Learning as a Machine: Crossovers between Humans and Machines

Mireille Hildebrandt
Law, Science, Technology, & Society Studies
Faculty of Law and Criminology
Vrije Universiteit Brussel
mireille.hildebrandt@vub.ac.be

ABSTRACT: This article is a revised version of the keynote presented at LAK '16 in Edinburgh. The article investigates some of the assumptions of learning analytics, notably those related to behaviourism. Building on the work of Ivan Pavlov, Herbert Simon, and James Gibson as ways of "learning as a machine," the article then develops two levels of investigation (processing of personal data and profiling based on machine learning) to assess how data driven education affects privacy, non-discrimination, and the presumption of innocence. Finally, the article discusses how data minimization and profile transparency will contribute to the methodological integrity of learning analytics, while protecting the fundamental rights and freedoms of human learners thus safeguarding the creativity, humour, and contestability of human learning.

Keywords: Data protection, data protection by design, privacy, machine learning, behaviourism, nudging, Pavlov, Simon, Gibson, capture, optimization, affordance

1   INTRODUCTION

Human beings learn from the moment they are born (Jarvis, 2005; Campbell & Dawson, 1995), and even before that. We learn because we are vulnerable, because we can suffer, because we have something to lose. We learn because we can die — because we will die, like all living organisms (or nearly all), and we want to delay that moment and flourish. We learn for the joy of learning, we are aware that we learn, and we help others to learn (our children, friends, colleagues, pupils). Machines learn because we configure them as such, to perform specific tasks for us, or even because we will learn from them about new solutions to known problems, about the problems that we missed, and we even configure machines to learn for the pure joy of it — as part of our creativity.[1] Machines also learn to contribute to our own learning processes.

The LAK '16 conference was (also) about privacy and learning analytics (LA),[2] making a pivotal point by investigating the pitfalls of data-driven learning as an issue that should inform the development of LA as a discipline, instead of framing such issues as side-effects to be dealt with at a later stage. Having studied the new *Journal of Learning Analytics,* I must admit to being deeply impressed with the high

---

[1] On the challenges of computational creativity, see Leijnen (2014).
[2] Special 2016 issue on "Ethics and Privacy in Learning Analytics," *Journal of Learning Analytics, 3*(1).

standards of research and the integration of methodological, ethical, and legal reflection.[3] In my keynote — and in this paper — I contribute to such reflection by raising the question of how machine learning will affect human learning, highlighting the need for effective redress of learners subjected to automated scoring practices based on LA. This is the introduction.

LA can be defined as data-driven applications that help students to improve their learning behaviours, and/or applications that help schools, colleges, and universities to manage, coordinate, and administrate information and communication regarding the progress of their students' performances (Baker & Inventado, 2014). Such progress is measured in machine-readable format, which implies that whichever part of the learning process cannot be made machine-readable will not be part of learning analytics and will not feed into the performance metrics. To the extent that these performance metrics will come to determine our learning processes, we need to acknowledge that whatever remains off the radar of LA will not visibly contribute to the flourishing of the students or the learning institutions we develop and create. This may have consequences for individual liberty and for the dignity of individuals insofar as their learning capabilities cannot be translated into machine-readable formats.[4] The famous legal philosopher Dworkin summed up democracy and the rule of law in terms of equal respect and concern for each individual person. This grounds democracy — one person, one vote — and the rule of law — human rights overrule majoritarian disrespect (Dworkin, 1991). To the extent that a person is reduced to inferences about her machine-readable behaviours, and targeted based on such inferences, her freedom from unreasonable constraints (liberty) to construct her identity (dignity) may be violated. This is not merely her private interest, but regards the capability of each and every citizen to develop an independent mind, which is the unspoken assumption of a viable democracy. Being tested by learning machines will lure learners into "thinking" like them, because they need to anticipate how these systems anticipate them. Learners may even start thinking in terms of their behaviours instead of their actions,[5] or be nudged into a preconceived learning process instead of being challenged to develop a critical distance.[6] The latter has consequences for our capability to engage in democratic participation, which depends on independent minds willing to challenge convention, power, and authority.

In this paper, I will investigate such potential implications in the light of data protection law. I will not, however, merely provide an analysis of current and upcoming legal rules. Rather, I hope to explain what is at stake when employing LA and how EU data protection law addresses potential threats. I will argue that to develop and integrate fair and sustainable LA into future proof learning institutions, we need to build effective protection and feedback into the relevant architectures, where feedback concerns the inner workings of the analytics and the foreseeable consequences of their scoring operations. This should enable the contestation of the analytics, whenever it significantly affects an individual learner. In section two, I will distinguish two levels of LA: 1) interventions at the level of identifiable students, and

---

[3] For excellent methodological reflection, see Berg, Mol, Kismihók, and Sclater (2016).
[4] One could rephrase by inquiring how LA scoring practices affect the agency and vulnerability of individual human learners, see saliently Prinsloo and Slade (2016).
[5] For a salient investigation of the difference between behaviour and action, see Vender (2011).
[6] On "the pre-emption of intent," see McStay (2011).

2) analysis and prediction based on aggregate data. Those concerned about data protection often get stuck at level 1, because this regards the processing of personal data. My main concern here is with level 2, because it enables personalized targeting based on machine learning. In section three, I will investigate how we should understand the idea of "learning as a machine," distinguishing 1) the Pavlov approach, 2) the Simon approach, and 3) the Gibson approach. This will enable us to better frame the threats that LA poses to privacy, non-discrimination, due process, and the presumption of innocence, to be discussed in section four at both levels. Finally, in section five, I will discuss the legal obligation to provide profile transparency, and how data protection by design could serve as an effective means to achieve "technological due process" within the settings of LA.

