

Data Science as Machinic Neoplatonism

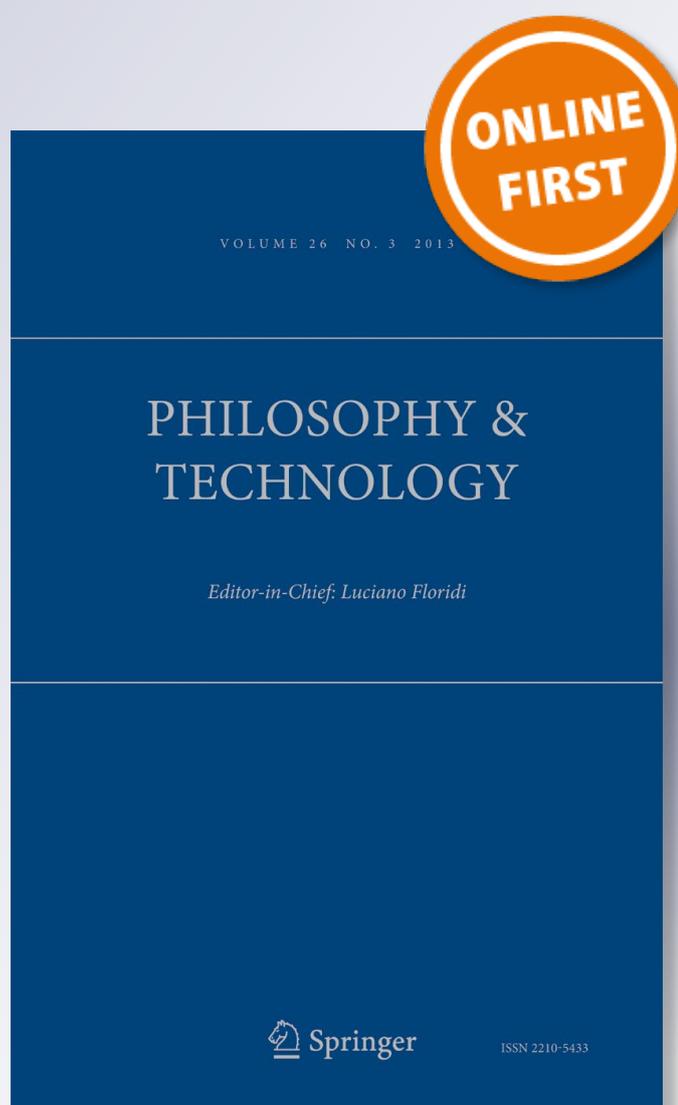
Dan McQuillan

Philosophy & Technology

ISSN 2210-5433

Philos. Technol.

DOI 10.1007/s13347-017-0273-3



Your article is published under the Creative Commons Attribution license which allows users to read, copy, distribute and make derivative works, as long as the author of the original work is cited. You may self-archive this article on your own website, an institutional repository or funder's repository and make it publicly available immediately.

Data Science as Machinic Neoplatonism

Dan McQuillan¹ 

Received: 22 October 2016 / Accepted: 2 August 2017
© The Author(s) 2017. This article is an open access publication

Abstract Data science is not simply a method but an organising idea. Commitment to the new paradigm overrides concerns caused by collateral damage, and only a counterculture can constitute an effective critique. Understanding data science requires an appreciation of what algorithms actually do; in particular, how machine learning learns. The resulting ‘insight through opacity’ drives the observable problems of algorithmic discrimination and the evasion of due process. But attempts to stem the tide have not grasped the nature of data science as both metaphysical and machinic. Data science strongly echoes the neoplatonism that informed the early science of Copernicus and Galileo. It appears to reveal a hidden mathematical order in the world that is superior to our direct experience. The new symmetry of these orderings is more compelling than the actual results. Data science does not only make possible a new way of knowing but acts directly on it; by converting predictions to pre-emptions, it becomes a machinic metaphysics. The people enrolled in this apparatus risk an abstraction of accountability and the production of ‘thoughtlessness’. Susceptibility to data science can be contested through critiques of science, especially standpoint theory, which opposes the ‘view from nowhere’ without abandoning the empirical methods. But a counterculture of data science must be material as well as discursive. Karen Barad’s idea of agential realism can reconfigure data science to produce both non-dualistic philosophy and participatory agency. An example of relevant praxis points to the real possibility of ‘machine learning for the people’.

Keywords Machine learning · Algorithms · Data science · Counterculture · Standpoint theory · Agential realism · Neo-platonism · Big data · Participation · Agency

✉ Dan McQuillan
d.mcquillan@gold.ac.uk

¹ Department of Computing, Goldsmiths, University of London, London, UK

1 Data Science as Organising Idea

Data science is not simply a method but an organising idea. That is, an underlying shift in perspective and practices of the kind that Kuhn called a paradigm (Kuhn 1996). This should be appreciated when trying to critique data science based on the occurrence of collateral damage. The commitment to adopting data science in more and more areas of life will not be constrained by limits such as data protection, because it is based on a new normativity. Only a countercultural critique can dissect the conditions that constitute the possibility of data science and propose an alternative.

To start the process of developing this counterculture, it is necessary to understand what data science actually does. The surge of popular interest in big data has also generated misinformation about its core concepts and practices. To grasp the essence of data science means examining the algorithmic methods that make data science possible, in particular, forms of machine learning. Shedding light on the practical strengths and weaknesses of data science allows us to illuminate real concerns about its operations in the world, ranging from algorithmic discrimination to the evasion of legal due process. But we cannot leave it at that. As Kuhn made clear, contrary to the assumptions of naive empiricism, reality is not directly accessible to us as ‘facts’ that can be recorded by suitable devices and rendered as theory. While we can access reality, we cannot do so without some level of meaning making. What we sense, through whatever device, already has meaning to us, and meaning is not an object of sensory perception. What is also present is a pattern of cognition which enables the ‘seeing’. As an organising idea, data science does not simply reorganise facts but transforms them.

Data science can be understood as an echo of the neo-platonism that informed early modern science in the work of Copernicus and Galileo. That is, it resonates with a belief in a hidden mathematical order that is ontologically superior to the one available to our everyday senses. Looking at how this defines the character of data science provides a skeleton key to understanding its likely consequences. It also helps explain the widespread commitment to data science in the face of actual results that fall far short of the promulgated vision. When characterising data science philosophically as a form of neo-platonism, it is important to understand the difference between data science and discursive philosophy. Data science does not affect by argument alone but acts directly in the world as a form of algorithmic force. It is machinic, that is, an assembly of flows and logic that enrolls humans and technology in a larger, purposeful structure. While algorithms and data are the bone and sinew of data science, its vital force comes from general computation. As computation becomes pervasive, capturing and reorganising human activity, data science exerts its philosophy directly as orderings, decisions and outcomes.

2 What is Data Science?

There is considerable variety in the way practitioners themselves describe data science. It is a lively term that functions both as a flag of convenience for statistics

and a genuinely new discipline (Quora 2014). As the latter, it embraces a grab bag of skills including programming (typically in R or Python), data munging (parsing, scraping and formatting data), statistics, linear algebra, multivariate calculus, SQL (structured query language), machine learning and data visualisation (Holtz 2014). The imaginary ideal data scientist is a Renaissance figure with a mastery of all these arts. The point of overlap of these skills occupied by any particular data scientist will usually reflect their personal background and the role they are recruited for. Typical routes into data science include someone with a background in statistics who learns to code, or someone strong in programming who has acquired an appreciation of analytical modelling and problem-solving. Hence, the much-retweeted saying: ‘Data Scientist (n.): Person who is better at statistics than any software engineer and better at software engineering than any statistician.’ (Wills 2012). The definition of data science as a practice is also malleable and is strongly shaped by context, which could be coding operational systems to claw basic metrics from a rising flood of data, or creating new data-driven products using sophisticated machine learning techniques at a level similar to academic research (O’Neil and Schutt 2013). In this swirl of recruitment and entrepreneurial pivoting, it can be hard to discern why data science has become so important, so quickly. The key dynamic is the encounter between the contemporary data flood and the forms of computation that can transform it into actionable statements.

