

## Doubtful Algorithms: Of Machine Learning Truths and Partial Accounts

**Louise Amoore**

During the war I worked on the atomic bomb. This result of science was obviously a very serious matter, it represented the destruction of people [...] Is there some evil involved in science? Put another way – what is the value of the science I had dedicated myself to – the thing I loved – when I saw what terrible things it could do? (Richard Feynman 1988: 239).

I remain a child of the Scientific revolution, the Enlightenment, and technoscience. My modest witness cannot ever be simply oppositional. Rather, s/he is suspicious, implicated, knowing, ignorant, worried, and hopeful. (Donna Haraway 1997: 3).

In a lecture delivered to the National Academy of Sciences in 1955, the physicist Richard Feynman reflected on a particular relationship between the practice of doing science and the value of a kind of orientation to *doubt*. As a graduate student in the 1940s, he worked on the physics of the atomic bomb at Los Alamos. Later, observing the terrible consequences of the weapons to which his physics had contributed, he began to doubt the value of science and its responsibility to society.<sup>1</sup> In the years that followed his work on the Manhattan Project, Feynman locates the practice of science in a particular method that permits a freedom to doubt, an animated curiosity for otherwise ordinary objects, and a sustained sense of encountering an unknown future

---

<sup>1</sup> While working on the Manhattan Project, Feynman and his colleagues calculated the energy release from the atomic bomb's implosion using IBM punch card machines. Feynman describes how they learned to compute "around the error" when the "machine made a mistake" (1992: 131). The use of such calculating machines allowed the scientists to conduct multiple calculations in parallel so that "a mistake made at some point in the cycle only affects the nearby numbers" (1992: 131). Whereas in their early attempts at the mathematics of the atomic bomb the linearity of the calculation meant that errors caused them to "go back and do it over again", the parallel processing of the punch cards allowed them to incorporate the error and continue to progress the program. In short, intrinsic to the science of the atomic bomb was a form of computation that could incorporate error and continue to a useable output. Some years later, after the events of Hiroshima and Nagasaki, Feynman reflected on the relationship between his partial and experimental computation and the science of the atomic bomb: "I looked out at the buildings and I began to think, you know, about how much the radius of the Hiroshima bomb damage was [...] How far from here was 34<sup>th</sup> Street? It's a terrible thing that we made." (1992: 136).

world. In his 1955 lecture, he describes how the scientist “must leave room for doubt”, proposes that “it is perfectly consistent to be unsure, it is possible to live and not know”, and makes a claim about responsibility: “permit us to question, to doubt, to not be sure, herein lies a responsibility to society” (1988: 245). For Feynman, there is a profound ambivalence within the idea of doubt, so that it expresses simultaneously “the openness of possibilities” into the future *and* the doubtfulness of science in the service of the state’s “monstrosities” of war (1988: 247).

In common with other physicists whose experimental fragments became lodged within a violent calculus, Feynman’s reflections signal something of what Rosi Braidotti has called the “anti-Humanism” of the “turbulent years after the Second World War” (2013: 160). Yet, as Donna Haraway acknowledges of her own subjectivity as “a child of the Scientific revolution, the Enlightenment, and techno-science”, what it means to be human is not separable from technoscience and its place in war (1997: 3). Though Feynman’s situated and partial accounts of the science of the atomic bomb do bear witness to the state’s enrolment of algorithmic calculation in the service of war, these are not accounts that stand outside of science in opposition. Rather, the partial account dwells uneasily within the practice of science itself.

Feynman’s ambivalent reflections on science, war, and doubt contain within them some elements which are of significance for our contemporary moment. Amid the pervasive twenty-first century political desire to incorporate doubt into calculation, algorithms are functioning today with the grain of doubt, allowing the uncertainties of data to become a means of learning and making decisions. In contrast to Feynman’s notion that science’s responsibility to society resides in leaving open the incalculability of the future, algorithms hold out the promise of securing against all possible future events (terrorism, irregular migration, financial crisis, climate change), via the analysis of data (Amoore, 2013; Fuller and Goffey, 2012). With contemporary machine learning algorithms, doubt becomes transformed into a malleable arrangement of weighted probabilities. Though this arrangement of probabilities contains within it a multiplicity of doubts in the model, the algorithm nonetheless condenses this multiplicity to a single output with a numeric value between 0 and 1. In short, the single output of the algorithm is rendered as a decision placed beyond doubt; a risk score or target that is to be actioned.

As doubt becomes the very terrain of algorithmic calculation, however, there are other forms of situated and embodied doubt that haunt the terrain. These partial and doubtful accounts seem to me to be present at all of the moments when an algorithm generates itself in relation to a corpus of data. For example, when a neural network algorithm doubts that it can recognise a face because the biometric datapoints were not present in training data, or when a desk analyst watching luminous screens of risk scores suspects that the model is “overfitting”: these are not moments of lack or an error, but a teeming plenitude of doubtfulness. What might it mean to invoke Feynman’s “room for doubt” or “to live and not to know” in the context of a techno-science driven to expose machine learning algorithms to data precisely in order to know and to act, indifferent to persistent doubt? In this essay I propose that in our contemporary culture, where the algorithm plays a major role in the calculability of doubts, the meaning of doubt should be reconsidered. The post Cartesian and posthuman form of doubt I envisage begins from embodied accounts of doubtfulness that decentre the liberal human subject.<sup>2</sup> This doubtful subject is not recognisable as a unitary individual, but is a composite subject in whom the doubts of human and non-human beings dwell together, opening onto an undecidable future, where one is permitted to ask new questions on the political landscape. As Katherine Hayles writes:

What embodiment secures is not the distinction between male and female or between humans who can think and machines which cannot. Rather, embodiment makes clear that thought is a much broader cognitive function... transforming the liberal subject, regarded as the model of the human since the Enlightenment, into the posthuman (1999: xiv).

Why might the revaluing of an embodied posthuman form of doubt matter in relation to an age where composites of humans and algorithms are making ethico-political worlds? In contrast with a critique of technoscience as opaque, blackboxed and unaccountable, I pose the problem differently, tracing the necessarily partial accounts of algorithms. When Haraway describes her “modest witness”, she is concerned with

---

<sup>2</sup> As Rosi Braidotti argues, the posthuman subject is neither “unitary or exclusively anthropocentric”, but its posthuman form must be a “site for political and ethical accountability” (2013: 103).