2   TWO LEVELS OF LA

2.1   First Level of LA: Processing Personal Data

First, LA concerns interventions at the level of an identifiable student, both when collecting her data, and when applying the results of LA to her, whether or not she is aware of this. Note that in terms of EU data protection law, identifiability includes both direct and indirect identification, while identification includes "being singled out."[7] The data collected may be data about previous grades, previously attended schools or universities, records of teachers' evaluations, grading, all types of tests, and even alumni data capable of linking school or college behaviours with employment achievements. This type of data processing concerns the automation of previously manual administration. LA, however, also refers to the collection of behavioural data, traced and tracked from new types of software that enable eLearning (Clark & Mayer, 2011). Such behavioural data may include keystroke and clickstream behaviours that trace reading habits and measure time spent on various tasks, or, for instance biometric behaviours such as eye movements that indicate boredom or loss of focus. It is pivotal to note that such behavioural data may also be derived from sources outside the educational setting of the student and institution. They may concern social media data, quantified-self metrics, or other types of data that may — for instance — be bought from data brokers. Such out-of-context data may be correlated with educational data to figure out how a student's background, lifestyle, or context correlates with her learning achievements. This may actually be legal or illegal, ethical or unethical. The point is that it is possible.

Next, LA concerns capturing student data: interventions at this level include feedback given to the student, such as real-time feedback during eLearning exercises or "early-warning" systems that support timely reconsideration of a student's capacities or learning strategies. The first level, however, also includes interventions regarding an identifiable student that she may or may not be aware of, such as

---

[7] Art. 4.1 General Data Protection Regulation R 2016/697 EU (GDPR) defines personal data as follows: "'personal data' means any information relating to an identified or identifiable natural person ("data subject"); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person." Recital 26 determines that "singling out" renders a person identifiable. The definition is nearly identical to the currently applicable art. 2(a) Data Protection Directive D 95/46/EC (DPD). The GDPR will come into force in May 2018 and have direct effect in all the Member States of the EU, replacing relevant national law.

real-time reconfiguration of her adaptive eLearning software, placement in a specific class or group, automated evaluation of tests, referral to a counsellor, selection for awards, grants, or being categorized as a potential drop-out, or failure to meet the standards of the course or educational institution (Fiaidhi, 2014; Kalampokis, Tambouris, & Tarabanis, 2013). It could eventually include decisions based on inferences concerning health risks, drug abuse, earning capacity, financial resources, or public security risks.

2.2   Second Level of LA: Analytics, Machine Learning, KDD

The second level of intervention is not about interaction with identifiable students, but concerns the analyses of the data that prepare potential interventions regarding the student. This data can be anonymous data or personal data; if it is personal data, it can be pseudonymous.[8] Anonymization would rule out the applicability of data protection law, as this implies that the data no longer qualifies as personal data. Pseudonymization implies that the "anonymization" is not irreversible, and therefore data protection law applies.[9] Pseudonymization, however, may qualify as proper protection of the data, meaning that it constitutes security by design or data protection by design. This depends on whether the processing of the data is fair and lawful and on the kinds of risks it entails for data subjects (the learner). The processing is unlawful if there is no legal ground or because the purpose of processing is not legitimate or has been exhausted. In that case, pseudonymization will not make the processing operations lawful.[10]

Analyzing large amounts of data can be done by means of, for instance, knowledge discovery in data basis (KDD) or machine learning (ML). KDD has been defined as a "nontrivial process of identifying valid, novel, potentially useful, and ultimately understandable patterns in data" (Piatetsky-Shapiro, 1996). ML has been defined as machines that "learn from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks T, as measured by P, improves with experience E" (Mitchell, 2006). The second level of intervention is about the detection of relevant correlations, association rules, and so forth. Basically, this level is all about pinpointing the key classifiers that will allow teachers and their institutions to gain insights into the data points deemed key to their students' learning capabilities and to their organizations' achievements. This can be done by means of various types of machine learning (Greller & Drachsler, 2012). It is the second level that forages patterns in big

---

[8] Art. 4(5) GDPR defines pseudonymous data as personal data: "'pseudonymization' means the processing of personal data in such a manner that the personal data can no longer be attributed to a specific data subject without the use of additional information, provided that such additional information is kept separately and is subject to technical and organizational measures to ensure that the personal data are not attributed to an identified or identifiable natural person."

[9] Recital 26 GDPR specifies: "To determine whether a natural person is identifiable, account should be taken of all the means reasonably likely to be used, such as singling out, either by the controller *or by another person* to identify the natural person directly or indirectly [my emphasis]." This means that even encryption does not anonymize, except if the decryption key cannot be accessed by anyone but the data subject. See also Opinion 5/2014 on Anonymization Techniques of the Art. 29 Working Party, available at http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2014/wp216_en.pdf

[10] Lawful processing requires one of six legal grounds, specified in art. 6 GDPR (art. 7 DPD), and compliance with the conditions of art. 5 GDPR (art. 6 DPD), notably that of purpose specification and use limitation. Recital 28 GDPR confirms that pseudonymization can contribute to compliance with the GDPR by reducing the risks for data subjects.

data to develop more adequate interventions at the first level, in order to improve the performance of the students and/or to advise them — or even to oblige them — "to change course." Such patterns may be about "when are students ready to move on the next topic? When are students falling behind in a course? When is a student at risk for not completing a course? What grade is a student likely to get without intervention? What is the best next course for a given student? Should a student be referred to a counsellor for help?" (Fiaidhi, 2014). This quotation is from of an editorial on learning analytics that nicely sums up some of the key objectives of LA. In the editorial, Fiaidhi (2014) actually proposes to invest in the analysis of so-called unstructured data to further the objectives of monitoring and improving the cognitive and intellectual development of "learners." He ends by commending a comprehensive learning architecture that integrates the analytics for structured and for unstructured data by means of a predictive engine, a content engine, an intervention engine, a feedback engine, and a measurement engine. The objectives and the various "engines" help to remind us of the close relationship of both levels of intervention, since the second level is instrumental for targeting students at the first level, either by providing them with feedback based on the knowledge mined, or by providing their teachers, institutional management, employers, or others that pay for their education with advice and relevant predictions. Such predictions enable covert targeting, excluding students from entering the next level of an application or from the next grade, from a school or college, or from remedial programs. This should be kept in mind when discussing privacy and other fundamental rights that may be infringed with large-scale implementation of LA. Though from a security and confidentiality perspective anonymization or pseudonymization should be a primary concern, there are major concerns that cannot be addressed by de-identifying the data used for LA. This concerns not merely privacy, but also non-discrimination, due process, and the presumption of innocence. Moreover, it concerns the affordances of learning institutions, and the kind of capabilities they enable for human learners.[11]