The most important methods that have been called forth by the presence of plentiful data and cheap computing infrastructure can be grouped under the heading of machine learning. These methods thrive on the volume and variety of their input and can thus be made to wrest meanings from big data. They do so by assuming a functional relationship between the input features, which can be any number of different measurable or categorisable aspects of the context under study, and the desired output, which can be anything from a prediction of future house prices to the likelihood of a tumour being malignant. Supervised machine learning algorithms are trained on data sets where the outcome is already known. During the process of training, the algorithm tries to force a fit between the selected features of the input data and the known output by varying the parameters of the assumed relationship. The forcing is mathematical, calculating an algorithmic ‘cost’ for the distance between the fit and the data.

For example, the method of logistic regression tries to find a boundary between two sets of input data, let us say between features that correlate with malignant or benign tumours. The set of features is selected from the available clinical data, along with the known diagnoses. The task of the algorithm is to find a decision boundary between the data for the cancerous tumours and the non-cancerous ones. If we imagine that there are only two key features (e.g. age and length of tumour) and these are used to plot the data on an x - y graph, the set of training data can be visualised as a cluster of green dots (benign) and a cluster of red dots (malignant) which intermingle a bit where the clusters overlap. The decision boundary would be a line that can be plausibly drawn between the points representing malignant tumours and the points representing benign tumours, such that the vast majority in each case fall clearly on one or other side of the line. (In reality, most machine learning involves a larger set of features and the vector space of features is multidimensional, so there is no easy way to visualise what is going on.). The boundary is created by assuming that the probability of a set of features mapping to

one of the outcomes takes the form of a sigmoid function,¹ that is, one that pushes minor differences quickly towards the asymptotic values of one or zero (malignant or benign) (Ng 2012). This is mathematically reasonable, but it is important to understand that it is not based on a causal or physical understanding of tumours. It is simply intended to force a mathematical fit.

The overall cost for a set of input data is calculated from all the individual differences between the actual data points and where they ‘should’ be according to the function being fitted. Iterating over this process finds a minimum cost for the fitted features, i.e. the compromise that best fits the predictions to the actual outcomes. The calculated cost is not a function of the input features but of parameters of the fit, that is, of the relative weights of the different features, so it is in effect deciding which of the features is important, and by how much. Having ‘learned’ how to discern the two kinds of data by processing thousands of known cases, the algorithm can make a rapid judgement on any future data by generating a prediction about whether that case is malignant or benign. In practice, most applications of machine learning will not restrict themselves to a couple of input features. Machine learning algorithms can work with hundreds or even thousands of different features, so they are well suited to a world where huge amounts of heterogeneous digital data have become available. However, the scale and complexity of the calculations also means that the algorithmic decision about which features are important is not necessarily reversible to human reasoning. While in some cases, it will be obvious why the algorithm picks a certain ratio of features, in other cases the algorithm’s ‘reasoning’ will be obscure. This obscurity is intensified in the case of neural networks.

Neural network algorithms are a modified form of machine learning that are becoming increasingly important (Hastie et al. 2003). They are suited to forms of input data that are hard to parameterise and complex to fit. For example, faces or handwritten letters come in many different forms; while humans learn from an early age to recognise them, it is tricky to write a specification that is precise enough for a machine yet flexible enough to deal with all the natural variations. The structure of a neural network applies the same starter logic as the logistic regression described earlier. There is a set of inputs, called nodes rather than features in this case, and a set of starting parameters which are intended to map the input nodes onto the target output. The difference is that this mapping goes via an additional hidden layer of nodes (Skymind 2016). Each of the initial nodes is mapped to each of the nodes in the hidden layer, and in turn, the hidden layer is mapped to the target (the desired outcome classification). The overall effect can be thought of as the hidden layer enabling the neural network to distil its own set of features, which it then uses to discriminate between ‘face’ and ‘not face’, or between letters or numbers. Given a large enough training set, the neural

¹ A sigmoid function is a mathematical function taking the form of a slanted ‘S’ shape. The middle of the S crosses the axis at $(x = 0, y = 0.5)$. For positive x , the top of the S quickly approaches one, which it finally reaches at infinity (i.e. asymptotically). Likewise, for negative x , the bottom of the S rapidly approaches zero, reaching it asymptotically at minus infinity. The effect of this slanted and stretched S shape is that as x increases slightly beyond zero, the corresponding value of y (which is a probability in our case) slides quickly towards one, and as x decreases slightly below zero, the value of y sinks quickly to zero. Thus, it tends to eliminate any ambiguous middle ground; for the vast majority of x , y (the probability) will either be effectively one or zero. The point here is that using the sigmoid function is a choice. It makes the machine learning effective, but it is not mandated by scientific causality.

network abstracts its own set of hidden features, which can be very effective for complex and messy input data.

But the nature of these features is also hidden from the operators. By definition, no human software engineer defines what these abstracted features are, and even if the contents of the hidden layer are examined, it is not necessarily possible to translate that back into comprehensible reasoning. In an animation of an early example of a neural network learning to recognise handwritten numerals, we can see recognisable though blurred representations of the numbers in the input layer (LeCun 1998). But, the contents of the hidden layer look like some kind of condensed barcode. Why this particular weighting? Why is this particular set of subtle combinatorics applied to the input data? We cannot necessarily tell. All we can say for sure is that, in many cases, it is surprisingly effective. Moreover, when we are talking about deep neural networks, where there are multiple hidden layers, the capacity for drawing predictions from messy data can be uncanny (Karpathy and Fei-Fei 2015). Deep neural networks are computationally demanding, but they have recently taken off because of a combination of supply and demand; the cheap distributed computing power has become available, and the need for methods that can handle huge levels of messy data has become urgent. Neural networks are emblematic of the machine learning's tendency to provide 'insight through opacity'. The opacity of machine learning is not only that of the black box. It is also a consequence of algorithms hidden behind the high walls of commercial secrecy. But it is also because they have a tendency to be opaque by nature.

3 The Problem with Data Science

Machine learning wrests apparent meaning from the streams of data that are the inevitable consequence of current digital conditions. It forces functional fits to draw tenuous connections between different phenomena that are otherwise inaccessible to human apprehension. Thus, data science appears to emulate science by transforming empirical data into patterns of regularity which have predictive power. In context, such as in a carefully parameterised support tool for clinical decisions, the abstracted insights of machine learning can add a lot of value. But the production of actionable results is deeply seductive, especially in contexts where risk or profit is at stake. Several problems have become apparent as early adopters rush to implement data-driven policies across the board. Some drawbacks are obvious; for example, data science and machine learning cannot transform bad input into good output. If the data itself carries embedded social bias or prejudice, then so will the nominally neutral output of the algorithms. This has become evident in predictive policing and parole software in the US criminal justice system, where there is an enthusiastic application of data science but where the underlying data reflects deeply rooted racial issues. A recent analysis of a recidivism algorithm used in sentencing in some US courts showed racial disparities in the predictions of future reoffending (Angwin et al. 2016).

