



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Machine behaviourism

Citation for published version:

Knox, J, Williamson, B & Bayne, S 2020, 'Machine behaviourism: Future visions of 'learnification' and 'datafication' across humans and digital technologies', *Learning, Media and Technology*, vol. 45, no. 1, pp. 31-45. <https://doi.org/10.1080/17439884.2019.1623251>

Digital Object Identifier (DOI):

[10.1080/17439884.2019.1623251](https://doi.org/10.1080/17439884.2019.1623251)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Learning, Media and Technology

Publisher Rights Statement:

This is an Accepted Manuscript of an article published by Taylor & Francis in Learning, Media and Technology on 10-06-2019, available online: <https://www.tandfonline.com/doi/full/10.1080/17439884.2019.1623251>.

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Special Issue: Education and technology into the 2020s: speculative futures

Machine Behaviourism: future visions of 'learnification' and 'datafication' across humans and digital technologies

Jeremy Knox, Ben Williamson, Sian Bayne
Centre for Research in Digital Education, University of Edinburgh

Machine Behaviourism: future visions of ‘learnification’ and ‘datafication’ across humans and digital technologies

Abstract This paper examines visions of ‘learning’ across humans and machines in a near-future of intensive data analytics. Building upon the concept of ‘learnification’, practices of ‘learning’ in emerging big data-driven environments are discussed in two significant ways: the *training* of machines, and the *nudging* of human decisions through digital choice architectures. Firstly, ‘machine learning’ is discussed as an important example of how data-driven technologies are beginning to influence educational activity, both through sophisticated technical expertise and a grounding in behavioural psychology. Secondly, we explore how educational software design informed by behavioural economics is increasingly intended to frame learner choices to influence and ‘nudge’ decisions towards optimal outcomes. Through the growing influence of ‘data science’ on education, behaviourist psychology is increasingly and powerfully invested in future educational practices. Finally, it is argued that future education may tend toward very specific forms of behavioural governance – a ‘machine behaviourism’ – entailing combinations of radical behaviourist theories and machine learning systems, that appear to work against notions of student autonomy and participation, seeking to intervene in educational conduct and shaping learner behaviour towards predefined aims.

Keywords behaviourism; behavioural economics; data science; nudging; machine learning; reinforcement learning

Intensive data processing increasingly pervades formal education activities, catalysing shifts in the discourses and educational practices associated with the concept of ‘learning’. The era of so-called ‘big data’ is ushering in important new changes in educational policy, pedagogical practice, and institutional strategy. Sociological studies have begun to establish critical perspectives on the rise of the ‘data revolution’ (Kitchin 2014), taking seriously the ‘social power’ of algorithmic systems (Beer 2017, 3) and their capacity to reproduce and amplify inequalities (O’Neil 2017; Eubanks 2017; Wachter-boettch 2017; Noble 2018), as well as considering implications for critical thinking in the era of machine learning (Mackenzie 2017). However, educational research has been slower to grasp the far-reaching effects of such ‘datafication’. There is a pressing need to develop such perspectives, given high-profile associations between ‘data science’ and broad educational agendas, from national policy initiatives, through research and development, to classroom practices ‘on the ground’. Importantly, these agendas demonstrate the increasing entanglement of sophisticated arrangements of software, infrastructure, and code with educational theory in ways that powerfully shape both the governance and day-to-day activities of teaching and learning in institutions.

The purpose of this paper is to analyse how methods associated with ‘data science’ are involved in (re)formulating, both conceptually and materially, what has become the central concern of contemporary education: ‘learning’. Building upon Biesta’s concept of ‘learnification’ (2005), which defines a consumerist and ‘student-centred’ orientation for education, we examine how the increasing ‘datafication’ of the sector both *follows from* the ‘marketised’ model, and *disrupts* some of its underlying assumptions about the learning process.

In particular, the paper considers specific techniques of ‘machine learning’ that describe software which is ‘trained’ to perform specific tasks, either through exposure to large data sets or with rewarding systems (Alpaydin 2016). These techniques are shown to frame learning, both conceptually and practically, in ways that are highly significant to understanding the future trajectory of the ‘datafied’ education sector, and the role of ‘learners’ within it. Firstly, machine learning exemplifies how data-driven technologies are beginning to influence educational activity, both through sophisticated technical expertise and a grounding in behavioural psychology. One of the most prominent areas is ‘reinforcement learning’, a technique involving the programming of an algorithm to react to an environment in ways that incur the greatest value reward. This technique can operate without any predetermined instructions for a given context, beyond the task of identifying and optimising behaviours that gain rewards.

Secondly, we examine how insights from behavioural economics—a branch of psychology that informs ‘behavioural public policy’ and ‘behaviour change’ initiatives (John 2018)—are increasingly utilized in educational software design to frame learner choices in ways that influence decisions towards optimal outcomes. The sentiments and behaviours of learners are thus ‘nudged’ in predictable ways, grounded in theoretical frameworks from psychology that assume irrational and unconscious decision-making. As behavioural science has merged with data science, new forms of ‘persuasive computing’ and ‘hypernudge’ techniques have become possible (Yeung 2017). Consequently, human learning appears susceptible to heightened forms of governance as digital environments become ever more connected, and education becomes subsumed into broader political agendas of behavioural governance. These techniques illustrate the escalating dominance of ‘data science’ in education, through which behaviourist psychology is powerfully invested in future educational practices.

Finally, this paper will suggest that future education may tend toward very specific forms of behavioural governance—a ‘machine behaviourism’—entailing potent combinations of radical behaviourist theories and powerful machine learning systems, that appear to work against notions of student autonomy and participation, seeking to intervene in educational conduct and shape learner behaviour towards predefined aims. As such, ‘learnification’ is being extended across humans and machinic systems of ‘datafication’.