3   THREE WAYS OF "LEARNING AS A MACHINE"

Before moving into fundamental rights infringements, we need to understand what machine learning does (as a tool to contribute to human learning) and how it may transform the practice and concept of human learning. This means investigating the affordances of machine learning applications in LA and how they interact with the capabilities of human learners targeted by LA. The investigation will inquire into three ways of "learning as a machine," which I will coin "the Pavlov approach," "the Simon approach," and "the Gibson approach." After that, I will investigate how second-level interventions (KDD and ML operations) may impact fundamental rights, and finally explain how data protection by design (DPbD) may help to prevent data obesitas, loss of reputation, and untrustworthy LA from destroying its potential as a tool for better education.

3.1   The Pavlov Approach: Manipulating Behaviour

At the beginning of the 20th century, psychologist Ivan Pavlov developed the so-called stimulus–

---

[11] On the concept of an affordance, see chapter five in Gibson (1986). On the concept of a capability, see Sen (1999) and Nussbaum (2011). Both concepts emphasize the relational nature of an individual person and her environment.

response theory on learning mechanisms (Pavlov, 1927). He based his theory on a series of experiments with a dog, allegedly demonstrating that behaviour is the result of a recurrent chain of stimuli (sensed objects and/or behaviours, such as food or signals associated with provision of food). These chains of stimuli supposedly trigger subsequent responses (actuated by means of specific behaviours, such as manipulating a lever, or producing saliva in anticipation of food). Interestingly, Pavlov was one of the first behaviourists, claiming to base his theoretical insights on observable behaviours without any reference to the human or animal mind, which he treated as a black box. This supposedly delivered objective, scientific fact-finding instead of the assumedly subjective theories of psychologists such as Wilhelm Wundt (Kim, 2014), who based their inquiries into the human mind on a methodology of introspection — aiming to observe what goes on in the mind of an individual person and trying to generalize from there. One of the assumptions of Pavlov's theory was the idea that behaviour can be explained by means of physiological reflexes that can be tested in an experimental setting. Because quite a few behavioural responses seem to depend on adaptation to the environment instead of innate reflexes, Pavlov developed experiments to "prove" that animal responses can be the result of acquired reflexes, based on a physiological association of specific events (the ringing of a bell) with other specific events (the provision of food). Pavlov's theory was further developed by Watson (1930) and Skinner (1976), and is actually closely connected with Turing's (1950) thought experiment and Dennett's (1989) "intentional stance." Both Turing and Dennett avoid the issue of consciousness, focusing on whether observable behaviour can help us to anticipate an organism's — or even a person's — next move. Watson (1930), for instance, suggested that:

> The interest of the behaviorist in man's doings is more than the interest of the spectator — he wants to control man's reactions as physical scientists want to control and manipulate other natural phenomena. It is the business of behavioristic psychology to be able to predict and control human activity. To do this it must gather scientific data by scientific methods. Only then can the trained behaviourist predict, given the stimulus, what reaction will take place; or, given the reaction, state what the situation or stimulus is that has caused the reaction.

Clearly, behaviourism has strong links with determinism and physicalism, assuming that mind is ultimately a matter of matter, while in fact the mind may not really matter once we can detect the mechanisms that rule behaviour, which is assumed to be subject to and defined by the laws of causality.[12]

Contemporary heirs of the behaviourist approach can be found in Pentland's (2014) "social physics," and Helbing's (2012) "computational social science." Their approaches to social science are akin to social engineering, aiming to unlock the trove of knowledge that can be mined from data-driven applications (Pentland) and/or computational simulations of human society by means of multi-agent systems (Helbing). Both seek the key to understand individuals as nodes in a network, hoping to contribute to the engineering of a fair and sustainable society. The idea is that if we know — statistically — how people learn (how they develop which conditional reflexes), we can reconfigure their environment (the set of relevant stimuli) to make sure they indeed learn what "we" think they should learn (in terms of

---

[12] On the idea that the mind is an accidental side-effect of human intelligence, see Robinson (2015).

behaviour). The question is who is "we" and whether we really want to manipulate people into learning what "we" think is best (for whom?). Though there may be noble intent behind such societal engineering, it entails the assumption that we "learn as a machine," and this easily leads to attempts to manipulate us into supposedly "good" behaviours as if we were mindless pawns in a game of chess (or even Go) (Thaler & Sunstein, 2008; Hausman & Welch, 2010). Such a view misses out on the complexity of societal networks, the plasticity of the human brain, and the vulnerability as well as the creativity of both individual humans and their society. This refers to the intractability of human learning, and some might suggest that increasing computational power will push the threshold of this intractability. But there is also the question of computability.[13] Our life world is built on the ambiguity of "natural" language that is generative of a highly dynamic web of meaning. The curious amalgam of uncertainty and creativity that springs from our language-saturated environment is what holds us together (as a people) and what sets us apart (from other animals). A mechanistic understanding of human learning that ignores our capabilities for generating meaning, and for giving reasons for our actions, has a number of highly problematic ethical and epistemological implications. This may be related to a move from teaching students, that is putting new knowledge in front of them, sharing acquired knowledge, explaining the reasons and causal connections for such knowledge, to inducing learning processes with students, that is manipulating their environment in a way that it induces them to acquire specific insights. One need not be radical to guess that this will corrupt their autonomy, their capability to reason and to make up their own mind about what to learn how (Slade & Prinsloo, 2013; Dalsgaard, 2006).