The less obvious problem here is the potential production of new forms of unrecognised prejudice. The whole point of data science is the analysis of a scale and complexity that is beyond direct comprehension by the people. Whereas critical observers are becoming alert to the machinic production of prejudice along race or gender lines (Social Media Collective 2016), the distant and multidimensional nature of

the machine learning's correlations may mean that subtler forms of discrimination go unnoticed. The algorithms may settle on a certain combination of features as an identifier that effectively minimises their cost functions, which amount to a real-world segmentation that no one expected or, given the scale of the data and the complexity of the algorithms, actually notices (McQuillan 2015). As the pace of data science adoption far outstrips the evolution of law and regulation, it is hard to know who will ensure fair play when, for example, 'Predikt's AI finds you candidates similar to your best hires' through an 'algorithmically curated pool of active and passive talent' (Predikt Inc 2016). Considering the way a deep convolutional neural network is trained for facial recognition, a cheerful introductory article aimed at practitioners remarks 'So what parts of the face are these 128 numbers measuring exactly? It turns out that we have no idea. It doesn't really matter to us. All that we care is that the network generates nearly the same numbers when looking at two different pictures of the same person' (Geitgey 2016). It may indeed be the case that metrics that seem important to humans, such as eye colour, are not that helpful to a computer analysing a histogram of pixel gradients and that deep learning does a better job than humans in figuring out which parts of the face to measure for machine learning. But what happens when the same, seemingly objective methodologies are carried over to the legal context? If the process of making these predictions is inherently opaque to human reasoning, we are abandoning the basic principle of due process. As an article in the *Stanford Law Review* put it, 'big data's power to enable a dangerous new philosophy of pre-emption' means 'the justification for a fundamental jurisprudential shift from our current *ex post facto* system of penalties and punishments to *ex ante* preventative measures that are increasingly being adopted across various sectors of society' (Earle and Kerr 2013). The potential is created for discrimination that evades due process. Moreover, the classifications of predictive algorithms may themselves change the people's behaviour in ways that the model did not learn about when it was trained (Mackenzie 2015) leading to a recursive reinforcement of the machine learning model as actual social practice (McQuillan 2016).

The problems of insight through opacity do not only occur at the level of individual applications. For some years, Silicon Valley luminaries have promoted ideas like algorithmic regulation, where social problems that drive aspects of government rulemaking are treated on a par with malware and spam (Howard 2012). According to Tim O'Reilly, for example, the idea of algorithmic regulation is core to the functioning of all internet platforms and should act as an inspiration for the design of a twenty-first century government. He valorises the data-driven automaticity of these systems over the alleged inefficacy of policy-making. As he puts it in a talk to the Long Now Foundation: 'If you look at, say, the way spam is regulated on the Internet, that's the beginnings of a kind of an immune system response to a pathogen and works a lot like biology: you recognise the signature of something new and hostile and you fix it.... You compare that to how government regulation works, and you go: 'It's just badly broken!' Somebody puts out some rules, and there's no method of enforcement' (Morozov 2013). Spam filter software, as a prosaic and practical application of machine learning, may decide to bin emails based on obscure combinations of apparently innocuous terms rather than using the clues that stand out to us, such as subject lines that shout about the chance to win millions of dollars (Burrell 2016). We might be prepared to tolerate some false positive from our spam filters, where a genuine email

ends up classified as junk. But it must be a matter for concern where a similar opacity leaks into governmental and legal systems.

Clearly, the traditional notion of data protection is of little relevance when it comes to data science. It is not viable to define a core set of private data when generating data is an innate function of so many essential systems, and any apparently innocuous data can be absorbed into the expansive correlations of the algorithms. Metadata alone is so powerful for surveillance that the NSA do not care about the actual content of our messages (Opsahl 2013), while logistic regression can turn Facebook friendships into predictions of hidden sexuality (Jernigan and Mistree 2009). Nevertheless, there are several emerging areas of activity that attempt to pre-empt a tide of bad outcomes from data science. Some demand the opening of corporate and governmental black boxes so that algorithms can be subject to examination, following the invocation of Judge Brandies that ‘Sunlight is said to be the best of disinfectants’ (Pasquale 2016). While having some effect, this does nothing to address the core opacities of methods like machine learning. Others attempt to probe the social consequences directly by porting methods from the social sciences, like the audit study (Sandvig et al. 2014), although this is most suited to the limited subset of algorithmic influence that presents itself through public interfaces. Finally, there is a small but growing number of computer scientists who are attempting to develop anti-discriminatory remedies at the level of data and algorithms (Feldman et al. 2014; Hajian and Domingo-Ferrer 2012). While having the merit of trying to correct data science from a perspective that understands the technicalities of its operations, it is constrained by seeing data science as an external set of methods rather than as a broader social apparatus in Foucault’s sense, that is ‘a thoroughly heterogeneous ensemble consisting of discourses, institutions, architectural forms, regulatory decisions, laws, administrative measures, scientific statements, philosophical, moral and philanthropic propositions’ (Foucault 1988). Humanist critiques of big data acknowledge that coming to the world through correlation has merit if the overall purpose is acting in the world, but point out that the cost is ‘a massive reduction of what it means to “know”’ (Bowker 2014). Moreover, the imaginaries that arise alongside big data place all their analytical value on the idea of anticipation, in part, because they are ‘shaped by positivist traditions that equate scientific value with predictive laws’ (Boellstorff 2013).

A broader framework for corrective action can be generated by seeing that data science is in fact more than the sum of its parts; that it represents a new way of structuring thought that draws allegiance from older historical currents and, as an organising idea, redefines observations and norms; and that it has a social momentum derived from both its metaphysical and machinic aspects.

These notions can be summed up by saying that data science is, or is becoming, an automated form of applied philosophy: a machinic neoplatonism

4 Neoplatonism

What would it mean to say that data science is neoplatonic? The philosophical school of platonism, as distinct from any arguments about what Plato himself did or did not believe, is committed to a two-world metaphysics (Yates 2002). Behind the world of the sensible, that which we experience through our senses, is the world of the Form or

the Idea. Experiences are the imperfect imprint of this perfect yet inaccessible second layer. As such, the world of the Idea is ontologically superior to the one we actually inhabit. There is no intelligibility in the world that we encounter through sense experience, and we can only come to true knowledge through contemplation of perfect Forms which are eternal and unchanging. Platonism has been a strong and continuing influence on many schools of later philosophy, so there are many forms of neoplatonism. For Plato and the neoplatonists, mathematics is the liminal realm between the imperfect and transitory world of the senses and the perfect and eternal world of pure spirit. Mathematical relations concerning triangles and circles, for example, are true independently of any particular triangle or circle. They are properties of pure triangularity or circularity which cannot be drawn as such. Yet, any triangle or circle that is drawn must reflect them imperfectly inasmuch as they are triangular and circular. Thus, each triangle or circle participates simultaneously both in the intelligible and the visible. (Plato 1998).

The kind of neoplatonism of most interest here emerged in the work of Copernicus and Galileo (Kuhn 1995). As a paradigm, it shaped the development of modern science, and it is resurfacing again in data science. Rather than being led to his beliefs by simple empirical observation, Copernicus 'took pains to read again the works of all the philosophers on whom I could lay hand' (Kuhn 1995, p. 142) and was influenced by his readings of older neoplatonism sources that contained the idea of a moving Earth and the central importance of the Sun in the universe. Copernicus was strongly influenced by the mathematical strand of neoplatonism and believed that the true order of things is a mathematical harmony consisting of arithmetical and geometric relationships. It is important to appreciate the way Copernicus, by his own account, was motivated by these geometric and mathematical symmetries. His dispute with the historically dominant system of Ptolemaic astronomy was not based on empirical observations but his perception that they had not 'been able thereby to discern or deduce the principle thing—namely, the shape of the Universe and the unchangeable symmetry of its parts' (Kuhn 1995, p. 139). In the Ptolemaic system, it is assumed that the Earth is stationary and the motion of the planets is circular, and in order to reconcile this with the observed complexity of planetary motion, it had, over time, been necessary to add an intricate system of epicycles and deferents. While the seven-circle system introduced early in Copernicus's work is both simpler and symmetric, and centred on the Sun, it is actually inferior to the Ptolemaic system in terms of accuracy of predictions. Copernicus had to introduce his own modifications of epicycles and eccentrics to make his version even match the older one for accuracy (Kuhn 1995, p. 154). The point is not to look back from our contemporary context in the knowledge that Copernicus was right, but to understand that the commitment to the Copernican view does not represent the triumph of empirical observation over inferior superstition. Rather, it was a commitment to carry forward a new organising idea that arose from a philosophical standpoint, a commitment which cannot be reduced to the scientific method as we understand it.

