

Articoli/1

Notes for a Political Epistemology of Algorithms

Massimiliano Badino  0000-0002-3521-9454

Articolo sottoposto a doppia blind peer review. Inviato l'11/11/2022. Accettato il 17/11/2022.

Algorithms are everywhere: they watch us, they give us advice, sometimes they take autonomous decisions. They are agents mediating between reality and us by virtue of their capability to treat huge amounts of data and to see patterns that are invisible to us. It is pretty clear that they are political actors by virtue of their epistemic features. In this article, I try to put together these two points and outline a political epistemology of algorithms. Firstly, I define political epistemology as concerned with epistemic performances that are essentially situated and therefore political all the way down. Secondly, I show that, in this sense, an epistemic analysis of algorithms cannot help being a piece of political epistemology.

Introduction

If there is a feature that really sets our age apart from other instalments of human culture is its reliance on artificial and automatic agents. Basically, every aspect of our present life is one way or another, influenced by artificial intelligence or other kinds of automation. Culture has not remained idle before this development. The last ten-odds years have witnessed a veritable boom of studies on various facets of artificial intelligence. Predictably, the lead has been taken by analyses of the ethical and political effects of the new technologies. I said predictably because, given their ubiquity, it is particularly urgent to develop an ethical, legal, and political framework to govern them.

However, most of this intellectual effort seems still to develop within disciplinary comfort zones. Arguably, the most striking feature of the new technologies is their ability to produce a certain kind of knowledge (more on this point later). This is an epistemic feature. Now, this ability has a number of ethical and political consequences. Although the connection between these two facets is strong and evident, surprisingly enough, scholars interested in the epistemic side are hardly interested in the political one and vice versa. Admittedly, combining the two features is not easy. Still, it is necessary to do so because an exclusively

epistemological focus runs the long-term risk of fostering the idea of neutrality of algorithms, while a politics-only focus might lack the epistemological flesh to have a real philosophical bite. In this article, I attempt to lay down some traits that I think a political epistemology of algorithms should display. My aim is purely programmatic. I do not think there is a real need to convince anyone that such a program should be carried out, but its contours are still very much open to debate.

Before starting, let me add a word of caution. In this article, I use the lump word ‘algorithm’ to indicate, quite generically, those AI-powered technologies which are able to extract knowledge from data, possibly by automatically learning how to do it, and to make decisions on that basis. Thus, I will speak of algorithms (a logic concept) to mean also machine learning (a computer science discipline), neural networks (a mathematical structure), recommendation systems, autonomous vehicles (technological artefacts), or artificial intelligence (a whole branch of science). I realize that, strictly speaking, this is an improper use of language, but improper as it might be, it has entered fairly regularly in common parlance. Thus, while bowing to simplicity, I want to stress that the meaning of the term has been deliberately stretched.

The article is organized as follows. In section 2, I discuss the concept of political epistemology, which may not be immediately transparent. In section 3, I clarify in which sense algorithms can be considered as performing epistemic action. I argue, then, that they are a specific kind of epistemic agent. In section 4, I present two cases particularly suitable to illustrate the notion of political epistemology of algorithms discussed in section 2. I conclude with some thoughts on future developments.

1. What is Political Epistemology?

In recent years there have been several attempts at enlarging the scope of epistemology. The traditional concerns with conceptual issues such as the definitions of knowledge and justification and the endless debate on internalism and externalism appear increasingly and hopelessly distant from the pressing questions of contemporary society. The first of these updates to the epistemological research agenda was social epistemology, which started in earnest at the turn of the century¹. The latest instantiation of this tendency is the so-called political epistemology, which has witnessed a battery of important publications coming out in the last couple of years². As it often happens when the adjective ‘epistemic’ or the substantive ‘epistemology’ are paired with another qualification, opinions

¹ A. I. Goldman, *Knowledge in a Social World*, Oxford 1999; M. Fricker, P. J. Graham, D. K. Henderson, N. Pedersen, *The Routledge Handbook of Social Epistemology*, London 2020.

² P. D. Omodeo, *Political Epistemology: The Problem of Ideology in Science Studies*, Berlin 2019; M. Hannon, J. de Ridder (eds.), *The Routledge Handbook of Political Epistemology*, New York 2021;

remain split – when they are uttered altogether – about the ambit of the new research field.

In late 2021, upon the COVID-constrained online presentation of *Political Epistemology*³, Regina Rini synthesized this issue with the effective label of ‘the spork problem’. A spork is a piece of cutlery that has a spoon at one end and a fork at the other end. While it might have multiple useful applications, it can be defined as both a spoon and a fork or neither of them. Analogously, while political epistemology aspires to be considered both political theory and epistemology, that is, to combine the best of both worlds, it runs the serious and ominous risk of ending up being neither good political theory nor substantial epistemology. Hence, the spork problem is how to think of such a hybrid field of study without winding up with an unmanageable academic hodgepodge. Instead of systematically discussing the answers given to this problem⁴, I will go straight to my proposal.