3.2   The Simon Approach: Optimal Optimization

In the second half of the 20th century, Herbert Simon (1983) wrote an intriguing article under the title "Why should machines learn?" This was the time when machine learning was hardly a success and notably "nerve net learning," as he calls it, achieved next to nothing. Simon (one of the patriarchs of artificial intelligence), however, makes a number of highly relevant observations about the differences between human and machine learning. One of his main points refers to what he calls "the tediousness of human learning," as this is a slow process that takes a lot of time. To achieve a certain level of knowledge, people spend many years in school, whereas once a computer has been programmed with specific knowledge this can easily be transferred from one computer to the next — no interant learning process is needed:

> [W]e should ask whether we really want to make the computer go through that tedious process, or whether machines should be programmed directly to perform tasks avoiding humanoid learning entirely!
>
> Only one computer would have to learn; not every one would have to go to school.

We may smile about the naïveté of the observation Simon seems to endorse, but I guess Simon was

---

[13] On computability, see Brey & Søraker (2009). On intractability, see Dean (2015).

using the rhetorical device of a tongue-in-cheek, making the point that such an approach would only work if the goal of learning were a matter of "doing the same task over and over again, but more efficiently." If, however, the goal of learning is to "acquire the ability to perform a wide range of tasks," or if learning refers to discovering new knowledge and finding a way to retain it, the tediousness of human learning may be the best way to optimize the process. According to Simon, this is indeed the case, because programs facing real-life problems are highly complex and suffer numerous bugs and computational artefacts. This results in them ending up as being highly inefficient and ineffective in the course of time. So, he says, instead of putting all hope in one program, human society consists of numerous independent programs that allow for new beginnings:

> Old programs do not learn, they simply fade away. So do human beings, their undebuggable
> programs replaced by younger, possibly less tangled ones in other human heads.

So, the death and the birth of individual humans serves to resolve the legacy problem for optimizing learning programs. Simon's ironic message is that the fact that my knowledge and insights can only be transferred to others by means of great effort on both sides is not a disadvantage but rather an incredibly smart way to ensure creative and robust adaptation. Anyone who claims to have invented or engineered the optimal learning machine for a particular set of problems must be put to the test by means of retrials and alternative machines. We could frame this in terms of Wolpert's famous "no free lunch theorem" that proves mathematically that specific optimization mechanisms that "work" for a specific type of problem do not necessarily work for other types of problems (Wolpert & Macready, 1997).

3.3   The Gibson Approach: The Ecological Niche

James Gibson has developed one of the most crucial concepts for our current age, that of "an affordance." It is firmly embedded in evolutionary theory, physical reality, psychological inquiry, and a deep understanding of the relationship between a living entity and its environment. It skips naïve attempts to divide the world into mind and matter, without discarding values or mental capabilities. Let me quote his own definition of an affordance:

> The affordances of the environment are what it offers the animal, what it provides or furnishes,
> either for good or ill. The verb to afford is found in the dictionary, but the noun affordance is not.
> I have made it up. I mean by it something that refers to both the environment and the animal in
> a way that no existing term does. It implies the complementarity of the animal and the
> environment. (Gibson, 1986)

When speaking of an affordance, Gibson is not interested in separating objective physical properties from the subjective experience of the perceiving organism. Though this is of course possible, Gibson is interested in the actionable "properties" of an environment, which depend on the agent that depends on the environment. Each type of agent thus has its own ecological niche: the set of affordances tuned to a specific type of agent that is in turn tuned to its ecological niche. For a human being, this for instance implies breathable air, walkable surfaces, graspable things, but also other humans capable of speaking or learning her language. For humans, affordances are — in terms of Pavlov — the set of

stimuli that enable and constrain them to behave one way or another, including enabling them to learn new behaviours or even to unlearn innate behaviours. Other than Pavlov's determinism, Gibson's keen attention to the mutual shaping of an agent and its ecological niche celebrates creativity and renounces the idea that either the agent or the niche is defined by their current properties. In terms of Simon, we can add, machines similarly depend on a particular environment that affords them to learn one thing and not another. As we know, many machines require an artificial environment geared to their behavioural potential, while also preventing them from causing harm to their human masters. In robotics this is called the envelop (Floridi, 2014), a controlled environment that makes sure the robot finds its way and is constraint in ways that enable its productivity. Since environments are sets of affordances, they are contingent upon the type of organism able to perceive and act upon them. Fish don't do well on land; humans don't do well underwater or high up in the air; surfaces that can sustain ants may not sustain elephants. An ecological niche, therefore, is not the habitat of a particular organism, but the actionability of the environment of that organism.

Humans have managed to reconfigure their ecological niche in myriad ways — with far-reaching consequences for other organisms, as we all know. Major transitions of the affordances of our ecological niche have been generated by the introduction of tools to support calculation and writing, including the printing press. Each of them afforded entirely novel learning mechanisms, including the retention of what has been learnt outside the body of an individual human (e.g., on a piece of paper). This has triggered the establishment of formal learning and teaching institutions such as schools and universities, as well as archives and libraries and other repositories of retained knowledge, such as computers and servers.

What fascinates me is how Gibson's approach enables us to raise the question of how eLearning and other types of LA will transform the affordances of our learning institutions, changing our ecological niche and thus also our selves. Whereas, so far, the other agents in our niche were living organisms such as animals, from cattle to pets, we are now increasingly confronted with mindless agents of our own making, capable of observing us and adapting their behaviour to the feedback they gain from ours. It is not merely that *we use these machines*, for instance to help us to learn faster or more efficiently, *they also use us* to improve their performance. And, finally, we may come to the point of acknowledging that we are actually *interacting* with them — e.g., teaching them how to make us learn. Training algorithms on a training set and checking the outcome on the test set is already a type of interaction, especially where knowing the precise operations of the software is not feasible (as with neural nets) or irrelevant for the outcome. When I teach a human person, I cannot access her brain and follow the neural operations — they would certainly be incomprehensible to me. But this has never stopped us from teaching or otherwise interacting with another — instead, it keeps us on our toes because of the uncertainty it implies and the creativity it requires. The same seems to go for teaching a neural net. However, other than living beings, cyber-physical systems that integrate ML have nothing to lose; they cannot suffer or fear their own death. They can probably simulate all that and more, based on synthetic emotions and affective computing; but simulation is not the same as what is simulated. To the extent that our learning environments become dependent on these systems, this may be cause for concern.

