

Politics and Ethics in the Age of Algorithms

The mathematical proposition has been given the stamp of incontestability. I.e.: “Dispute about other things; this is immovable—it is a hinge on which your dispute can turn.”

—Ludwig Wittgenstein, *On Certainty*

“A Hinge on Which Your Dispute Can Turn”

It is March 2016, and I am seated in a London auditorium, the gray curve of the river Thames visible from the windows. A tech start-up business, specializing in developing machine learning algorithms for anomaly detection, is presenting its latest algorithmic innovations to the assembled government and corporate clients. The projection screen displays a “protest monitoring dashboard” as it outputs risk scores of “upcoming threats of civil unrest” in cities around the globe, their names scrolling: Chicago, London, Paris, Cairo, Lahore, Islamabad, Karachi. The score that the analyst reads from the dashboard is the singular output of deep neural network algorithms that have been trained to recognize the attributes of urban public life, the norms and anomalies latent in the data streams extracted from multiple sources, from Twitter and Facebook to government databases. As the presenter explains to the audience of national security, policing, and border officials, “We train our algorithm to understand what a protest is and is not,” and “it gets better,” “adapting day by day,” as it iteratively learns with humans and other algorithms.¹ The process of learning “what a protest is” from the clustered attributes in data and modifying the model continues when the algorithm is later deployed in the city or at the border: “We give you the code,” he pledges, “so that you can edit it.” How does an algorithm learn to recognize what a protest is? What does it mean to cluster data according to the attributes and propensities of humans to gather in protest or in solidarity? At the London event, as the presenter displays a still

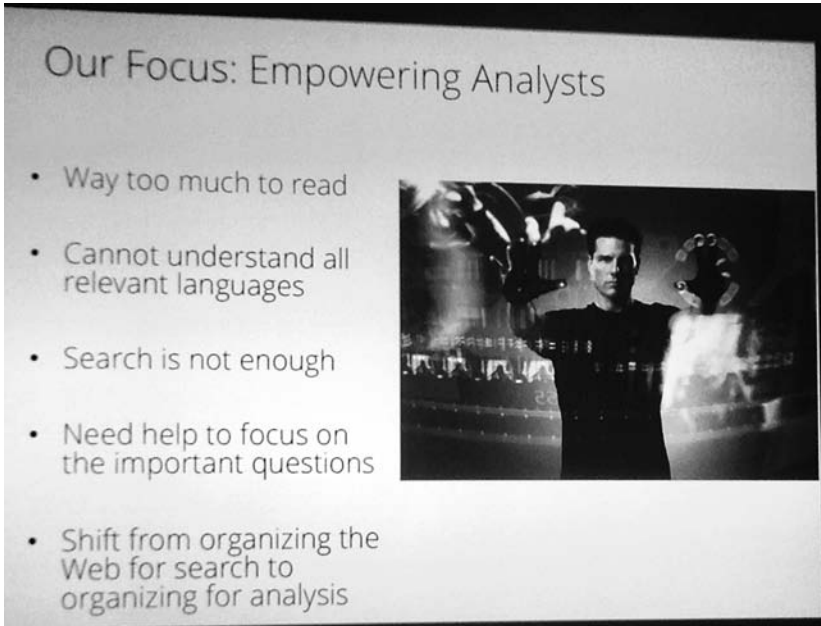



Figure I.1 An image from Stephen Spielberg's film adaptation of Philip K. Dick's novel *Minority Report* appears in a technology company's presentation to government and corporate analysts. Author's photograph.

from *Minority Report* (figure I.1), my thoughts turn to the protests that took place one year earlier, in the US city of Baltimore.