My fundamental claim is that political epistemology is primarily a form of epistemology and only derivatively a form of political theory. In other words, I inscribe political epistemology into the old and venerable tradition of the philosophical study of knowledge in its multiple varieties. One immediate objection that springs to mind against this claim is that political epistemology would sit awkwardly within the epistemological tradition. From Descartes up to contemporary analytical philosophy, the study of knowledge has been deliberately and ostensibly separated from the worldly matters of politics. Its very same tools of inquiry, introspection, logic, and conceptual analysis, have next to nothing to do with mundane social transactions, which, on the contrary, seem to require empirical methods. This objection would make political epistemology, at best, an inferior form of epistemology, an unclear hybridization still in need of definition. Ultimately, the above-mentioned spork problem would reiterate itself.

My answer to this objection is that it rests on an unduly reification of the historical tradition. It is a contingent fact, not a conceptual truth that epistemology developed as the study of disembodied, individual processes of cognition and knowledge production. Several historical reasons, not least Galilei’s clash with the Church, advised Descartes to set the investigation of knowledge and certainty as a purely academic and abstract endeavour to be carried out in perfect solitude before a stove rather than in the square before the people. However, it could have gone another way. Descartes could have understood epistemological investigations as political all the way down in the first place. And even if Descartes could have not, that would have been again for purely historical reasons. There is nothing in the concept of epistemology

³ E. Edenberg, M. Hannon (eds.), *Political Epistemology*, Oxford 2021.

⁴ To mention only one example, in the above mentioned talk, Regina Rini proposed to tentatively view political epistemology as the study of politically relevant doxastic practices. See also R. Rini, *Weaponized Skepticism. An Analysis of Social Media Deception as Applied Political Epistemology*, in E. Edenberg, M. Hannon (eds.), *Political Epistemology*, cit., pp. 31-48.

that forces us to disentangle the analysis of knowledge from its specific contexts. So much so that other writers developed epistemology precisely in this way⁵. Hence, my point is that we should not take a historically determined way of doing epistemology for the royal road and all others as suboptimal options. The legitimacy of political epistemology as a genuine form of epistemology is an empirical question to be adjudicated by working on this option.

Having disposed of this preliminary objection, it's about time I discuss what I mean by political epistemology. My proposed working definition reads as follows: *political epistemology is an essentially and inherently situated inquiry into the relationship between knowledge and action*. Likely, the only perspicuous term is 'inquiry', so let me unearth what I mean by this formulation. By stating that political epistemology is 'essentially situated', I intend to defend the view that it concerns epistemic questions that emerge *only* when the agent is embedded in a social and cultural context. The implicit assumption is that there are, in general, questions that are *genuinely and irreducibly contextual*, that is to say: they only arise when the agent deals with other agents. For example, any issue connected with the establishment of power relations falls within this category. Power relations can subsist only if we have at least two agents: they simply make no sense with a single entity. Hence, any issue derived by the fact that the agents involved are related by some sort of power relation is essentially situated.

It is not trivial to claim that this thesis also holds for epistemic questions. An authoritative tradition in social epistemology seems to rest on the implicit assumption that epistemic issues in a social context can be treated by scaling up the usual individualistic epistemology⁶. For example, the way in which some writers treat the epistemology of testimony does not depend in any relevant fashion on the fact that evidence is obtained by a human agent who might be subject to social or political constraints. Until the appearance of Miranda Fricker's work on the concept of epistemic injustice⁷, testimony was largely regarded as the acquisition of knowledge from the point of view of the receiver alone. Another way to phrase the point of the essential situatedness of political epistemology is that I am interested in epistemological issues that do not simply depend on the fact that there is more than one agent involved but depend on the specific social and political relations that link them. In this sense, a political context cannot be reduced to a sum of individuals.

⁵ One example is the tradition of historical epistemology, which counts several authors. For an introduction see M. Badino, I. Gerardo, P. D. Omodeo, *Epistemologia Storica*, Roma 2022.

⁶ One prominent author in this line of thought is Alvin Goldman (see A. I. Goldman, *Knowledge in a Social World*, Oxford 1999). Let me add a caveat. I don't want to overstretch this point and claim that there is a cleavage in the way in which social epistemology is practiced. It has been argued, however, that social epistemology comes in two varieties (D. K. Henderson, *On the Background of Social Epistemology*, in *The Routledge Handbook of Social Epistemology*, eds. M. Fricker, P. J. Graham, D. K. Henderson and Nikolaj J. L. L. Pedersen, London 2020, pp. 3-9), one which emanates directly from the tradition of individualistic epistemology and one that spins off the tradition of philosophy of science. My point holds for the former.