What should concern us here is the fact that in order to interact with machines, people need to anticipate how these machines learn (Christian, 2011). This goes for those that develop the LA, e.g., when they train the algorithms for eLearning, or write the code for administrative decision-making. But it also goes for the students, e.g., when they interact with eLearning systems that can only respond to their machine-readable behaviours, or when they figure out which of their behaviours is taken into account when decisions are made on their eligibility for grants, admission to selected courses, or education institutions. As a result, our ecological niche may increasingly be shaped in a way that accommodates "learning as a machine" by humans. This is partly due to the unfathomable plasticity of the human brain. Though we are not born with a reading brain we somehow manage to reconfigure the organization, the morphology, and the behaviour of our brains to afford reading, as the combination of experimental psychology and brain science has demonstrated (Wolf & Stoodley, 2008). Once we learned to read and write, our ecological niche was hugely extended with the affordances of written and printed text. There is no reason to think that "ordinary people" cannot develop a brain more geared to mathematics, computation, and calculation if — as a society — we decide to make the effort (Sousa, 2014). This might indeed fast-track their interactions with LA systems, as they gain a more intuitive understanding of how these machines operate or even "think." But it may also have transformative effects on what it means to be human.

4 PRIVACY, NON-DISCRIMINATION, DUE PROCESS, AND THE PRESUMPTION OF INNOCENCE

4.1 First Level Issues: Identification and Targeting

As discussed above, the first level concerns interventions with identifiable individual learners. It is about collection of their personal data, whether from their admission forms or Blackboard, from teacher input such as grades, or from eLearning applications that capture behavioural data (or even from data troves gathered outside the educational setting, such as social networks and data brokers). It is also about applying inferences based on LA to an individual student, whether she is aware of this or not. This may involve automated or semi-automated decision-systems that categorize students and condition access to extra facilities, counselling, rewards, or grants. Results of LA may even condition sharing data with the police or intelligence agencies. I am speaking of affordances, not — yet — about whether this is ethical or legal.

At this level, privacy may be concerned whenever data is processed (for instance, shared) in a way that violates legitimate expectations, notably when data is shared out of context (Nissenbaum, 2010). This is at stake when access to the data within the learning institution is not appropriately secured (unreliable authentication), which results in teaching and other staff being able to look up how a particular student is doing — without any necessity. This is also at stake if publishers of electronic textbooks or eLearning software gain access to the data of individual students to improve the functionality of their software, or if grant providers demand an abundance of personal data to conduct fraud detection. As Khaliah Barnes, director of the student privacy project at the Electronic Privacy Information Center notes, "Rampant data collection is not only destroying student privacy, it also threatens students' intellectual freedom"

(Barnes, 2014; Drachsler et al., 2016). The tax authority, the police, or the intelligence services may demand access — probably accompanied by a prohibition to communicate this access with the relevant person or anyone else. This implies that undetected privacy infringements could easily occur, depending on the whims of whoever is in charge. Unless a system of safeguards is built into the architecture of the technical systems and the organizational design, privacy will be up for grabs. The point here is not bad intentions, but a far greater threat: the transformation of the affordances of learning institutions.

When LA inferences are applied to an individual based on knowledge that is not shared with her, or when such applications confront her with things about herself that she was not aware of, this may amount to another type of privacy violation (Hudson, 2005). Inferences may indicate health problems, notably concerning mental health, or they may show that students with a certain religious or ethnic background lack specific coping strategies, resulting in failure to finish the course or obtain the relevant degree. This is not only about privacy, but also about non-discrimination. It may be great to know that people from a certain background (ethnic, religious, gender, economic, geographic) have specific learning disabilities that can be remedied by means of specific interventions, based on extensive AB test designs.[14] That implies, however, that lots of students will be used as guinea pigs, while their sensitive data are employed for scientific research, to secure the policy goals of the educational institution, or — more mundane — to increase the profits of a service provider.

To the extent that students are not aware of all this, their due process rights may be at stake (Citron & Pasquale, 2014). They could be placed in certain categories based on clustering techniques that result in classifiers that discriminate (which is not necessarily prohibited) based on certain features that their "nearest neighbours" have, though they may not have these features (the profiles inferred will often be non-distributive).[15] To the extent that they are not aware of this, they cannot object. Indeed, even if they knew about this, they might have a problem arguing their case, because the categorization may be presented as objective and applicable and could only be countered by means of software verification or training alternative algorithms on the same data set. Another problem is that the arguments available to object to their being profiled may be restricted because objections must be submitted online and are formatted in a specific way. This probably entails that arguments that don't fit the format cannot be registered. Finally, those involved in dealing with the appeals may have formatted responses, e.g., stating that the system does not allow the argumentation. This is already often the case when objecting to fines from the taxation office or complaining about telecom providers. "Computer says no" is a popular phrase for good reason (Dean, 2016). This, finally, touches upon the presumption of innocence. In administrative law, unlike in criminal law, it is assumed that the government got it right and the citizen needs to prove her innocence. So, the student may have to prove that the system got it wrong, in terms dictated by the system. Kafka is around the corner — or amongst us (Solove, 2004). This is not science fiction and it is not even about bad intentions; it is about the consequences of entering another ecological niche with a different set of affordances.

---

[14] For an early precursor of AB research design in an educational setting, see Foster, Watson, Meeks, and Young (2002).
[15] On non-distributive profiles, see Vedder (1999).