“telling the truth, giving testimony, guaranteeing important things” in a way that is “optically dense” (1997: 22-24). The truthfulness of the account is thus not grounded in the objective sight of the liberal reasoning subject, but precisely in the very impossibility of a clear-sighted and transparent account. As Judith Butler reminds us, “my account of myself is partial, haunted by that for which I have no definitive story”, so that “a certain opacity persists” in a partial account that is precisely the condition of ethics (2003: 29). In short, it is not the case that algorithms bring new problems of opacity and partiality but precisely that they illuminate the already present problem of locating a clear-sighted account in a knowable human subject. Here we begin to find the possibilities of an alternative orientation to doubt, one in which the subject necessarily doubts the grounds of their claims to know.

Thus, the political decision to detain someone at the border, to authorize a drone strike, to refuse credit or employment, takes place in what Donna Haraway describes as an “enlarged community” of posthuman knowledge, comprising an amalgam of humans and algorithms as “knowing subjects” (Haraway, 1997; Braidotti, 2013: 11). Inspired in part by Feynman’s foregrounding of doubt in science, and placing it into a conversation of curiosity with Haraway’s enlarged posthuman community, I am interested in how posthuman ethics might begin from a doubtful account, or from the impossibility of giving a coherent account of things. This form of doubt places in question the unity of the subject so that the “avowing subject”, as Michel Foucault describes it, “loosens its hold” on itself (2014: 200). This loosening of the hold on subjectivity is precisely a relinquishing of the idea of a grounds from which the subject speaks. As we see also in Foucault’s late lectures on the historical practices of truth-telling, or what he calls “risky speech” or “parrhesia”, an assertion of the truth is necessarily ethico-political because it forges relations to oneself and to others (Foucault 2010: 43; 2011: 6). What kind of relation to self and others is entailed by the algorithm’s particular claims to the truth? Could ethical relations between technoscience and society begin from the plural and posthuman doubts that grow and flourish when the boundaries of human and algorithm, always arbitrary, “highly permeable and eminently revisable”, are relinquished (Haraway 1997: 11)?

## **Ground Truths**

As algorithms become increasingly pervasive in supplying solutions for state decision-making, from the risk management of borders and immigration to decisions on outcomes in the criminal justice system, they hold out the promise of a particular claim to truth. In public forums, when algorithmic decision making is interrogated, the form of truth is commonly said to be an efficient outputting of an optimized decision (Science and Technology Select Committee 2018). In fact, though, the mode of truth-telling of contemporary algorithms pertains to the “ground truth”: a labelled set of training data from which the algorithm generates its model of the world. In a process of supervised machine learning, the algorithms learn from a ground truth model of data labelled by humans. When a new set of features is extracted from the input data, these features are weighted in relation to the ground truth data. For example, a facial recognition algorithm used in urban policing is able to identify a face because of its exposure to a *ground truth* dataset of images labelled by outsourced humans via Mechanical Turk.<sup>3</sup> Increasingly, as deep learning algorithms derive their own ground truth by clustering raw unlabelled data, a model of what is ‘normal’ in the data is generated by the algorithms. The claim to truth made by machine learning algorithms, then, is not one that can be opposed to error or falsity (Parisi, 2013). Rather, the algorithm learns from the degree of probabilistic similarity with a ground truth, itself often generated by algorithms. So, when neural network algorithms at the border or in policing or immigration reach a decision, this is a decision derived from an output signal that is entirely contingent on weighted probabilities (Alpaydin, 2016: 116). The architectures of neural network algorithms can contain multiple hidden layers, hundreds of millions of weights, and billions of potential connections between neurons (Bottou 2015). Put simply, the mode of truth of the algorithm is entirely contingent on a particular notion of ground truth in the data.

It is this malleable relation to a ground truth that supplies the algorithm with the capacity to work with the grain of doubt and uncertainty. The output of the algorithm

---

<sup>3</sup> Though it is commonly considered that the secret or proprietary nature of algorithms renders them impossible to study, in fact most algorithms are openly disclosed and discussed in computer science journals when they are at an experimental stage. It can be many years before they are commercialized and become proprietary. Methodologically, the computer science and engineering journals are a means of tracing the emergence of particular algorithms that go on to have roles in public and private decisions. In my study of the notion of a “ground truth” in machine learning, see for example Yang, Zhou, Balasubramana, and Sastry (2013) and Hinton, Osindero and Teh (2006).

places a decision beyond doubt in the sense that it always already embodies the truth-telling of the ground truth data. One could contest the output of a recidivism algorithm for being “false” in the sense that it wrongly assigns someone a high probability of re-offending, for example, but its degree of truth will always remain intact in its relations to data. To be clear, the algorithm does not eradicate doubt, but neither does it only productively incorporate doubt, as has been observed in the methods of scenarios, catastrophe modelling, and pre-emption in geopolitics (Anderson 2010; Aradau and Van Munster 2011; de Goede 2012; Amoore 2013). Though this is a science that can hold together multiple possible versions of events simultaneously, each possibility weighted as a layer of computation in the algorithm, it cannot live with doubt as such. That is to say, the machine learning algorithm must reduce the vast multiplicity of possible pathways to a single output. At the instant of the actualization of an output signal, the multiplicity of potentials is rendered as one, that moment of decision is placed beyond doubt.

In the sections that follow I address the practice of scientific truth telling in today’s machine learning algorithms, and I seek to give some revitalized life to the embodied doubt that is always already present within the science. In one sense this is a form of critique of vision-dominated objectivity that claims to have a truth beyond doubt, but more precisely it seeks what Donna Haraway has called an “embodied objectivity” of “partial perspective” (Haraway, 1988: 581; see also Daston and Galison 2007), and N. Katherine Hayles captures as an “embodied actuality” (1999: 287). This different orientation to doubt begins from the embodied doubts inhering within all notions of subjectivity and objectivity, doubting oneself and one’s capacity to know. With this orientation to doubt one could reopen the contingencies of the ground truth of data, giving life to the fallibilities of what the algorithm has learned about the world across its billions of parameters. Running against the grain of the Cartesian doubt of the fallibility of data derived from the senses, doubt in this alternative register is felt, lived, and sensed as embodied actuality in the process of an algorithm learning through its relations with the world.