⁷ M. Fricker, *Epistemic Injustice. Power and the Ethics of Knowing*, Oxford 2007.

Let's move on to inherent situatedness. Partly, this is a consequence of my previous point. For, if political epistemology inquiries phenomena that depend essentially on the social and political context, then its questions are inherently different both from classical epistemology and standard political theory. But this is only part of the story. Also, on the methodological front, political epistemology must set for itself new agendas and new tools. To mention only one example, the classical approach of seeking definitions of epistemic concepts such as knowledge and justification in terms of necessary and sufficient conditions seems wholly inadequate to capture the complexity of the matter. Perhaps more importantly, research in cultural studies, social studies, and history, to mention only a few, become relevant for epistemological inquiry. Thus, far from being a unified field of study, political epistemology is a cluster of different approaches, concepts, methods, and practices kept together by the common project of understanding how knowledge dynamics influences and is influenced by politically salient factors⁸.

A final word concerning knowledge and action. The interplay between knowledge and action is not a stranger in contemporary epistemology. The revival of studies on the ethics of belief or the emergence of the concept of pragmatic encroachment demonstrates a sustained interest in this problem. However, I want to stress that, typically, contemporary epistemology understands this interplay in terms of how and to what extent the action dimension impacts the usual criteria of adjudicating knowledge, justification, and the like. What is usually left out of the equation is the situatedness of the problem, more specifically, the agent. In my view, political epistemology should not be concerned merely with knowledge and action in general, but with the interconnection between types of knowledge with types of actions perpetrated by types of agents. This means, on the one hand, to repeat, that one should abandon the project of providing a necessary-and-sufficient-conditions characterization of knowledge, on the other, that conceptual resources from the philosophy of action, social psychology, and cognitive sciences (among others) should be mobilized to team up with epistemology.

2. Algorithms as Epistemic Performers

This article deals with a particular form of political epistemology, i.e., that one concerned with algorithms. Taking for granted, for the sake of argument, the points made in the previous section, two questions are still in need of an answer, two questions that, as it turns out, are closely related. I tackle them both in this section, and I continue with more specific examples in section 4.

⁸ It's fair to remember that this line of thought was anticipated by Joe Rouse back in the late 1980s (see J. Rouse, *Knowledge and Power*. Ithaca-NY 1987; Id., *Engaging Science: How to Understand Its Practices Philosophically*, Ithaca-NY 1996). Rouse especially insisted on the role of practice as an intermediate between epistemology and situated action and power.

The first question aims at, so to speak, potentially debunking my argument so far. One obvious and well-known sense in which the production, dissemination, appropriation, and modification of knowledge can matter for politics comes straight from the thesis of the non-neutrality of knowledge (more specifically, scientific knowledge). A long tradition of thinkers, which gathers both analytical and continental philosophers, sociologists, cultural theorists, and political scientists, have forcefully held that knowledge, by embodying concepts such as objectivity and truth, might become a tool (if not a weapon) of political action⁹. Michel Foucault, arguably one of the most vocal representatives of this tradition, often insisted that the alleged objectivity of scientific knowledge can be used to justify any sort of political project. Also, John Dewey famously warned against the perils of technocracy and the 20th-century characteristic tendency to turn political and social problems into scientific ones¹⁰. So, one can legitimately ask what is the difference, if any, between this clearly political view of knowledge and the one I have outlined above.

First of all, I have to remark, that I am not denying the relevance of this tradition of looking at knowledge production and dissemination as a matter of political import. On the contrary, this form of interplay between knowledge and power must remain one key component of political epistemology in general. But it is just one component. In particular, it stresses the possible political *uses* of scientific knowledge, but it also assumes that such knowledge might be developed in a neutral environment. That this is the case follows easily by reflecting on a simple fact: while scientific knowledge can be deployed for oppression and political control, it can also be deployed for human flourishing. Without becoming exceedingly optimistic or naive, it is difficult to overlook the great material and spiritual gains that scientific progress brought to humankind. Thus, there is a sense according to which initially politically and ethically neutral knowledge can take one path or the other. Or, to put the same point in other words, there is a sense in which ideally neutral knowledge can be used for good or evil. Granted, it might be practically difficult to separate out the content of knowledge from its deployment, but, at least as an analytical tool, this difference is well defined.

Thus, whereas I do not deny that scientific knowledge might be put to the service of vicious political projects, I see political epistemology to also deal with forms of knowledge production and appropriation that are *constitutively* and not only *derivatively* political. This view is clearly stronger than the former. For, I am holding that there is a sense in which the separation between producing knowledge and acting politically is not possible even as an analytical idealization.

