products and services so they seamlessly share data across platforms' (AWS 2020a). It also strongly promotes its 'Machine Learning for Education' services to 'identify at-risk students and target interventions', 'improve teacher efficiency and impact with personalised content and AI-enabled teaching assistants and tutors', and 'improve efficiency of assessments and grading' (AWS 2020b). These examples indicate how huge global technology companies are in competition for structural dominance over the digital infrastructures of education, providing cloud systems that are capable of integrating various platforms and interoperability programs for enabling seamless and frictionless data flow, aggregation and analysis.

In these ways, some of the world's most powerful technology companies are seeking to insert themselves into education as back-end infrastructure providers, as well as front-end suppliers or vendors of specific educational services such as G Suite or Microsoft education products. McStay (2020) in this issue, for example, highlights the role of the 'affective computing' company Affectiva as a supplier of 'emotional AI' software that can be embedded in other educational technologies. Affectiva originally emerged from the MIT Media Lab's affective computing centre, and from collaborative research on 'affect-sensitive autotutor' technologies conducted within the emerging field of emotional learning analytics in the early 2000s (D'Mello, Picard & Graesser 2007). Affectiva is now one of the world's leading emotional AI companies, having amassed a databank of billions of images of human faces that can then be used as the basis for real-time emotion expression analysis. But as McStay (2020) shows, in its current role, rather than taking a centre-stage role in edtech, Affectiva prefers a back-stage role, enabling it to scale out its software and its mass collection of facial images through other products and brands. Affectiva inserts itself into education as emotional AI infrastructure.

These brief examples highlight the importance of approaching AI in education as at least partly the genealogical result of historically situated technology sector efforts to extend AI infrastructures into a wide variety of other industries and practices. Importantly, they demonstrate how education is increasingly fused to transnational private technology companies and the business logics of platform capitalism. They show how AI enters education through mundane back-end AI-as-a-service plug-ins, rather than in the more spectacular guise of automated pedagogic agents or tutoring systems. This genealogical development of AI in education is quite distinct from the efforts of either academic AIed specialists or the dedicated edtech sector. Its significance as yet remains unclear, but it seems likely that schools and universities will increasingly rely on these infrastructural arrangements to handle the flows and analysis of data required for institutional performance management, or that are demanded for processes of governance and accountability.

AI policy and governance

AI has also become an emerging concern in education policy and governance, as part of a long history of governance by numbers (Piattoeva & Boden 2020). In particular, from the 1990s onwards, concerns with educational accountability and evidence-based policy drove the implementation of data systems that could be used to record progress towards performance targets and improvement goals (Lingard, Martino & Rezai-Rashti, 2013). The use of accountability data as a key source of education policy and governance was enabled by large-scale information infrastructures for collecting and processing the data (Anagnostopoulos, Rutledge & Jacobsen, 2013), with interest growing over the subsequent decades in analytics packages and data dashboards for analysing, interpreting and displaying up-to-date data:

> shrinking fiscal resources and the expansion of a global competitive education market have fueled this increasing pressure for educational accountability. The offshoot of these economic drivers has been the development in the education sector of standardized scalable, real-time indicators of teaching and learning outcomes. (Lockyer, Heathcote & Dawson, 2013, 1439)

The nascent fields of Educational Data Mining and Learning Analytics, from around 2005, opened up new opportunities for systematic quantitative analysis in the field of education

(Agasisti & Bowers 2017). Edu-businesses such as Pearson began promoting the idea that these forms of data mining and analytics could be used to transform policy processes, by focusing on the real-time performance of institutions and even individuals rather than relying on temporally periodic assessment events as insights for policymaking (Hill & Barber 2014). In this context, processes of educational governance have become increasingly reliant upon data stored in student information systems and learning management systems, and by the proliferation of analytics packages for generating predictive insights into institutional and individual-level performance. Emerging AI-enabled data infrastructures that can perform intelligent analytics of vast quantities of educational information represent the next instantiation of governance by numbers.

In situated practice, new forms of digital, data-led governance have emerged in governmental departments of education and associated agencies. For example, the UK Behavioural Insights Team ('Nudge Unit') collaborated with the schools inspectorate Ofsted (Office for Standards in Education) to develop a 'school inspection algorithm' that could automatically evaluate a school's data records and then highlight areas of concern for the embodied human inspector to observe and report on. Moreover, policy-influencing organizations including the OECD have begun to firmly advocate for the use of AI to measure and improve learning. An OECD report, for example, advocates 'the use of Big Data, Artificial Intelligence algorithms, education data mining and learning analytics … to improve learning and education' through transmitting scientific evidence of learning into 'real-world education practice and policy' (Kuhl et al 2019, 13-14). As these examples indicate, governance by numbers in education has mutated into digital, data-led governance through algorithms, justified increasingly explicitly through discourses of scientific, AI-based policymaking.

