

Journal Pre-proof

On the Paradigms of Learning Analytics: Machine Learning Meets Epistemology

Shayan Doroudi

PII: S2666-920X(23)00071-1
DOI: <https://doi.org/10.1016/j.caeai.2023.100192>
Reference: CAEAI 100192


To appear in: *Computers and Education: Artificial Intelligence*

Received date: 3 December 2022
Revised date: 4 December 2023
Accepted date: 4 December 2023

Please cite this article as: S. Doroudi, On the Paradigms of Learning Analytics: Machine Learning Meets Epistemology, *Computers and Education: Artificial Intelligence*, 100192, doi: <https://doi.org/10.1016/j.caeai.2023.100192>.

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2023 Published by Elsevier.



On the Paradigms of Learning Analytics: Machine Learning Meets Epistemology

Shayan Doroudi^{a*}

^aSchool of Education, University of California, Irvine, Irvine, USA

*Email Address: doroudis@uci.edu

Mailing Address: 401 E. Peltason Drive, Suite 3200, Irvine, CA 92617

Phone Number: (949) 824-6896

On the Paradigms of Learning Analytics: Machine Learning Meets Epistemology

Abstract

Baker, Gašević, and Karumbaiah (2021) recently proposed using a philosophical framework to classify learning analytics research in terms of four paradigms. Here I build on their theme of reflecting on philosophical differences in different approaches to learning analytics. I first present two limitations of their classification, which raise questions for how to best classify different approaches in learning analytics. In an attempt to resolve these questions, I draw upon the bias-variance tradeoff from machine learning and show how different learning analytics approaches can be viewed in terms of their positions on the tradeoff. However, I claim that this is not enough, as we must also be attuned to the underlying epistemologies behind different approaches. I claim a constructivist epistemology for learning analytics has been missing, which could, in part, explain Baker et al.'s (2021) observation that constructivist work has been relatively absent in established learning analytics research communities. Drawing on prior work from different fields, I present a sketch of what a constructivist data science philosophy might look like and how it could help advance learning analytics. Sitting at the nexus of the learning sciences and machine learning, the field of learning analytics is in a unique position to theorize about philosophy and epistemology; this paper encourages us to pursue more work in such a direction.

Keywords: research paradigms; learning analytics; epistemology; constructivism; bias-variance tradeoff

On the Paradigms of Learning Analytics: Machine Learning Meets Epistemology

Although it is often seen as an applied field, learning analytics is actually in a special position to both draw on philosophical and epistemological discussions and perhaps even contribute to them. This is because learning analytics is built on two fields, which inherently intertwine with epistemology:

1. **Learning sciences** are naturally concerned with how people come to know, a central question of epistemology.
2. **Machine learning** is concerned with how machines can learn and ultimately what can and cannot be learned in theory, which is again entangled with epistemology.

A particularly ripe territory for philosophizing about learning analytics might be seeing how concepts from one of these fields could be brought to bear on the other. Recent scholarship has begun to do this. Hennig (2002, 2003, 2010) takes a philosophy that was developed in education and the social sciences, namely that of radical constructivism, and discusses how it can be used to develop a new philosophy of data science. On the other hand, Doroudi (2020) shows that the bias-variance tradeoff—a conceptual and theoretical principle in machine learning—can be extended to make sense of ongoing debates in education. As such, both of these tools, namely constructivism and the bias-variance tradeoff, may have something to offer philosophical theorizing in learning analytics.

Despite its position to draw on diverse philosophical discourses, to date, the field of learning analytics has seemingly not devoted much attention to its philosophical and epistemological foundations. Fortunately, a step in this direction was taken when the 1st Workshop on Philosophy of Learning Analytics (“Towards a Philosophical Framework for Learning Analytics”) was hosted at the LAK 2021 conference. At this workshop, Baker and Gasevic presented an abstract of their work on the four philosophical paradigms of learning analytics. Although this was seemingly the first paper on this framework, Baker had already

been referring to this framework for several years in his course on “Learning Analytics: Process and Theory.” (Hearing about this framework had been influential in my own thinking about epistemological and methodological differences in the field.) Since the workshop, Baker et al. (2021) have taken an important and commendable step in philosophically grounding methodological differences in the field of learning analytics (and educational data science more broadly) by presenting a clear exposition of these four paradigms of learning analytics in a publication that is now widely accessible to researchers in the field. However, a clear exposition of a framework also opens doors to further scrutinize the framework—to critique some aspects of it and expand upon others. That is the aim of the present work.

In particular, I will begin by discussing two limitations of the present framework by drawing on historical examples of researchers from related fields. First, I claim their description of the “essentialist” paradigm is misleading. Second, I claim that the four paradigms largely lie on a single axis, suggesting the possibility of there being only two broad paradigms or approaches. Third, I show how the bias-variance tradeoff can help provide a theoretical rationale for the existence of these two distinct approaches. I then show that while the bias-variance tradeoff provides some clarity, to further explain distinctions between different research approaches, we must look towards epistemological differences; it is here that constructivism comes into play. Finally, I end with thoughts on why learning analytics should welcome a constructivist philosophy of data science. In the end, this paper presents a new framework for understanding philosophical and methodological differences in learning analytics and related fields. Although the framework I propose differs in important ways from that of Baker et al. (2021), we ultimately work towards the same goal:

We call for a better understanding, within each of us, of how our philosophical stances impact our research and practice. By understanding ourselves, we can understand the deep perspectives that lead to the specific choices we make, and seek analogies in the history of other work coming from that philosophical tradition. By understanding the philosophical stances that our colleagues in our broader community bring to bear on their own research and practice, we can better understand why they make the choices they make, and how these choices emerge from deep (if sometimes unspoken) philosophical commitments rather than from ignorance, laziness, or foolishness. Perhaps even more important, by understanding our colleagues’ philosophical stances, we may be able to better see what they may be able to see that we cannot see; we may be able to better learn from them; and we may be more able to craft research projects that take full advantage of what we each have to contribute to our field and to the learners we as communities strive to serve. (p. 7)

1 Baker et al.'s (2021) Four Paradigms

Baker et al. (2021) describe four paradigms of learning analytics research, by drawing on a framework of four different approaches to philosophy set forth by the philosopher Richard McKeon (McKeon, 1990). These paradigms are the entitative, ontological, existentialist, and essentialist schools of thought. Baker et al. (2021) also refer to the entitative school as reductionist and ontological school as dialectical. As they rightly suggest, McKeon is “incomprehensible to many,” and as such my discussion of these four paradigms is based on how Baker et al. (2021) have presented them, which may or may not exactly align with McKeon’s own description.¹ Below, I briefly provide the most relevant parts of Baker et al.’s (2021) description of the four philosophies for the present discussion:

Reductionism can be viewed as an approach towards understanding phenomena that consists of breaking down those phenomena into their constituent components and then analyzing the relationship between those components.

Dialectic is the key method of the ontological school. This school adopts the goal of understanding phenomena as wholes, or understanding systems as systems, where components cannot be properly understood without understanding the whole system...

Existentialism views reality as fundamentally individually constructed and therefore asserts that phenomena should be understood as the participants themselves understand them and that these understandings are irreducibly valid (terms such as philosophical constructionism and phenomenological understanding are sometimes seen)...

