A Review of Gulson, Sellar, and Webb's *Algorithms of Education*

Daniel Shussett, Villanova University, dshusset@villanova.edu

———————————————

Algorithms of Education: How Datafication and Artificial Intelligence Shape Policy
Kalervo N. Gulson, Sam Sellar, and P. Taylor Webb
University of Minnesota Press, 2022
190 pp.

In a time when the layperson theorizes at length on social media about the impact of artificial intelligence (such as ChatGPT, in recent months) on our daily lives, an academic treatment of the impact of artificial intelligence (AI) on education policy comes as a welcome respite from a deluge of banal "hot takes." *Algorithms of Education* does well to show its readers the vast range of possibilities, positive and negative, that may occur from the introduction of artificial intelligence into the educational space, extending far beyond the dime-a-dozen post about how students may write essays with the assistance of ChatGPT or how every teacher will soon be replaced by a robot. It turns out, thanks to the work of these authors, that our fears ought to be far deeper, and our optimism far greater than merely intuitive, knee-jerk responses to news media's depiction of current events.

The authors accomplish this by careful conceptual work, on both what it is to "govern" in the space of education policy and what new terrain is opened by artificial intelligence. Following this theoretical landscaping, the authors turn to three case studies, examining what AI has done to education policy *already*. While there are criticisms to be made about the authors' choice of theoretical tools and methodological framework, the book's value comes when it turns to applying theory to empirical work. The reader may be left yearning for another book on the same topic, by the same authors, but this is only due to a greater awareness of the issues, a better understanding of the technical aspects of AI's integration into education policy, and the complications of recent events (i.e., what *would* they say about ChatGPT?) than when they began reading.

The Epistemological

For the frequent reader of *SERRC* there is much truly epistemological work in the book. Some of this is done through French-inflected critical philosophy, some occurs organically through case studies, but both give any reader wondering about the impact of AI on how knowledge is socially produced and evaluated much to chew on.
One stand-out example of such work comes during the authors' discussion of an oft-mentioned "fix" for issues related to bias in machine learning's (ML) application to facial recognition, here as used in China for attendance-taking and attention-tracking. This solution is referred to as "human in the loop," where it is assumed that the involvement of a person in developing an algorithm may help to "break open the black box"[1] and reduce biased outcomes. The authors show that machine learning obfuscates the distinction between human and algorithmic input (a claim that leads the authors to the term "synthetic thought," where human and machine think together, inextricably). Thus, a human in the loop does not have the privileged epistemic vantage point that would open black boxes. They continue to

---

[1] A black box in general (not necessarily regarding AI) refers to a process where the relation between inputs and outputs of a system is opaque.

show that, even if this were true, the inclusion of a human does not *necessarily* reduce bias, given unequal demographic representation in AI research and development.

The book is similarly strong when it comes to problematizing the datafication that is inherent to AI/ML approaches to education policy. Data that often could not have been derived without non-human forms of "thought"[2] (and often, cannot be understood by human thought) is converted into more digestible representations, which the authors note is always political, and then itself used to make governance decisions. This theme is closely tied to the authors' claim that many AI/ML applications in education governance deal not with the quantification of current student performance but instead predict future student behavior, aiming to extend governance in the classroom to events which have not yet happened. Issues like these appear throughout the book, giving the epistemologically-minded reader much to learn and consider further.

One major epistemological theme of the book is the "uncertainty" that the introduction of machinic thought into education governance reveals. The authors argue that the uncertainty created by AI's entry into "evidence-based policy-making" is due to the use of abduction rather than induction (122-123). Citing Luciana Parisi, we read that "abductive reasoning is ignorance preserving," this ignorance itself giving rise to "indeterminacy" that "potentially open[s] new possibilities for decision-making and understanding policy problems" (123). For these authors, using probability to justify predictions is a more full-throated acceptance of the uncertainty inherent to prediction than we get from other methods.

Troublingly, the authors at one point claim that "the limits of prediction" that give rise to this uncertainty are "clear to data scientists," but "can run up against political desires for prediction" (125). Thus, the authors view the data science enabled by AI/ML as somehow qualitatively different than other types of predictive approaches,[3] but they also believe that the data scientists, perhaps those most desperately in search of certainty, are best equipped to recognize the shortcomings of their method. It's possible that we could excuse a lack of clarity on exactly what "uncertainty" is (compared to the contingency that generally characterizes our world) in this context if it weren't for the fact that it is precisely this uncertainty that gives the authors reason for optimism as well as pessimism for the future.