The learnification of education

Through a number of works, Gert Biesta’s concept of ‘learnification’ has identified broad changes in the social, political, and theoretical framing of education in the 21st century, which

have worked to reorder the sector and its activities around the concept of ‘learning’ (2005; 2006; 2012; 2013). Where terms such as ‘teaching’ might have been privileged in the past, Biesta argues contemporary discourse tends to frame educational activity as ‘the provision of learning opportunities or learning experiences’ (2005, 55). While often emphasising the rise of ‘the language of learning’ (2005, 54), Biesta’s analysis is critically concerned with the concrete social infrastructures that enable discourse, and the very real conditions and practices that are represented and shaped by the ways education is discussed. Central to this work has been the development of critical perspectives around the figuring of the self-directing and autonomous ‘learner’, which Biesta attributes to a transactional framing of education, where:

‘(i) the learner is the (potential) consumer, the one who has certain needs, in which (ii) the teacher, the educator, or the educational institution becomes the provider, that is, the one who is there to meet the needs of the learner, and where (iii) education itself becomes a commodity to be provided or delivered by the teacher or educational institution and to be consumed by the learner’ (58)

Importantly, Biesta (2005) identifies ‘learnification’ not as the result of a direct or intentional agenda, but rather as the consequence of a range of occurrences, situated across education, wider society, and the political sphere. Biesta specifies the acceptance of constructivist and socio-cultural theories in the field of the psychology of learning, which foreground the idea that ‘knowledge and understanding are actively constructed by the learner, often in co-operation with fellow learners’ (2005, 56). He also identifies the broader academic interest in postmodernist theory with its underlying critique of the project of education and its assumed role in the Enlightenment ideals of emancipation and liberation. Shifting to sociological trends, Biesta outlines the significant rise of informal learning activities outside of traditional education institutions, underlining the ‘individualistic and individualised’ character of these activities, where ‘learners are primarily struggling with themselves...with their body, their identity, and their relationships’ (2005, 57). Not only is the figure of the learner placed at the centre of the educational arrangement, but the individual becomes the site of learning. Socio-economic and political aspects of learnification are also identified in ‘the erosion of the welfare state and the rise of neo-liberalism’ (Biesta 2005, 57), which have intensified the commodification of educational activity, and the culture of accountability accompanying it.

Biesta’s overarching critique demonstrates how a language of learning tends to generate two problematic assumptions. Firstly, he challenges ‘the underlying assumption that learners come to education with a clear understanding of what their needs are’ (Biesta 2005, 59) by arguing that education exists to provide the very means by which learners become *able* to understand those needs. In this way, he criticises the ‘learnified’ institution as one which shuns responsibility for defining educational aims. For Biesta, where the predefined ‘needs’ of the learner begin to provide the core justifications for education, the role of the educational institution and its teachers becomes merely responsive, one in which the institution exists to supply educational ‘services’ in response to learner demand. Secondly, and as a consequence of the transactional structure of the educational relationship, the assumption that ‘the only

questions that can meaningfully be asked about education are technical questions, that is, questions about the efficiency and the effectiveness of the educational process' (2005, 59). Learnification is portrayed as blind to broader questions about the role and purpose of education in wider society, about how particular educational aims and aspirations are negotiated and established, and what kind of power structures underpin these processes.

The concept of 'learnification' is important for this current paper in two ways. Firstly, his critique of the marketization of education demonstrates some of the ways in which recent political and conceptual shifts have made the educational landscape more amenable to data-intensive practices. In particular, by highlighting the optimistic figuration of the self-directing learner habitually used to rationalise the neoliberal revisioning of the sector (Busch, 2016), it reveals underlying commitments to a particular kind of human condition, with which powerful data-intensive technologies are being designed to intervene. Secondly, the concept of 'learnification' is important in order to account for the increasing prevalence of such technologies across the education sector, and the ways these reshape shared understandings of learning while also structuring how it is measured and advanced.

The 'datafication' of education

A number of important factors can be identified in the increasing 'datafication' of education, many aligning with 'learnification'. It is perhaps the term 'learning analytics' that is used most frequently to describe the application of 'machine learning' (see below) in education, and developments here serve as a useful view of how 'datafication' follows from the 'learnified' state of the sector, while also influencing the forms and categories of learning taking place within it. Learning analytics has garnered significant attention in education, following high-profile research (see Lang *et al.* 2017) and policy agendas (see Tsai & Gašević 2017), as well as consistent predictions of its impending disruption of formal education (for example Johnson *et al.* 2016; Adams Becker *et al.* 2017). However, as part of 'learnification', learning analytics can be better understood not as straightforwardly 'disruptive' but rather as deeply entwined with broader shifts in the political economy of the educational landscape. Nonetheless, the concept of 'learnification' is not enough to straightforwardly explain the 'datafication' of education, and important (re)turns to particular educational theories will also be discussed as significant deviations from the core ideas associated with 'learner-centred' practices.

A significant feature of the early literature has involved attempts to define learning analytics according to 'essential characteristics' (Cooper 2012), and align emerging approaches with existing techniques from 'business intelligence' and 'data mining' (Ferguson 2012). While subsequent work has called for more connections with established educational research (Gasevic *et al.* 2015), further collaborative and interdisciplinary exchange (Ferguson & Clow 2017), and the establishing of ethical perspectives (Slade & Prinsloo 2013; Rubel & Jones 2016), learning analytics approaches are clearly grounded in the notion of a technical 'discipline' (Siemens 2013), as signified by the term 'education data science' (Cope & Kalantzis 2016). The discursive alignment of education with 'data science' reflects the much broader narrative of 'Silicon Valley solutionism' (Morozov 2013), which tends to frame data-

work as an outside force of revolutionary disruption, capable of radically enhancing a particular sector or social practice with technology-fuelled efficiency and precision. However, Biesta's work on 'learnification' suggests an educational landscape already oriented towards the kind of transactional arrangements in which 'datafication' might thrive. It is useful, therefore, to understand 'business intelligence', not simply as a coincidental precursor to the advent of convenient learning analytic systems, but as a set of practices and technologies that became relevant to education partly as a *result* of the increasing commodification of educational practices and materials.

Similarly, the interest in learning analytics—which is concerned with the 'measurement, collection, analysis and reporting of data about student progress and how the curriculum is delivered' (JISC 2018)—only becomes plausible within an education system already oriented around a centrally important 'learner'. The promotion of learning analytics is often premised upon its ability to reveal insights about learning unobtainable without the collection and analysis of learner-data. As a kind of precision science for education, it is attributed with the ability to 'penetrate the fog' (Long and Siemens 2011) of educational activity, in which teachers and institutions have been 'driving blind' (Pea 2014, 16). However, the idea that we are blinded by a 'fog' is to assume that there is a blind spot in our understanding as educators that can be 'made clear' by data. Such an assumption is itself a product of recent orientations towards 'learning' as a central concern, and a subsequent marginalisation of teacher expertise. Biesta's 'learnification' helps to explain the existing conditions in which questions about the scientific interrogation of educational practices can be asked, and perceived as necessary. Where learning analytics is habitually positioned on the 'horizon' (Johnson *et al.* 2016; Adams Becker *et al.* 2017), it should be noted that the educational territory is already one in which creating 'a model of the learner' (Pea 2014, 24) is perceived as a coherent and sensible task.