### **Doubtful science**

To consider the doubtfulness of partial perspectives is not the same as casting doubt on the algorithm as such. Indeed, the point is that one could never stand outside of

the algorithm to adjudicate on its veracity since one is always already implicated in the data that makes the algorithm's adjudication of truth. Instead, doubtfulness here is an absence of grounds, expressed in the plenitude of teeming doubts in the experimental algorithmic model. In order to explain what might be at stake in this reorientation to doubt, I will turn to an example drawn from a set of interviews with the designers of deep learning algorithms for border and immigration control. The significance of this example is that a specific neural network algorithm emerges partially and experimentally via posthuman iterations. One of the computer scientists comments that he "plays with" his neural networks, taking his experimental model to the uniformed border operations team in the adjoining building to test it against the specified target outputs. His model was generated from a training data set extracted from past travel and immigration data and it is modified further as its output is calibrated against a target output. When the algorithm designers tune or adjust their algorithm, this space of play involves experiments with the proximity between a specified target value and the actual output signals from their model. They make adjustments of the weights in one of the hidden layers of the algorithm, observing how the numeric risk scores generated by their model diverge or converge on the target. My point is that the computations are infinitely malleable and contingent on plural interactions of humans and algorithms – a small change in the weighting of probabilities in the model will transform the output signal and, therefore, the decision.

Understood as a practice that is partial, iterative and experimental, the neural network algorithm is not only doubtful in the sense that it supplies a contingent probability for an absolute decision, but moreover because it actively generates thresholds of normality and abnormality. Put simply, the algorithm does not need to eradicate doubt or establish certainty for the decision because it *generates the parameters* against which uncertainty will be adjudicated. As computer scientists Ted Dunning and Ellen Friedman ask of their machine learning practice, "what is normal?; how far is far, if something is to be considered anomalous?" (2014: 14). In their account of machine learning for anomaly detection, they describe precisely the kind of experimental playfulness with thresholds I encountered in the borders laboratory:

You must experiment to determine at what sensitivity you want your model to flag data as anomalous. If it is set too sensitively, random noise will get flagged

and it will be essentially impossible to find anything useful beyond all the noise. Even if you've adjusted the sensitivity to a coarser resolution such that your model is automatically flagging actual outliers, you still have a choice to make about *the level of detection that is useful* to you. There are always trade-offs between finding everything that is out of the ordinary and getting alarms at a rate for which you can handle making a response. (2014: 15-16; *my emphasis*).

What is happening here is that the algorithms are working with the grain of doubt, learning to recognise and redefine what is normal and anomalous at each parse of the data. In all cases where machine learning algorithms are generating targets of interest – from *Cambridge Analytica's* targeting of the uncertain voter to *Palantir's* targeting of the migrant of uncertain status – decisions are made in the context of profound doubts.<sup>4</sup> What is the sensitivity of the model? What is the optimal threshold of false positives to false negatives? How should the error be distributed? The response to such questions is never authored by a clearly identifiable human, but rather from a composite of algorithm designers, frontline security officers, the experimental models of the mathematical and physical sciences, a training dataset, and the generative capacities of machine-learning classifiers working on entities and events. In the algorithmic systems that emerge, destined for apps-based desktops from counter-terrorism and border security to finance and life insurance, the composite fragments of human and machine elements are impossible to map onto a unitary body. As Gilles Deleuze writes, “our condition condemns us to live among badly analysed composites, and to be badly analysed composites ourselves” (1991: 28), so that “arbitrarily grouped” elements make up both human and machine (1991: 18). It is precisely the arbitrary groupings and attributes of the algorithm that become effaced when it becomes a technoscience for resolving political difficulty.

---

<sup>4</sup> In the UK, the Information Commissioner is conducting an inquiry into the use of targeted social media in the EU referendum, with a focus on *Cambridge Analytica's* machine learning algorithms. See <https://www.theguardian.com/technology/2017/may/17/inquiry-launched-into-how-uk-parties-target-voters-through-social-media> In the US *Palantir's* algorithms are generating the targets for Trump's ICE deportation programme <https://theintercept.com/2017/03/02/palantir-provides-the-engine-for-donald-trumps-deportation-machine/>



There is a great deal at stake politically in the erasure of the doubtfulness of algorithms at the point of decision. Though the making of a neural net algorithm is a fraught, political and doubtful practice, the computational structure of such algorithms dictates that the final output must be a single numeric value between 0 and 1. In effect, the output of a neural network is itself a numeric probability, a single value distilled from a teeming multiplicity of potential pathways, values, weightings and thresholds. It is this process of condensation and reduction to one from many that allows algorithmic decision systems to retain doubt within computation and yet to place the decision beyond doubt. In the liberal humanist account of algorithms, it is at this point of output as decision that the oft-cited “human in the loop” is invoked as the ethical guarantor (Intelligence and Security Select Committee 2015; National Academy of Sciences, 2015).<sup>5</sup> This human in the loop is, though, an impossible figure who can never meaningfully engage the plurality of posthuman doubts lodged within the calculus. What would it mean to be able to express a posthuman doubtfulness in this context? How does one speak against the grain of the single output generated from millions of potential parameters? Is it possible to locate and amplify the doubtfulness dwelling within the partial fragments of the science itself?

### **Partial accounts**

In order to detail how what Donna Haraway calls “partial perspective” might advance a different orientation to doubt (1988: 584), I turn here to a historical moment when the output of a risk algorithm led to a catastrophic decision. This is a moment when a more conventional Cartesian doubt was expressed that inaccurate data led to a catastrophic decision. However, when the event is understood as a series of partial fragments, what comes to matter is not whether something could be calculated accurately from data, but rather how partial probabilities became a singular calculus for a lethal decision. On January 28<sup>th</sup> 1986, 73 seconds after its launch at 11.25 am, the NASA Space Shuttle Challenger broke apart over the Atlantic Ocean, killing all 7 crew members, including Christa McAuliffe, NASA’s ‘teacher in space’. By February, President Ronald Reagan had established the Rogers Commission to “review the circumstances surrounding the

---

<sup>5</sup> As machine learning algorithms are being rolled out in the policing of urban spaces, the human governing of the algorithm’s output has become the sole locus of ethical recourse. When a series of algorithms are generating the targets for facial recognition decisions, for example, it is the human operator who is said to make a meaningful decision to stop and question a person in the street. The process of surfacing the target of interest (weights, thresholds, bias..) is not perceptible by the operator.

accident and establish the probable cause or causes” and to “develop recommendations for corrective or other action” (House of Representatives, 1986).

Among the 14 commissioners appointed to the investigation was the physicist Richard Feynman, persuaded, by his own account, in spite of “having a principle of going nowhere near government” after his experience of working on the atomic programme at Los Alamos (1988: 113). More specifically, he was wary of bureaucratic reason and the governing of what he thought of as an unruly science by bureaucratic rules and protocols. The Cold War cybernetic rise of “algorithmic rules that could be executed by any computer, human or otherwise” with “no authority to deviate from them” had extended a particular kind of “arithmetic into the realm of decision” (Erickson, Klein, Daston, Lemov, Sturtevant and Gordin, 2013: 29-30). Such an entanglement of mathematics and decision was of real concern to Feynman, whose letters and diaries reveal a disdain for algorithmic decision procedures and axiomatic formulae, and a propensity to ask questions that ran against the grain of mathematical rules. “Doing it by algebra was a set of rules”, writes Feynman of his early encounters with mathematics, “which, if you followed them blindly, could produce the answer” (1988: 17). For Feynman, algebra was a means of imposing an axiomatic set of rules on a puzzle that could otherwise intuitively work towards an unknown solution.