⁹ For a comprehensive historical account see L. Daston, P. Galison, *Objectivity*, Princeton 2010. See also the discussion of this point in A. I. Goldman, *Knowledge in a Social World*, Oxford 1999.

¹⁰ For an interesting analysis on this issue, which compares Dewey and the tradition of political theology, see M. Chun, *Beyond Technocracy and Political Theology: John Dewey and the Authority of Truth*, «Philosophy & Social Criticism», February, 2022, pp. 1-27.

There is no conceptual separation between the political and the epistemic act: some forms of epistemic performances (or, more precisely, some epistemic agents) are *ipso facto* also political performances.

Thus, the second question is: how do algorithms enter this picture? In the remainder of this article, I want to defend the claim that algorithms are precisely the kind of agents for which epistemic performances tend to be constitutively political. My argument consists of a general analysis of the notion of epistemic performance and in the discussion of two major examples: the problem of algorithmic fairness and the issue of AI carbon footprint. Let us proceed with the first step.

By epistemic performance, I mean, very plainly, a process or an action that produces information or knowledge as its final result.¹¹ This definition is deceptively simple. Firstly, it is ambiguous because the term ‘performance’ can refer both to the process itself and to the outcome of the process. This difference is immaterial to my argument. The adjective ‘epistemic’ characterizes the performance as information- or knowledge-oriented, that is to say, as a process that eventually produces a piece of information or a piece of knowledge. Clearly enough, a performance is to be assessed both on the basis of the sequence of actions that lead to that particular outcome and on the basis of the quality of the outcome itself. Thus, to take a very common example in analytical epistemology, one might be right by a pure stroke of luck and thus reach an optimal outcome in spite of a poor process¹². By contrast, it might happen that perfectly reliable methods fail us, and thus a disappointing outcome follows a more than respectable process. Depending on the situation, one will be more important than the other (or they will be equally important), but in the following, I will speak indifferently of epistemic performance, including both the process and the outcome. Secondly, the previous definition alludes to information and knowledge. Once again, these two concepts should be kept separate. Typically, information has to do with our expectations given a set of data (syntactic information) or the semantic content that we see in the data themselves (semantic information)¹³. By contrast, knowledge traditionally refers to belief, truth, justification, and, more importantly, understanding. Still, to the extent that there is a difference between these notions, I do not need to enter their intricate relations here. I can

¹¹ For simplicity's sake, I will confine myself to public information and knowledge. I will, therefore, exclude the production on private knowledge, that is learning or understanding.

¹² Epistemologists used to amuse themselves with the ‘barn problem’, which was originally conceived by Alvin Goodman. During a drive in the countryside, a man ends up in an area full of barns. In fact, that area is an old cinematographic set and almost all barns are fake, except one. Now, the man is luckily standing in front of the only genuine barn and forms the belief that he's facing a barn. He is right, so his belief is true (the best epistemic performance possible), but he is so by pure luck.

¹³ For a good introduction see L. Floridi, *Information: A Very Short Introduction*, Oxford 2010. A more ambitious and comprehensive treatment is given in L. Floridi, *The Philosophy of Information*, Oxford 2011.

be content with assuming that an epistemic performance sometimes produces information, and sometimes produces genuine knowledge.

The only thing that it really does matter is that algorithms are capable of epistemic performances, they are, hence, epistemic performers. To fix the ideas, let us consider a deep neural network used for face recognition. Typically, such a system is previously trained with many examples, and during the training phase, it spots statistical correlations between groups of data, in this case, pixels on the screens. In more anthropomorphic parlance, it finds patterns, although they are patterns that, more often than not, mean absolutely nothing to us. But they mean to the neural network and enable it to sort the image out, say, in men and women. I submit that this, simple as it may look, is an epistemic performance. For sure, it would be called so if performed by a human being and I do not see any reason, as far as the outcome is concerned, to deny the same status to a machine.¹⁴ But if we, for any reason, feel uncomfortable with using the comparison with humans to attribute a feature to a machine, we can move to super-human performances. A chess software based on machine learning can spot a winning variation much faster and more effectively than any human player, a feat that, I submit, should be called epistemic¹⁵.

3. Algorithms Going Political

Now that I have clarified in which way algorithms can be viewed as epistemic performers, it is about time to discuss in which sense these performances are constitutively political. To do this, I submit two examples of algorithmic epistemic performances that I deem as illustrative of my main point.