The supposed 'datafication' of education also, though, entails significant shifts away from the 'learnification' of the sector. Firstly, one might consider the kind of technology used to collect learner-data, and ultimately 'model' the learner: the software platform. Just as one can discern a shift in priority from the naming conventions of such technologies—from 'course management system (or CMS), suggestive of institutional or pedagogical supervision and control, to 'virtual learning environment' (or VLE), emphasising a 'learner centred' arrangement¹—the 'dashboard' now features prominently on educational software platforms. As such, these systems are now more overtly cast as devices for the harvesting of student data, and the presentation of 'actionable intelligence' (Campbell *et al.* 2007) to teachers and administrators. This is a crucial shift. While Biesta (2005) writes against the idea of the student-consumer, demanding learning experiences at will and without purpose, a tangible new prioritisation, for teachers and students alike, is detectable in the encouragement of 'data-driven decision making' (van Barneveld *et al.* 2012). Here, learners are assumed to respond directly to what the dashboard reveals, rather than evoking some kind of consumerist

¹ What better cultural indicator of such a shift than, at the time of writing, Wikipedia redirecting 'course management system' links to the 'virtual learning environment' page.

desire. Significantly, this appears to suggest a reassertion of the kind of centralised control that advocates of ‘learning’ sought to redress (Biesta 2012). Moreover, the role of the educator is not simply one of the ‘facilitation’ or ‘delivery’ of ‘learning opportunities’ as critiqued by Biesta, but a practice now initiated and determined by systems of data analysis.

Evoking the well-rehearsed adage—often attributed to Douglas Rushkoff (2010)—that the social media ‘user’ is the *product* rather than the consumer of the service, the shift from the ‘VLE’ to ‘data dashboard’ signals a very different view of the ‘learner’. As ‘student data’ becomes more necessary to provide substance to the analytic engines of education (discussed further in the next section on machine learning), learner behaviour within the software platform becomes a valuable commodity in itself. Mackenzie (2017, 29, n.2) notes that machine learning models learn ‘to optimize their predictions on the basis of “experience”/data’ and that, when employed in learning analytics programs, students’ own ‘learning is learned by machine learners’. Thus, the business of analytics becomes concerned with developing more accurate models of user behaviour, for which increased user activity is necessary. This gestures towards an important change in educational prioritisation that appears to counter Biesta’s framing of learner-consumers. As analytic systems become increasingly important to educational practices, learners are positioned as prized products, from which valuable behaviours can be extracted and consumed by ever-improving algorithmic systems. In these ways, consideration of the educational software platform highlights important aspects of learning analytics that entail significantly different ideas about learning and the learner than those assumed by a ‘student-centred’ orientation: not only is data positioned before the desires of the learner as the authoritative source for educational action, but the role of the learner itself is also recast as the product of consumerist analytic technologies.

It is also important to note the expansion of learning analytics beyond the mere software platform, to incorporate so-called ‘Internet of Things’ devices (for example Cheng & Liao 2012) and wearable technologies (for example Di Mitri *et al.* 2017; Yu *et al.* 2017). These technologies present the means not only to expand analytics procedures to the very ‘real’ spaces and infrastructure of the campus—in addition to the ‘online’—but also to the bodies of learners (as discussed more later). This ‘datafication’ of the educational environment would seem to constitute another significant turn away from notions of a passive surround, pliant to the needs of the learner—an environment ‘grounded in constructivist theory of learning’ (Biesta 2012, 38)—towards the learning environment as an active producer of data with its own ‘smartness’. While it is important to acknowledge that the software platform already establishes a data collection routine around the profiling of the individual, wearables offer technologies capable of extending data capture beyond the usual kind of interface to incorporate bodily events, such as facial expressions, biophysiological responses, or neural signals. Significantly, this indicates a shift towards the desire to detect and measure ‘learning’, not in the intentional expressions or decisions of a self-directing, autonomous learner, but in activity and behaviour located in ‘non-conscious’ mental processes and autonomic ‘bodily’ functions. This highlights a crucial final consideration for the relationships between the so-called ‘learnification’ and ‘datafication’ of education: the

(re)turn to behaviourism (a specific form of this combination will be elaborated below in the section on ‘nudging’).

Alongside the grounding in technical expertise from computer science disciplines, learning analytics are usually aligned with psychology, and the so-called ‘learning sciences’ (Gasevic *et al.* 2015). Therefore, accompanying the ‘re-emergence of empiricism’ ushered in by the ‘data revolution’ (Kitchin 2014), has been *radical* behaviourism, with which the work of B.F. Skinner ‘has been making a comeback’ (Friesen 2018, 3). According to Zuboff (2019, 296), the emphasis on Skinnerian behaviourism in reinforcement learning, which Skinner sought to operationalize through a pervasive ‘technology of behaviour’, now underpins many of the psychological techniques used by data scientists and social media designers for online ‘behaviour modification’ and ‘behaviour engineering’. Similar forms of ‘behaviour design’ underpin many recent data-processing educational technologies (Watters 2017). It is important to clarify particularities of the school of behaviourism favoured by the proponents of ‘data science’. It does not refer to a notion of ‘behaviour’ simply as a proxy for ‘learning’, as one might detect in the framing of data-intensive computation as merely providing predictive ‘insight’, or indeed being straightforwardly ‘analytic’, as might be associated with ‘classical conditioning’. ‘Behaviour’, in its radical form, must be understood in an interventionist sense, where the task is not just to discern learning through observation of behaviour, but to actively modify and determine it through direct and intentional control of the environment. Particularly significantly for educational concerns here is that, as a field of research, education has largely rejected radical behaviourism as an explanatory theory (Friesen 2018). Skinner’s ‘operant conditioning’ appears as a rather bleak view of the learner as passive *tabula rasa*, in comparison with what has become a common sense of active and self-knowing individuals in constructivist theory. To attempt to control behaviour through positive and negative reinforcements, or indeed punishments, sounds clinical and detached, if not decidedly primitive, in contemporary understandings of teaching as ‘facilitation’ and learning as individual ‘empowerment’. Nevertheless, the resurfacing of radical behaviourism is occurring in two important areas that have profound implications for understanding ‘learning’ in a ‘datafied’ educational system. We turn to these next.

The learnification of machines

One of Biesta’s concerns with the discourse of learning is that it is portrayed as ‘something inevitable, something we have to do and cannot *not* do’ (Biesta 2013, 4). As such, learnification is naturalised beyond the boundaries of the institution and throughout our lives. However, in the era of ‘datafication’, learning continues to expand, beyond the boundaries of individual humans and social relations and further into the domain of machines. Significant attention has been given to ‘machine learning’ in recent years, referring to software that can be ‘trained’ with large data sets to perform specific tasks such as recommendation, prediction, image or voice recognition, and autonomous driving. Describing machine learning as ‘the new A.I.’, Alpaydin (2016) portrays the emergence of these techniques as a kind of natural progression of digital technology as hardware, and in particular graphics processing units, increase in power (also see Peters 2018), as well as various media pervading ever more aspects of social life, and in so doing, generating ever more data. As such, machine learning

tends to be understood as a technical process through which computer systems learn from data without the requirement of formal structuring by a human programmer (Berry 2017).