Essentialism... states that meaning is inherent in the universe. This viewpoint underpins ideas such as a common and universal core of mathematics. It is seen also in perspectives that argue for the “unreasonable effectiveness of data” as justification for rejecting interpretable modeling methods (Halevy et al., 2009), where direct modeling of reality is seen as sufficient and no attempt at theory or explanation is needed (or, indeed, desired). (p. 4)

Importantly, these four schools are conceptualized on two axes, with the entitative/reductionist and ontological/dialectical schools being on opposite ends of one axis and existentialism and essentialism being on opposite ends of the other axis. Baker et al. (2021) then describe how different communities of relevance to learning analytics tend to represent one or more of these positions. According to their characterization,

¹Indeed, in McKeon’s (1990) framework each mode of thought has an associated type of interpretation and a method of inquiry, but McKeon admits that philosophers often use interpretations and methods from different modes of thought. For example, McKeon labels Hegel as using the dialectical method but using an entitative interpretation. This is counterintuitive since Hegel is regarded as being anti-reductionist (Kabeshkin, 2021), and so it challenges Baker et al.’s (2021) notion that dialectical is interchangeable with ontological and that reductionist is interchangeable with entitative.

at least in the early days of these conferences, Learning Analytics and Knowledge (LAK) was primarily represented by the ontological school, the Educational Data Mining (EDM) conference was primarily represented by the reductionist school, the Learning @ Scale (L@S) conference was represented by essentialism, and the International Conference on Quantitative Ethnography (ICQE) was represented by existentialism. Baker et al. (2021) use the term “learning analytics” both to refer to the narrower learning analytics research community (as represented by the Society for Learning Analytics Research and its LAK conference) as well as the broader set of research communities (including LAK, EDM, L@S, and ICQE) that focus on applying data science to understand learning and improve education. To avoid ambiguity, I have chosen to use the term learning analytics in the broader sense² and I use LAK to refer to the narrower community.

2 An “Essential” Question: Where Does Deep Learning Fit?

One issue with the four paradigms as presented above is that it seems to conflate different phenomena. In this section, we investigate the paradigm of essentialism in particular. Baker et al. (2021) mention one manifestation of the essentialist position being “This viewpoint underpins ideas such as a common and universal core of mathematics.” This is the essentialist philosophy in education. However, this brand of essentialism does not necessarily seem to particularly inform any of the four research communities that Baker et al. (2021) focus on, including the Learning @ Scale community, which they claim aligns with essentialism. Baker et al. (2021) further claim that essentialism is also manifested “in perspectives that argue for the “unreasonable effectiveness of data” as justification for rejecting interpretable modeling methods (Halevy et al., 2009).” It is this style of work from the L@S community that they primarily focus on in establishing its connection to essentialism. However, it is not clear why black-box machine learning approaches are characteristic of essentialism. For example, what does the use of atheoretical machine learning algorithms have to do with educational essentialism?

In fact, educational essentialism emphasizes the use of curricula that are rooted in certain principles, theories, or values about what should be taught rather than data-driven evidence. A “common and universal core” is also not personalized or individualized for different learners. This is in stark contrast to data-driven approaches that are atheoretical and afford the ability for personalization and differentiation. In fact, one of the affordances of using machine learning in education is for personalization based on data.

Moreover, Peter Norvig, the second author of the paper on the “unreasonable effectiveness of data,”

²I could instead have used the term “educational data science,” but this refers to an even broader movement to use data science to advance education research, much of which is not directly concerned with learning. While much of this paper applies to educational data science and data science in general, I have chosen to use the term learning analytics to ground the conversation in research directly concerned with studying learners and learning.

is a strong proponent of data-driven theories to deal with the complexity of the real-world in place of the simple, elegant models which have sufficed us in earlier centuries (Halevy, Norvig, & Pereira, 2009; Norvig, 2012). In the context of computational linguistics, Norvig places value in large statistical models that can learn patterns in language (and which have become increasingly popular in recent times); he contrasts his view with that of the famous linguist Noam Chomsky, who believes language should be explained by relatively simple theoretical models (Norvig, 2012).³ Interestingly enough, Chomsky has been termed “the intellectual ancestor of Essentialism,” a school of thought in linguistics. Of course, names can be misleading and essentialism in linguistics may not be the same one that McKeon is referring to. Indeed, Chomsky’s essentialism seems to have roots in Plato (Norvig, 2012), who is often also associated with the term “essentialism” (Politis, 2021), despite the fact that McKeon (1990) considers Plato to be the progenitor of the ontological school⁴. Regardless, I believe it is easier to make the argument that Chomsky is an essentialist—by most definitions—than Norvig.

Indeed, Šekrst and Skansi (2022) point out that while essentialism may align with certain kinds of machine learning or feature engineering that utilize a (small) number of essential features to a prediction task, this is not true of deep learning: “deep-learning feature engineering does not have to correspond to some natural kinds or essential properties: it is not really essentialism, but a certain kind of accidentalism” (p. 187). Moreover, while it might seem like deep learning does away with the need for human interpretation (in contrast to existentialism), it is really transferring the job of interpretation from people to neural networks that have their own (not necessarily well-understood) interpretative lenses. Thus, fitting two different neural network architectures to the same dataset can result in two models that make different predictions in many cases and even fitting the same neural network architecture to two different datasets sampled from the same underlying distribution may result in models that make different predictions. Thus deep learning does not really get at “meaning [that] is inherent in the universe” even if some of its adherents think that it does.

So now we return to the question. How would we characterize deep learning in terms of the four paradigms? I am not sure if there is a clear answer according to McKeon (who was primarily concerned with classifying philosophers), but regardless, I think we can take a different approach. Before answering this question, I turn to another issue with the framework; in the subsequent section, I present a solution to answer both dilemmas.

³Interestingly, this distinction not only applies to how researchers model language, but also how people acquire language; Chomsky believes the mind has a universal, and largely innate, language acquisition device, while Norvig and others put stock in the idea that the mind can gradually learn the rules of language by being exposed to it in a similar way to complex machine learning algorithms.

⁴Compare this with Pepper (1942)’s categorization of four major philosophical schools of thought where he considers Plato and Aristotle as belonging to the same school, “formism,” which could be seen as analogous to essentialism.