4.2   Second Level Issues: Analytics and Prediction

Let's now check second level issues. These are entirely different, because the inferences drawn need not be personal data. Instead they will be patterns in the data set, linking specific attributes (features, behaviours, background, context) to specific performance metrics to create hindsight (what kind of attributes correspond with what level of performance?), foresight (what can be expected *ceteris paribus*?) and insight (which interventions could improve future student performance, depending on what circumstances?). The issue usually highlighted at this level is privacy in the narrow sense of hiding personal data and/or making sure that the data is protected against unauthorized disclosure. There are now a number of techniques meant to enable big data analytics in compliance with data protection law, based on pseudonymization (as irreversible anonymization is usually not an option). Frameworks such as MIT OpenPDS with its SafeAnswers module,[16] Max Planck's AIRCLOAK with its Insights AQR,[17] and the so-called PEP framework (polymorphous encryption & pseudonymization) developed at Radboud University (Verheul, 2015; Verheul, Jacobs, Meijer, Hildebrandt, & Ruiter, 2016), provide sophisticated ways of combining encryption with pseudonymous sharing or enable running code on the raw data without ever sharing the data (zero knowledge protocols). These are highly relevant and provide important means of engaging in data protection by default (data minimization). They do not, however, go to the heart of the matter. Once a service provider, grant provider, insurance company, potential employer, or educational institution gets ahold of the results of the analytics, a new set of affordances comes into play.

As discussed under first level issues, as soon as the inferences are applied to individual students, the infringements of privacy, non-discrimination, due process, and the presumption of innocence return at full speed. That is why I devote the final part of this article to the implications of LA for the fundamental core of human learning, i.e., our ability to reflect on our own behaviours. This "ability" may be an affordance of written text that is not necessarily an affordance of LA.[18]

5   PROFILING AND DATA PROTECTION BY DESIGN

5.1   The Relevance of the GDPR for First Level LA: Data Minimization

The upcoming General Data Protection Regulation consolidates and reinforces the core principles of the current Data Protection Directive. Nevertheless, it will be a game-changer for those who process personal data — notably for those who conduct additional processing, engaging in secondary usage, and re-purposing of personal data. Compared to the DPD, the GDPR will "do" three things (amongst many others). *First*, it will create a level playing field for stakeholders that base their business case on the processing of personal data. This level playing field is the result of deterrence: high fines, data breach notification, transparency requirements with high impact on reputation, and tort liability.[19] This will

---

[16] http://openpds.media.mit.edu/#architecture.

[17] https://www.aircloak.com/downloads/Aircloak-One-Pager.pdf.

[18] This has been extensively argued by Hildebrandt (2015).

[19] Art. 83 GDPR stipulates that supervisory authorities can punish violations of the GDPR with fines of up to 4% of global

institute countervailing powers capable of mitigating some of the problematic affordances of LA — speaking law to power. In the end, this level playing field will enable companies to act ethically without being pushed out of the market. *Second*, it will require a data protection impact assessment whenever high risk is expected for the rights and freedoms of natural persons (DPIA, art. 35 GDPR), which is related to the obligation to implement data protection by default and by design (DPbD, art. 25 GDPR). The latter requires that data minimization and other legal obligations be embedded into the architecture of data processing systems. These legal obligations will reconstruct the ecological niche; they will reconfigure the set of affordances that emerge in the wake of large-scale application of LA. This is not merely speaking law to power, but architecting trustworthy infrastructure that reduces risks to privacy and other fundamental rights by default, based on the technical specifications of LA systems. *Third*, it will privilege the processing of pseudonymized data — subject to stringent conditions — to enable big data analysis.[20] In combination with the previous points (creating a level playing field and integrating data protection into the hardwired and softwared LA systems), this will provide for an incentive structure that favours systems that enable secondary use for a number of specific purposes while reducing the risk of identification. Again, this will have a transformative impact on the affordances of LA, as it will simultaneously constrain the use of data that enables direct identification, while enabling big data analytics on pseudonymized data. An excellent example of this type of Data Protection by Design and Default can be found in the DELICATE checklist, developed by Drachsler and Greller (2016), highlighting the need to develop a sustainable framework for safe and agile LA that safeguards the data protection rights of learners.

5.2   The Relevance of the GDPR for Second Level LA: Profile Transparency

DPbD, however, also applies to the obligation of profile transparency, which is already part of the current legal framework but further clarified in the GDPR. I dare say that profile transparency is the only legal constraint that directly targets the issue of due process and the presumption of innocence in the case of data-driven applications that involve personal data. One could say that profile transparency is a form of what Citron and Pasquale (2014) termed "technological due process" in their groundbreaking article on the "scored society."[21]

Profiling refers to the patterns that result from ML operations, notably the analysis or prediction of a person's "performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements" (art. 4.4 GDPR).[22] Profile transparency requires providing information about three types of information: 1) the existence of profiling, 2) meaningful information

---

turnover. Art. 33 GDPR imposes a duty to notify the supervisory authority within 72 hours of a data breach, and art. 34 GDPR imposes a similar duty to notify the data subject whose data have been leaked, "[w]hen the personal data breach is likely to result in a high risk to the rights and freedoms of natural persons."

[20] Art. 6.4.e GDPR mitigates the duty to comply with purpose limitation in the case of re-use and re-purposing, qualifying pseudonymization as a safeguard that may contribute to legitimize additional processing. Art. 25 (DPbD) qualifies pseudonymization as an appropriate organizational and technical measure to comply with data minimization. Art. 32 (security of personal data) qualifies pseudonymization as a form of security by design.

[21] See also "procedural data due process" in Crawford and Schultz (2014).

[22] Still highly relevant: Hildebrandt and Gutwirth (2008), co-authored by computer scientists, sociologists, and lawyers.

about the logic involved, and 3) information about the envisaged consequences (art. 13.2.f and 14.2.g GDPR, 12.a DPD). Enabling individuals to exercise this right will be the real challenge, as it implies taking learners seriously, engaging them on their own terms, and helping them to critically assess the learning algorithms that define their progress (Hildebrandt, 2012). This is crucial in the case of automated decision-making, especially when such decisions have legal effect (fail or pass an exam, be admitted to a learning institution or a job position) or similarly significantly affect the concerned person. The GDPR provides data subjects with a right not to be subject to such decisions (art. 22.1), and prohibits them if they are based on so-called sensitive data (ethnicity, religion, etc.; art. 22.4). Here we see how and why the fundamental right to data protection is not restricted to privacy but explicitly addresses non-discrimination, due process, and even the presumption of innocence in its broadest sense. Even though exceptions apply, it should be clear that the GDPR is ready to address the harms that result from unwarranted automated application of LA. Learning institutions and App developers that work to integrate LA into the way we learn should take to heart that human learning is not in the first place about the management of learning behaviours, but about one of the most salient characteristics of being human. The challenge will be to make sure that learners are treated as individual persons worthy of equal respect and concern, by creating systems that afford capabilities such as intellectual independence, critical distance, and human flourishing. Having the right not to be subject to the outcome of fully automated LA decision machines may, therefore, be the most important fundamental right in the era of data-driven education.