Confronted with the Rogers’ Commission’s setting of procedural steps in the investigation of the Challenger disaster, Feynman worked to reinstate the doubtfulness present at each link in NASA’s chain of events. He travelled to meet with the engineers, avionics scientists and physicists whose data on particular components had made up the aggregate risk calculation on which NASA had based their launch decision. What he discovered was not a catastrophic departure from the normal rules, nor a “human error” or failure as such. Instead, the launch decision belonged properly to a posthuman composite of algorithms, where the steps of a normalized risk calculation protocol had been followed beyond the limits of the calculable. Rather as the sociologist Diane Vaughan proposes in her account of the Challenger disaster, “harmful actions” can be “banal”, and they can “occur when people follow all the rules” (1996: xiv-xv). Though Vaughan’s famous account emphasises the role of the situated culture of NASA in the errors made, however, Feynman’s account foregrounds the doubt already present within the science of the programme. Where NASA failed,

according to Feynman, was in its tendency to “aggregate out” the multiple small fractures and failings that existed at the level of the most mundane of components and instruments. Understood in this way, Feynman’s critical scientific method does not seek to correct the inaccuracies of a calculus but instead brings to the surface the doubts already present within each of the fragmented elements so as to open up the breaches in their calculative arrangement.

There are two aspects of Feynman’s method I reflect on here for the possibilities they may offer to a critical method of partial accounts. The first is a particular *reinstatement of doubt* in data as it is given. The method of reinstatement is intended specifically to *reinstate*, or to give something back a position it had lost. This is of some significance, for it is not the case that doubt functions to cast uncertainty on data that was heretofore settled and certain. Feynman’s approach to doubt begins from the position that all scientific data is contingent, uncertain, and full of doubt. Indeed, among Feynman’s major contributions to quantum physics was his “sum over histories” method, in which the calculation of particle interactions must “take into account a weighted tally of every possible chronology” (Halpern, 2017: 114). In this sense Feynman’s physics theorised how a calculation might keep open the possibilities of multiple incalculable pathways.<sup>6</sup> As Karen Barad suggests in her account of the intimacies of feminist science studies, “life is not an unfolding algorithm” so that “electrons, molecules, brittlestars, jellyfish, coral reefs, dogs, rocks, icebergs, plants, asteroids, snowflakes, and bees stray from all calculable paths” (2012: 207). In my refiguring of posthuman doubt I envisage the unfolding algorithm itself as straying from incalculable paths. The hesitant and non-linear temporality of the etymology of doubt, from the Latin *dubitare*, suggests precisely a straying from all calculable paths. To be doubtful could be to be full of doubt, in the sense of a fullness and a plenitude of other possible incalculable pathways.

What might a method of reinstating doubt, or giving doubt a presence that it had lost, look like? Returning to Feynman’s investigations into the Challenger disaster, the

---

<sup>6</sup> In his 1961-64 lectures at Caltech, Feynman explained Werner Heisenberg’s uncertainty principle as a means of “protecting quantum mechanics” because “if it were possible to measure the momentum and the position simultaneously with greater accuracy, the quantum mechanics would collapse”. As he described to his students, “nobody could figure out a way to measure the position and the momentum of anything – a screen, an electron, a billiard ball, anything – with any greater accuracy” (2011: 1-11).

seals of the solid fuel rocket boosters of the Shuttle were secured by the use of pairs of rubber 'O rings'. It was known to NASA that, during some launches, the hot gases from the boosters could leak from the seals, causing what was called "blowby" over the liquid oxygen and nitrogen of the orbiter's engines. Though the engineers observed corrosion of the O-rings over time, they assumed that the heating of the rubber O rings during launch caused small expansions to close the gaps in the seal. At the fatal launch the temperature was 28 Fahrenheit (-2 celsius), the coldest launch recorded prior to this had been 53 Fahrenheit. On the evening of January 27<sup>th</sup>, the night before the launch, the engineers at Thiokol, who manufactured the seals, warned NASA that the launch should not take place if the temperature was below 53 degrees. Yet, the public record shows that NASA, "took a management decision" on the basis that "the evidence was incomplete", that "blowby and erosion had been documented above 53 degrees", so "temperature data should be excluded in the decision" (Feynman 1988: 135). As Feynman writes, NASA testified that "the analysis of existing data indicated that it is safe to continue flying" (137). In effect, the aggregate data of 24 past flights without a mission failure had placed the launch decision beyond the doubt that would otherwise have been reinstated by attentiveness to the components, and to the inscribed traces they embodied. A sense of doubtfulness that might be considered to be fully posthuman – dwelling in the material marks of blowby and erosion, and in the touch of engineers on fissures and cracks – was aggregated out in the calculation of an output.

In historical instances such as Feynman's account of Challenger, we can see some of the fallacies of an objective risk-based or data-driven decision placed beyond doubt. Remarking on the computer model used by NASA in support of their launch decision, Feynman writes that: "It was a computer model with various assumptions. You know the danger of computers, it's called GIGO: garbage in, garbage out" (1988: 107). In order to question the assumptions of the computer model, Feynman began to focus on the fragments of data and their associated probabilities, reinstating doubt within each element. Understood in this way, the probabilistic calculation of risk gives way to what Donna Haraway has called "partial, locatable knowledges", in which there is the "possibility of new webs of connections" (Haraway 1988: 584). NASA had testified to the Commission that the probability of failure of a mission was calculated to be 1 in 100,000. "Did you say 1 in 100,000?", queried Feynman, "that means you could fly the

shuttle every day for 300 years between accidents?” (Feynman, 1988: 180). Opening a breach in the calculation, Feynman gathered together the scientists who had worked on the various different components of Challenger, asking them to estimate the probability of failure of the Shuttle, and to write this probability onto a slip of paper. When he collected the fragments of paper with their pieces of data, he found that the calculus reflected the particular and situated relationship to a material component and its properties. Thus, for the seal engineers the probability of failure was felt to be 1 in 25, for the orbiter’s engines 1 in 200, with none of the data elements showing NASA’s aggregate probability of 1 in 100,000. Feynman’s point was not to correct the inaccuracy but rather to dramatize its incalculability. From the paper fragments of each embodied likelihood, gathered together as incongruous scraps, there emerges the possibility of new webs of connections.