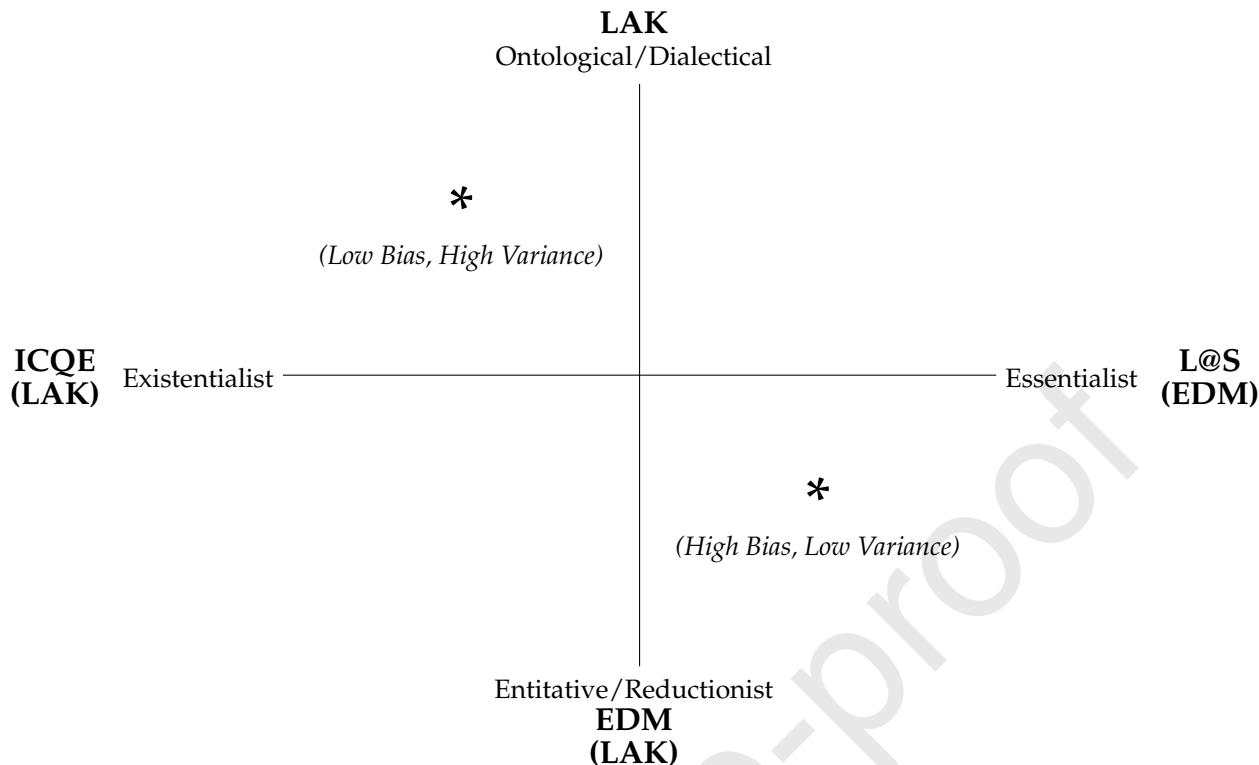


Figure 1: A reproduction of Baker et al.'s (2021) diagram depicting four paradigms of learning analytics on two axes. The text is taken directly from Figure 1 of their paper except for the asterisks and italicized text in the upper left and bottom right quadrants.

3 Two Axes or One?

The framework that Baker et al. (2021) adopt from McKeon is composed of two axes: (1) reductionism vs. holism, and (2) whether knowledge is primarily a subjective construction by the individual (existentialism) or whether it is an objective property of the world (essentialism). Their two-axis representation is reproduced in Figure 1 (ignoring the asterisks and italicized text which are my additions and explained below). The use of two axes seems to suggest that an individual researcher or research community would adopt a position in this two-dimensional space (e.g., I can be a moderate existentialist that exclusively uses reductionistic methods or I can be an extreme essentialist that uses a mix of reductionistic and holistic methods). However, the framework they present suggests each individual typically “prefers to work in one, or perhaps two, of these perspectives” (p. 4). Moreover, Baker et al. (2021) align each of the four learning analytics research communities (at least when they started) with primarily one of the four paradigms, situating the research communities at the end points of the axes rather than anywhere else in the quadrants. This raises the question of whether it is meaningful to have two axes.

I suggest that researchers will actually *tend to be* in one of two quadrants in the two-dimensional space.

Specifically, I argue below that researchers will tend to gravitate towards either the top left quadrant or the bottom right quadrant, as indicated by the asterisks in Figure 1. That is, there is a strong correlation between reductionist methods and essentialism and there is a strong correlation between dialectic methods and existentialism. In fact, there are already clues of this in Figure 1. As shown in parentheses underneath the primary research community identified with each paradigm, Baker et al. (2021) identify that over time, LAK has secondarily adopted the existentialist paradigm and EDM has secondarily adopted the essentialist paradigm (although this could be misleading based on the classification of deep learning as essentialist). They also suggest that LAK has recently adopted more entitative approaches, which goes against the clustering I have suggested, but this is attributed to “many researchers who had previously published at AIED or EDM beginning to publish their work at LAK” (p. 5).

To give concrete evidence for the correlation between the different paradigms, I turn towards intellectual ancestors of some of the research approaches used in learning analytics. Baker et al. (2021) interpret existentialism as asserting that individuals construct knowledge (or reality) for themselves. This idea is rooted in the work of Jean Piaget, the founder of constructivism in psychology. What might be less known is that Piaget was very much influenced by a holistic way of thinking (Burman, 2020). In his autobiography, Piaget (1952) claims that when he was around 18 years old, he reached an important conclusion that would later pervade his work on children’s thinking:

In all fields of life (organic, mental, social) there exist “totalities” qualitatively distinct from their parts and imposing on them an organization. Therefore there exist no isolated “elements”; elementary reality is necessarily dependent on a whole which pervades it. (p. 242)

Hence, Piaget’s work can be simultaneously described as following the existentialist paradigm and following the ontological (holistic) paradigm. Indeed, dialectics is an important aspect of Piaget’s theory of equilibration, and Piaget is often compared to Hegel (Kitchener, 1980).

Moreover, Seymour Papert, a student of Piaget and founder of constructionism, not only emphasized an individual’s knowledge construction, but also gave importance to the social and cultural context in which knowledge construction takes place and the embodied nature of learning (Papert, 1976, 1980). Situated and embodied cognition, which advocate against reducing cognition to a set of elements or operations located in the mind, are associated with the ontological paradigm. Papert’s approach to artificial intelligence has also been labeled as a type of “emergent AI,” which as the name suggests takes a more holistic, rather than reductionistic, approach to modeling AI (Turkle, 1991).

To turn to a scholar in the field of learning analytics, David Williamson Shaffer, who was a student of Papert, developed the technique of epistemic network analysis with his colleagues (Shaffer, Collier, &

Ruis, 2016; Shaffer et al., 2009). Baker et al. (2021) categorize epistemic network analysis as an existentialist method, presumably due to the influence of Papert and Piaget on Shaffer's work. However, like the work of Papert and Piaget, epistemic network analysis would also fall under the ontological paradigm: "learning is defined not by the *possession* of isolated bits of knowledge and other competencies but by the *structure of connections among them*" (Shaffer et al., 2016, p. 11).

Turning to the debate between Anderson, Reder, and Simon (1997) and Greeno (1997) which Baker et al. (2021) referred to as representing entitative vs. ontological positions, it is interesting to note that Anderson, Reder, Simon, Ericsson, and Glaser (1998) also debated against constructivist approaches to education research. Furthermore, Anderson et al. (1998) explicitly noted connections between situated learning and constructivism, which they found surprising on some level:

The alliance between situated learning and radical constructivism is somewhat peculiar, as situated learning emphasizes that knowledge is maintained in the external, social world; constructivism argues that knowledge resides in an individual's internal state, perhaps unknowable to anyone else. However, both schools share the general philosophical position that knowledge cannot be decomposed or decontextualized for purposes of either research or instruction. (p. 235)

The last sentence brings us back to learning analytics research and can also help see the link between the entitative and essentialist paradigms. A popular approach in EDM, specifically in constructing predictive models of student learning, is to decompose knowledge into component parts. This typically reflects both the entitative and essentialist paradigms as it supposes there is a core body of knowledge or curriculum that all students must learn, that students generally learn that knowledge in the same way, and that knowledge can be neatly decomposed into components. On the other hand, the ontological situativists and existentialist constructivists push back against these assumptions by claiming students construct their own knowledge in idiosyncratic ways, a student's knowledge depends on the socio-cultural context in which it was learned, and the knowledge that a student constructs is a complex interconnected whole that cannot be readily reduced into components (see e.g., Greeno, 1997; Shepard, 1991).

Thus it appears that at least in many cases, including in the work of some important education researchers and psychologists, the ontological and existentialist perspectives are aligned and the entitative and essentialist positions are aligned. But why is this the case? As Anderson et al. (1998) noted, "The alliance between situated learning and radical constructivism is somewhat peculiar" because they seemingly have different emphases. To understand this, we now turn to the bias-variance tradeoff.