This commitment was shared by Galileo, and he further developed the Copernican model. In doing so, he established some fundamental tenets of modern science. Rather than being deterred by the counter evidence of our senses, that the motion of the Earth according to Copernicus should be apparent through the experience of strong wind and the fact that objects in the air would be left behind by the rotating surface, Galileo inverted the problem. Instead of doubting the motion of the planet, he developed a new

physics of motion to explain the way bodies move as if the Earth was at rest; ‘the crucial thing is being able to move the Earth without causing a thousand inconveniences’ (Galileo 2016). His basic idea can be described as ‘indifference’, that is, a body is indifferent to its state of motion in general. If a body is indifferent to its state of motion, it can have several motions at the same time without them interfering with each other. In other words, the net motion can be seen as the superposition of analytical components. Thus, all bodies have the state of motion of the Earth while appearing to us the same as if the Earth was at rest. The idea that movement was a state was a radical shift from the older idea that motion was something involving the essential nature of the body; a necessary feature, a ‘becoming’ of the body itself. It laid the foundations for the later idea of inertial motion and the physics of Newton. This breakthrough was founded on a belief that truth could be discerned by going against the direct experience of the senses. The metaphysics that Copernicus and Galileo bequeathed to science was a belief in hidden layer of reality which is ontologically superior, expressed mathematically and apprehended by going against direct experience.

5 Neoplatonic Data Science

As a method for revealing a hidden mathematical order in the world, data science strongly echoes this neoplatonic project. For the data scientist, computation plays the role of the intermediary between the imperfect world of data and the pure function that relates the features to the target. While the scientific project required a mathematisation of the world, data science requires the datafication of the world. The scientific requirement that empirical facts must be measurable led to the division of qualities into primary and secondary. Primary qualities such as number, magnitude, position and extension can be expressed mathematically, whereas aspects which seem to us an inseparable part of phenomena are relegated to secondary qualities, mere sensory echoes. Hence, Newton replaced colour with ‘degrees of refrangibility’. For data science, the primary qualities are those that can be expressed as data. Rather than drawing on the first person view of reality, it follows the scientific pattern of standing outside, registering events from an external perspective. Events in data science are constituted not from experiences but from those traces of experience which can be datafied. The consequence is the same as it was for science: a displacement of significance away from direct apprehension. Data science echoes neoplatonism by moving to a point of view ‘against experience’. Moreover, the scale of operations with data makes the processes inaccessible to us directly. Thus, data science prioritises data over the phenomenological and uses this to reveal mathematical orderings.

But does data science actually propose that the revealed order is ontologically superior? Data science as such does not claim to reveal causal relationships. In fact, it substitutes correlation at any cost for causal mechanisms and is not constrained by any wider framework of consistency, unlike the physical sciences. And yet, at the same time, it is increasingly enrolled as a justification for action in the world. How can a method which simply reveals patterning become so influential in terms of decision-making authority? This comes from the continuation through data science of what critics of science would call ‘onlooker consciousness’. We perceive ourselves to be standing outside of a reality which we observe and manipulate. This is, in fact, the

constituting condition for the possibility of scientific experiment. Nature is organised into a set of concepts which can be represented quantitatively, and the scientist works with the organisation of these conceptual representations. In our scientific culture, these are the preconditions of superior truth claims. Data science enables us to stand outside a mathematised and manipulable context. Data science seems to fulfil the same criteria as science and, thus, by extension, accrues a similar authority. In effect, the pronouncements of data science are being treated as ontologically superior without actually having to make that claim.

The neoplatonic character of data science makes it hard to constrain. It creates the structural conditions not only for specific injustices caused by bad data or false positives but also the elevation of epistemic injustice, where data science has more sway than the testimony of the subject, or where a community is unable to contest the data science because they lack the capacity to express their knowing in the same way (Fricker 2009). Where data science provides ‘insights’, the testimony or participatory understanding of individuals or groups without access to this insight becomes devalued, even where they are the central subjects of inquiry. Inverting the traditional slogan of the disability movement, data science seeks to know ‘everything about me, without me’. Data science as an organising idea rekindles a commitment to a new way of seeing, and this commitment can transcend contradictory or disappointing results in the short term. The new paradigm redefines ‘the facts on the ground’, because, as both Kuhn and Feyerabend pointed out, the very idea of what constitutes facts can change with a shift in the overall pattern of thought (Kuhn 1996). Traditional safeguards and civic protections become ineffective, because the ground they stand on is modified by a new neoplatonism. The force of the change comes in part from this paradigmatic reframing, but also because this is a worldview that is at the same time its own enactment. Unlike previous forms of metaphysics, neoplatonic data science attenuates the world directly because it is also machinic.

6 Machinic Neo-Platonism

Data science is machinic, because it is an apparatus that not only makes possible a certain way of knowing but also acts directly on the knowledge produced. In that sense, it is very different to science, which seeks to distance itself from implementation in order to retain the veil of neutrality.

The apparatus of data science extends beyond the moment of calculation to include the networked infrastructures that generate the data and the mechanisms that actuate the algorithmic judgements. The action may simply be the reordering of updates in your Facebook feed, and this consequences a minor change in your mood (Kramer et al. 2014). But the cumulative effect of predictions that becomes pre-emptions must be the foreclosure of life chances. As Agamben said about the state of exception, it has the force of law without being of the law (Agamben 2005). Data science always has a target, in the same sense that Husserl characterised consciousness as intentional, that is, always a consciousness of something (Husserl et al. 2001). As a targeting machine, it raises complex questions about accountability. But this is different to traditional AI (artificial intelligence) debates about whether machines are capable of moral reasoning. If the algorithmic part of data science is AI, then it is AI without cognition or

apprehension. It is simply savant at scale, a narrow and limited form of ‘intelligence’ that only provides intelligence in the military sense, that is, targeting. Where machine learning makes reasoning inaccessible, and where the computation itself is subject to errors which cannot be pinned down, accountability for mistakes acquires a core obscurity.

However, this assemblage still includes human agency at most junctures. Decisions for drone strikes are not yet taken by autonomous weapons systems, as far as we know, and even Facebook’s algorithmic filtering may involve a somewhat blurred amount of human intervention (Chaykowski 2016). With data science, we have moved from metadata to metaphysics; it is an embedded, even weaponised, philosophy. Where humans are part of the data science apparatus, what can be said about the effect on human agency of data science as an organising idea? By providing actionable numbers with the aura of authority, the algorithmic predictions become forceful at a human level. The potential exists to sideline ethical concerns or amplify pre-existing biases. Consider the way the police spokesperson defended the targeting of black suspects by the Chicago ‘heat list’ algorithm (Gomer 2013). It could not be racist, because it was algorithmic. Most people charged with reacting to a data science prediction are unlikely to have the benefit of time for reflection. The social worker given a predictive score for the likelihood of parents committing child abuse cannot retreat into academic critique (‘Government Halts Abuse Prediction Study’ 2015). The risk is that pervasive data science at the level of the social will give rise to more of what Hannah Arendt described as ‘thoughtlessness’ (Arendt 2006). Arendt developed this concept through her efforts to comprehend Eichmann and his actions. She used thoughtlessness to characterise the ability of functionaries in the bureaucratic machine to participate in a genocidal process. Of course, we are not concerned here with fascism per se. But thoughtlessness, which is not a simple lack of awareness, is also a useful way to assess the operation of algorithmic governance with respect to the people enrolled in its activities. In wrestling with the legal basis of the trial that she was observing, Arendt argued that the ability to judge is a necessary condition of justice: that legal judgement is founded on the fact that the sentence pronounced is one the accused would pass upon herself if she were prepared to view the matter from the perspective of the community of which she is a member. If we are unable to understand the judgement of the algorithms, which are opaque to us, we are in some way released from categories of intent or accountability. The result is an apparent indifference to the consequences of following a programme of action mandated by an abstracted authority.