The Methodological

While *Algorithms of Education* is at its best during its three case studies, there are still issues to be had with the methodological approach. To begin with, while the reader is given ample evidence for the authors' claims in the empirical chapters, there is too little evidence offered in the beginning chapters where the conceptual work occurs. Thus, the reader is left wondering if any of the claims are correct or tendencies are actual. These chapters play off the sense that a major "shift" is coming to education policy and the interacting worlds of governance and technology more broadly, but it's not always clear what the shift consists of.

---

[2] For those philosophers of mind who may be reading, no account is given of how machinic thought is similar to, or differs from, other forms of cognition. This is understandable given the authors' approach, but someone who is in search of clarity on this might want to look elsewhere.

[3] The limitations of our inferential capacities have long been noted, whether in regard to abduction or other forms. One notable example is the "problem of induction" from Nelson Goodman's *Fact, Fiction, and Forecast* (1983).

The reader is often left questioning if the changes to education policy in light of AI/ML are actually changes in kind, not just magnitude.

Further, one of the authors' favorite methods for the case studies is the interview. One might wonder what such a subjective approach yields in the context of something as technical as AI and ML, where engineers themselves often have no thorough conception of what it is their creations do or are capable of. The authors claim this approach is taken because their forebears in infrastructure studies focus on the technical side while under-exploring how infrastructure is "represented and interpreted in stories" (65). While this is well-taken, it is concerning when the authors suggest that computer scientists who are involved in the development of AI systems are best positioned to open black boxes.

You may recall my review's earlier praise of the authors' recognition of the bias that AI engineers may harbor. But here, prior to that claim, the authors instead incorporate that very same bias into their methodological approach. In the empirical chapter, the authors cite an interview with a software developer from a large company who claims that their company actively works to represent the views of smaller companies that cannot afford to participate in pilot projects (87). Are we supposed to take this developer at their word, despite their possible naiveté and/or profit motive? Or, in more epistemological terms, why think these insiders have a privileged epistemic position at all, particularly regarding the black boxes they themselves had a hand in building, instead of an "outsider-within," following Patricia Hill Collins?[4]

The possible methodological concerns continue when one considers the case studies that the authors have selected. First, two of the case studies come from Australia and one from China. How representative are these cases of the education policy space writ-large? For the Australian cases, we are given a decent amount of context on the history of education policy in Australia, but the same cannot be said for the case study on China. In fact, in the case of China's use of facial recognition, the authors write that "a focus on facial recognition takes us in an admittedly speculative direction due to its currently limited use in education" (97). Why use this case study at all, if the goal is to make theoretical claims based on empirical work? A similar issue occurs when the authors note that the focus of their third case study is a "leader… among Australian education departments and potentially among education systems globally" (115). They continue to argue that this makes the program "an ideal site for an exploratory case study," without considering that this atypicality might make their findings non-generalizable and therefore not widely applicable.

The Theoretical

The authors are likely aware of what criticism comes next, given the amount of preemptive defense of the framework throughout the book. *Algorithms of Education* makes extensive use of the approach to philosophical inquiry (among other things) called "accelerationism." To

---

[4] Patricia Hill Collins. 1986. "Learning from the Outsider Within: The Sociological Significance of Black Feminist Thought." *Social Problems* 33 (6): s14–s32.

give the authors their due, their summary of accelerationism is perhaps the best and easiest to digest out there, especially if you usually dismiss such a theory (or are new to it). The reader is given two main reasons for this choice: first, that accelerationism captures the ubiquitous and pervasive influence of technology in every element of life, and second, that the ability of humans to predict and control the outcomes of this technological influence is extremely limited. The authors openly eschew Actor-Network Theory and the New Materialisms (37), but it is left unclear why they do so, especially when both approaches can accomplish similar views of the interweaving of humanity and technology. The inability to predict or control the future need not be described in one theoretical framework; this ought to be a readily understood fact for anyone who accepts that contingency is a general feature of the human experience. Is accelerationism a necessary starting point to argue that the combination of human and nonhuman agents leads to unpredictable and unintended consequences?