4 The Bias-Variance Tradeoff

In machine learning (specifically supervised learning), we try to estimate a function using data sampled from the true function. The quality of the estimated function will depend on the particular sampled dataset that we have. However, to assess the quality of a machine learning algorithm, we can look at how well that algorithm will perform on average, over the randomness of the sampled data. Any machine learning algorithm is susceptible to two kinds of errors that affect our ability to derive a good estimate.⁵ *Bias* refers to how far our estimated function is on average from the true function. *Variance* refers to how much the estimated function varies as a function of the sampled data (e.g., an algorithm that always outputs a similar function, regardless of the dataset, would be low in variance while an algorithm whose output is highly dependent on the particularities of the sampled data would be high variance). The bias-variance decomposition is a mathematical theorem that states that the mean squared error in an algorithm's ability to fit a function is the sum of the algorithm's bias squared and the algorithm's variance. Typically as machine learning algorithms increase in complexity—that is, the complexity of the functions they use to fit the data—the bias of the algorithm decreases but the variance increases; this is referred to as the bias-variance tradeoff.⁶

Doroudi (2020) showed how the bias-variance tradeoff can be formally extended to apply to different approaches and theories in education research. For example, the cognitivist or information-processing approach tends to give precise (often computational) theories of cognition. Because these theories are precise, they tend to be low variance. However, this approach also tends to ignore or downplay individual differences, the importance of context, the role of the body in cognition, etc. Hence, at least its opponents would claim it is biased. On the other hand, situated and constructivist approaches try to model the variation in human cognition and learning across time, place, and individuals. These theories are therefore often less precise and harder to generalize; hence they tend to exhibit high variance. However, they try to make up for the increase in variance by modeling the complexities of cognition and learning in the real world, which reduces bias.

Interestingly, the bias-variance decomposition applies to any situation where one is trying to learn some

⁵There are actually three kinds of error when there is random noise in the generated data. However, true noise is a property of the phenomenon being studied, not the methods we use to estimate it. What we often call noise is actually our inability to fully understand the complexity of the phenomenon being modeled, which may result in variance. Variance is actually sometimes referred to as noise (Kahneman, Sibony, & Sunstein, 2021).

⁶Recent machine learning research has shown that, in some cases, for complex models like neural networks that are overparameterized (i.e., have more parameters than the number of data points), as the number of parameters increases, the variance actually decreases (Belkin, Hsu, Ma, & Mandal, 2019; d'Ascoli, Refinetti, Biroli, & Krzakala, 2020). This suggests that highly complex neural networks actually do not always have high variance, contradicting the classical machine learning paradigm. Nonetheless, it is worth noting that this phenomenon is still being actively researched and tradeoffs may still exist in terms of other notions of model complexity and a tradeoff still exists when the noise-to-sample ratio is high enough (Xue, Whitecross, & Mirzasoleiman, 2022), which I expect is the case in many learning analytics applications. Therefore, in this paper I will consider complex models like neural networks as relatively high variance approaches.

phenomenon to the best of one's ability. Hence, it applies to how researchers try to develop theories about learning (e.g., in whether we apply a quantitative or qualitative method or whether we use a simple or complex machine learning model). But it also applies to learners as they learn, meaning that both how people learn and the models we use to describe how people learn tend to either be relatively biased or relatively high variance. For example, Chomsky's language acquisition device assumes human learning is relatively biased but low variance while a neural network model of language acquisition assumes human learning has relatively low bias but is high variance. Therefore, in the field of learning analytics in particular, the bias-variance tradeoff appears both in the *process* of studying learning (i.e., our research methods) and in the *product* of studying learning (i.e., the properties of the models we develop).⁷

Thus, the bias-variance tradeoff can simplify our distinctions between different paradigms by situating them on a single axis between two extreme poles of high bias and high variance. As indicated by the italicized text in Figure 1, the entitative and essentialist paradigms tend to be biased but have low variance; the ontological and existentialist paradigms tend to exhibit low bias but high variance.⁸ Approaches can be high in both bias and variance, but such approaches are bad and would be less likely to be pursued. Ideally, researchers should aim for approaches that are low in both bias and variance—and some methods might be better attuned to achieving that than others—but as the term “tradeoff” suggests, it is typically necessary to make choices about whether to welcome more bias or more variance. This tradeoff gives rise to the different prominent approaches in education research (in general) and in learning analytics (in particular).

We now return to the question of where deep learning fits into this picture. This is easy to answer now as deep learning is generally considered a high-variance approach in machine learning. This would suggest that deep learning is aligned with the ontological and existential approaches, and in opposition to essentialism. This may seem at odds with how deep learning is used today in a largely atheoretical way; many constructivists and situativists might be wary of machine learning engineers trying to use deep learning to definitively answer educational questions. However, there have been historical connections between connectionism, the intellectual predecessor to deep learning, and situated and constructivist theories (Bereiter, 1991; Quartz, 1999; Winograd, 2006). In an article in *Educational Researcher*, Bereiter (1991) made this connection, and its relevance to education research, quite clear:

At least since Rousseau, there has been a strain of educational thought opposed to the clas-

⁷I contend that the two are correlated; that is researchers who use high bias approaches tend to develop models of learning that have high bias while researchers who use high variance approaches tend to develop theories of learning that have high variance. However, this need not necessarily be the case. For example, a researcher can intuit (without any empirical data) a very complex, high-variance model for how people learn. This researcher used a very biased research methodology (based solely on their intuition) but developed a high-variance theory. The interplay between the bias-variance tradeoff in the process and product of developing learning theories would be an intriguing area of further inquiry.

⁸I claim that this statement can be formalized to apply to these four paradigms broadly, whether in the context of philosophy, psychology, science, or learning analytics; however, that is beyond the scope of this paper. For some insights on how to formalize this claim, see Doroudi (2020).

sical, rule-based view of learning and cognition. It has often appealed to biological concepts of growth, emergence, and organicism or to social and cultural concepts and has emphasized imagination, spontaneity, feeling, and the wholistic character of understanding. One of its hallmarks has been opposition to the teaching of rules. . . . Grounded in conviction and experience, however, it has lacked intellectual tools for recognizing and solving its internal problems. Connectionism could provide the missing scientific basis for this approach to education. On the one hand, it would provide a way of formulating and perhaps testing intuitions about the organic character of human cognition. On the other hand, it would permit educators to find a way, within an organic view of cognition, to deal pragmatically with the use of rules in teaching, as suggested in the preceding points. In the absence of such a scientific basis, there is a tendency for valid intuitions about the nature of cognition to be generalized into dogmatic opposition to rule teaching. (p. 15)

Bereiter's (1991) suggestion to use connectionism as a scientific or computational basis for the situated and embodied cognition movements seemingly never came to fruition. Today, the use of deep learning in a largely atheoretical fashion may be high variance but it is not obviously linked to ontological or existentialist approaches to studying learning. Nonetheless, it seems that the learning analytics community, given its interdisciplinary background, might be in a particularly strong position to consider the use of connectionist techniques to model situated, embodied, and constructivist theories of learning.