On April 12, 2015, Freddie Gray, a twenty-five-year-old African American man, sustained fatal injuries in the custody of the Baltimore Police Department. The profound violence of Gray's murder is an all-too-familiar event in the racialized architectures of our contemporary cities. During the days that followed his death, however, as people gathered on Baltimore's streets to protest the killing, the violence of the act extended into the plural actions of a set of machine learning algorithms that had been supplied to the Baltimore Police Department and the US Department of Homeland Security by the tech company Geofeedia. With the use of techniques similar to those described in the London protest-monitoring software, the Geofeedia algorithms had been trained on social media data, analyzing the inputs of Twitter, Facebook, YouTube, Flickr, and Instagram and producing scored output of the incipient propensities of the assembled people protesting Gray's murder. "Several known sovereign citizens have begun to post social media attempting to rally per-

sons to demonstrate,” recorded the Baltimore Police Department in a memo that promised to “continue to evaluate threat streams and follow all actionable leads.”² Indeed, Geofeedia went on to market its algorithms to other states on the basis of a Baltimore “case study” (figure 1.2) in which Freddie Gray is said to have “passed away,” the city to have “braced itself for imminent protests,” and the police to have seized “opportunities” to analyze “increased chatter from high school kids who planned to walk out of class.”³

During those days in April, terabytes of images, video, audio, text, and biometric and geospatial data from the protests of the people of Baltimore were rendered as inputs to the deep learning algorithms. Even the written text embedded within social media images—such as the “police terror” placards carried aloft and captured on Instagram—was extracted by a neural network and became features in the algorithm.⁴ People were arrested and detained based on the outputs of a series of algorithms that had—as the London scene also proposed—learned how to recognize what a protest is, what a gathering of


CASE STUDY: BALTIMORE COUNTY PD

Baltimore County Police Department and Geofeedia Partner to Protect the Public During Freddie Gray Riots

BACKGROUND

When Freddie Gray passed away in Baltimore on April 25, 2015 from injuries allegedly sustained during his arrest by the City of Baltimore Police, Detective Sergeant Andrew Vaccaro with the Baltimore County Police Department's Criminal Intelligence Unit knew trouble was brewing. With Ferguson's Michael Brown still fresh in the nation's mind and racial tensions running high, Baltimore braced itself for the imminent and expected protests.

OPPORTUNITY

"The Freddie Gray incident was a watershed moment for the City of Baltimore police," Vaccaro said. "The minute his death was announced, we knew we needed to monitor social media data at key locations where protesting was likely, especially at the local police precinct where Gray had been arrested."

In a stroke of luck, the Baltimore County Police Department had renewed their Geofeedia contract a week before the trouble began. The Criminal Intelligence Unit had experienced the tool's power first-hand before, and they didn't hesitate to call in reinforcements when trouble arose.

When an event at Camden Yards on April 25 turned violent, a ten-minute police nightmare was set into motion. It was the Criminal Intelligence Unit's job to identify and track the protesters, and to set up perimeters around key locations, set up automated alerts, and forward real-time information directly to Vaccaro's team via email




Figure 1.2 Geofeedia’s account of the Baltimore protests in the marketing of software analyzing social media data for the detection of incipient public protests. American Civil Liberties Union, 2016.

people in the city might mean. As Simone Browne has argued in her compelling account of the “digital epidermalization” of biometric algorithms, what is at stake is the recognizability of a body as human, as fully political.⁵ Among Baltimore’s arrests and detentions were forty-nine children, with groups of high school students prevented from boarding buses downtown because the output of the algorithm had adjudicated on the high risk they posed in the crowd.⁶ Based on the so-called ground truth of features that the algorithms had learned in the training data, the algorithms clustered the new input data of people and objects in the city, grouping them according to their attributes and generating a numeric scored output.⁷

The profound violence of the killing of one man, and the residue of all the past moments of claims made in his name, and in the name of others before him (note that the names Freddie Gray and Michael Brown persist in the training of subsequent algorithms to arbitrate protest), becomes lodged within the algorithms that will continue to identify other faces, texts, and signs in future crowds. Understood as the principal architecture of what N. Katherine Hayles calls the “computational regime,” what matters to the algorithm, and what the algorithm makes matter, is the capacity to generate an actionable output from a set of attributes.⁸ What kind of new political claim, not yet registered as claimable, could ever be made if its attributes are recognizable in advance? The very capacity to make a political claim on the future—even to board a bus to make that claim—is effaced by algorithms that condense multiple potential futures to a single output.