Although machine learning has a long history in statistics, cybernetics, data mining and computer science, it has grown rapidly over the last two decades (Kitchin 2014). In his ‘archaeology’ of the field, Mackenzie (2017) has documented the proliferation of machine learning as an ‘accumulation’ of data practices, demonstrating its spread beyond computer science, statistics and engineering to almost every domain of contemporary science, as well as to health, humanities and social science fields. Machine learning is ‘laying claim to the apparatus of knowledge production’ (Mackenzie 2017, 13), partly as a result of a massive industry in textbooks, how-to manuals, online courses, code libraries, demonstrator models and other pedagogic materials. As a field, machine learning has produced ‘learning algorithms’—or just ‘learners’—that can process vast databases of historic and real-time data in order to generate predictions and probabilities, and thereby to ‘mediate future-oriented decisions,’ or ‘rule out some and reinforce other futures’ (Mackenzie 2017, 7-8). Machine learning is concerned with analysing data to classify and predict events or identify associations among people, things, and processes, often for scientific, commercial or governmental purposes. In fact, as Mackenzie (2017: 216) elaborates, machine learning cannot be reduced wholly to an automated computational process—in contrast, perhaps, to learning by a conscious human subject—but should rather be understood as the outcome of a ‘human-machine assemblage’. He describes ‘machine learners’ as ‘both humans and machines or human-machine relations’ (6) that also include code, databases, infrastructures, platforms and interfaces, new technical settings and human experts, scientific and commercial settings. It also involves a vast proliferation of techniques, methods and practices that have emerged from more than a century of work in mathematics, statistics and computer science, as well as from psychology, cybernetics, cognitive science, genetics and neuroscience research.

As such, machine learning evades simple definition, especially when it comes to the question of how ‘learning’ is conceptualized. As Mackenzie (2017: 48) puts it, ‘to understand what machines can learn, we need to look at how they have been drawn, designed, or formalized.’ Berry (2017) notes that the ‘learning’ in machine learning ‘is very specific and technical in its deployment, and relates to the ability to undertake skills or tasks, not to wider humanistic connotations of learning as understanding, interpretation, etc.’ At its core, the learning of machine learning refers to the finding of appropriate mathematical functions that can transform data into classifications and predictions (Mackenzie 2015). Machine learners learn by finding and optimizing functions that can produce knowledge, and which also feed back into how the machine learner operates in an iterative fashion.

Three particular forms of learning can be found in machine learning: supervised, unsupervised and reinforcement learning (Alpaydin 2016). Supervised learning involves the construction of a predictive model based on a sample of ‘training data’, which can then be tested out on unseen test data; unsupervised learning looks for well-characterized patterns in data and produces its own algorithms and programs to undertake analysis ‘in the wild’;

reinforcement learning proceeds by ‘rewarding’ itself as it performs actions (discussed further below). As such, different forms of machine learning learn through supervised direct instruction, unsupervised experiences, or through behavioural rewards, in ways which emulate different branches of psychological learning theory. In other words, machine learning needs to be apprehended as a complex and shifting combination of mathematical functions, algorithmic processes, underlying models of the things or processes to be classified and predicted, as well as human operators, training sets and code, all of which have to be coordinated in disparate but connected ways in order for machine learning to ‘learn’.

Reinforcement learning, currently one of the most prominent areas of machine learning, is the technique underpinning the highly publicised AlphaGo² programme developed by Deep Mind, and an approach which draws explicitly from behavioural psychology. Alongside neural network architecture, a cutting-edge ‘self-play reinforcement learning algorithm’—‘an algorithm that learns ... superhuman proficiency’ as its science team described it in *Nature*—is AlphaGo Zero’s primary technical innovation (Silver et al 2017, 354). It is ‘trained solely by self-play reinforcement learning, starting from random play, without any supervision or use of human data,’ and as such, ‘AlphaGo becomes its own teacher’ (Silver et al 2017, 354). Reinforcement learning systems are therefore not reliant on human behavioural data, as described previously, and learn exclusively from their own attempts to seek rewards. Furthermore, it is claimed that AlphaGo’s reinforcement learning systems ‘are trained from their own experience, in principle allowing them to exceed human capabilities, and to operate in domains where human expertise is lacking’ (Silver et al 2017, 354). As the algorithm processes its own activity in the game, it is ‘rewarded’ and ‘reinforced’ by the wins it achieves, in order ‘to train to superhuman level,’ and even to ‘learn’ new moves that it can then ‘teach’ to human players (Silver et al 2017, 358). The implication, in other words, is that powerful self-reinforcing machine learners could be put to the task of training better humans, or even of outperforming humans to solve real-world problems.

While the influence of reinforcement learning on formal education is as yet unclear, it is associated with very particular ideas about teaching and learning processes that potentially stand to shape developments in artificial intelligence technologies for education. For example, Alpaydin describes reinforcement learning as ‘learning with a critic’ where ‘[a] critic differs from a teacher in that they do not tell us what to do, but only how well we have been doing in the past’ (Alpaydin 2016, 127). This reveals a very different notion of learning to the autonomous neoliberal consumer defined by Biesta (2005), one that appears to strictly condition behaviour according to predefined values. Furthermore, as explained by Alpaydin (2016), such feedback is necessarily summative; feedback on individual actions within a larger interdependent sequence have no value or purpose, only ‘winning’ or ‘losing’ counts. This stark vision of learning may be able to train machines towards ‘superhuman’ capabilities, but there is certainly no suggestion that such a brutal regime would be suitable for human learners. Indeed, the area of psychology from which reinforcement learning derived was concerned, not with humans at all, but with animals. In their seminal text on

² <https://deepmind.com/research/alphago/>

reinforcement learning, Sutton and Barto (first published in 1998) are explicit about the connection with ‘learning theories developed to explain how animals like rats, pigeons, and rabbits learn in controlled laboratory experiments’ (Sutton & Barto 2018, 343), and acknowledge that the field of psychology itself had largely shifted interest towards cognition. Nonetheless, their text demonstrates the influence of psychological models of learning on algorithm design, which they expect to develop further as ‘development of reinforcement learning theory and algorithms will likely exploit links to many other features of animal learning’ (Sutton & Barto 2018, 370). Such connections are important ‘because they expose computational principles important to learning, whether it is learning by artificial or by natural systems’ (Sutton and Barto 2018, 343). Moreover, they note ‘mounting evidence from neuroscience that the nervous systems of humans and many other animals implement algorithms that correspond in striking ways to reinforcement learning algorithms’, in particular involving dopamine, the chemical involved in reward processing in the brains of mammals (Sutton & Barto 2018, 377).