Categorizing different communities adjacent to learning analytics in terms of their position on the bias-variance tradeoff does not suggest that it is not useful to also view these approaches in terms of multiple paradigms as suggested by Baker et al. (2021). However, dividing up approaches using a philosophical framework like McKeon's can be difficult and possibly result in mischaracterizing some approaches as highlighted above. On the other hand, the bias-variance tradeoff gives a very clear and formal way of delineating different approaches by relating the approaches to concepts in a mathematical theorem; see Doroudi (2020) for more details. Moreover, the bias-variance tradeoff gives us a rationale for why these different approaches exist and why "both sides" of different debates might be reasonable: each tries to minimize a real source of "error." But the bias-variance decomposition also points to the ideal of minimizing both bias and variance; learning analytics researchers could thus look for creative ways of pragmatically combining approaches that are traditionally associated with different paradigms towards this end.

One conundrum still remains. Given that deep learning can be used in ways that are not very compatible with ontological or existentialist approaches, one might ask whether the bias-variance tradeoff as a single axis is a useful enough heuristic to distinguish between different approaches to learning analytics and adjacent

research communities. I claim there *is* something missing in the bias-variance tradeoff, and that lies in different epistemologies in learning analytics, which we turn to next.

5 Epistemological Differences

Although various theories (e.g., situated learning, embodied learning, and constructivism) and methods (e.g., deep learning and qualitative methods) can all be regarded as high variance, it would be an oversimplification to claim that all of these theories and methods are always compatible. A key distinction that we have glossed over lies in epistemology. Here, a brief survey of different epistemological traditions may be useful. While the discussion here is primarily centered on the epistemology of researchers, which in turn guides research methods; I reiterate that epistemology is also central to learning analytics in that to study learning we must make epistemic assumptions about the source of *learners'* knowledge as well. To be self-consistent, it makes sense for a learning analytics researcher who makes certain epistemological assumptions about learners to use those same assumptions to guide their choice of research methods, but this may not always be the case. I hope that the following discussion can help some learning analytics researchers better see how they can align their epistemic assumptions about learners with their own research epistemology.

Two competing epistemological traditions that were prominent prior to the twentieth century were *rationalism* and *empiricism*. Rationalists claim that our knowledge primarily comes from the human faculty of reason, while empiricists argue that our knowledge is primarily rooted in sense experiences (Markie & Folescu, 2023; Phillips & Burbules, 2000). Rooted in empiricism, *positivism* is an epistemological position that asserts that we can only know the truth by verifying ideas using empirical evidence (i.e, our senses) and logical deduction. Positivists (such as the Auguste Comte, the nineteenth-century philosopher who is often credited as originating this position) argued that “the method of science (the “positive” method) was *the* method of arriving at knowledge” (Phillips & Burbules, 2000, p. 8, emphasis in original). Hence positivists rejected metaphysical claims or any claims that could not be verified with empirical data as meaningless (Phillips & Burbules, 2000). While rationalism and empiricism/positivism are often seen as perpetually in conflict, as Phillips and Burbules (2000) describe it, both epistemologies are *foundationalist* and share similar epistemological difficulties that other epistemologies, which we discuss below, attempt to address.

Returning to the debate between Chomsky and Norvig, I characterized their positions as opposing one another and thus falling under competing paradigms in McKeon’s framework. Indeed, they also differ in terms of epistemology as Chomsky is often described as a rationalist (Markie & Folescu, 2023; Norvig, 2012) while Norvig’s view can be characterized as empiricist (Childers, Hvorecký, & Majer, 2021). However, in

another sense, the two positions may not be so radically different in terms of epistemology:

Even though they represent antagonistic views within the area of artificial intelligence, discourse on research paradigms... allows the assumption that both positions are positivistic. One claims for modeling the world objectively a priori - Chomsky's notable work on the principled modeling of linguistic structures - and the other assumes that prediction based on statistics is explanation - Norvig's pleading to accept the success of the probabilistic approaches that purely rely on large and ever growing amounts of raw digital data. (Luczak-Rösch, 2013, p. 3)

As per the discussion above, I would argue that it is more apropos to label the two positions as foundationalist rather than positivist. However, Luczak-Rösch's (2013) characterization is indeed consistent with discussions of research paradigms in education research and social sciences, which tend to gloss over the distinction between rationalism and empiricism, clumping foundationalist positions together under the label of positivism (Guba & Lincoln, 1994).

In the twentieth century, many rejected some of the tenets of positivism (and foundationalism more generally), leading to *post-positivism*, the most famous rendition of which is Karl Popper's *falsificationism* (Phillips & Burbules, 2000). Popper claims that scientific theories are always tenuous and hence cannot be verified; rather they can always be falsified (and hence a hypothesis that is not falsifiable should be discarded). Science progresses by finding theories compatible with evidence, but allowing for the fact that those theories may eventually be falsified with further evidence, leading to better theories over time. Most scientists likely hold onto a positivist or postpositivist epistemology (whether they know it or not!).

On the other hand, existential theories lend themselves to a *constructivist* epistemology. Constructivism as an epistemology—which is often called “radical constructivism” (von Glasersfeld, 1991)—is not the same as constructivism as a theory of learning or a pedagogical approach, but the three are related to one another. All of them are rooted in the work of Jean Piaget, although a constructivist epistemology can also be found in much earlier philosophical writings (von Glasersfeld, 1991). As an epistemology, radical constructivism states that since every individual must construct their own knowledge, we each construct our own subjective reality that does not necessarily “represent” any external reality. As von Glasersfeld and Cobb (1983) have described this position:

In our habitual way of thinking and speaking, “to know something” is intended to mean that one possesses a conceptual structure that matches some part or aspect of something that is considered ontologically real. From the constructivist perspective, this is an impossibility, and we therefore replace the notion of match with the notion of fit... From the radical constructivist perspective, “knowledge” fits reality in much the same way that a key fits a lock that it is able

to open. The fit describes a capacity of the key, not a property of the lock. When we face a novel problem, we are in much the same position as the burglar who wishes to enter a house. The “key” with which he successfully opens the door might be a paper clip, a bobby pin, a credit card, or a skillfully crafted skeleton key. All that matters is that it fits within the constraints of the particular lock and allows the burglar to get in. (p. 220)

This epistemology is largely a consequence of adopting the existentialist paradigm, but since such a paradigm is closely connected to the ontological approach, ontological researchers often also adopt such an epistemology. Some ontological researchers choose to adopt a pragmatic approach, which ignores epistemological differences; this is also true of Cobb, who shifted from being a radical constructivist to a pragmatist (Cobb, 2007). As evident from the burglar example above, a constructivist epistemology lends itself to pragmatism.

So where does deep learning fit in terms of epistemology? While deep learning and constructivist research are both high-variance approaches, deep learning is typically not applied to problems in a way that is compatible with a constructivist epistemology. In fact, philosophers have recently linked deep learning with empiricism (Buckner, 2018; Childers et al., 2021; Skansi & Kardum, 2021). Childers et al. (2021) explain how modern deep learning can be situated in a long line of empiricist thought about the human mind dating back to B. F. Skinner and Willard Van Orman Quine, both of whom also engaged in debates with Chomsky (decades before his debate with Norvig).⁹ While data scientists may not self-identify as positivists, as I argue below, most data science research is not compatible with a constructivist epistemology. This could explain Baker et al.’s (2021) observation that “Existential researchers have until this point generally not found a strong home in any of the existing learning analytics communities” (p. 5). The authors then state that the International Conference on Quantitative Ethnography formed to give a home to some existential learning analytics researchers; however, they note that *constructionist* research (as a specific kind of *constructivist* research) was largely absent from the first ICQE conference, perhaps

due in part to the methodological focus of quantitative ethnography — the use of epistemic network analysis and related methods — which represents one take on how to do existentialist work with large-scale educational data but is clearly not the only such method. (p. 5)

Indeed, while epistemic network analysis provides one research method that could be seen as being compatible with a constructivist epistemology, it is not the right method for many problems that existential

⁹As an aside, to the long list of people who have engaged in debates with Chomsky we can also add Piaget and Papert (Piattelli-Palmarini, 1980). Interestingly, while Papert’s relationship with neural networks is complicated (Papert, 1988), he actually used the perceptron (a predecessor to modern neural networks) to argue against Chomsky’s claim that learning mechanisms could not account for language acquisition (Piattelli-Palmarini, 1980). Moreover, this very debate between Papert and Chomsky inspired Yann LeCun, one of the pioneers of modern deep learning, to begin pursuing research in this area (Fisher, 2016). Hence, we have another connection between constructivism and deep learning, although not in terms of research epistemology.

researchers, including constructionists, tackle.