As Webb, Gulson and Sellar (2020) argue in this issue, data-led forms of governance have now begun mutating into 'anticipatory' forms of AI-enhanced governance. AI and learning analytics platforms, they suggest, create new conceptions of temporality in education, where past actions captured in huge comparative datasets can be used to make predictions at the level of the individual student in 'real time'. Such platforms are designed to identify students 'at-risk' and intervene computationally in ways that ideally ameliorate deleterious learning outcomes. This involves the modelling of probable futures from continuously generated and analyzed datastreams, providing synchronous diagnoses and interventions which model multiple temporal trajectories in order to anticipate student futures.

Digital regulation is essential consideration when exploring the complex array of factors that are shaping these data-led forms of governance. In their paper for this issue, Bernedt, Littlejohn and Blakemore's (2020) examine the benefits and risks of AI in education in relation to fundamental human rights. They demonstrate that current laws, such as the European Union's data protection law (GDPR) introduced in 2018 that introduced legally enforceable standards relating to privacy and data protection, are insufficient. The authors argue that human rights need to be protected at a transnational level, as national approaches are unlikely to work particularly given the global reach of many AI systems. The importance of human rights, and the rights of the child is also an important theme in McStay's (2020) contribution that questions the extent to which emotional AI and EdTech are serving the public good. Such arguments chime with the discussion outside of education about the needs for children's rights to be firmly integrated into the agendas of global debates about ethics and data science (Livingstone et al., 2015; Berman and Albright, 2017; Lupton and Williamson, 2018). Yet despite this interest and investment, the necessary legal frameworks are not fit for purpose.

Missing links

Our intention with the genealogical notes in the previous sections was to illustrate how we understand AIed as a complex set of relations and confluences among many factors and historical pathways. In these final parts of this introduction we would like to turn our attention to the potential for future work in this emerging area, focusing first on the missing links. We suggest three areas of focus: ethnographies of AI in use, links with the philosophy of education, and more meaningful interactions with other academic communities working in AIed.

As can be seen from the discussion above, this special issue covers an array of rich perspectives to promote the critical study of AIed. The lens of political economy (Knox 2020), science and technology studies (Perrotta and Selwyn 2020), socio-legal studies (McStay 2020; Berendt, Littlejohn, Blakemore 2020), education governance and time (Webb, Sellar and Gulson 2020), and new materialist and Black feminist thought (Dixon-Román, Nichols and Nyame-Mensah 2020) all feature. Many of the papers also effectively move across micro and macro levels of analysis and across disciplines to bring together their argument. Such a multi-disciplinary mix of critical perspectives is central to ensuring the development of this area of study.

Like the papers in this special issue, much of the critical work on AIed to date is based on rich analyses of AI products, mapping power structures, raising questions of ethics and making visible systems of infrastructures, the role of the commercial sector and new forms of digital governance. All of these are crucial, yet there is far less data from a critical perspective of what happens when these systems are used on a daily basis in varied educational contexts. We know very little, for example, about how learners and teachers really use AI systems, and how AI is embedding (or not) into the everyday workings of schools, colleges and other sites of education and learning. As many of the authors in this special issue have highlighted, there are likely key risks of such systems. However, we need to better understand how these happen in practice. We now have a growing and theoretically rich account of the risks and possibilities of AIed but fewer accounts (beyond more instrumental studies of effectiveness) of what really happens in the classroom.

More ethnographic case-study work of AI in different educational contexts would add an additional perspective to enrich the important themes emerging from work in this area. Such an approach fits well with what makes for critical study of education and technology, where studies of the 'state of the actual' and investigations into how the use of AIed plays out in varied contexts are central (Selwyn 2010).

A common call across the papers in this special issue is that education is a special part of society that requires specific understanding and attention. Given the uniqueness of education, it is clear that while authors exploring AIed can certainly learn from and inform work in other spheres of social life (e.g. policing, health, employment) it is also important to work with a consideration of the purposes and values of education and make this central to the debate. One example of this raised by many of the contributions to this special issue centre around questions of power: the power of systems to reinforce or exacerbate bias and discrimination, the powers of the commercial sector, and the absence of power for students and teachers. These power dynamics combine to lead to a context where students (and indeed teachers) are not able to give or withhold) meaningful consent to use AI products as part of their education. This is not the same as in other areas of social life, where individuals have some (although often admittedly limited) power, in terms of the decision whether or not, or how to use, a certain system (Berendt, Littlejohn, Blakemore 2020; Zeide 2017).