It is important here not to view reinforcement learning as an isolated technique, both in the sense that in machine learning research it tends to form just one part of a broader range of integrated, and exploratory, approaches, as well as often requiring monitoring from a human agent. Returning to Mackenzie’s (2017) notion of a ‘human-machine assemblage’ is useful here to highlight the extent of technologies, techniques, methods, and practices that make reinforcement learning possible. Further, while there is no necessary suggestion that reinforcement learning is set to instil notions of animalistic behaviour as the new standard for learning in the ‘datafied’ institution, the disciplinary origins of this particular technique are important to acknowledge. Behaviour, it seems, continues to be viewed as a key part of understanding and developing the learning process, both in humans and machines. The next section turns to specific ways in which the theories of reinforcement learning are gaining traction in education through an increasing interest in behavioural intervention. It is important, therefore, not to suggest that reinforcement learning, or indeed machine learning in general, be equated with ‘human learning’ – researchers in machine learning are certainly cognisant of the differences between pattern recognition and human intelligence (see Lake *et al.* 2017). Rather the key point of these sections is to underscore how particular ideas about ‘learning’ – significantly, those involving behaviour - appear to be traversing computer science and educational disciplines through the use of data-intensive technologies.

Learning as ‘nudging’

Behavioural science and its variants (behavioural economics, behaviour change theory, and the psychology of persuasion) have become key intellectual vectors of governmental authority in the first decades of the twenty-first century. As Jones, Pykett and Whitehead (2013) have shown in their studies of behavioural science as a form of ‘psychological governance’, a key insight of the field is that most human decision-making is inherently irrational, habitual, and predictable. Overturning classical economic models of humans as rational decision-makers which have long dominated political thought, and seeming to underpin Biesta’s ‘learnification’ (2005) as discussed previously, behavioural science instead insists that people do not act in their own long-term best interests or make socio-economic or

personal decisions based on available information. Instead, people tend to favour immediate need and gratification over future planning, with their decisions and behaviours involving emotional responses, habits, social norms, and the automatic, unconscious and involuntary within human action, which might nonetheless be predicted, regulated, enhanced and exploited (Feitsma 2018a).

As such, ‘human nature’ has been reconceptualised in psychological terms by behavioural science, and made amenable to governmental and commercial programs designed to persuade people to make better choices regarding their own personal well-being and the wider socio-economic well-being too (John 2018). Feitsma (2018b, n.p.) notes that the application of behavioural insights by governments appeals to a modernist ideal ‘to manage society through hard fact-finding’ but also to ‘a neoliberal agenda, seeking to responsabilize citizens to alter their problematic behaviours rather than change the underlying socio-economic, political, and institutional structures that underpin such behaviours’. New disciplines such as behavioural design, cognitive engineering and persuasive computing have emerged which focus on the idea that the environment has an enduring impact on human behaviours and habits. As such, behavioural designers have set themselves the challenge of designing ‘choice architectures’—the physical, socio-cultural and administrative environments in which choices are framed—that can influence people to make favourable decisions, while persuasive computing engineers create software programs in order to stimulate users to behave in ways deemed desirable (Yeung 2017; Zuboff 2019). A key aspect of behavioural economics relates to neuroscientific insights into dopamine and reward-processing in the brain, which itself has been conceptualized through ‘the convergence of computational reinforcement learning and results of neuroscience experiments’ (Sutton & Bardo 2018, 377). As such, computational theories of reinforcement learning from the domain of machine learning have begun to shape the scientific understanding of human reward-seeking, and these insights have been taken up in governmental strategies to ‘nudge’ people to make ‘rewarding’ decisions (Whitehead et al 2018).

The field of education has become a site for behavioural intervention and nudging (Bradbury, McGimpsey & Santori 2013). The most influential text in the behavioural science literature, *Nudge* by Thaler & Sunstein (2008) explicitly notes that the shaping of educational options for students counts as a choice architecture. The UK government’s Behavioural Insights Team (informally known as the ‘Nudge Unit’) has partnered with the education business Pearson on a report about using behavioural insights in formal education (O’Reilly et al 2017). The UK higher education regulator, the Office for Students, has also adopted aspects of behavioural design to inform how it presents data to prospective university students - thus nudging them to make favourable choices about which degrees to study and which institutions to choose for application - while the Department for Education’s new ‘Education Lab’ positions behavioural science as a key source of scientific expertise in policy design (Knight 2019). These ‘educative nudges’ are behavioural interventions designed to promote a significant degree of human reflection in the long-term decisions that affect them (John 2018).

The most significant instantiation of educational nudging, however, relates to the use of digital technologies, often involving machine learning systems, that are designed to shape students' choices and decisions based on constant tracking and predicting of their behaviours, emotions and actions (Nemorin 2017). As Whitehead et al (2018) have noted in relation to the enlarged scope of behavioural science more broadly, the behavioural potentials of digital technologies and social media that track data about users' consumption and behavioural habits makes the potential for behavioural modification and experimentation ever greater. In the new 'guinea pig economy' connected devices such as smartphones and wearables, combined with data analytics software make it possible for companies to continuously test their application of behavioural insights on users' habits, and to refine their products to produce optimal outcomes. Davies (2016, 222) argues that the combination of big data, the sharing of thoughts and feelings on social media, and more emotionally intelligent machines that can read an individual's feelings through sentiment analysis of text, or facial analysis of the body, face and behaviour, is opening up unprecedented opportunities for 'psychological surveillance' and behaviour modification. For Zuboff (2019, 282), this 'new frontier of rendition trained ... on your emotional life' is concerned with the mass-scale prediction and modification of behaviour under 'surveillance capitalism'. Mass digitization and data analytics offers a mode of 'psychological audit,' with 'the truth of our emotions' made plain 'once researchers have decoded our brains, faces and unintentional sentiments', with the result that 'the real facts of how to influence decision-making may come to light' and 'society becomes designed and governed as a vast laboratory' intended to 'make life easier, healthier and happier for all' (Davies 2016, 223-24).

As part of current educational policy preoccupations with 'social and emotional learning', students are increasingly being approached as irrational subjects whose behaviours can be nudged towards particular 'desirable' or 'preferable' outcomes by carefully-designed persuasive computing technologies. The 'non-cognitive' and 'socio-emotional' aspects of students' learning have become of interest to behavioural science since it is assumed that many students do not engage sufficiently with academic curriculum demands out of rational decision-making about their long-term best interests; instead, students are understood to be behaviourally shaped by emotions, habits and other noncognitive processes which need to be 'optimized' for academic learning (Lavecchia et al 2014). As John (2018, 14) notes, emerging forms of behavioural governance are concerned with changing long-term behaviours and recalibrating individuals' goals, aims which are linked to psychological concepts such as 'grit'. 'Grit' is one of the core concepts of the social-emotional learning movement along with related concepts of 'growth mindset' and 'character', which, as pedagogic implementations of behavioural science, focus on identifying and then intervening in students' behavioural habits in order to lead to better long-term outcomes. The OECD is developing a large-scale international test of social and emotional skills as concern for measuring these qualities of learning becomes more widespread, amid rising policy concern about how social-emotional learning can be improved (OECD 2017).