The second aspect of Feynman’s method I will discuss here is the *affordance of capacities* to algorithmic devices – in particular, an affordance of the capacity to give accounts of themselves and their limits. One approach to such affordances would be to say that, put simply, when something fails, it also speaks to us and tells us something of its limit points (Bennett 2010). Yet, this is not merely a question of how the “non-human” thing, in the Heideggerian (1954) tradition, has a capacity to gather a community of other human and non-human beings to it. Instead, as Katherine Hayles reminds us, “cognition is much broader than human thinking”, so that cognitive capacities are afforded to “other animals as well as technical devices” within a “rich ecology of collaborating, reinforcing, contesting and conflicting interpretations (Hayles 2016: 29). Within this broad ecology of cognition, the always present ethico-political difficulty of humans giving an account must necessarily extend to the partial accounts of other beings.

When Feynman found that the Rogers Commission was unwilling or unable to hear his critical account, he brought with him part of his material scientific community, inviting the material device of the O-ring to give an account in the public forum. As the government officials and press gathered for a meeting with the panel, Feynman placed a piece of rubber O-ring in a clamp bought in a downtown hardware store. Requesting a glass of iced water, he dropped in the materials and waited for them to cool. When the rubber was removed from the glass, brittle and compressed, it broke apart before

the assembled audience. One reading of the event of the public accountability of the O-ring is that Feynman follows the “structure of heroic action” in science, speaking objectively and “self-invisibly” through the “clarity and purity of objects” (Haraway 1997: 24, 34). Certainly the now famous New York Times report of the O-ring event, citing Feynman’s few words: “there is no resilience in this material when it is at a temperature of 32 degrees”, bears a form of witness close to that of Haraway’s critique. However, I locate a different form of account here, one in which there is no unified authorial source of truth, but rather a distributed and oblique account of the impossibility of resolving truthfulness before the public. Refusing accountability as such, the different mode of giving account is closer to what Karen Barad describes as the “condensation of responsibility” in matter, wherein the multiple past decisions are lodged within the object, engaging us in “a felt sense of causality” (2012: 208). In this way, a material limit that was ordinary in the embodied experiences of engineers and materials scientists – familiar to their touch – could be expressed as a claim that could be heard in the world. It is not the case that Feynman’s claim expressed the voice of “the victor with the strongest story” (Haraway 1997: 35), but rather his embodied, tactile and local account was derided and rejected by the Commission. At the specific moment I describe, Feynman speaks against the grain of Cartesian method and instantiates a mode of doubt more commonly feminized and annexed with subjectivity. It is precisely this mode of intuitive causality and embodied doubtfulness that I am seeking as a resistant and critical form of responsibility.

### **Risky speech**

To speak of the excesses and limits of technoscience in our contemporary present has become extraordinarily difficult. Every output of the machine learning algorithm, even where it leads to wrongful detention or racialized false positives, is reincorporated into the adjustment or “tuning” of the weights of a future model. However, the method of reinstating doubtfulness within algorithmic arrangements does have the potential to yield a kind of “risky speech” or “parrhesia”, as Michel Foucault describes, where an account is given that places itself at risk (2011: 79). For Foucault, the defining quality of parrhesia, for Foucault, is that the parrhesiast is bound to the truth of their claim, but they place themselves at risk in so doing, because they speak against the grain of prevailing thought. As Judith Butler comments on Foucault’s mode of criticism, parrhesia “involves putting oneself at risk”, “imperilling the very possibility of being

recognized by others”, and risking unrecognizability as a subject” (2003: 19). What marks out parrhesia from other forms of discursive and performative speech is its capacity to open onto an indeterminate future:

There is a major and crucial difference. In a performative utterance, the given elements of the situation are such that when the utterance is made, the effect which follows is known and ordered in advance, it is codified [...] In parrhesia, on the other hand, the irruption determines an open situation, or rather opens the situation and makes possible effects which are, precisely not known. *Parrhesia* does not produce a codified effect; it opens up an unspecified risk (Foucault 2010: 62, *emphasis in original*).

Understood thus, parrhesiatic speech opens up a breach in a situation and makes possible effects that could not be ordered in advance. It is this opening of the situation beyond codified effects that I consider to be so very necessary in an age of machine learning outputs. Once could be doubtful of the claims of a human guarantor of ethics, and of the bias and discrimination that could be excised from the algorithm, for example, and begin instead from the unspecified risks emerging between the algorithm and the data corpus from which it learns.

Ultimately, on discovering that the Rogers Commission report contained a previously undisclosed “tenth recommendation”, sheltering the risk algorithm within “error” and commending NASA’s work and mission to the nation, Richard Feynman refused to add his signature to the report. His fellow commissioners sought to persuade him, suggesting that the tenth recommendation was “only motherhood and apple pie”, and that “we must say something for the president” (Feynman 1988). The tenth recommendation, which would authorize a particular decided future of continued NASA missions, was demarcated from the real science of objective findings, and rendered maternal and nation-building.

In response, Feynman asked what kind of truth could be spoken to the government, “why can I not tell the truth about my science to the President?” In the event, his own report appeared as an appendix to the published document, titled “personal observations on the reliability of the shuttle”, as though his merely personal situated

account could only be subjective and without grounds. Feynman's appendix is risky speech in Foucault's sense, an argument that places him at risk, runs against the grain of the Commission report, and challenges the form of truth-telling of the public report. In this sense, to be *doubtful* is to open onto the contingency of a situation and to be responsive to unspecified effects. The fullness of an embodied doubt I envisage here is a critical cousin to Lauren Berlant's notion of a "cruel optimism", where a "cluster of promises could seem embedded in a person, a thing, an institution, a text, a norm, a bunch of cells, smells, a good idea" (2011: 23). The cruelty Berlant so vividly describes is a "relation of attachment to compromised conditions of possibility whose realization is discovered to be impossible, sheer fantasy, or too possible, and toxic" (p.24). Though Berlant does not probe directly what it might mean politically to be doubtful of the promises as they embed in a person, a thing or a technology, her work does envisage the making of claims on the present so that there remains "the possibility of the event taking shape otherwise" (2011: 262). Understood thus, to be doubtful could be to experience a fullness or multiplicity of the present moment and the many ways it might unfold, such that the cruelly optimistic promises of technoscience do not cling so tightly to ideas of the "optimal" or the "good idea". The optimism of the algorithm is founded on its optimization, that is to say on its capacity to reduce the fullness of the present moment to an output to be actioned in the future.