Hence, I claim what is missing from learning analytics and what may have resulted in less constructivist work in the field, is the lack of a coherent constructivist philosophy of data science, one that could support a variety of data science methods.

6 A Constructivist Data Science Philosophy

In a typical data science task (particularly in supervised learning), our goal is to maximize the accuracy of a model, or how well that model fits the data. A model with higher accuracy is deemed better than a model with lower accuracy. A positivist approach to data science research might presuppose that we aim to find the *correct* model, and once we have a model that fits the data sufficiently well, we have found the correct model. A more nuanced post-positivist (or falsificationist) position would aim to progressively refine our models; as we find models with higher accuracy we falsify our previous models, leading us to closer and closer approximations of reality.

I suggest that much of the research in data science in general, and learning analytics in particular, follows an epistemology that lies between positivism and post-positivism. To be fair, a lot of data scientists are *pragmatic*, in that they do not care if their model accurately captures reality; rather they just care if their model has sufficient predictive accuracy to be useful, but even then solely focusing on accuracy can be misleading. Indeed, it seems unlikely that many researchers would take the extreme positivist stance and claim they have found the correct model, given the innumerable number of methods and parameters that can be used to fit data-driven models and given that the research community seems to always find ways to incrementally improve the accuracy of models in competitive tasks. Nonetheless, although researchers may not explicitly state that they are using a “correct” model, they may implicitly assume this or make inferences as though their model were true. For example, when using a linear regression model, a researcher may make a statement such as “an increase of 1 hour in time spent using a tutoring system is associated with an increase of 0.2 on post-test scores.” While this hypothetical researcher was careful to avoid making a causal claim (which is not true of all researchers!), their statement still suggests that there is a linear relationship between the variables, even though this was simply a modeling assumption. (In fact, we *know* the relationship cannot be linear in this case because post-test scores must be bounded by the minimum and maximum possible scores.)

As another example, when discussing fitting knowledge component (KC) models, Alevan and Koedinger (2013) seem to suggest that models with a high enough predictive accuracy have psychological reality:

for a KC model to have psychological reality... it means the model can be used to make accu-

rate predictions of a given student's performance on future problems based on that student's performance on past problems. Put differently, it means that the transfer predictions that are implied by the model are actually observed in data about student performance, typically, tutor log data (p. 167)

Finally, Doroudi and Brunskill (2017) have shown that when a student model is misspecified, interpreting the parameters of the model might lead to misleading conclusions. For example, many researchers have noticed that Bayesian knowledge tracing (BKT) models often have high slip probabilities (i.e., the probability that a student who knows a KC would "slip" and answer it incorrectly). The most immediate interpretation of this is that students slip up often. However, Doroudi and Brunskill (2017) showed using simulations that high slip probabilities could actually result from model misspecification (e.g., when a student can have varying degrees of knowing the KC, rather than a binary knowledge state).

It may be for such reasons that Box (1979) famously wrote "all models are wrong but some are useful" (p. 202, capitalization removed). A constructivist epistemology requires shifting from model *correctness* to model *usefulness* or *robustness*. Although we may still optimize model fit, we only value the models insofar as they make predictions and inferences that are useful without assuming the models are correct. Just as a burglar can get away with any key that fits the lock, perhaps we can get by with a model that is robust to our assumptions about learning and reality. Of course, if a model suffers from bad accuracy, we should generally not expect it to be useful and should not use it in practice; but model accuracy (beyond some threshold or in comparison to other models) is not sufficient to guarantee a useful model. Even though many learning analytics researchers may agree with this, I suspect few researchers in learning analytics and related fields adopt a fully constructivist epistemology. Indeed, few researchers even appear to explicitly reflect on or state their epistemological stance when conducting data science research.

Hennig (2002, 2003, 2010) proposed such a constructivist epistemology for data science in general, but the number of citations to his work in this area suggests that his call was not heeded. Perhaps learning analytics researchers who already take a constructivist approach to studying learning are in a strong position to adopt such an epistemology in data science. Hennig (2003) points out some thought-provoking consequences of not taking this constructivist view of statistical modeling:

It is possible that individuals or social systems are influenced so much by the formalized discourse and the repercussion of the models, that they reduce their reality to the formalized aspects and, in the most extreme case, that they are no longer able or willing to observe deviations. That is, models can match an observer's reality, but this does not say that they fit any observer-independent reality. Instead, it says something about the reduction of the perceptions of the

observers. To formulate it provoking: While wrong models may be useful, “correct” models are dangerous . . . There will be modeled and non-modeled aspects. The act of modeling highlights the modeled aspects and weights down the others. The non-modeled aspects are always in danger of vanishing from the scientific discourse. Thus, conscious ignorance is crucial for a reasonable work with models. By this I mean that the non-modeled aspects (including deviations between model and the researchers’ perceived realities) either should be kept explicitly in the discussion, or that a clear and conscious decision is made that these aspects are not important with respect to the problem at hand. (pp. 238-239)

For example, an entitative researcher using a constructivist epistemology might be explicit about the fact that their model ignores social interactions or context-specific variables; they can then argue either why they believe those variables are not needed to make the kinds of inferences they care about *or* what limitations there might be as a result of not including those variables. On the other hand, an ontological researcher using a constructivist epistemology might be explicit about how their social network analysis does not model the precise low-level mechanisms of learning and reason about whether different cognitive mechanisms would impact the reliability of their findings.

As mentioned, seeking robust models or inferences may be a useful objective for constructivist data scientists. One strategy for doing so was articulated long ago by the quantitative sociologist, Duncan (1975):

Analysis of specification error relates to a rhetorical strategy in which we suggest a model as the “true” one for sake of argument, determine how our working model differs from it and what the consequences of the difference(s) are, and thereby get some sense of how important the mistakes we will inevitably make may be. (pp. 101-102)

For example, such a strategy was used by Doroudi and Brunskill (2017) to investigate what might happen when a student model is misspecified.

A version of this strategy was also articulated by Levins (1966), an ontological researcher in biology:

We attempt to treat the same problem with several alternative models each with different simplifications but with a common...assumption. Then, if these models, despite their different assumptions, lead to similar results, we have what we can call a robust theorem that is relatively free of the details of the model. Hence, our truth is the intersection of independent lies. (p. 423)

Interestingly, this quote resonates strongly with a strategy that Minsky and Papert (1971) took in AI:

We are dependent on having simple but highly developed models of many phenomena. Each model—or “micro-world” as we shall call it—is very schematic. . . we talk about a fairyland in

which things are so simplified that almost every statement about them would be literally false if asserted about the real world. Nevertheless, we feel they are so important that we plan to assign a large portion of our effort to developing a collection of these micro-worlds and finding how to embed their suggestive and predictive powers in larger systems without being misled by their incompatibility with literal truth.