Data analytics techniques and apps to capture measures of aspects of social-emotional learning have begun to proliferate as a way of identifying students in need of behavioural

intervention and ‘nudges’. In primary schools, apps such as ClassDojo are used by teachers to monitor students’ observable behaviours that correlate with categories of grit and growth mindset (<https://www.classdojo.com/>). For high schools, school-wide data analytics packages are available such as Panorama Education, which has been designed for schools to help ‘measure and understand social-emotional learning ... with research-backed measures and actionable data reports’ (<https://www.panoramaed.com/social-emotional-learning>). Explicitly rooted in theories and practices relating to grit and growth mindset, Panorama generates ‘early warning indicators’ which flag students in need of behavioural improvement. These technologies enable teachers and school leaders to capture long-term data relating to students’ social-emotional learning, to view the data on data visualization dashboards, and to target individuals and groups for behavioural intervention and support. As with other forms of educative nudging, their aim is to change long-term behaviours and individual goal-setting, but they also serve up data-based cues and prompts to teachers where need-for-intervention is signalled in the data.

Other social-emotional learning technologies include wearable biometric and facial vision devices capable of capturing real-time data about students’ engagement, attention and mood during educational activities. McStay (2017, 9) terms such devices ‘emotional AI’ that can ‘feel-into’ human emotional life, and refers to them as ‘automated industrial psychology’. Academic research centres dedicated to exploring ‘emotional learning analytics’ and ‘wearable-enhanced learning’ have generated insights into how to detect emotions from bodily signals that have begun to inspire commercial development (D’Mello 2018). The Affective Computing group of the MIT Media Lab, for example, has spawned spinout companies to support wearable biometric wristbands and facial vision analytics that are considered appropriate for use in educational settings (Santiago et al 2015; McDuff et al 2016):

Wearable devices can also provide a minute-by-minute record of someone’s emotional state, potentially helping to build self-awareness and even empathy, both of which are critical components of social and emotional skills. The Embrace watch, from Empatica, is a wearable device that tracks physiological stress and activity. It can be programmed to vibrate when stress reaches a specific level, giving someone time to switch to a more positive response before stress gets out of control. (WEF 2016, 14)

Through wearables, the ‘nudge’ is made haptic as it touches the body itself. Another Affective Computing product uses a webcam and algorithms to capture, identify and analyse human emotions and reactions to external stimuli, using facial data to compare a user’s expressions with a database of more than 1 billion facial expressions to assess differences between emotions such as happiness, fear, surprise and confusion. Through such emerging technical innovations ‘artificial intelligence and multimodal social computing could help improve cognitive, social and emotional skills’ (WEF 2016: 15). These emotion-detecting devices originally from the MIT Affective Computing lab collect data from ‘unconscious’ flickers of affect in facial expressions and other biophysiological responses, in order to ‘render both conscious and unconscious emotion as observable behaviour for coding and

calculation' (Zuboff 2019, 286). As such, behaviourist forms of observation have moved to machine learning platforms.

A further outgrowth of digital nudging is the effort by global commercial companies Google and Amazon to teach children how to be 'polite' to smart home assistants such as Amazon's Alexa and Echo and Google's Home devices. Through in-built politeness features, these devices are designed to encourage children to treat artificial intelligence with courtesy, as if the machines have feelings and emotions. The politeness feature of Google Home, for example, 'uses positive reinforcement and reciprocal niceties to encourage manners in children', thus offering 'subtle nudges toward politeness', while also potentially 'creating a lifelong intuition in children that software has feelings that can be hurt, that it's an intelligent being to be respected—or even an authority to be obeyed' (Elgan 2018). The logic of such devices is that they can read mood from the student's voice, body, brainwaves or face and deliver feedback which is supposed to prompt the student to shift to a more positive and rewarding state. These technologies are pedagogic variants of the 'hypernudge' techniques employed by social media and web companies to influence users and maintain consumer engagement (Yeung 2017).

As these examples demonstrate, various machine learning technologies of nudging, persuading and reinforcing positive or preferable emotional conduct are currently flowing into educational spaces. Rooted in behavioural science and its view of human nature as irrational, social-emotional learning technologies seek to measure feelings, moods and emotions by rendering them as observable biophysical behaviours. These developments are giving rise to machine learning systems that are themselves 'learning' human affect, and may even be able to generate appropriate emotional responses themselves in ways that might persuade or influence the human subject (Rose, Aicardi & Reinsborough 2016), such as by delivering haptic 'nudges' to the skin or adapting the 'choice architecture' in a digital learning environment. In this way, learning itself is reconceptualised in terms of psychologically quantifiable affective characteristics which are both detectable as autonomic bodily signals and amenable to being changed and modified in line with particular theories about what constitutes the 'correct', 'preferable', or 'desirable' behaviours for learning. Psychologists of grit, growth mindset and character have supplied the intellectual grounding for the advance of behaviour change and nudge programs in education, inspiring developers of analytics packages and apps to embed behavioural design approaches in their products, and to create emotionally-sensitive and potentially persuasive machine learning systems.

Conclusions

This paper has explored some of the ways 'learning' is being defined, promoted, and practiced in educational activity through the use of contemporary data-driven technologies. It may be useful to term these shifts in the conceptual and functional understandings of learning 'machine behaviourism': forms of learning that are shaped and conditioned by a combination of complex machine learning technologies with (often radical) behaviourist psychology and behavioural economics. The significance of this shift in the framing of learning has been emphasised through the discussion of Gert Biesta's concept of 'learnification' (2005)—a