6   CONCLUSION

Human learning cannot be reduced to Pavlov's stimulus–response hypothesis. Even animals are driven crazy by his experiments that disrespect their animal dignity, while the limited set of responses that defines their laboratory situation says little about the repertoire they develop in their "normal" ecological niche (Buytendijk & Plessner, 1936; de Waal, 2016). Simon nicely demonstrated the benefit of individual learning processes, clarifying how the tediousness of human learning may be the outcome of an optimization strategy in view of the unfathomable complexity of human society. Taking Gibson seriously, we need to acknowledge that this complexity is connected with our language capabilities and the institutional dynamics they require and produce. Our ecological niche has myriad affordances that are continuously renegotiated and attuned to the capabilities they afford while they — in turn — reshape the niche.

Core to human learning are creativity, humour, and reflection, corresponding with art and ethics, judgement, politics, and law. Creativity and humour combine, detecting the unexpected with the need to face both life and death. Reflection entails an externalized awareness of what we think we learnt, instead of succumbing entirely to unconscious learning processes. That is why teaching will remain crucial to democracy. Teaching refers to *explicit presentation and explanation* of knowledge, such that it can be the object of debate and discussion. LA may be a great way to *induce* learning processes, as it were, behind a person's back and, although this is not necessarily wrong or bad, LA must be (re-) configured in a way that allows for critical reflection *on what and on how* we learn. Therefore,

pseudonymization, encryption, and discrimination-aware data mining must be complemented with profile transparency and the effective means to object to being profiled as a specific type of learner. This is what US legal scholars call "technological due process"; as Europeans, we should start the concrete implementation of profile transparency and the right to object to automated decision-making as cutting-edge and pertinent elements of the Rule of Law.

REFERENCES

Baker, R. S., & Inventado, P. S. (2014). Educational data mining and learning analytics. In J. A. Larusson & B. White, (Eds.), *Learning analytics: From research to practice* (pp. 61–75). Springer New York. http://dx.doi.org/10.1007/978-1-4614-3305-7_4

Barnes, K. (2014). Big data in the classroom is out of control. *The New York Times,* 19 December 2014. Retrieved from http://www.nytimes.com/roomfordebate/2014/09/24/protecting-student-privacy-in-online-learning/student-data-collection-is-out-of-control

Berg, A. M., Mol, S. T., Kismihók, G., & Sclater, N. (2016). The role of a reference synthetic data generator within the field of learning analytics. *Journal of Learning Analytics*, *3*(1), 107–128. http://dx.doi.org/10.18608/jla.2016.31.7

Brey, P., & Søraker, J. H. (2009). Philosophy of computing and information technology. In A. Meijers (Ed.) *Philosophy of Technology and Engineering Sciences*, Vol. 14. D. Gabbay, P. Thagard, & J. Woods (Gen. Eds.) *The Handbook for Philosophy of Science.* Elsevier. https://dx.doi.org/10.1016/B978-0-444-51667-1.50051-3

Buytendijk, F. J., & Plessner, H. (1936). Die physiologische Erklärung des Verhaltens. *Acta Biotheoretica*, *1*(3), 151–172. http://dx.doi.org/10.1007/BF02147637

Campbell, S., & Dawson, A. J. (1995). Learning as embodied action. In R. Sutherland & J. Mason (Eds.), *Exploiting mental imagery with computers in mathematics education* (pp. 233–249). Springer Berlin Heidelberg. http://dx.doi.org/10.1007/978-3-642-57771-0_16

Christian, B. (2011). *The most human human: What talking with computers teaches us about what it means to be alive*. New York: Doubleday.

Citron, D. K., & Pasquale, F. A. (2014). The scored society: Due process for automated predictions. *Washington Law Review*, *89*(1), 1–33.

Clark, R. C., & Mayer, R. E. (2011). *E-Learning and the science of instruction: Proven guidelines for consumers and designers of multimedia learning*. San Francisco, CA: Pffeifer.

Crawford, K., & Schultz, J. (2014). Big data and due process: Toward a framework to redress predictive privacy harms. *Boston College Law Review*, *55*, 93–128.

Dalsgaard, C. (2006). Social software: E-learning beyond learning management systems. *European Journal of Open, Distance and E-Learning, 9*(2). Retrieved from http://www.eurodl.org/materials/contrib/2006/Christian_Dalsgaard.htm

de Waal, F. (2016). *Are we smart enough to know how smart animals are?* New York: W. W. Norton & Co.

Dean, W. (2015). Computational complexity theory. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/archives/fall2015/entries/computational-complexity/

Dean, J. (2016). Computer says no: Go champion loses in man v machine. *The Times* (London), 16 March 2016. Retrieved from http://www.thetimes.co.uk/tto/technology/article4709094.ece

Dennett, D. C. (1989). *The intentional stance*. Cambridge, UK: Bradford Book.

Drachsler, H., Hoel, T., Cooper, A., Kismihók, G., Berg, A., Scheffel, M., Chen, W., & Ferguson, R. (2016). Ethical and privacy issues in the design of learning analytics applications. *Proceedings of the 6th International Conference on Learning Analytics and Knowledge* (LAK '16), 25–29 April 2016, Edinburgh, UK (pp. 492–493). New York: ACM. https://dx.doi.org/10.1145/2883851.2883933

Drachsler, H., & Greller, W. (2016). Privacy and analytics: It's a DELICATE issue a checklist for trusted learning analytics. *Proceedings of the 6th International Conference on Learning Analytics and Knowledge* (LAK '16), 25–29 April 2016, Edinburgh, UK (pp. 89–98). New York: ACM. https://dx.doi.org/10.1145/2883851.2883893

Dworkin, R. (1991). *Law's Empire*. Glasgow, UK: Fontana.