7 Critiques of Science

What are we to do when faced with an essentially aesthetic epistemology that masquerades as empirical, asserting superior insight while remaining essentially blind to the prior concepts that constitute it as a possibility? If the historical roots of modern science have made us susceptible to neoplatonic data science, we can look to critiques of science to help us develop an alternative. Of particular relevance are feminist and post-colonial critiques of technoscience. They confront head on the idea that there is only one valid form of science, whose superiority is a product of its internal features, i.e. the scientific method, the use of mathematics to represent natural laws and the

scientific idea of objectivity. The main target of their ire is the assertion that there is nothing culturally specific in the representations of nature that are produced, and they especially focus on those elements in science which can be traced to patriarchal or colonial influence. Standpoint theory says that the scientific method and its ideas about objectivity do not immunise science against these influences. It accepts that the scientific method is good at removing individual bias or problematic experimental results. But while some sexist, racist or distorted elements in scientific research come from not following proper scientific method, others stem from inadequacies in the way those methods and norms are conceptualised. As Sandra Harding makes clear, a central weakness in scientific thinking is the understanding of objectivity. Prevailing standards for objectivity are too weak to identify culture-wide assumptions that shape selection of specific scientific procedures as good ones in the first place (Harding 1998). Standpoint theory is concerned with the way that assumptions, discursive frameworks and conceptual schemes generated by certain ways of life shape the way dominant groups think about both the natural world, and about social relations, and the way those assumptions get hard coded into the way everyone else gets to understand the world. These critiques are not saying that science just makes things up, but that any particular form of science is modulated by the social order in which it develops. Objectivity is strengthened by dispensing with claims to neutrality that hide its social history.

Without strong objectivity, science can indeed come unstuck. Reardon recounts the downfall of the Human Genome Diversity Project, an attempt to sample and archive the world's human genetic diversity whose main protagonists were some of biology's most socially progressive scientists. Despite their good intentions, the project was halted by outrage from an alliance of indigenous advocacy groups and anthropologists (Reardon 2011, p. 322). Lacking any traction on the entanglement of forms of knowing and forms of governance, the scientists were caught off guard by questions about power, especially regarding who gets to make authoritative claims about human diversity. One aspect of Reardon's analysis which is particularly relevant to current developments with machine learning is the way 'categories used to classify human diversity in nature and those used to order relevant aspects of social practice' in turn 'loop back' on each other 'to produce new societal arrangements' (Reardon 2011, p. 328). Regarding the Human Genome Diversity Project, she concludes that, despite trying to provide a scientific basis for the famous UNESCO statements which debunked race as a scientific category (UNESCO 1952), it failed because it excluded too many people from the debate 'whose knowledge might have provided important insights into what it meant to interpret and define human diversity using the tools of scientific (genetic) experts', in other words, by lacking the traits of standpoint theory.

By contrast, a domain of scientific practice where Harding's ideas have been applied in practice is molecular biology. As one researcher in reproductive neuroendocrinology said, 'I realized that...it was no longer sufficient for me to simply engage in feminist critiques of science. I needed to formulate a concrete feminist model of scientific inquiry that spoke directly to my experience' (Roy 2004). For her, the concrete difference was not in the method at the level of the lab bench, as 'there is not a feminist way to pipette, centrifuge, or run a statistical test', but to draw on the work of Harding and others in her approach to epistemology and methodology. This altered the direction of her research into the effect of melatonin on reproduction at the level of gonadotropin-releasing hormone (GnRH) neurons of the brain. Whereas some clinical trials could

have justified high doses of melatonin as a contraceptive, applying standpoint theory revealed an underlying gendered bias in the research programme which led the researcher to focus on other effects of melatonin on the brain and on cellular components such as energy production by mitochondria and to build a case that it should not be used as a contraceptive. Another area of molecular biology which has seen the application of Harding's work is the debate about the contested identity of the HeLa cell line. The HeLa cell type is the original so-called immortal cell line, first isolated in 1951: human cells that do not quickly die off outside of the body but can be cultured in the lab for medical research. Decades later, research suggested that due to their many years of growing in culture, these cells had diverged sufficiently to become a new species of their own: a regression from human to protist cells that in turn have shown themselves to be 'aggressive in invading tissue cultures and extending their biogeographic range' (Strathmann 1991). The fact that the uninformed and unconsenting donor of the original HeLa cells was Henrietta Lacks, a black woman from Baltimore with cervical cancer (Haider 2017) was the starting point for a practising molecular biologist working with these cells to use standpoint theory to highlight the racial and gendered bias in this new scientific proposal. Tackling the way 'metaphors of proliferation and miscegenation enter into and intersect with categories of race and gender in microscopic discourse', she deconstructed both the historical scientific debate about race and the assumption that the presence of papillomavirus type 18 DNA in HeLa cells mean that Henrietta Lacks 'slept around' (Weasel 2004).

Despite the incremental acknowledgement of standpoint theory within some scientific domains, the intersection of genomics and data science highlights the tendency of the latter to pull in the other direction. The multiplicity of variables at play in the algorithmic pattern finding allows other measures to be used as proxies for race, while the social complexities of data construction (e.g. how the racial group of a swab sample is assigned in the first place) are glossed over. A study that combined analysis of scientific papers on genomics and interviews with academic and biotech practitioners found 'a striking trend back toward racial realism in the social shaping of genome technologies' (Chow-White and Green 2013). The authors conclude 'that these new forms of knowledge production are producing rational discrimination', referring to data analysis that generates the identification and classification of groups based on the quantification of risk. To understand the broader social implications, they draw on the work of Gandy (2009) and in particular the notion of cumulative disadvantage that 'reinforces and reproduces disparities in the quality of life' (Gandy 2009, p. 55, quoted from Chow-White and Green 2013). The striking contrast between data science and Harding's ideas has been highlighted before, especially in terms of the big data's distorting positivism and its distancing effect from notions of race, gender and class (Jurgenson 2014). The claims of the pure data school must be contested 'to underscore that there is no Archimedean point of pure data outside conceptual worlds' (Boellstorff 2013). But a wariness of the data science's '30,000 ft view' (Boyd and Crawford 2011) does not provide enough substance to contest its successes. The aim here is that a trenchant theorising and historicising of data science as neoplatonic can provide the traction for standpoint theory to tackle it at every level.