### **Posthuman doubt**

If the signalling of multiple errors and faults in contemporary machine learning algorithms cannot meaningfully provide a ground for critique, then can there be risky speech amid algorithmic techniques? I note here that throughout the process of following algorithm designers, there has been a curious mood of twinned optimism and doubt: optimism that a "good enough" model can always be found, and doubt that elements of "junk" data in social media or image files could ever be adequately prepared or labelled. If, however, the very many misidentified faces in biometric systems, or wrongly seized assets in an algorithmic immigration system merely "harness dysfunctions, errors, and crisis", as Luciana Parisi suggests, in order to "script uncertainties within the programming of relations" (2013: 96), then doubting the veracity or accuracy of the data could never be sufficient. The faults, errors and inaccuracies that would have been considered impediments to programming in past statistical modes, have become important stimuli for the machine learning algorithm



to train, learn and develop. Put differently, Feynman excavated and amplified the doubt within analogue data elements by painstakingly reconstructing the plural probabilities of each finite piece. Yet, if this method allows us to pause and dwell within the doubtfulness of the calculus, can we imagine today an equivalent method of asking the computer scientists we research to please “write the probability of the failure of your piece of the machine-learning software on this piece of paper”? And how would the imagination of such probabilities extend to the future actions of one algorithm on another? How would it account for errors that are productively reabsorbed into the model? If one seeks a method to try to reinstate the doubt already lodged within the data, within its weighting and thresholds, then what kind of truth-telling practice might this be?

When Richard Feynman asked the scientists on the Shuttle programme why the fragments of their doubts had not found their way into the future calculations of risk assembled by NASA, he was informed that “it is better if they don’t hear the truth, then they cannot be in a position to lie to Congress” (1988: 214). The problem of truth-telling in this instance lies less with science as such, and more with the various ways in which embodied doubtfulness in science does not give political accounts of itself. Similarly, when the US Director of National Intelligence James Clapper was found to have lied to Congress on the analysis of data on US citizens in March 2013, he later argued that he had provided the “least untruthful answer possible in a public hearing” (Harding 2014; Greenwald 2013). To speak against the grain of the promise algorithms make in the name of geopolitical security is to bring to the public assembly a kind of contemporary heresy. Following the UK terror attacks in Manchester and London in 2017, for example, the official inquiry by David Anderson QC has concluded that what is required is a more effective algorithm for the identification of high risk individuals among 20,000 so-called closed subjects of interest.<sup>7</sup> In effect, this implies training an algorithm to recognise the attributes of known individuals in the dataset in order to identify similar attributes among unknown future subjects. The weighted probabilities among the 20,000 will be used to generate target outputs for security action. To signal

---

<sup>7</sup>[https://www.gov.uk/government/uploads/system/uploads/attachment\\_data/file/664682/Attacks\\_in\\_London\\_and\\_Manchester\\_Open\\_Report.pdf](https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/664682/Attacks_in_London_and_Manchester_Open_Report.pdf) Last accessed December 2017.

the doubtfulness and of the arrangement of probabilities of such an algorithm is to risk annulling the promise of a securable future.

Let us turn to the truth telling of our times and why it has become so difficult to reinstate the doubt which would be necessary for a meaningful politics. The key doubt here is the doubtfulness of the grounds on which one makes a political claim. As Thomas Keenan writes, “what we call ethics and politics only come into being or have any force and meaning thanks to their ungroundedness”, so that “we have politics” precisely because we have no “reliable standpoints” (1997: 3). If a future political claim is to be possible, then “its difficulty and its persistence stem from its terrifying, challenging, removal of guarantees” (1997: 42). If there are to be future claims to human rights not yet registered or heard, for example, then there must always also be doubts haunting the possibility of rights, doubts that acknowledge the irresolvability of all political claims. Here, doubt pervades what Keenan describes as a “darkened frontier” where the essential difficulties of politics must take place, and where all decision must, following Derrida, “act in the night of non-knowledge and non-rule” (Keenan 1997: 12; Derrida 1992: 26). In holding out the promise of resolving political difficulty – coming out of the darkness with the condensed output of a machine learning algorithm – this technoscience harms the very terrain of the political. As I have argued elsewhere, the algorithmic technique promises to resolve the political by holding together a multiplicity of possibilities, whilst carving out a space for a singular decision (Amoore 2013; 2014). This is not to say that the decision is made without doubt. On the contrary, it is doubtful, full of doubt and yet it decides with indifference.

Seated in the audience in a technology company’s auditorium, I watched the pitch made to government departments in the name of resolving political difficulty with machine learning. *Recorded Future* explained how their “threat centaur” machine learning algorithms would “scrape the web for geopolitical events” and to send “real time alerts” on “political protests, terrorist attacks, and civil unrest”.<sup>8</sup> As the algorithm designers describe the system, it deploys a combination of rule-based and deep learning algorithms. The rule-based algorithms are said to be based on “human intuition” about whether “an entity is associated with some kind of risk”, while the

---

<sup>8</sup> <https://www.recordedfuture.com/artificial-threat-intelligence/> Last accessed December 2017.

machine learning algorithms are trained on a “large data set, using trusted threat list sources as ground truth for what constitutes a malicious entity” (ibid: 4). The figure of the “centaur” in the algorithmic system is a posthuman body with a capacity to reach an unsupervised algorithmic judgement (“this event is critical”) but also to display “a human readable motivation for that judgement” so that action can be authorized against a threat. [insert figures 1 and 2 here]

At one level the algorithms of the threat detection centaurs appear to run counter to the form of risk calculation observed by Feynman at NASA. The smallest of infractions or the subtlest of signals of a component failure or a change in online clickstream patterns, for example, is said to yield a threat alert to the centaurs. Like the Rolls Royce jet engines that yield real-time data on the performance of aircraft as they fly, the twenty-first century O-ring would almost certainly be given a data voice. And yet, at the level of practices of truth-telling, there is something of significance here. The Challenger launch decision teemed with a cacophony of embodied doubt – on the part of the engineers, physicists and mathematical modellers, in communion with their non-human partners – and yet, the decision itself was placed beyond doubt. What we are seeing emerge in contemporary algorithmic technoscience is an orientation to truth-telling that says that all geopolitical decisions can be placed beyond doubt.