Developing stronger connections with arguments around the philosophy of education could be a useful way to articulate and promote the uniqueness of education as a site for critical study. Such an approach would bring attention back to the purposes and values we have and aspire to for our education systems, something that often gets lost in more common instrumental discussions of learning effectiveness and efficiency that are very characteristic in debates about AIed (Biesta 2015; 2019).

As noted in the early parts of this introduction, there are varied communities and stakeholders engaged in this domain, each with their own ways of conceptualising AIed. Within this complex field it is important to make a distinction (yet also note the complex relationship between) academics from the learning and cognitive sciences and those who work in the commercial sector. Perspectives from learning and cognitive scientists are retold by the authors in this special issue – but there is no one from this community who represents themselves. Although highly challenging, it would be interesting to promote more dialogue across the different academic communities focused on AIed. Elsewhere there have been some moves in this direction (Buckingham Shum and Luckin, 2019; Selwyn, 2019) and it would be valuable for these to continue.

More meaningful debate across these communities may lead to some interesting ways forward in future work, and are important in developing meaningful understandings of the motives, understandings and objectives of researchers working in related fields of education. It may also provide stronger arguments and ways to address the strong commercial presence in AIed in all its forms.

Future directions

The papers in this special issue have highlighted how the current and future imaginaries of AIed are problematic in multiple ways. To change it requires the engagement and interaction across multiple relevant social groups (Bijker 2010). No one trajectory is inevitable or fixed, and the emerging dominant ways of thinking about AIed can be reconfigured to support the more democratic vision of the future of education envisioned by many papers in this special issue (Eynon and Young, 2020).

We would suggest that one important way forward could be the use of more participative approaches, enabling the development of insight and changes in practice simultaneously. A long standing and well established approach in critical education, human computer interaction, and related fields, it is attracting growing attention by critical scholars of AI. For example, D'Ignazio and Klein draw on feminist theory to argue for the importance of multiple actors, particularly those who are most marginalised, to be involved in data work to facilitate positive change. They argue that enabling everyone to be part of data projects, particularly those who are likely to be most impacted by such systems, could encourage an array of responses from using data to make inequalities visible, to developing a consensus around local projects or to support community storytelling. As the authors note, data scientists alone are unlikely to have the same impact (D'Ignazio and Klein, 2019). Costanza-Chock calls for Design Justice, where marginalised communities lead the process of designing new technologies to facilitate positive social change (Costanza-Chock, 2020). In a similar vein, scholars of learning analytics and educational data mining have called for an "ethics by design" approach in developing and implementing data centric policies and practices in Education (Harel Ben Shahar, 2017; Gray and Boling, 2016).

Developing a broader academic community who could both use and critique the development of AI is also an important aspect for future work. Many of the authors in this special issue have made visible how AI systems work and in doing so the choices, assumptions and flaws that underpin these systems. Using AI as a research tool could add to that understanding, and also provide new ways to engage a range of stakeholders in debates about the future of AIed. AI could be used to assist in mapping the hidden networks of actors promoting AIed over time, it could aid in the analysis of the discourses around AIed and make visible structural inequalities in education systems. Such approaches may add to understanding, offer new ways to demonstrate the issues and further understanding and informed critique of what such methods can, and cannot offer the critical edtech community (Williamson, Potter & Eynon, 2019).

Expanding methodologies and approaches to research in these ways may be an important strategy to protect against the current pattern of events where, despite the long and complex history of AIed we have shown above it seems an area very prone to the challenges of hype cycles, short termism and forgotten histories. These ahistorical approaches are a continued problem in many studies of edtech, and enable those with limited expertise but significant (often commercial) power to take centre stage in shaping and investing in educational futures. This special issue showcases a collection of work that takes an important step towards resisting such trends through the use and development of theory and an awareness of history. An important additional trajectory of research is for academics to develop ways to more directly intervene in shaping the future imaginaries of AIed.

References

Alkhatlan, A. & Kalita, J. 2018. Intelligent tutoring systems: A comprehensive historical survey with recent developments." arXiv: http://arxiv.org/abs/1812.09628

Agasisti, T. and Bowers, A.J., 2017. Data analytics and decision making in education: towards the educational data scientist as a key actor in schools and higher education institutions. In Handbook of contemporary education economics. Edward Elgar Publishing.