prominent critique of the marketization of education through notions of ‘student-centred’ learning. This paper has attempted to demonstrate that many of the assumptions embodied in the idea of learnification—concerning the autonomy and empowerment of learners that continue to pervade educational discourse—are being eroded by the intervention of powerful data-driven machine learning systems that are being employed to shape and intervene in learner behaviour. While Biesta’s ‘learnification’ critiqued the shift towards a market-driven model of the education, where learner-consumers and their rational decision-making were placed at the centre of responsive institutions and constructivist pedagogies, the era of data-intensive technology is necessitating renewed attention to the ways ‘learning’ is being conceived and practiced. The ‘learner’ is now an irrational and emotional subject whose behaviours and actions are understood to be both machine-readable by learning algorithms and modifiable by digital hypernudge platforms. ‘Machine behaviourism’ marks a (re)turn to the influence of behaviourist psychology on educational practice, and appears to usher in new powerful regimes of centralised control, in which ‘correct’ forms of performance and conduct have already been decided, and learners are increasingly ‘nudged’ towards predefined modes of participation and behaviour. Across this shift, this paper has attempted to emphasise the entanglement, rather than the simplistic ‘disruption’, of educational technology within the political economy of education, drawing attention to the ways data processing not only intensifies the marketised and performance-driven vision of the institution, but also offers heightened forms of governance and surveillance. It is also highly significant that, in their need for data to drive machine learning systems, a substantial interest in new education technology development appears to be directed towards ‘bodily’ and ‘emotional’ data, as a source of supposedly new accuracy and authenticity in understanding human learning. This appears to overlook entirely the kind of rational, ‘conscious’, and constructive participation from students that has, despite the critique of marketization, tended to characterise the era of ‘learnification’.

Learning analytics, machine learning, and ‘nudging’ are certainly not mainstream aspects of education presently, and are probably better understood as experimental, and somewhat marginal, uses of technology within the sector. Nevertheless, the powerful discursive regimes that underlie the global interest in machine learning and so-called ‘big data’ provide a seemingly seductive narrative of efficiency and precision that may be increasingly hard for educational institutions to resist. In a future where not just humans but also machines are tasked with learning, attention needs to be given to the ways learning itself is theorised, modelled, encoded, and exchanged between students and progressively more ‘intelligent’, ‘affective’, and interventionist educational technologies. More interdisciplinary work across education and computer science is needed, not only to broaden the scope of learning beyond narrow measures of performance or delimited methods of behaviour management, but also to do more to develop and define the position and purpose of education within a wider political economy increasingly characterised by data-intensive practices. This work could draw from the growing critical interest in the societal effects of algorithms, which, as a result of the embedding of such systems in key social institutions, have been shown to disproportionately affect the poor and marginalised (O’Neil 2017; Eubanks 2017; Noble 2018). Given the pervasive discourses of efficiency-savings that accompany the promotion of data-intensive

technologies in education, it seems likely that less financially stable educational institutions will be those more inclined to adopt systems that offer cost-effective management of student cohorts. Further, as Wilson *et al.* (2017) discuss in relation to learning analytics, the significant emphasis on retention in the development of such systems, appear to suggest contexts with low teacher-to-student ratios, where reduced contact time makes it difficult to identify 'at risk' student populations without technical assistance. In these ways, the increasing use of machine learning in education, not only appears more likely to proliferate in low socio-economic or developing contexts, but also to disproportionately impact those without access to elite educational institutions.

References

Adams Becker, S., Cummins, M., Davis, A., Freeman, A., Hall, Giesinger, C., and Ananthanarayanan, V. 2017. *NMC Horizon Report: 2017 Higher Education Edition*. Austin, Texas: The New Media Consortium.

Alpaydin, E. 2016. *Machine Learning: the new AI*. Cambridge: MIT Press.

Beer, D. 2017. The social power of algorithms. *Information, Communication & Society* 20(1), 1-13

Berry, D.M. 2017. Prolegomenon to a Media Theory of Machine Learning: Compute-Computing and Compute-Computed. *Media Theory* 1(1): 74-87.

Biesta, Gert. 2005. Against Learning. Reclaiming a Language for Education in an age of Learning. *Nordisk Pedagogik* 25(1): 54-66.

Biesta, Gert. 2006. *Beyond Learning. Democratic Education for a Human Future*. Boulder, CO: Paradigm Publishers.

Biesta, Gert. 2012. Giving Teaching Back to Education: Responding to the Disappearance of the Teacher. *Phenomenology & Practice* 6(2): 35-49.

Biesta, Gert. 2013. Interrupting the Politics of Learning. *Power and Education* 5(1): 4-15.

Bradbury, A., McGimpsey, I. & Santori, D. 2013. Revising rationality: the use of 'nudge' approaches in neoliberal education policy. *Journal of Education Policy* 28 (2): 247-67.

Campbell, John P., Peter B. De Blois, and Diana G. Oblinger. 2007. Academic Analytics: A New Tool for a New Era. *Educause Review* 42(4): 40-57.

Cheng, H.C. & Liao, W.W. 2012. Establishing an lifelong learning environment using IOT and learning analytics. 14th International Conference on Advanced Communication Technology (ICACT), PyeongChang, 2012: 1178-1183.

Cooper, A. 2012. What Is Analytics? Definition and Essential Characteristics. *CETIS Analytics Series* 1(5): 1-10.

Cope, B, Kalantzis, M (2016) Big data comes to school: Implications for learning, assessment and research. *AERA Open* 2(2): 1-19.

Davies, W. 2016. *The Happiness Industry: How the Government and Big Business Sold Us Well-Being*. London: Verso.

D'Mello, S. 2017. Emotional learning analytics. In Lang, C., Siemens, G., Wise, A. & Gasevic, D. (eds) *Handbook of Learning Analytics*, 115-127. Society of Learning Analytics Research.

Di Mitri, D., Scheffel, M., Drachler, H., Börner, D., Ternier, S., Specht, M. 2017. Learning pulse: using wearable biosensors and learning analytics to investigate and predict learning success in self-regulated learning. Proceedings of the Seventh International Learning Analytics & Knowledge Conference, Vancouver, British Columbia, Canada, March 13 – 17: 188-197

Elgan, M. 2018. The case against teaching kids to be polite to Alexa. *Fast Company*, 24 June: <https://www.fastcompany.com/40588020/the-case-against-teaching-kids-to-be-polite-to-alexa>

Eubanks, V. 2018. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. Macmillan.

Feitsma, J. 2018a. Brokering behaviour change: the work of behavioural insights experts in government. *Policy & Politics*: <https://doi.org/10.1332/030557318X15174915040678>

Feitsma, J. 2018b. The Rise of the Nudge Expert. *Discover Society*, 3 April: <https://discoversociety.org/2018/04/03/policy-politics-the-rise-of-the-nudge-expert/>

Ferguson, R. (2012). Learning analytics: drivers, developments and challenges. *International Journal of Technology Enhanced Learning* 4(5/6): 304–317.

Friesen, N. 2018. Personalized Learning Technology and the New Behaviorism: Beyond Freedom and Dignity. Available: https://www.academia.edu/36111979/Personal_learning_new_behaviorism2.docx

Gasevic D, Dawson S, Siemens G. 2015. Let's not forget: Learning analytics are about learning. *TechTrends* 59(1): 64.