How might one intervene in models such as the threat centaur in ways that reinstate the doubtfulness of the algorithm? It is necessary to engage the political traces of machine learning algorithms, to be curious about them, to doubt them, but also to listen to the doubt they themselves express in the world.<sup>9</sup> In engaging research methods that follow elements of an algorithmic solution as they travel and have onward life, I am trying to be a receptive listener when the piece of data appears to say “you do know that I cannot really be cleaned, I remain in there, muddying the calculation”; or when the biometric matching algorithm says “in 14% of cases, where ambient lighting exceeds these parameters, I cannot be read”; or when the pattern recognition algorithm says “the training data set I have been exposed to has established what looks normal or abnormal to me”. For there is doubt, it proliferates everywhere, at

---

<sup>9</sup> Writing on the method of “learning from machine learning” Adrian Mackenzie proposes that understanding how particular algorithms, such as random forest decision trees, “order differences” could substantially “change how we relate to what we see, feel, sense, hear, or think” (2017: 11).

every step in the application of algorithmic solutions to political questions, doubt lives and thrives and multiplies. And, where there is an ineradicable doubt, there is the darkened frontier and also the potential for politics. Inspired by Richard Feynman's method of reinstating doubt in the data points, and being an interlocutor for the people and things 'in the engine' who actively give accounts of limit points of all kinds, I urge that we too experiment, like the algorithmic models do, iteratively and recursively giving *doubtful accounts* of the output of a calculation.

### **Conclusions: the open channel**

To speak in praise of an embodied form of posthuman doubt is to draw together a number of things which do not sit together easily in terms of critique. The form of doubt I propose is to be distinguished from Cartesian skepticism, for it does not seek to annexe fallible sensory doubts in the pursuit of foundational truths. As posthuman life forms, the doubtful subjects I depict here dwell uneasily within partial and situated knowledges and, when they make a claim in the world, they do so in ways that stray from calculable paths. I include within this category of doubtful subjects all those whose science bears a responsibility for an apparently risk-free, error tolerant political decision, from drone strikes to border controls and from voter targeting to immigration decisions. In so doing, I follow Foucault's sense that the scientist and the geometer do not stand outside of the capacity for parrhesiatic or risky speech, alongside Karen Barad's caution that we are materially immersed in and inseparable from science. Doubtfulness expresses the many ways in which algorithms dwell within us, just as we too dwell as data within their layers, so that we could not stand apart from this science even if we wanted to. The doubtfulness of our relations to ourselves and to algorithmic systems means that critique will always involve "putting oneself at risk", as Judith Butler describes it, and risking "unrecognizability as a subject" (2003: 19, 45). In this way, the unrecognizability of the subject extends into the opacity of the algorithm, so that new forms of composite subject emerge. The next time one hears that the "black box" of the algorithm should be opened, one might usefully reflect on the obscured unrecognizability of all forms of self and ask how such an opening could be possible.

To express doubt in these terms is to try to find a means to respond to algorithms that incorporate doubt into computation, converting doubts into weighted probabilities that will yield a condensed output that can be actioned. A critical response cannot merely

doubt the algorithm itself, pointing to its errors and contingencies, for it is precisely through these variations that the algorithm learns what to do. Instead, a posthuman orientation to doubt decentres the authoring subject, even and especially where this author is a machine, so that the grounds – and the “ground truths” – are called into question. Understood in this way, to doubt is also to reopen a “decision worthy of the name”, as Derrida notes, or to open onto “a politics of difficulty”, as Thomas Keenan describes, where the “break with the humanist paradigm” does not close the ethico-political but provides its very starting point (Derrida, 1995: 24; Keenan, 1997: 2).

What could it mean to value doubtfulness as a posthuman critical faculty, to reinstate doubt into the composite creature that is automated algorithmic calculation? It is not the same thing at all as to call for an ethics of human decision in machinic security (putting the human back in the loop), or to say science at the service of geopolitics is intrinsically malevolent. Instead, it expands the space for doubt beyond what Deleuze, following Bergson, calls “the decisive turn”, where lines diverge according to their differences in kind (Deleuze 2011: 29). Understood in this way, the action precipitated by the output of the algorithm is never placed beyond the darkness of doubt, for it carries doubt within, it does not know what is around the decisive turn. Though contemporary algorithms reduce the multiplicity to one at the point of output signal, they do not leave behind the knots of doubt that dwelled within the multiple hidden layers. As Karen Barad writes on Feynman’s quantum field theory, where the electron absorbs its own proton and “there is something immoral about that”, the action of the body is never fully beyond doubt (2012: 212).

In so many ways, of course, Richard Feynman is a curious figure to invite into conversation with Haraway, Braidotti, Hayles, and Barad on science and the ethico-political, particularly given his role in the material violence of the nuclear age. Yet, the point is that we none of us can stand outside of this science and judge it to be good or evil, to say “things should not be this way”. We are it and it is us, such that critique cannot begin from an outside to the algorithm. As Barad suggests, “ethicality entails noncoincidence with oneself” (2010: 265) so that to give an ethical account is to resist anchoring the claim. Doubt does not stand outside and pronounce judgement, doubt is also interior, essential to subjectivity, for we doubt ourselves. For Feynman, doubt is the open channel in politics and in science:

If we take everything into account... then I think we must frankly admit that we *do not know*. But, in admitting this, we have probably found the open channel [...] The openness of possibilities was an opportunity, and doubt and discussion were essential [...] We must leave the door to the unknown ajar [...] It is our responsibility to leave the people of the future a free hand (Feynman, 1992: 247).

To find the open channel and to leave the door to the unknown ajar is to resist the promises made that the future is resolvable through the optimized output of algorithmic decision engines. To leave the door ajar is to reopen the very many points in the algorithm where another future was possible, where the hinge does not quite demarcate the axis of possible movement. Whilst the enfolded doubts of the algorithm's feedback loop and back propagation deny the impossibility and difficulty of decision, an embodied doubtfulness pauses with the undecidability of alternative pathways and their contingent probabilities. The claim to a ground truth in data that pervades our contemporary political imagination precisely closes the door to the future, offering algorithmic solutions to resolve the difficulties of decision. To reinstate embodied doubt within the algorithm, and to allow the components of the badly formed composite to speak of their limits, is to seek to leave the door ajar for the making of future political claims. For it is these future unknown claims that can never appear in the clusters, attributes, and thresholds of the algorithm.

### **Acknowledgements**

This article was first prepared as a lecture at the Holtz Center for Science and Technology Studies, University of Wisconsin-Madison, April 2015. My thanks to all at the Center for their generous engagement with the work. The paper benefitted greatly from a presentation of the ideas at the Goldsmiths University of London Centre for Research Architecture, February 2016. The anonymous reviewers for TCS have contributed with extraordinary generosity to my thinking on doubt, thank you. The research from which this article is drawn has been funded by ESRC (Securing Against Future Events, ES/K000276/1) and Leverhulme Trust (Ethics of Algorithms, MRF-2015-164).