The first example is generally known under the label: algorithmic fairness. By algorithmic fairness, it is commonly meant the whole gamut of ethical and political problems that can arise from the use of algorithms and machine

¹⁴ It won't escape the attentive reader that this argument rests on an analogous of the parity principle introduced by Clark and Chalmers (see A. Clark and D. Chalmers, *The Extended Mind*, «Analysis», 58 (1), 1998, pp. 7-19). This impression is correct. I am in fact suggesting that the category of 'epistemic performer' can be attributed to any entity capable of a human-like epistemic performance. In other words, I am evoking the following assumption: if the entity X performs in a way that, had the same performance being carried out by a human being we would not hesitate to call it an epistemic performance, then X's performance is an epistemic performance and X itself is an epistemic performer. This assumption, in turn, rests on a sort of epistemic functionalism: a performance can be called epistemic independently of the entity carrying it out, that is to say, being human is a sufficient, but not necessary condition to be an epistemic performer.

¹⁵ This does not imply that machines can think letting alone being able to think in human terms. I have suspended judgement to the nature of the process precisely to allow the adjective 'epistemic' to be related with knowledge and information only. A winning variation is a relevant piece of knowledge extracted from a state of affair, which, in my view, is enough to qualify the relevant process as epistemic.

learning technology for classification, profiling, and prediction.¹⁶ One of the most important applications of algorithms is in recommendation systems. We encounter them every time we visit the Amazon website, or we choose what to watch on Netflix. Algorithms can analyse tons of data about our behaviour to predict what better suits our needs, desires, or momentary inclinations. Innocent as this application might appear, it can get very tricky. For one, in the case of Amazon or Netflix there is a problem with privacy and data ownership. But let's put these issues aside for the moment and concentrate on more problematic cases of recommendations.

One such case is predictive police. Algorithms can be profitably deployed for a number of security-related tasks. For example, they have been employed to predict the recidivism probability in parole cases or to determine which neighbourhood criminal activities are more likely. At first sight, this looks like a huge help in making police action faster and more effective. But several analyses of PrePol algorithms (as they are called), show that they are heavily racially biased¹⁷. A particularly outrageous case is COMPAS (Correctional Offender Management Profiling for Alternative Sanctions), a software developed to assess the risk of recidivism and used in the courts of New York, Wisconsin, and Florida, among others. As it was revealed in 2016 by the independent press agency Pro Publica¹⁸ the algorithm was in fact, deeply biased against black people. In particular, while «blacks are almost twice as likely as whites to be labelled a higher risk but not actually re-offend», COMPAS «makes the opposite mistake among whites: They are much more likely than blacks to be labelled lower-risk but go on to commit other crimes»¹⁹.

Analogous problems plague basically any application of a machine learning system to decision problems, even apparently innocent ones such as language recognition²⁰ or online advertising²¹. Examples could easily be multiplied, but my point here is more general. The systematic presence of bias in recommendation

¹⁶ The literature on algorithmic fairness is growing exponentially. A general treatment of the issue is given in M. Coeckelbergh, *AI Ethics*. Cambridge MA 2020 and in M. Kearns and A. Roth, *The Ethical Algorithm: The Science of Socially Aware Algorithm Design*, Oxford 2020. For a specific focus on the ethical and legal impact of fairness see also W. Hinch, *Differences That Make a Difference: Computational Profiling and Fairness to Individuals*, in *The Cambridge Handbook of Responsible Artificial Intelligence* (eds. S. Voeneke, Ph. Kellmeyer, O. Mueller and W. Burgard), Cambridge 2022, pp. 229-51, and A. Von Ungern-Sternberg, *Discriminatory AI and the Law: Legal Standards for Algorithmic Profiling*, in *The Cambridge Handbook of Responsible Artificial Intelligence* (eds. S. Voeneke, Ph. Kellmeyer, O. Mueller, W. Burgard), Cambridge 2022, pp. 252-78.

¹⁷ For a telling example, see K. Lum, and W. Isaac, *To Predict and Serve?*, «Significance», 13 (5), 2016, pp. 14-19.

¹⁸ J. Angwin, J. Larson, *Machine Bias*, «ProPublica», <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>, 23.5.2016.

¹⁹ *Ibidem*.

²⁰ A. Caliskan, J. Bryson, A. Narayanan, *Semantics Derived Automatically from Language Corpora Contain Human-like Biases*, «Science», 356 (6334), 2017, pp. 183-86.