Like Levins, Minsky and Papert acknowledged that simple models are false, but their constructivist attitude did not push them away from utilizing such models. By acknowledging that such models are false while not “being misled by their incompatibility with literal truth,” we can carefully gain insights about the phenomena we are investigating. Of course, this does not give us a direct recipe for how to develop constructivist research methods in learning analytics, but it does give some hope that constructivists (especially constructionists whose work is intellectually rooted in Papert’s) can help articulate such methods in ways that are consistent with a constructivist epistemology. To provide a concrete example, agent-based modeling is one methodology that aligns with this approach but is not currently popular in learning analytics. Agent-based models are a popular method for constructionists (Abrahamson, Blikstein, & Wilensky, 2007) as they are useful both as models that scientists can use to study phenomena (including learning) and as models that learners can use as they construct their own understanding of scientific phenomena. Indeed, Papert (1980) himself proposed a simple agent-based model to explain a phenomenon in child development. Consistent with a constructivist epistemology, Papert acknowledged that the “model is absurdly oversimplified . . . Dozens or hundreds [of agents] are needed to account for the complexity of the real process. But, despite its simplicity, the model accurately conveys some of the principles of the theory” (pp. 168-169).

Moreover, a constructivist approach of aiming for model robustness can be useful for researchers from different paradigms in learning analytics, even if they do not wholly prescribe to a constructivist epistemology. (It is especially well-suited to those who see themselves as pragmatists.) Ultimately, a constructivist approach to learning analytics should aim for models that can be used for robust decision-making that meets the needs of the context in which the model would be used. For example, learning analytics researchers may apply such a strategy to the following scenarios to have more confidence about the decisions they make:

- To make sure that we make reliable decisions about when to move students on to another topic (when implementing mastery learning or personalized learning), we can see if those decisions are reasonable under various plausible assumptions about how students learn.
- To see if we want to make decisions based on a model that predicts when students are likely to strug-

gle or fail a course, we can make sure the predictions are equally valid for students from different demographic groups. (Such techniques are growing in popularity under fair machine learning, but here we draw attention to the fact that they are a kind of robustness check for our models.)

- To make decisions about how to modify instructional activities based on learning curves, we can make sure our decisions are robust across various KC models that are qualitatively different but have similar levels of accuracy.
- To see if a particular variable is important in predicting a good self-regulation strategy, we might fit multiple black box models and use multiple explainable AI techniques to see if they agree on the relative importance of that feature (even if they differ on other aspects of predictions).

Notice that in each of these cases, we have to move beyond simply measuring the accuracy of our models. We use models to make meaningful inferences and useful decisions, but we do not make strong truth claims about any particular model. I end this section with another quote that nicely sums up this epistemology of model fitting, here offered by McClelland (2009), one of the pioneers of connectionism:

I argue that we should think of models as tools for exploring the implications of ideas. They can teach us things about the consequences of particular ways of construing the processes that take place when humans engage in particular kinds of cognitive tasks, with sometimes surprising consequences... A good fit never means that a model can be declared to provide the true explanation for the observed data; a poor fit likewise does not necessarily show that the core principles embodied in a model are necessarily the source of the misfit. (p. 12)

7 Conclusion

The foregoing analysis suggests a new philosophical framework for categorizing different approaches to learning analytics: learning analytics research can be categorized based on (1) its position in terms of the bias-variance tradeoff and (2) its underlying epistemology. However, I reiterate that such an analysis does not suggest it is not of value to categorize learning analytics approaches into different paradigms and to look for connections between those paradigms and research communities. If we do want to make use of the paradigms that Baker et al. (2021) proposed, we should take further care to ensure the paradigms are accurate labels for the research they represent. As suggested above, we need to be clear on what kind of learning analytics research would be accurately classified as essentialist, what paradigm(s) deep learning research falls under, and whether researchers tend to gravitate towards one or two paradigms.

Moreover, if we do seek to categorize learning analytics research in terms of four or more paradigms, McKeon's categorization is not the only approach. Pepper's (1942) systematic philosophical categorization of "world hypotheses" offers an alternative approach that boils down different approaches to making sense of the world (whether from the lens of philosophy, science, religion, history, art, etc.) to the root metaphors that underlie those approaches. He suggests there are four world hypotheses that are all defensible but mutually exclusive. Researchers have applied this framework to categorizing different psychological theories (Hayes, Hayes, & Reese, 1988). An analysis of learning analytics research in terms of their underlying world hypotheses—and the relationship between these hypotheses and the bias-variance tradeoff—could be an interesting area to explore in future work.

The approach taken in this paper has also helped uncover new potential directions for learning analytics. First, seeing the relationship between connectionist, situated, and constructivist approaches can suggest a new area of inquiry. Can neural networks (whether state-of-the-art deep learning or a more theoretically-informed exploration of neural networks) offer new analytics methodologies that support situated and constructivist theories of learning? Second, we have presented a preliminary formulation of a constructivist philosophy of data science that might be of value to learning analytics. One impetus for such a philosophy of data science is helping constructivist researchers see how their epistemological stance is not necessarily mutually exclusive with data science research. But even if such a data science epistemology does not broaden participation in learning analytics, it suggests useful methodological practices for learning analytics researchers in general. Some learning analytics researchers may already implicitly adopt such an epistemology in their work; however, formally theorizing about what data science research looks like under a constructivist epistemology can help ensure that our methods are consistent with such an epistemology.

Finally, this paper started by acknowledging that learning analytics is in a unique place for philosophical theorizing, especially around epistemology, because it combines the study of how people learn and the study of how machines learn. Although learning analytics primarily views machine learning from an applied perspective, researchers can take a step back to see how insights about human learning can guide their data science methodology and epistemology and to see how theorizing about data science might inspire new insights on how people learn. The application of the bias-variance tradeoff and a constructivist epistemology are but two examples of such an approach. By furthering this approach, learning analytics researchers might serve to inform the learning sciences, machine learning, and even philosophy.

References

- Abrahamson, D., Blikstein, P., & Wilensky, U. (2007). Classroom model, model classroom: Computer-supported methodology for investigating collaborative-learning pedagogy. In *Proceedings of the 8th international conference on computer supported collaborative learning* (pp. 49–58).
- Aleven, V., & Koedinger, K. R. (2013). Knowledge component (KC) approaches to learner modeling. *Design Recommendations for Intelligent Tutoring Systems, 1*, 165–182.
- Anderson, J. R., Reder, L. M., & Simon, H. A. (1997). Situative versus cognitive perspectives: Form versus substance. *Educational researcher*, 26(1), 18–21. doi: <https://doi.org/10.3102/0013189X026001018>
- Anderson, J. R., Reder, L. M., Simon, H. A., Ericsson, K. A., & Glaser, R. (1998). Radical constructivism and cognitive psychology. *Brookings Papers on Education Policy*(1), 227–278. Retrieved from <http://www.jstor.org/stable/20067198>
- Baker, R. S., Gašević, D., & Karumbaiah, S. (2021). Four paradigms in learning analytics: Why paradigm convergence matters. *Computers and Education: Artificial Intelligence, 2*, 100021.
- Belkin, M., Hsu, D., Ma, S., & Mandal, S. (2019). Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32), 15849–15854.
- Bereiter, C. (1991). Implications of connectionism for thinking about rules. *Educational researcher*, 20(3), 10–16.
- Box, G. E. (1979). Robustness in the strategy of scientific model building. In R. L. Launer & G. N. Wilkinson (Eds.), *Robustness in statistics* (pp. 201–236). Elsevier.
- Buckner, C. (2018). Empiricism without magic: Transformational abstraction in deep convolutional neural networks. *Synthese*, 195(12), 5339–5372.
- Burman, J. T. (2020). On kuhn’s case, and piaget’s: A critical two-sided hauntology (or, on impact without reference). *History of the Human Sciences*, 33(3-4), 129–159.
- Childers, T., Hvorecký, J., & Majer, O. (2021). Empiricism in the foundations of cognition. *AI & SOCIETY*, 1–21.
- Cobb, P. (2007). Putting philosophy to work. *Second handbook of research on mathematics teaching and learning: A project of the National Council of Teachers of Mathematics*, 1(1).
- Doroudi, S. (2020). The bias-variance tradeoff: How data science can inform educational debates. *AERA Open*, 6(4), 2332858420977208.
- Doroudi, S., & Brunskill, E. (2017). The misidentified identifiability problem of Bayesian Knowledge Tracing. In *Educational data mining* (pp. 143–149).
- Duncan, O. D. (1975). *Introduction to structural equation models*. Elsevier.