Standpoint theory can act as a counter to the neoplatonic vision of data science. Firstly, standpoint theory can be used to question data science at the level of metaphysics and not just at the level of consequences. Secondly, it does not promote the idea

of abandoning empirical methods but of strengthening them by putting them into dialogue with plural perspectives. A counterculture of data science refuses to throw the baby out with the bathwater; it does not abandon the idea that empirical and mathematical methods of data science can generate valid propositions about the world. But, like standpoint theory, an alternative form of data science must also tackle the question of objectivity. As we have seen, data science is slippery on this point; while not claiming to discern objective reality, it operates through forms of mathematical and computational objectivity. Combined with a dualistic metaphysics, this results in the production of an apparently neutral and external authority with the tendency to encourage thoughtlessness at the point where its judgements are applied. This encourages the scientific perspective which Donna Haraway calls the ‘view from nowhere’: the objective and neutral view which is by its own definition above, outside of, unlocated and therefore cannot be held to account (Haraway 1988). She calls instead for situated and embodied knowledges as the grounding for rational knowledge claims. This would mean that the way out of a machinic metaphysics that eludes accountability is to find a form of operating that takes embodied responsibility. Moreover, this embodiment should start at the ‘edges’. Standpoint theory proposes that positions of social and political disadvantage can actually become sites of analytical advantage, because they can challenge hegemonic assumptions while owning their own perspective. So we can search for a way to develop a counterculture of machine learning by starting from the perspective of those who are disadvantaged by the current construction of data science. While Harding’s point that ‘abstractness and formality express distinctive cultural features not the absence of all culture’ may seem to condemn machine learning as a colonialist project, she is also keen to point out that cultural influence is an inevitable and essential part of developing forms of science. ‘This co-evolution of sciences and the rest of their social orders turns out not to just limit the growth of knowledge, as it always does in some way, but also simultaneously to be a resource for its growth, enabling different cultures, and different historical eras in the same culture, to detect yet more aspects of nature’s order.’ The task is not to seek a fairness that relies on a neutral definition of what is fair, by maximising standardisation, impersonality or some other quality assumed to contribute to fairness. A counterculture of data science is a creolisation of machine learning.

8 Counterculture of Data Science

From what has been said so far, it may seem tempting to dispute data science as just another form of rhetoric. That is, as a form of persuasive argumentation that acts in the world. Datafication itself is a rhetorical move, because it is saying that the important aspects of reality are ones that can be expressed as data. The specific algorithms of machine learning seek to persuade us that a relationship between features can be determined by a particular algorithm, that the cost function can be constructed from a particular probability distribution and so on. Perhaps the reconsideration of machine learning as rhetoric could point the way to its democratic assimilation. Machine learning as another form of proposition becomes amenable to the discourse of peers. Seeing data science as a form of rhetoric rather than a way to X-ray reality would allow its propositions to be returned to their proper place, as basically political statements that

need to be debated. And yet, data science is not simply social constructivism by computational means. It is a material process that participates in social production.

Data science is powerful, because it is an apparatus in the sense that Foucault sets out: a specific set of material and conceptual techniques that coerce by means of observation, ‘an apparatus in which the techniques that make it possible to see induce the effects of power, and in which, conversely, the means of coercion make those on whom they are applied clearly visible’ (Foucault 1988). Data science is an apparatus engaged in the production of subjectivity. While its claims to ontological authority are unsound, a retreat to purely discursive critique loses the power of performativity and drops the material aspect of the philosophy. We need a way to work with the materiality of data science with a different effect. We seek to mobilise the specific constraints and opportunities in a way that extends participation and agency instead of reinforcing dualism and hegemony. We can retain a materialist understanding by viewing data science through Karen Barad’s idea of agential realism (Barad 2007).

Agential realism draws both from Foucault and from the quantum philosophy of Niels Bohr to articulate the idea of material-discursive phenomena as the objective referent for any concept of measurement. In other words, the productive nature of power in co-constituting the subject on which it acts and the non-dual nature of observer and observed revealed by quantum physics, are recast as mutually reinforcing descriptions of a holistic social-material philosophy.

Bohr’s analysis of quantum experiments led him to reject basic assumptions of orthodox science: that the world is made of determinate objects with well-defined properties independent of specific experimental practices and that measurements of these properties can be properly assigned to the object as separate from the agencies of observation. In other words, the stuff that we characterise through experiments cannot be said to exist in that defined form between, and independent of, us measuring it (Barad 2007, p. 196). This breaks with classical, representational science. Instead, Bohr talks about ‘phenomena’ as particular instances of wholeness: the inseparable object measurement event.

In Barad’s account, phenomena are these inseparable physical-conceptual interactions. She introduces the term ‘intra-action’ to signify the mutual constitution of objects and agencies of observation with the phenomena. In Bohr’s understanding, concepts are determined by the circumstances required for their measurement—they are specific material arrangements. A specific arrangement introduces a cut between object and observation that materialises a specific set of properties while excluding others. Likewise, Foucault proposes that the objects (subjects) of knowledge do not exist beforehand but emerge through discursive practices involving apparatuses (Foucault 1988). Barad assimilates Foucault’s ideas by expanding Bohr’s analysis from physical-conceptual devices of observation to the notion of the material-discursive apparatus. Phenomena are produced by the agential intra-actions of material-discursive apparatuses, which are not just measuring instruments but boundary-drawing practices.

Phenomena are specific material configurations, not social constructions, but neither are they independent of human practices. Humans, according to Barad, are part of the ongoing reconfiguration of the world: ‘humans (like other parts of nature) are of the world, not in the world, and surely not outside of it looking in. Humans are intra-actively (re)constituted as part of the world’s becoming. Which is not to say that humans are mere effect but neither are they/we the sole cause, of the world’s

becoming'. Human practices are 'agentive participants' as phenomena are 'sedimented out' of this ongoing process. Agential realism is the notion of material-discursive practices that produce the world through a process of sedimentation, that is, the iterative layering of phenomena that produces 'subject' and 'object' rather than taking them as pre-existing entities.

Barad draws on Judith Butler's ideas about performativity to explain the way agency emerges from the iterative production of reality (Butler 2011). Rather than a deterministic causality, we have constraints, within which there is the space for new possibilities. Moreover, the agency comes through acting rather than being seen as an attribute of subjects or objects. It is a 'matter of making iterative changes to particular practices, in refiguring boundary articulations and exclusions'. Agential realism reinscribes participation rather than reinforcing dualism. 'If our descriptive characterisations do not refer to properties of abstract objects or observation-independent beings, but rather through their material instantiation in particular practices contribute to the production of agential reality, then what is being described by our theories is not nature itself but our participation with nature'.

Agential realism does not presume specifics about the world prior to the enactment of material-discursive practices. Considering data science in this way brings two key benefits: a non-dualism that contrasts starkly with the current neoplatonism and the possibility for participatory agency.

By dispensing with onlooker consciousness, the non-dual perspective counters the ethical split that runs through the neoplatonism of both science and data science. This ethical split has allowed some aspects of the world to be labelled as object, as opposed to subject, and therefore open to instrumental manipulation without any consideration of whether intrinsic harm is being caused. Agential realism sees the production of the real through participation in material-discursive practices that are constrained but not deterministic. As a productive machinic process, data science is open to a participatory reworking.

The paradoxical result of reforming data science as agential realism is to take it more seriously than it takes itself at the moment. That is, to see it not as a description of a hidden layer of reality, but to understand it as part of the production of reality. Agential realism suggests that the world is 'sedimented out of the process of making the world intelligible through certain practices and not others' and data science itself is a prime example of such a material-discursive practice. Understanding data science through agential realism is both to dispute its objective knowledge claims while recognising it as an apparatus whose role in 'sedimenting reality' is open to participatory reworking. Agential realism is a guidebook for developing a countercultural data science as praxis.

The problem with countercultures of orthodox science, such as standpoint theory, is that they stay largely at the level of critique. While indigenous forms of knowledge production cling on wherever marginalised cultures are able to survive, the net impact of standpoint theory has yet to touch the vast core of modern scientific practice. A counterculture of data science, however, can be a critique that also becomes its own practice. In other words, a machinic form of praxis. Praxis is more than simply reflective practice, because it also contains the idea of the good, that is, an overall goal of human flourishing. Instead of *techne*, a way of being concerned with making things and with what things can make, praxis is political action as a mode of togetherness (Arendt 1998, p. 175 ff). A participatory counterculture of data science can develop

praxis by engaging with social justice. As a form of ‘standpoint activism’, the first task of a renewed data science is to actively involve outside perspectives instead of relying on data about them. That is, to recast machine learning as a critical pedagogy where people and communities are involved in both setting the questions and determining the meaning of what is found. The demand of a new data science, in line with agential realism, is the refusal of separation between observer and observed.