²¹ L. Sweeney, *Discrimination in Online Ad Delivery: Google Ads, Black Names and White Names, Racial Discrimination, and Click Advertising*, «Queue», 11 (3), 2013, pp. 10-29.

systems results from two concomitant causes. The first cause is epistemological in nature. Machine learning systems are fuelled by data (a huge amount thereof), which are then manipulated by means of statistical techniques. Algorithms are, in fact, recipes to elaborate data according to well-established statistical principles. The final recommendation is the conclusion that the algorithm draws from data. Statistics is full of ingenious tools to extract information from data, but there are two drawbacks. First, information does not come for free. A famous theorem shows that, when measured over all possible populations, no algorithm is better than random choice.²² This means that algorithms, to work at all, need to be put on the right track with some preliminary assumptions on the type of solution we expect. And this leads us to the second drawback. Recommendation systems need a bias, they cannot work, so to speak, in the dark. More importantly, the bias can be reduced only at the expense of the vagueness of the final recommendation, which statistically is called variance. In other words, algorithms work by balancing a delicate trade-off between the initial bias and the variance of the prediction.²³

Hence, far from being something to steer away from, bias is one of the conditions for algorithms to perform epistemically. Algorithms can produce knowledge only by making some sort of stereotypical assumptions about the initial data. Here comes the second cause, which concerns the deep structure of society. For, it is a reasonable move to base the initial bias on the data at disposal. But the data are not neutral: in fact, they encode a variety of structural inequalities sedimented over decades. Historically, ethnic minorities have been segregated and pushed toward a harsh life. Social data reflect this history. It is, therefore, clear that using those data to predict recidivism or the capability to pay back a mortgage means contributing to sustaining and propagating these structural inequalities.

Thus, the interplay between epistemological and political factors emerges very strongly here. Machine learning systems rely on social statistical inference (a very dynamic version of it, but statistical inference anyway), and it is common wisdom that social statistics is politically loaded. It begins with a quantification of some social features to produce a conflated snapshot of society that is completely oblivious to a historical perspective. To use a slightly more philosophically precise language, it reifies historical differences and turns them into natural kinds. But these differences express political relations and dynamics, which get, therefore, completely obliterated by the statistical snapshot.

²²This result is called the no free lunch theorem, see D. H. Wolpert, *The Lack of A Priori Distinctions Between Learning Algorithms*, «Neural Computation», 8 (7), 1996, pp. 1341-90; D. H. Wolpert, and W. G. Macready, *No Free Lunch Theorems for Optimization*, «IEEE Transactions on Evolutionary Computation», 1 (1), 1997, pp. 67-82.

²³On the key concept of bias/variance tradeoff in machine learning see S. Geman, E. Bienenstock, R. Doursat, *Neural Networks and the Bias/Variance Dilemma*, «Neural Computation», 4 (1), 1992, pp. 1-58.

Thus, the troublesome issues elicited by fairness and discrimination, as well as the discussions related to the concepts of equality, justice, and reparation applied to algorithms, do not concern the mere implementation of neutral machines. On the contrary, these issues are part and parcel of the way in which algorithms produce knowledge, i.e., their epistemic structure. Here we see an example of the point I made back in section 2: there are forms of knowledge production that are political all the way down, and for them, any epistemological analysis is *ipso facto* a politico-epistemological analysis.²⁴

My second example illustrates another case of this dynamics, with somewhat of an ironic twist. Among the contributions that machine learning can bring to our present global concerns, one is mentioned more and more frequently: the fight against climate change. Notoriously, climate change is caused mainly by our political and economic system, which uses energy irresponsibly and looks at nature as an infinitely capable reservoir where the wastes of our production chain can be externalized. In tackling this problem, machine learning systems can be twofold helpful. First, they can contribute to improving the efficiency of our supply chain. By fast elaborating huge amounts of data, algorithms are able to direct the collection of raw material, organize its elaboration, and to govern the distribution of the final product much more efficiently than any human manager. Optimizing resources is not only good for the economy: it also means reducing energy emissions in the environment. Thus, algorithms might magically harmonize the requirements of the environment and those of economic growth. A second way in which algorithms can come to rescue us from global warming disaster is by shedding light upon the way in which climate evolves. Famously, the climate is fantastically complicated to describe because it depends on a huge number of local parameters. Humans are not able to find their way in such a maze, but algorithms are. In other words, using machine learning can give us unprecedented knowledge about the underlying dynamics of climate and hence help us mitigate dangerous effects.

It is important to stress that both these contributions rest on two major features of algorithms. One is, as I have shown, their characteristic epistemic performance. Algorithms can see correlations among data, which are completely invisible to us. Furthermore, they can elaborate big chunks of information at a computational speed incomparably higher than the human brain. Algorithms happen to be great at optimizing and finding patterns. The second feature is the perceived immateriality of algorithms. Technically, algorithms are recipes to solve problems. Can you imagine anything more immaterial than this? They work with data, which, in turn, are immaterial information stored in the 'cloud'. Their action is immaterial too. Algorithms, we just said, optimize and find

²⁴To be sure, there is another factor that one should add, that is the ubiquity of algorithms. Their political import also derives from the fact that an increasing portion of our processes is relinquished to artificial intelligence. However, this state of affair certainly contributes to the *scale* of the phenomenon, but the political *nature* of the phenomenon itself depends on epistemic factor.

patterns. This means that they merely re-arrange or spot what is already there without any need to bring in more resources or matters.