Alpaydin, E., 2016. *Machine Learning: The New AI*. Cambridge, MA: MIT Press.

Anagnostopoulos, D., S.Rutledge, and R.Jacobsen, eds. 2013. *The Infrastructure of Accountability: Data Use and the Transformation of American Education*. Cambridge, MA: Harvard Education Press.

AWS (Amazon Web Services). 2020a. K12 and Primary Education. Amazon. https://aws.amazon.com/education/K12-primary-ed/

AWS. 2020b. Machine learning in education. Amazon. https://aws.amazon.com/education/ml-in-education/

Beer, D. 2019. *The Data Gaze: Capitalism, power and perception*. London: Sage.

Benjamin, R. 2016. Catching Our Breath: Critical Race STS and the Carceral Imagination. *Engaging Science, Technology, and Society*, 2, 145-156.

Berman, G. and Albright, K. (2017). Children and the Data Cycle: Rights and Ethics in a Big Data World, *Innocenti Working Papers* no. 2017-05, UNICEF Office of Research - Innocenti, Florence

Biesta, G. 2015. Good Education in an Age of Measurement: Ethics, Politics, Democracy. London: Routledge.

Biesta, G. 2019. What Kind of Society Does the School Need? Redefining the Democratic Work of Education in Impatient Times. Studies in the Philosophy of Education, 38 (6), 657-668.

Bijker, W.E. 2010. How is technology made? – That is the question! Cambridge Journal of Economics, 34 (1), 63–76.

Boninger, F., Molnar, A., & Saldaña, C. 2020. Big Claims, Little Evidence, Lots of Money: The Reality Behind the Summit Learning Program and the Push to Adopt Digital Personalized Learning Platforms. Boulder, CO: National Education Policy Center. https://nepc.colorado.edu/publication/summit-2020

Buckingham Shum, S. & Luckin, R. 2019. Learning analytics and AI: Politics, pedagogy and practices. *British Journal of Educational Technology*, 50(6), 2785-2793.

Bulger, M., 2016. Personalized learning: The conversations we're not having. *Data and Society*, *22*(1).

Chamuah, A. & Ghildiyal, H. 2020. AI and education in India. Tandem Research, 16 April: https://tandemresearch.org/publications/ai-and-education-in-india

Costanza-Chock, S. 2020. Design justice: Community-led practices to build the worlds we need. Cambridge: MIT Press.

Dai, S. 2020. Online education start-up Yuanfudao raises US$1bn as sector heats up on coronavirus impact. *South China Morning Post*, 1 April: https://www.scmp.com/tech/start-ups/article/3077915/online-education-startyuanfudao-raises-us1bn-sector-hots

D'Ignazio, C. & Klein, L. 2019. Data Feminism. Cambridge: MIT Press.

D'Mello, S., Picard, R. & Graesser, A. 2007. Toward an Affect-Sensitive AutoTutor. *IEEE Intelligent Systems*, 22(4), 53-61.

du Boulay, B. 2019. Escape from the Skinner Box: The case for contemporary intelligent learning environments. *British Journal of Educational Technology*, 50(6), 2902–2919.

Eynon, R. & Young, E. 2020. Methodology, Legend, and Rhetoric: The Constructions of AI by Academia, Industry, and Policy Groups for Lifelong Learning. Science, Technology, & Human Values. https://doi.org/10.1177/0162243920906475

Fischer, C. et al. 2020. Mining Big Data in Education: Affordances and Challenges. Review of Research in Education, 44(1), 130-160.

Garland, D., 2014. What is a "history of the present"? On Foucault's genealogies and their critical preconditions. *Punishment & society*, 16(4), 365-384.

Gilliard, C. & Culik, H. Filtering content is often done with good intent, but filtering can also create equity and privacy issues. Common Sense Education, 24 May: https://www.commonsense.org/education/articles/digital-redlining-access-and-privacy

Gray, C.M. & Boling, E. 2016. Inscribing ethics and values in designs for learning: a problematic. Education Technology Research and Development 64, 969–1001.

Greene, D., Hoffman, A.L. & Stark, L. 2019. Better, Nicer, Clearer, Fairer: A Critical Assessment of the Movement for Ethical Artificial Intelligence and Machine Learning. *Proceedings of the 52nd Hawaii International Conference on System Sciences*: https://scholarspace.manoa.hawaii.edu/bitstream/10125/59651/0211.pdf

Harel Ben Shahar, T. 2017. Educational justice and big data. Theory & Research in Education, 15(3), 306–320.