The authors create a four-part typology of the branches of accelerationism, where they align themselves with "problematization" (49). This position argues for both nihilism and creativity in light of the runaway effects of AI/ML. My argument is that the uncertainty that prompts this nihilism and creativity is always a part of the human experience, AI/ML becoming widely influential or not. Further, and returning to methodology, the authors acknowledge that not every AI/ML intervention can be explained using accelerationism, writing that "positive feedback" in one of their case studies "describes an accelerationist dynamic," but then continuing to say that "of course, education technology markets do not always grow in this way" (81). This critique doesn't mean the authors ought not use accelerationism at all, but it shouldn't be their exclusive theoretical approach if it fails to explain certain scenarios.

My contention with the authors' use of accelerationism as their primary theoretical framework, to the exclusion of others, largely centers around the resulting reduction of human agency. The authors take the acceleration of technology to be characterized by uncertainty and uncontrollability when AI/ML is introduced into the human world of education governance. They describe a "cascading process of automation" that is somehow not meant to be read as technological determinism (99). Ultimately, the approach of synthetic thought and synthetic governance that the authors advance ends up not being so much a binding of human and nonhuman agency (as they argue) but instead an overestimation of non-human agency and an erosion of human agency.

The Binding of Human and Machine

This reduction of human agency might not be such a bad thing if it were limited to a theoretical playing field. However, there are real consequences that come with starting with the assumption that humans can do little to control the contingent futures opened up by the maturation of AI. The authors write that "we do not discuss issues of privacy or data ownership; rather, we investigate what gets done when machines begin to do education governance" (96). The analysis cares not for what humans can do to address the influence of AI on the education of humans, but instead cares for what happens if/when machines become the main agents in human affairs. While the latter is admittedly important, it seems the authors choose accelerationism as a framework in part because they do not consider what humans can do in the face of technological advancement.

At another point, the authors describe the binding of the human and machinic in thought/governance as "technological somnambulism" which they then define as "willingly sleepwalking" (133). Their argument is that we cannot be certain of the outcomes that open up by sleepwalking into the AI governance of education. My counterargument is that the *willing* of their "willingly sleepwalking" shows us the path forward—there is something humans can do to create a more certain future, if only we have the strength of will. An example of this divergence comes in the conclusion, where the authors write that their "focus… is not on proposals to regulate algorithms or machines as new policy actors—while recognizing that such proposals are vitally important—but on thinking about the developments mapped in this book from the perspective of what can be done with the convergence between automated technologies and human thought and practices" (136).

Perhaps the reason for our sleepwalking is that skilled theorists with empirical evidence to support them, like the authors of this book, choose to do the latter rather than the former. Granted both approaches are important, but there is nothing inevitable about humanity's decreasing agency, just as there is nothing certain about the outcomes of technological advancement.

Further evidence of underestimated human agency comes shortly after the above cited passage, when the authors write that "infrastructure creates new sites of control, such that 'far removed from legislative processes, dynamic systems of space, information, and power generate de-facto forms of polity faster than even quasi-official forms of governance can legislate them'" (136). Here too I would argue that this need not be the case. Governments *can* intervene, even if always on the back foot, but often don't in any meaningful way. Sometimes governments may even welcome this shift, as the book shows with Australia and China, but once again this is merely a contingency of the governments we allow ourselves to have.

Conclusions

Despite some issues with the methodology used in the empirical, case-study sections, and despite over-reliance upon accelerationism as a theoretical framework, there is much to learn from and enjoy about *Algorithms of Education*. The book only occasionally falls short because of its great aspirations, which are to use pre-existing theory from a wide range of sources to help the reader come to terms with the revolution underway in the world of technology—a revolution that has expanded the world of AI/ML technology into our policy spaces, such as education, and even further into the ways communities come to know things. This expansion is rife with uncertainty, and the authors do a remarkable job of hedging their bets between optimism and pessimism, which is quite rare in the AI fervor of today.
I highly recommend the book to academics of any stripe who are curious about how AI affects real-world issues. I further recommend the book to the world beyond academia, but mostly to those interested in either AI, at a non-technical level, or education policy.

References

Goodman, Nelson. 1983. *Fact, Fiction, and Forecast: Fourth Edition.* 4th edition. Cambridge, MA: Harvard University Press.

Hill Collins, Patricia. 1986. "Learning from the Outsider Within: The Sociological Significance of Black Feminist Thought." *Social Problems* 33 (6): s14–s32.