As it turns out, algorithms are, in fact, very material. Two begin with, data are not in the cloud but in huge buildings occupying thousands of square meters and called data centers. More importantly, machine learning systems need to be trained on big data before they can perform, and training devours a lot of energy resulting in a huge carbon footprint. To take just one example, consider a natural processing language system (NLP) whose architecture is based on a deep neural network. It has been estimated that the training phase of such a system has an environmental impact equal to that of five cars during their entire lifecycle.²⁵ Here one can see the ironic entrenchment of epistemology and politics: the very same epistemic features that enable machine learning to contribute to mitigating climate change make it a remarkable player in the same game. This fact raises huge political questions related to the geographical distribution of data centres, the access to those systems, the resources mobilized for their maintenance, and so forth, whose implications are still investigated by scholars²⁶.

Conclusion

The foregoing sections have driven home a major point: algorithms, in the broad sense outlined in the introduction, are subject matter for political epistemology discussed in section 2. I am by no means claiming that all applications of algorithms have this characteristic. I cannot see a political implication in automatic theorem proving, for one. But there are numerous and popular applications, perhaps the majority of useful applications, in which epistemic performance and political effect are flipped sides of the same coin. In section 4, I have discussed two examples, but they can easily be multiplied.

A question remains open: How should such a political epistemology of algorithms develop? This is clearly a theme for the future. In recent years cultural analyses on several aspects of artificial intelligence have skyrocketed. For sure, political epistemology requires a combination of multiple expertise. As I have stressed, the political impact of algorithms is closely related to their epistemic performance, that is to say, one needs to understand deeply how they produce knowledge. But, once again, this understanding is not enough. The examples discussed above show that a political-epistemological analysis might

²⁵ See, for example, E. Strubell, A. Ganesh, A. McCallum, *Energy and Policy Considerations for Deep Learning in NLP*, *arXiv*. <http://arxiv.org/abs/1906.02243>, 2019.

²⁶ Very useful thoughts on this issue can be found in C. Freitag *et al*, *The Real Climate and Transformative Impact of ICT: A Critique of Estimates, Trends, and Regulations*, «Patterns», 2 (9), 2021, p. 100340; F. Lucivero, *Big Data, Big Waste? A Reflection on the Environmental Sustainability of Big Data Initiatives*, «Science and Engineering Ethics», 26 (2), 2020, pp. 1009-30; G. Tamburrini, *The AI Carbon Footprint and Responsibilities of AI Scientists*, «Philosophies», 7 (1), 2022, p. 4.

require expertise in statistics, moral philosophy, climate science, and industrial technology, to mention just a few.

Massimiliano Badino
Università degli studi di Verona
✉ massimiliano.badino@univr.it

References

- Alfano, M., Emily, S. 2021. *Online Trust and Distrust*, in *The Routledge Handbook of Political Epistemology* (ed. M. Hannon and J. De Ridder), New York: Routledge, pp. 480-91.
- Angwin, J., Larson, J. 2016. *Machine Bias*, «ProPublica», <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>. (23.5.2016).
- Badino, M., Gerardo I., Omodeo, P. D., 2022. *Epistemologia Storica*. Rome: Carocci.
- Binns, R. 2018. *Fairness in Machine Learning: Lessons from Political Philosophy*, «Journal of Machine Learning Research», 81, pp. 1-11.
- Caliskan, A., Bryson J. and Arvind Narayanan, A., 2017. *Semantics Derived Automatically from Language Corpora Contain Human-like Biases*, «Science», 356 (6334), pp. 183-86.
- Chun, M., 2022. *Beyond Technocracy and Political Theology: John Dewey and the Authority of Truth*, «Philosophy & Social Criticism», February, pp. 1-27.
- Clark, A., and Chalmers, D., 1998. *The Extended Mind*, «Analysis», 58 (1), pp. 7-19.
- Daston, L., and Galison, P., 2010. *Objectivity*, Princeton: Zone Books.
- Coeckelbergh, M. 2020. *AI Ethics*. Cambridge, MA: The MIT Press.
- Coeckelbergh, M. 2022. *The Political Philosophy of AI*. Cambridge: Polity Press.
- Edenberg, E., and Hannon, M. (eds.) 2021. *Political Epistemology*. Oxford: Oxford University Press.
- Floridi, L. 2010. *Information: A Very Short Introduction*. Oxford: Oxford University Press.
- Floridi, L. 2011. *The Philosophy of Information*. Oxford: Oxford University Press.
- Freitag, C., Mike Berners-Lee, M., Widdicks, K., Knowles, B., Blair, G. S., and Friday, A., 2021. *The Real Climate and Transformative Impact of ICT: A Critique of Estimates, Trends, and Regulations*. «Patterns», 2 (9), p. 100340.
- Fricker, M. 2007. *Epistemic Injustice. Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Fricker, M., Graham, P. J., Henderson, D. K., and Pedersen, N., 2020. *The Routledge Handbook of Social Epistemology*. London: Routledge.