Goodfellow, I., Bengio, Y., & Courville, A. 2016. *Deep Learning*. London: MIT Press.

JISC. 2018. Learning Analytics. Designing learning and assessment in a digital age. Available: <https://www.jisc.ac.uk/guides/designing-learning-and-assessment-in-a-digital-age/learning-analytics#>

John, P. 2018. *How far to nudge? Assessing behavioural public policy*. Cheltenham: Edward Elgar.

Johnson, L., Adams Becker, S., Cummins, M., Estrada, V., Freeman, A., and Hall, C. 2016. *NMC Horizon Report: 2016 Higher Education Edition*. Austin, Texas: NMC.

Johnson, S. 2017. This Company Wants to Gather Student Brainwave Data to Measure 'Engagement'. *EdSurge*, 26 October: <https://www.edsurge.com/news/2017-10-26-this-company-wants-to-gather-student-brainwave-data-to-measure-engagement>

Jones, R., Pykett, J. & Whitehead, M. 2013. *Changing Behaviours: On the rise of the psychological state*. Cheltenham: Edward Elgar.

Kitchin, R. 2014. *The Data Revolution: Big Data, Open Data, Data Infrastructures and Their Consequences*. London: Sage.

Knight, A. 2019. Putting people at the heart of policymaking. DfE Digital and Transformation blog, 1 April: <https://dfedigital.blog.gov.uk/2019/04/01/about-user-centred-research-programme/>

Lake, B.M., Ullman, T.D., Tenenbaum, J.B. and Gershman, S.J., 2017. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40.

Lang, C., Siemens, G., Wise, A., & Gašević, D. 2017. Handbook of Learning Analytics. SoLAR. DOI: 10.18608/hla17.

Lavecchia, A.M., Liu, H. & Oreopoulos, P. 2014. *Behavioral Economics of Education: Progress and Possibilities*. National Bureau of Economic Research, Working Paper 20609: <http://www.nber.org/papers/w20609.pdf>

Mackenzie, A. (2015). The production of prediction: What does machine learning want? *European Journal of Cultural Studies* 18(4-5): 429-445.

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

McDuff, D., Amr, M., Mahmoud, A., Turcot, J., Mavadati, M. & el Kaliouby, R. 2016. AFFDEX SDK: A Cross-Platform Real-Time Multi-Face Expression Recognition Toolkit. CHI'16 Extended Abstracts, May 07-12, 2016, San Jose, CA, USA. https://www.affectiva.com/wp-content/uploads/2017/03/McDuff_2016_Affdex.pdf

McStay, A. 2017. *Emotional AI: The rise of empathic media*. London: Sage.

Morozov, E. 2013. *To save everything click here: Technology, solutionism and the urge to x problems that don't exist*. London: Penguin.

Nemrin, S. 2017. Affective capture in digital school spaces and the modulation of student subjectivities. *Emotion, Society and Space* 24: 11-18.

Noble, S.U. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press.

OECD. 2017. *Social and Emotional Skills: Well-being, connectedness and success*. Paris: OECD.

O'Neil, C. 2017. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. London: Penguin.

O'Reilly, F., Chande, R., Groot, B., Sanders, M. & Soon, Z. 2017. *Behavioural Insights for Education: A practical guide for parents, teachers and school leaders*. London: Pearson.

Pea, R. (2014). *The learning analytics workgroup: A report on building the field of learning analytics for personalized learning at scale*. Palo Alto: Stanford University.

Rose, N., Aicardi, C. & Reinsborough, M. 2016. *Future computing and robotics: A foresight report from the Human Brain Project lab*. London: King's College London.

Rubel, A. & Jones, K.M.L. (2016) Student privacy in learning analytics: An information ethics perspective. *The Information Society* 32(2): 143-159.

Rushkoff, D. 2010. You Are Not Facebook's Customer. Available: <http://www.rushkoff.com/you-are-not-facebooks-customer/>

- Santiago, J., Tholen, L., Nguyen, M. & Nophlin, A. 2015. Empatica impacts education, 20 April. Retrieved from:
http://auapps.american.edu/alberto/www/itr/ITRFiles/Competition_2015S/Group7_ITR_EmpaticaE4_NewsArticle.pdf
- Siemens, G. (2013). Learning Analytics: The Emergence of a Discipline. *American Behavioral Scientist* 57(10): 1380–400
- Silver, D., Schrittwieser, J., Simonyan, K. et al. 2017. Mastering the game of Go without human knowledge. *Nature* 550: doi:10.1038/nature24270
- Slade, S. and Prinsloo, P. (2013). Learning analytics: ethical issues and dilemmas. *American Behavioral Scientist* 57(10): 1509–1528.
- Sutton and Barto. 2018. *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge: MIT Press. Open Access online version available: <https://drive.google.com/file/d/1xeUDVGGWGUUv1-ccUMAZHJLej2C7aAFWY/view>
- Thaler, R. & Sunstein, C. 2008. *Nudge: Improving decisions about health, wealth and happiness*. London: Penguin.
- Tsai, Y. S., & Gasevic, D. (2017). Learning analytics in higher education challenges and policies: a review of eight learning analytics policies. In Proceedings of the Seventh International Learning Analytics & Knowledge Conference: 233-242. ACM.
- van Barneveld, Angela, Kimberly E. Arnold, and John P. Campbell. 2012. Analytics in Higher Education: Establishing a Common Language. ELI Paper 2012 (1): 1–11.
- Wachter-boettch, S. 2017. *Technically Wrong: Sexist Apps, Biased Algorithms, and Other Threats of Toxic Tech*. Norton and Co.: New York.
- Watters, A. 2017. Education technology and the new behaviourism. Hack Education, 23 December: <http://hackededucation.com/2017/12/23/top-ed-tech-trends-social-emotional-learning>
- Wilson, A., Watson, C., Thompson, T.L., Drew, V. & Doyle, S. 2017. Learning Analytics: challenges and limitations. *Teaching in Higher Education*, 22:8, 991-1007, DOI: 10.1080/13562517.2017.1332026
- Whitehead, M., Jones, R., Lilley, R., Pykett, J. & Howell, R. 2018. *Neoliberalism: Behavioural government in the twenty-first century*. London: Routledge.
- Yeung, K. 2017. Hypernudge: Big Data as a mode of regulation by design. *Information, Communication & Society* 20(1): 118-136.
- Yu L., Sen Z., Zhiqiang Z., Wendong X., and Shengquan Y. 2017. A Framework for Learning Analytics Using Commodity Wearable Devices. *Sensors*. 17(6): 1382. doi: 10.3390/s17061382
- Zuboff, S. 2019. *The Age of Surveillance Capitalism: The fight for a human future at the new frontier of power*. London: Profile Books,