9 Conclusions

Data science has been described here as the operation of machinic metaphysics that travels like a resonant wave through the medium of our scientific culture. Constructing a sufficient counterculture means countering its claims at the level of concepts and attempting a deliberate paradigm shift to a more participatory ontology. But a counterculture is not only a set of concepts. An effective counterculture is one that can not only rebut the claims of the dominant culture but also repurpose its artefacts to construct something novel. When Theodore Roszak invented the term counterculture to describe the fusion of hippies and the New Left, he was highlighting the way actual social formations were enacting a vital critique of the technoscientific worldview (Roszak 1969). It may be harder to conceive a counterculture of data science arising under contemporary conditions, when ‘everything is a target’ (Gharavi 2014) and the most visible challenge to the neo-liberal order is a resurgent right wing. So, it is interesting to look at a parallel example where an abstract and potentially alienating mathematical-computational method has been assimilated by a progressive social movement.

The mathematics of graph theory is arguably as abstract as anything in the field of machine learning. It abstracts the context under study to a set of elements (nodes) and their connections (edges). In its applied form as social network analysis, it has, like data science, found eager application in technology platforms whose business consists of networked relations. As a way of understanding social behaviour, it can be as alienating as anything that data science can produce, in the sense that it is also an instantiation of onlooker consciousness with an overt mission to manipulate (Pentland 2014). And yet, the social struggle against a right wing and Islamist regime in Turkey has produced a project called Graph Commons that repurposes network analysis as a collaborative community-led activity (Graph Commons 2016). Graph Commons aims to ‘empower people and organisations to transform their data into interactive maps and untangle complex relations that impact them and their communities’. The initiative gained initial momentum by creating ‘Networks of Dispossession’, a network mapping of the complex political-commercial connections behind the destruction of Gezi Park in Istanbul (Grant 2016). It seems that the idea of mapping networks makes sense to the participants as a form of critical pedagogy: as a way to help reshape shared understandings in the context of an active social struggle. The ongoing form of Graph Commons is conceptual, practical and aesthetic. The core of the project is a technical platform that makes the computation and visualisation of network mappings accessible in a way that does not rely on mathematical ability (Arıkan 2016). But there is an equal emphasis on outreach to diverse potential user communities through participatory hackathons (Graph Commons Hackathon

2016). Without overvalorising this single example, we can count it as an attempt to develop a praxis using abstract conceptual-computational means.

The example of Graph Commons supports the contention of standpoint theory as applied to data science. The development of a relevant praxis of countercultural data science must also reach out to the social edges, however they are defined. As Harding says, ‘to get a critical perspective on...conceptual frameworks, research must begin from the ‘outside’. Standpoint projects do this by starting research from the daily lives of social groups that are not well served by dominant institutions’ (Harding 2010). Bolstered by an agential realism that resonates with the technical means to hand, it is possible for movements in data science to emulate initiatives like Science for the People (Science for the People 2013) or science shops (Wachelder 2003) and develop authentic forms of ‘machine learning for the people’. The metaphysical will meet the machinic at the point of relevance to social struggles. It is plausible that machine learning can find this common ground, given its characteristic of making connections. By exploring and positing new forms of correlation and association, and doing so freely without any claim to superior knowledge, machine learning could become part of an apparatus that promotes mutuality and interdependence. Not so much data science, as data solidarity.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Agamben, G. (2005). *State of exception*. (K. Attell, Trans.) (1 edn). Chicago: University Of Chicago Press.
- Angwin, J., Larson, J., Mattu, S., Kirchner, L. (2016). Machine bias: there’s software used across the country to predict future criminals. And it’s biased against blacks. - ProPublica. May 23. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.
- Arendt, H. (1998). *The Human Condition*. Chicago & London: University of Chicago Press.
- Arendt, H. (2006). *Eichmann in Jerusalem: a report on the banality of evil* (1st ed.). New York, N.Y: Penguin Classics.
- Arikan, B. (2016). Analyzing data networks. *The Graph Commons Journal*. April 26. <http://blog.graphcommons.com/analyzing-data-networks/>.
- Barad, K. (2007). *Meeting the universe halfway: quantum physics and the entanglement of matter and meaning*. Durham: Duke University Press Books.
- Boellstorff, T. (2013). Making big data, in theory. *First Monday*, 18(10). <http://firstmonday.org/ojs/index.php/fm/article/view/4869>.
- Boyd, D., & Crawford, K. (2011). *Six provocations for big data*. SSRN scholarly paper ID 1926431. Rochester, NY: Social Science Research Network. <http://papers.ssrn.com/abstract=1926431>.
- Bowker, G. C. (2014). Big Data, Big Questions| The Theory/Data Thing. *International Journal of Communication*, 8(0), 5.
- Burrell, J. (2016). How the machine ‘thinks’: understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1), 2053951715622512. doi:10.1177/2053951715622512.
- Butler, J. (2011). *Bodies that matter: on the discursive limits of sex* (1st ed.). Abingdon: Routledge.
- Chaykowski, K. (2016). Facebook backtracks after removing iconic Vietnam War photo. *Forbes*. September 9. <http://www.forbes.com/sites/kathleenchaykowski/2016/09/09/facebook-backtracks-after-removing-iconic-vietnam-war-photo-for-nudity/>.
- Chow-White, P. A., & Green Jr., S. (2013). Data mining difference in the age of big data: communication and the social shaping of genome technologies from 1998 to 2007. *International Journal of Communication*, 7, 28.