- Geman, S., Bienenstock, E., and Doursat, R., 1992. *Neural Networks and the Bias/Variance Dilemma*, «Neural Computation», 4 (1), pp. 1-58.
- Goldman, A. I. 1999. *Knowledge in a Social World*. Oxford: Oxford University Press.
- Hannon, M., and de Ridder, J., eds. 2021. *The Routledge Handbook of Political Epistemology*. Routledge Handbooks in Philosophy. New York: Routledge.
- Henderson, D. K. 2020. *On the Background of Social Epistemology*. In *The Routledge Handbook of Social Epistemology*, eds. M. Fricker, P. J. Graham, D. K. Henderson and Nikolaj J. L. L. Pedersen, London: Routledge, pp. 3-9.
- Hinsch, W. 2022. *Differences That Make a Difference: Computational Profiling and Fairness to Individuals*. in *The Cambridge Handbook of Responsible Artificial Intelligence*, eds. S. Voenekey, Ph. Kellmeyer, O. Mueller and W. Burgard, Cambridge University Press, pp. 229-51.
- Hogan, M. 2018. *Big Data Ecologies*. «Ephemera: Theory & Politics in Organizations Patterns», 18 (3), pp. 631-57.
- Kearns, M., and Roth, A., 2020. *The Ethical Algorithm: The Science of Socially Aware Algorithm Design*. Oxford: Oxford University Press.
- Lucivero, F., 2020. *Big Data, Big Waste? A Reflection on the Environmental Sustainability of Big Data Initiatives*. «Science and Engineering Ethics», 26 (2), pp. 1009-30.
- Lum, K., and Isaac, W., 2016. *To Predict and Serve?*, «Significance», 13 (5), pp. 14-19.
- Omodeo, P. D., 2019. *Political Epistemology: The Problem of Ideology in Science Studies*. Berlin: Springer International Publishing.
- Rini, R., 2021. *Weaponized Skepticism. An Analysis of Social Media Deception as Applied Political Epistemology*, in *Political Epistemology*, eds. by E. Edenberg and M. Hannon, Oxford: Oxford University Press, pp. 31-48.
- Rolnick, D., Donti, P. L., Kaack, L., Kochanski, K., Lacoste, A., Sankaran, K., Slavin A., et al. 2019. *Tackling Climate Change with Machine Learning*. *arXiv*. <http://arxiv.org/abs/1906.05433>.
- Rouse, J. 1987. *Knowledge and Power*. Cornell University Press.
- Rouse, J. 1996. *Engaging Science: How to Understand Its Practices Philosophically*. Ithaca, NY: Cornell University Press.
- Strubell, E., Ganesh A., and McCallum, A., 2019. *Energy and Policy Considerations for Deep Learning in NLP*, *arXiv*. <http://arxiv.org/abs/1906.02243>.
- Sweeney, L., 2013. *Discrimination in Online Ad Delivery: Google Ads, Black Names and White Names, Racial Discrimination, and Click Advertising*, «Queue», 11 (3): 10–29.
- Tamburrini, G., 2022. *The AI Carbon Footprint and Responsibilities of AI Scientists*, «Philosophies», 7 (1), p. 4.
- Tsamados, A., Aggarwal, N., Cows, J., Morley, J., Roberts, H., Taddeo, M., and Floridi, L., 2021. *The Ethics of Algorithms: Key Problems and Solutions*. «AI & Society», February.

- Von Ungern-Sternberg, A., 2022. *Discriminatory AI and the Law: Legal Standards for Algorithmic Profiling*. In *The Cambridge Handbook of Responsible Artificial Intelligence*, eds. S. Voenekey, Ph. Kellmeyer, O. Mueller and W. Burgard, Cambridge University Press, pp. 252-78.
- Wolpert, D. H. 1996. *The Lack of A Priori Distinctions Between Learning Algorithms*, «Neural Computation», 8 (7), pp. 1341-90.
- Wolpert, D. H., and Macready, W. G., 1997. *No Free Lunch Theorems for Optimization*. «IEEE Transactions on Evolutionary Computation», 1 (1): 67–82.