## References

- Alpaydin, Ethem (2016) *Machine Learning*, Cambridge, MA: MIT Press.
- Amoore, Louise (2013) *The Politics of Possibility: Risk and Security Beyond Probability*, Durham NC: Duke University Press.
- Anderson, Ben (2010) 'Preemption, Precaution, Preparedness: Anticipatory Action and Future Geographies', *Progress in Human Geography*, 34: 6.
- Aradau, Claudia and Rens Van Munster (2011) *Politics of Catastrophe: Genealogies of the Unknown*, London: Routledge.
- Barad, Karen (2007) *Meeting the Universe Halfway*, Durham NC: Duke University Press.
- Barad, Karen (2012) 'On Touching: The Inhuman That Therefore I Am', *Differences: A Journal of Feminist Cultural Studies* 23:3, pp.206-223.
- Bennett, Jane (2010) *Vibrant Matter*, Durham NC.: Duke University Press.
- Berlant, Lauren (2011) *Cruel Optimism*, Durham NC: Duke University Press.
- Braidotti, Rosi (2006) 'Posthuman, All Too Human: Towards a New Process Ontology', *Theory, Culture & Society* 23:7-8, pp.197-208.
- Braidotti, Rosi (2013) *The Posthuman*, Cambridge: Polity.
- Butler, Judith (2003) *Giving an Account of Oneself*, Fordham Press.
- Daston, Lorraine and Peter Galison (2007) *Objectivity*, MIT Press.
- de Goede, Marieke (2012) *Speculative Security: The Politics of Pursuing Terrorist Monies*, Minneapolis: University of Minnesota Press.
- Deleuze, Gilles (1991) *Bergsonism*, New York: Zone Books.
- Derrida, Jacques (1992) 'Force of Law: The Mystical Foundations of Authority', in *Deconstruction and the Possibility of Justice*, Ed. D. Cornell, M. Rosenfeld, and D. G. Carlson, London: Routledge.
- Derrida, Jacques (1995) *The Gift of Death*, Chicago: Chicago University Press.
- Dunning, Ted and Ellen Friedman (2014) *Practical Machine Learning*, Sebastopol CA: O'Reilly.
- Erickson, Paul, Judy L. Klein, Lorraine Daston, Rebecca Lemov, Thomas Sturm, and Michael D. Gordin (eds) (2013) *How Reason Almost Lost Its Mind: The Strange Career of Cold War Reasoning*, Chicago: Chicago University Press.
- Feynman, Richard (1992) *Surely You're Joking Mr Feynman!: Adventures of a Curious Character*, New York: W.W Norton.

Feynman, Richard (1988) *What do you Care What Other People Think?: Further Adventures of a Curious Character*, New York: W.W Norton.

Feynman, Richard, Robert Leighton and Matthew Sands (2011) *The Feynman Lectures on Physics: Volume III Quantum Mechanics*, New York: Basic Books.

Foucault, Michel (1994) 'The Masked Philosopher', *The Essential Foucault: 1954-1984*, Ed. Paul Rabinow and Nikolas Rose, New York: New Press.

Foucault, Michel (2010) *The Government of Self and Others, Lectures at the Collège de France 1982-83*, Basingstoke: Palgrave.

Foucault, Michel (2011) *The Courage of Truth, Lectures at the Collège de France 1983-84*, Basingstoke: Palgrave.

Foucault, Michel (2014) *Wrong-Doing, Truth-Telling: The Function of Avowal in Justice*, Ed. Fabienne Bron and Bernard Harcourt, Chicago: Chicago University Press.

Fuller, Matthew and Andrew Goffey (2012) *Evil Media*, MIT Press.

Greenwald, Glenn (2014) *No Place to Hide: Edward Snowden, the NSA and the Surveillance State*, Random House.

Halpern, Paul (2017) *The Quantum Labryrith: How Richard Feynman and John Wheeler Revolutionized Time and Reality*, New York: Basic Books.

Harding, Luke (2014) *The Snowden Files*, Faber and Faber.

Haraway, Donna (1988) 'Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective', *Feminist Studies* 14:3, pp.575-599.

Haraway, Donna (1997) *Modest\_Witness@Second\_Millennium. FemaleMan@\_Meets\_ Oncomouse* London, New York: Routledge.

Hayles, N. Katherine (1999) *How we Became Posthuman: Virtual Bodies in Cybernetics, Literature and Informatics*, Chicago: Chicago University Press.

Hayles, N. Katherine (2010) 'How we Became Posthuman: Ten Years On An Interview with N. Katherine Hayles', *Paragraph* 33:3, pp318-330.

Hayles, N. Katherine (2012) *How we Think: Digital Media and Contemporary Technogenesis*, Chicago: Chicago University Press.

Hayles, N. Katherine (2017) *Unthought: The Power of the Cognitive Nonconscious*, Chicago: Chicago University Press.

Heidegger, Martin (1954) 'Bauen, Wohnen, Denken', *Vorträge und Aufsätze* (Pfullingen: Neske).

Hinton, Geoffrey E., Simon Osindero, and Yee-Whye Teh (2006) 'A fast learning algorithm for deep belief networks', *Neural Computation* 18:7, pp.1527-54.



House of Representatives (1986) *Investigation of the Challenger Accident*, available: <https://www.gpo.gov/fdsys/pkg/GPO-CRPT-99hrpt1016/pdf/GPO-CRPT-99hrpt1016.pdf>

Intelligence and Security Select Committee of Parliament (2015) *Privacy and Security: A Modern and Transparent Legal Framework*, London HMSO.

Keenan, Thomas (1997) *Fables of Responsibility: Aberrations and Predicaments in Ethics and Politics*, Stanford CA: Stanford University Press.

Mackenzie, Adrian (2017) *Machine Learners: Archaeology of a Data Practice*, Cambridge MA: MIT Press.

Mackenzie, Donald (2011) 'How to make money in microseconds', *London Review of Books*,33:10, pp.16-18.

National Academy of Sciences (2015) *Bulk Collection of Signals Intelligence: Technical Options*, Washington DC: National Academies Press.

Parisi, Luciana (2013) *Contagious Architecture: Computation, Aesthetics and Space*, MIT Press.

Vaughan, Diane (1996) *The Challenger Launch Decision*, Chicago: University of Chicago Press.

Yang, Allen, Zihan Zhou, Arvind Ganesh Balasubramana, S. Shankar Sastry, and Yi Ma (2013) 'Minimization algorithms for robust facial recognition', *Transactions on Image Processing* 22: 8, pp.3234-3246.