Fiaidhi, J. (2014). The next step for learning analytics. *IT Professional, 16*(5), 4–8. https://dx.doi.org/10.1109/MITP.2014.78

Floridi, L. (2014). Future/the green gambit. *The fourth revolution: How the infosphere is reshaping human reality*, Chapter 7. Oxford, UK: Oxford University Press.

Foster, L. H., Watson, T. S., Meeks, C., & Young, J. S. (2002). Single-subject research design for school counselors: Becoming an applied researcher. *Professional School Counseling*, *6*(2), 146–154.

Gibson, J. J. (1986). *The ecological approach to visual perception*. Mahwah, NJ: Lawrence Erlbaum.

Greller, W., & Drachsler, H. (2012). Translating learning into numbers: A generic framework for learning analytics. *Educational Technology & Society*, *15*(3), 42–57.

Hausman, D. M., & Welch, B. (2010). Debate: To nudge or not to nudge. *Journal of Political Philosophy, 18*(1), 123–136. http://dx.doi.org/10.1111/j.1467-9760.2009.00351.x

Helbing, D. (2012). *Social self-organization: Agent-based simulations and experiments to study emergent social behavior*. Springer Berlin Heidelberg.

Hildebrandt, M. (2012). The dawn of a critical transparency right for the profiling era. *Digital Enlightenment Yearbook 2012*, 41–56.

Hildebrandt, M. (2015). *Smart technologies and the end(s) of law: Novel entanglements of law and technology*. Cheltenham, UK: Edward Elgar Publishing.

Hildebrandt, M., & Gutwirth, S. (2008). *Profiling the European citizen: Cross-disciplinary perspectives*. Dordrecht: Springer.

Hudson, B. (2005). Secrets of self: Punishment and the right to privacy. In E. Claes & A. Duff (Eds.), *Privacy and the criminal law* (pp. 137–162). Antwerp/Oxford: Intersentia.

Jarvis, P. (2005). *Towards a comprehensive theory of human learning*. London: Routledge.

Kalampokis, E., Tambouris, E., & Tarabanis, K. (2013). Understanding the predictive power of social media. *Internet Research, 23*(5), 544–559. http://dx.doi.org/10.1108/IntR-06-2012-0114

Kim, A. (2014). Wilhelm Maximilian Wundt. In E N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/archives/win2014/entries/wilhelm-wundt/

Leijnen, S. (2014). *Creativity and constraint in artificial systems*. http://www.leijnen.org/

McStay, A. (2011). *The mood of information: A critique of online behavioural advertising*. New York:

Continuum.

Mitchell, T. M. (2006). *The discipline of machine learning* (Vol. 9). Carnegie Mellon University, School of Computer Science, Machine Learning Department. http://www-cgi.cs.cmu.edu/~tom/pubs/MachineLearningTR.pdf

Nissenbaum, H. (2010). *Privacy in context: Technology, policy, and the integrity of social life*. Stanford Law Books.

Nussbaum, M. C. (2011). *Creating capabilities*. Harvard University Press.

Pavlov, I. P. (1927). *Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex*. Translated and edited by G. V. Anrep. Oxford, UK: Oxford University Press.

Pentland, A. (2014). *Social physics: How good ideas spread: The lessons from a new science*. New York: Penguin Press.

Piatetsky-Shapiro, G. (1996). *Advances in knowledge discovery and data mining* (Vol. 21). U. M. Fayyad, P. Smyth, & R. Uthurusamy (Eds.). Menlo Park, CA: AAAI Press.

Prinsloo, P., & Slade, S. (2016). Student vulnerability, agency and learning analytics: An exploration. *Journal of Learning Analytics*, *3*(1), 159–182. http://dx.doi.org/10.18608/jla.2016.31.10

Robinson, W. (2015). Epiphenomenalism. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/archives/fall2015/entries/epiphenomenalism/

Sen, A. (1999). *Commodities and capabilities*. New York: Oxford University Press.

Simon, H. A. (1983). Why should machines learn? In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning* (pp. 25–37). Springer Berlin Heidelberg. http://dx.doi.org/10.1007/978-3-662-12405-5_2

Skinner, B. F. (1976). *About behaviorism*. New York: Vintage.

Slade, S., & Prinsloo, P. (2013). Learning analytics: Ethical issues and dilemmas. *American Behavioral Scientist*, *57*(10), 1510–1529. https://dx.doi.org/10.1177/0002764213479366

Solove, D. J. (2004). *The digital person: Technology and privacy in the information age*. New York University Press.

Sousa, D. A. (Ed.). (2014). *Mind, brain, & education: Neuroscience implications for the classroom*. Bloomington, IN: Solution Tree Press.

Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.

Turing, A. M. (1950). Computing machinery and intelligence. *Mind, 59*(236), 433–460.

Vedder, A. (1999). KDD: The challenge to individualism. *Ethics and Information Technology*, *1*(4), 275–281. http://dx.doi.org/10.1023/A:1010016102284

Vender, D. (2011). Is balancing emblematic of action? Two or three pointers from Reid and Pierce. *Humana: Mente Journal of Philosophical Studies*, *1*(15), 251–270.

Verheul, E. R. (2015). Privacy protection in electronic education based on polymorphic pseudonymization. *IACR Cryptology ePrint Archive*, *2015*, 1228.

Verheul, E. R., Jacobs, B., Meijer, C., Hildebrandt, M., & de Ruiter, J. (2016). Polymorphic encryption and pseudonymisation for personalised healthcare. *IACR Cryptology ePrint Archive*, *2016*, 411.

Watson, J. B. (1930). *Behaviorism*. University of Chicago Press.

Wolf, M., & Stoodley, C. J. (2008). *Proust and the squid: The story and science of the reading brain*. Cambridge, UK: Icon.

Wolpert, D. H., & Macready, W. G. (1997). No free lunch theorems for optimization. *IEEE transactions on evolutionary computation*, *1*(1), 67–82. https://dx.doi.org/10.1109/4235.585893