Hill, P.W. and Barber, M., 2014. Preparing for a Renaissance in Assessment. London: Pearson.

Holmes, W., Bialik, M. & Fadel, C. 2019. *Artificial Intelligence in Education: Promises and Implications for Teaching and Learning*. Boston: Centre for Curriculum Redesign.

Knox, J., Williamson, B., & Bayne, S. 2020. Machine behaviourism: future visions of ‘learnification’and ‘datafication’ across humans and digital technologies. *Learning, Media and Technology*, 45(1), 31-45.

Lingard, B., Martino, W. and Rezai-Rashti, G., 2013. Testing regimes, accountabilities and education policy: Commensurate global and national developments. *Journal of Education Policy*, 28(5), 539-556

Liu, Y. 2020. The future of the classroom? China’s experience of AI in education. Nesta, 18 May: https://www.nesta.org.uk/report/the-future-of-the-classroom/#content

Livingstone S., Carr J. & Byrne, J. 2015. *One in Three: Internet Governance and Children’s Rights*. Global Commission on Internet Governance Paper Series. Oct; No. 22.

Lockyer, L., Heathcote, E. and Dawson, S., 2013. Informing pedagogical action: Aligning learning analytics with learning design. American Behavioral Scientist, 57(10), 1439-1459.

Luckin, R. Holmes, W., Griffiths, M. and Forcier, L. B. 2016. I*ntelligence Unleashed: An argument for AI in Education*. Pearson Education, London.

Luckin, R. 2020. AI in education will help us understand how we think. The Financial Times, 10 March: https://www.ft.com/content/4f24adca-5186-11ea-8841-482eed0038b1

Lupton, D., & Williamson, B. 2017. The datafied child: The dataveillance of children and implications for their rights. *New Media & Society*, *19*(5), 780–794.

McArthur, D., Lewis, M. & Bishay, M. 1995. The roles of artificial intelligence in education: Current progress and future prospects. *i-manager’s Journal of Educational Technology*, 1(4), 42-80.

Mackenzie, A. 2017. *Machine Learners: Archaeology of a Data Practice*. Cambridge, MA: MIT Press.

Microsoft Education Team. 2020. New IDC report shows big opportunities to transform higher education through AI. Microsoft Education Blog, 3 March: https://educationblog.microsoft.com/en-us/2020/03/new-report-shows-big-opportunities-to-transform-higher-education-through-ai/

Piattoeva, N. & Boden, R. 2020. Escaping numbers? The ambiguities of the governance of education through data. *International Studies in Sociology of Education*, 29(1-2), 1-18.

Rieder, B. 2020. Engines of Order: A mechanology of algorithmic techniques. Amsterdam: Amsterdam University Press.

Schank, R.C. & Edelson, D.J. 1989. A role for AI in education: Using technology to reshape education. International Journal of Artificial Intelligence in Education, 1(2), 3-20.

Schmidhuber, J. 2019. Deep Learning: Our Miraculous Year 1990-1991. arXiv: https://arxiv.org/abs/2005.05744

Schmidhuber, J. 2020. The 2010s: Our Decade of Deep Learning / Outlook on the 2020s. 20 Feb: http://people.idsia.ch/~juergen/2010s-our-decade-of-deep-learning.html

Selwyn, N. 2010. Looking beyond learning: Notes towards the critical study of educational technology. Journal of computer assisted learning, 26(1), 65-73.

Selwyn, N. 2019. What's the Problem with Learning Analytics? Journal of Learning Analytics, 6(3), 11–19.

Srnicek, N. 2017. *Platform Capitalism*. Cambridge: Polity.

Tuomi, I. 2018. The Impact of Artificial Intelligence on Learning, Teaching, and Education: Policies for the future. Luxembourg: Publications Office of the European Union.

Whittaker, M., Crawford, K., Dobbe, R., Fried, G., Kaziunas, E., Mathur, V., … Schwartz, O. 2018. AI Now Report 2018, AI Now Institute, New York University, NY. https://ainowinstitute.org/AI_Now_2018_Report.pdf

Williamson, B., Potter J. & Eynon, R. 2019. New research problems and agendas in learning, media and technology: the editors' wishlist, Learning, Media and Technology, 44(2), 87-91.

Zeide, E. 2017. The structural consequences of big data-driven education. Big Data, 5(2), 164-172.