- Earle, J., & Kerr, I. (2013). Prediction, preemption, presumption: how big data threatens big picture privacy. *Stanford Law Review Online*, 66, 65.
- Feldman, M., Friedler, S., Moeller, J., Scheidegger, C., Venkatasubramanian, S.. 2014. Certifying and removing disparate impact. *arXiv:1412.3756 [Cs, Stat]*, December. <http://arxiv.org/abs/1412.3756>.
- Foucault, Michel. 1988. *Power/knowledge: selected interviews and other writings, 1972–1977*. 1st American Ed edition. New York: Random House USA Inc.
- Fricker, M. (2009). *Epistemic injustice: power and the ethics of knowing*. Oxford: Oxford University Press.
- Galileo. (2016). Dialogue concerning the two chief world systems. Accessed 17 Oct. <http://law2.umkc.edu/faculty/projects/frtrial/galileo/dialogue2.html>.
- Gandy, Jr., O. H. (2009). *Coming to terms with chance: Engaging rational discrimination and cumulative disadvantage*. Burlington: VT: Ashgate.
- Geitgey, A. (2016). Machine learning is fun! Part 4: modern face recognition with deep learning. *Medium*. July 24. <https://medium.com/@ageitgey/machine-learning-is-fun-part-4-modern-face-recognition-with-deep-learning-c3cfc121d78#pq11walzo>.
- Gharavi, M. M. (2014). Everything is a target—full text of interview with Peter Galison. *The New Inquiry*. January 4. <http://thenewinquiry.com/blogs/southsouth/everything-is-a-target-full-text-of-interview-with-peter-galison/>.
- Gomer, J. (2013). Chicago police use heat list as strategy to prevent violence. *Chicago Tribune*. August 21. http://articles.chicagotribune.com/2013-08-21/news/ct-met-heat-list-20130821_1_chicago-police-commander-andrew-papachristos-heat-list.
- Government Halts Abuse Prediction Study. (2015). *From Nine to Noon*. Radio New Zealand. <http://www.radionz.co.nz/national/programmes/ninetoonoon/audio/201764456/govt-halts-abuse-prediction-study>.
- Grant, C. (2016). Turkey's attempted coup and social media. *Click - BBC World Service*. <http://www.bbc.co.uk/programmes/p04165nq>.
- Graph Commons. (2016). Graph Commons—map networks together. Accessed 22 Oct. <https://graphcommons.com/>.
- Graph Commons Hackathon. (2016). Hackathon documentation: creative use of complex networks, Istanbul. January. <http://graphcommons.github.io/hackathons/2016/02/15/istanbul-creative-use-of-complex-networks-documentation/>.
- Haider, F. (2017). The immortal cells of Henrietta Lacks, Witness - BBC World Service'. BBC. Accessed 9 Mar. <http://www.bbc.co.uk/programmes/p04trbdc>.
- Hajian, S., & J. Domingo-Ferrer. 2012. A study on the impact of data anonymization on anti-discrimination. In *2012 I.E. 12th International Conference on Data Mining Workshops*, pp. 352–359. doi:10.1109/ICDMW.2012.19.
- Haraway, D. (1988). Situated knowledges: the science question in feminism and the privilege of partial perspective. *Feminist Studies*, 14(3), 575–599. doi:10.2307/3178066.
- Harding, S. (1998). *Is science multicultural?: Postcolonialisms, feminisms, and epistemologies* (1st ed.). Bloomington: Indiana University Press.
- Harding, S. (2010). Standpoint methodologies and epistemologies: a logic of scientific inquiry for people. In *World social science report, 2010: knowledge divides*, pp. 173–75. France: UNESCO and International Social Science Council.
- Hastie, T., Tibshirani, R., Friedman, J., et al. (2003). *The elements of statistical learning: data mining, inference, and prediction*. 1st ed. 2001. Corr. 3rd printing edition. New York: Springer.
- Holtz, D. (2014, November 7). 8 Skills You Need to Be a Data Scientist. Retrieved 28 September 2016, from <http://blog.udacity.com/2014/11/data-science-job-skills.html>.
- Howard, A. (2012). Rethinking regulatory reform in the Internet age—O'Reilly Radar. July 25. <http://radar.oreilly.com/2012/07/rethinking-regulatory-reform-in-the-internet-age.html>.
- Husserl, E., Moran, D., Dummett, S. M., & Findlay, J. N. (2001). *Logical investigations volume 2: vol 2 (new edition)*. London: Routledge.
- Jemigan, C., & Mistree, B. F. T. (2009). Gaydar: Facebook friendships expose sexual orientation. *First Monday*, 14(10). <http://firstmonday.org/ojs/index.php/fm/article/view/2611>.
- Jurgenson, N. (2014). View from nowhere. *The New Inquiry*. October 9. <https://thenewinquiry.com/essays/view-from-nowhere/>.
- Karpathy, A., & Fei-Fei, L.. (2015). Deep visual-semantic alignments for generating image descriptions. In pp. 3128–3137. http://www.cv-foundation.org/openaccess/content_cvpr_2015/html/Karpathy_Deep_Visual_Semantic_Alignments_2015_CVPR_paper.html.
- Kramer, A. D. I., Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111(24), 8788–8790. doi:10.1073/pnas.1320040111.

- Kuhn, T. S. (1995). *The Copernican revolution: planetary astronomy in the development of Western thought*. Cambridge: Harvard University Press.
- Kuhn, Thomas S. 1996. *The structure of scientific revolutions*. New ed of 3 Revised ed edition. Chicago: University of Chicago Press.
- LeCun, Y. (1998). MNIST demos. <http://yann.lecun.com/exdb/lenet/index.html>.
- Mackenzie, A. (2015). The production of prediction: what does machine learning want? *European Journal of Cultural Studies*, 18(4–5), 429–445. doi:10.1177/1367549415577384.
- McQuillan, D. (2015). Algorithmic states of exception. *European Journal of Cultural Studies*, 18(4–5), 564–576. doi:10.1177/1367549415577389.
- McQuillan, D. (2016). Algorithmic paranoia and the convivial alternative. *Big Data & Society*, 3(2), 2053951716671340. doi:10.1177/2053951716671340.
- Morozov, E. (2013). The Meme Hustler. *The Baffler*. <http://thebaffler.com/salvos/the-meme-hustler>.
- Ng, A. (2012). CS229 lecture notes part II: classification and logistic regression. Stanford University. <http://cs229.stanford.edu/notes/cs229-notes1.pdf>.
- O’Neil, C., & Schutt, R. (2013). *Doing Data Science: Straight Talk from the Frontline* (1 edition). O’Reilly Media.
- Opsahl, K. (2013). Why metadata matters. *Electronic Frontier Foundation*. June 7. <https://www.eff.org/deeplinks/2013/06/why-metadata-matters>.
- Pasquale, F. (2016). *The black box society: the secret algorithms that control money and information*. Reprint edition. Harvard University Press.
- Pentland, Alex. 2014. *Social physics: how good ideas spread—the lessons from a new science*. First Edition Used edition. New York: Penguin Press.
- Plato (1998). *The Republic*. Translated by Benjamin Jowett. <http://www.gutenberg.org/ebooks/1497>.
- Predikt Inc. (2016). Predictive hiring software. Instantly find and source top talent. *Predikt*. <https://www.predikt.co/>.
- Quora. (2014). What is data science? Retrieved 28 September 2016, from <https://www.quora.com/What-is-data-science>.
- Reardon, J. (2011). The Human Genome Diversity Project: what went wrong? In S. Harding (Ed.), *The postcolonial science and technology studies reader*. Durham: Duke University Press Books.
- Roszak, T. (1969). *The making of a counter culture*. Garden City: Anchor Books/Doubleday & Co, Inc..
- Roy, D. (2004). Feminist theory in science: working toward a practical transformation. *Hypatia*, 19(1), 255–279.
- Sandvig, C., Hamilton, K., Karahalios, K., & Langbort, C.. (2014). Auditing algorithms: research methods for detecting discrimination on internet platforms. *Data and Discrimination: Converting Critical Concerns into Productive Inquiry*.
- Science for the People. (2013). *Science for the people magazine*. <http://science-for-the-people.org/sftp-resources/magazine/>.
- Skymind. (2016). Introduction to deep neural networks - Deeplearning4j: Open-source, distributed deep learning for the JVM. <https://deeplearning4j.org/neuralnet-overview.html>.
- Social Media Collective. (2016). Critical algorithm studies: a reading list. *Social Media Collective Research Blog*. Accessed 30 Aug. <https://socialmediacollective.org/reading-lists/critical-algorithm-studies/>.
- Strathmann, R. R. (1991). From metazoan to protist via competition among cell lineages. *Evolutionary Theory*, 10, 67–70.
- UNESCO. (1952). *The race concept: results of an inquiry*. Paris: UNESCO.
- Wachelder, J. (2003). Democratizing science: various routes and visions of Dutch science shops. *Science, Technology & Human Values*, 28(2), 244–273. doi:10.1177/0162243902250906.
- Weasel, L. H. (2004). Feminist intersections in science: race, gender and sexuality through the microscope. *Hypatia*, 19(1), 183–193.
- Wills, J. (2012, May 3). Data Scientist (n.): Person who is better at statistics than any software engineer and better at software engineering than any statistician. [microblog]. Retrieved 27 September 2016, from https://twitter.com/josh_wills/status/198093512149958656.
- Yates, F. A. (2002). *Giordano Bruno and the Hermetic tradition*. Taylor & Francis Ltd.