- d'Ascoli, S., Refinetti, M., Biroli, G., & Krzakala, F. (2020). Double trouble in double descent: Bias and variance (s) in the lazy regime. In *International conference on machine learning* (pp. 2280–2290).
- Fisher, L. M. (2016, Aug). *In memoriam: Seymour papert 1928-2016*. ACM. Retrieved from <https://cacm.acm.org/news/205495-in-memoriam-seymour-papert-1928-2016/fulltext>
- Greeno, J. G. (1997). On claims that answer the wrong questions. *Educational Researcher*, 26(1), 5–17. doi: <https://doi.org/10.3102/0013189X026001005>
- Guba, E. G., & Lincoln, Y. S. (1994). Competing paradigms in qualitative research. In N. K. D. . Y. S. Lincoln (Ed.), *Handbook of qualitative research* (Vol. 2, p. 105). Sage Publications.
- Halevy, A., Norvig, P., & Pereira, F. (2009). The unreasonable effectiveness of data. *IEEE intelligent systems*, 24(2), 8–12.
- Hayes, S. C., Hayes, L. J., & Reese, H. W. (1988). Finding the philosophical core: A review of stephen c. pepper's world hypotheses: A study in evidence. *Journal of the experimental analysis of behavior*, 50(1), 97.
- Hennig, C. (2002). Confronting data analysis with constructivist philosophy. In *Classification, clustering, and data analysis* (pp. 235–243). Springer.
- Hennig, C. (2003). How wrong models become useful—and correct models become dangerous. In *Between data science and applied data analysis* (pp. 235–243). Springer.
- Hennig, C. (2010). Mathematical models and reality: A constructivist perspective. *Foundations of Science*, 15(1), 29–48.
- Kabeshkin, A. (2021). Hegel's anti-reductionist account of organic nature. *Intellectual History Review*, 31(3), 479–494.
- Kahneman, D., Sibony, O., & Sunstein, C. R. (2021). *Noise: A flaw in human judgment*. Little, Brown.
- Kitchener, R. F. (1980). Piaget's genetic epistemology. *International Philosophical Quarterly*, 20(4), 377–405.
- Levins, R. (1966). The strategy of model building in population biology. *American scientist*, 54(4), 421–431.
- Luczak-Rösch, M. (2013). Beyond positivism in computer science.
- Markie, P., & Folescu, M. (2023). Rationalism vs. Empiricism. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford encyclopedia of philosophy* (Spring 2023 ed.). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2023/entries/rationalism-empiricism/>.
- McClelland, J. L. (2009). The place of modeling in cognitive science. *Topics in Cognitive Science*, 1(1), 11–38.
- McKeon, R. (1990). Philosophic semantics and philosophic inquiry. *Freedom and History and Other Essays: An Introduction to the Thought of Richard McKeon*, 242–256.
- Minsky, M., & Papert, S. (1971). *Artificial intelligence progress report* (Tech. Rep.). MIT Artificial Intelligence

- Laboratory. Retrieved from <https://web.media.mit.edu/~minsky/papers/PR1971.html>
- Norvig, P. (2012). Colorless green ideas learn furiously: Chomsky and the two cultures of statistical learning. *Significance*, 9(4), 30–33.
- Papert, S. (1976). *Some poetic and social criteria for education design* (Tech. Rep.). Retrieved from <https://dspace.mit.edu/bitstream/handle/1721.1/6250/AIM-373.pdf>
- Papert, S. (1980). *Mindstorms: Children, computers, and powerful ideas*. Basic Books, Inc.
- Papert, S. (1988). One ai or many? *Daedalus*, 1–14.
- Pepper, S. C. (1942). *World hypotheses: A study in evidence*. University of California Press.
- Phillips, D. C., & Burbules, N. C. (2000). *Postpositivism and educational research*. Rowman & Littlefield.
- Piaget, J. (1952). Jean Piaget. In E. G. Boring, H. Werner, H. S. Langfeld, & R. M. Yerkes (Eds.), *A history of psychology in autobiography* (Vol. 4, p. 237–256). Clark University Press. doi: <https://doi.org/10.1037/11154-011>
- Piattelli-Palmarini, M. (Ed.). (1980). *Language and learning: The debate between jean piaget and noam chomsky*. Harvard University Press.
- Politis, V. (2021). *Plato's essentialism: Reinterpreting the theory of forms*. Cambridge University Press.
- Quartz, S. R. (1999). The constructivist brain. *Trends in cognitive sciences*, 3(2), 48–57.
- Šekrst, K., & Skansi, S. (2022). Machine learning and essentialism. *Philosophical Problems in Science (Zagadnienia Filozoficzne w Nauce)*(73), 171–196.
- Shaffer, D. W., Collier, W., & Ruis, A. R. (2016). A tutorial on epistemic network analysis: Analyzing the structure of connections in cognitive, social, and interaction data. *Journal of Learning Analytics*, 3(3), 9–45.
- Shaffer, D. W., Hatfield, D., Svarovsky, G. N., Nash, P., Nulty, A., Bagley, E., ... Mislevy, R. (2009). Epistemic network analysis: A prototype for 21st-century assessment of learning. *International Journal of Learning and Media*, 1(2).
- Shepard, L. A. (1991). Psychometricians' beliefs about learning. *Educational Researcher*, 20(7), 2–16.
- Skansi, S., & Kardum, M. (2021). A prolegomenon on the philosophical foundations of deep learning as theory of (artificial) intelligence. *Disputatio philosophica: international journal on philosophy and religion*, 23(1), 89–99.
- Turkle, S. (1991). Romantic reactions: Paradoxical responses to the computer presence. In J. J. Sheehan & M. Sosna (Eds.), *The boundaries of humanity: Humans, animals, machines* (pp. 224–252). University of California Press. doi: <https://doi.org/10.1525/9780520313118-014>
- von Glasersfeld, E. (1991). An exposition of constructivism: Why some like it radical. In *Facets of systems science* (pp. 229–238). Springer.

- von Glasersfeld, E., & Cobb, P. (1983). Knowledge as environmental fit. *Man-Environment Systems*, 13(5), 216–224.
- Winograd, T. (2006). Shifting viewpoints: Artificial intelligence and human–computer interaction. *Artificial intelligence*, 170(18), 1256–1258.
- Xue, Y., Whitecross, K., & Mirzasoleiman, B. (2022). The final ascent: When bigger models generalize worse on noisy-labeled data. *arXiv preprint arXiv:2208.08003*.

Journal Pre-proof

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Journal Pre-proof