At the level of the algorithm, it scarcely matters whether the clustered attributes are used to define the propensities of consumers, voters, DNA sequences, financial borrowers, or people gathering in public space to make a political claim.⁹ Thus, when in 2016 Cambridge Analytica deployed its deep learning algorithms to cluster the attributes of voters in the UK EU referendum and the US presidential election, or when Palantir’s neural networks supply the targets for the US ICE deportation regime, what is at stake ethically is not only the predictive power of algorithms to undermine the democratic process, determine the outcomes of elections, decide police deployments, or make financial, employment, or immigration decisions. Of greater significance than these manifest harms, and at the heart of the concerns of this book, algorithms are generating the bounded conditions of what a democracy, a border crossing, a social movement, an election, or a public protest could be in the world.

Ethics of Algorithms

At first sight, the potential for violent harm precipitated by algorithms that learn to recognize human propensities appears to be a self-evident matter for critique. Surely, one could say, the ethical terrain of the algorithm resides in the broader political landscape of rights and wrongs, good and evil. After all, one could readily identify a set of rights, already apparently registered as belonging to rights-bearing subjects, that has been contravened by algorithms that generate targets, adjudicating which people may peaceably assemble, or which people are worthy of credit or employment, and on what terms. Indeed, on this terrain of delineating the rights and wrongs of algorithmic actions is precisely where many critical voices on the harms of the algorithm have been heard. Writing in the *New York Times*, for example, Kate Crawford identifies machine learning's "white guy problem," arguing that "we need to be vigilant about how we design and train machine learning systems."¹⁰ The dominant critical perspectives on algorithmic decisions have thus argued for removing the "bias" or the "value judgements" of the algorithm, and for regulating harmful and damaging mathematical models.¹¹ Within each of these critical calls, the ethical problem is thought to dwell in the opacity of the algorithm and in its inscrutability, so that what Frank Pasquale has called the "black box society" is addressed through remedies of transparency and accountability.¹² In sum, the rise of algorithmic power in society has been overwhelmingly understood as a problem of opaque and illegible algorithms infringing or undercutting a precisely legible world of rights belonging to human subjects. In such a framing, there is an outside to the algorithm—an accountable human subject who is the locus of responsibility, the source of a code of conduct with which algorithms must comply. To call for the opening of the black box, for transparency and accountability, then, is to seek to institute arrangements that are good, ethical, and normal, and to prevent the transgression of societal norms by the algorithm.

Yet, when people gathered to protest on Baltimore streets, or when Facebook users' data fueled the political and commercial models of Cambridge Analytica (figure I.3), legible rights to peaceable assembly or to electoral due process were not violated primarily by illegible algorithms. Rather, the means by which people could appear in a political forum, the conditions of their appearance, and the capacities they had to make a recognizable political claim in the world were subject to algorithmic regimes of what Michel Foucault calls truth telling and wrongdoing.¹³ In short, what matters is not primarily the identification and regulation of algorithmic wrongs, but more significantly how algo-



Cambridge
Analytica

Data drives all we do.

Cambridge Analytica uses data to change audience behavior. Visit our Commercial or Political divisions to see how we can help you.

Figure 1.3 Cambridge Analytica advertises how “data drives all we do” in the fused commercial and political pursuit of ways “to change audience behavior.” Screenshot archived by the author in May 2018, when the firm ceased trading under that name.

rithms are implicated in new regimes of verification, new forms of identifying a wrong or of truth telling in the world. Understood in these terms, the algorithm already presents itself as an ethicopolitical arrangement of values, assumptions, and propositions about the world. One does not need to look beyond the algorithm for an outside that is properly political and recognizably of ethics. Indeed, there can be no legible human outside the algorithm and underwriting its conduct, for as John Cheney-Lippold reminds us, we are enmeshed in the data that produce each “freshly minted algorithmic truth.”¹⁴ One cannot sustain a search for codes of ethics that instill the good, the lawful, or the normal into the algorithm. Contemporary algorithms are not so much transgressing settled societal norms as establishing new patterns of good and bad, new thresholds of normality and abnormality, against which actions are calibrated.

Actions one might consider harmful, as William Connolly notes, are not merely “actions by immoral agents who freely transgress the moral law” but are “arbitrary cruelty installed in regular institutional arrangements taken to embody the Law, the Good, and the Normal.”¹⁵ Amid the widespread search for new ethical arrangements for the scrutiny and regulation of algorithms,

what becomes of the arbitrary harms lodged within embodied algorithmic arrangements? One could imagine a world in which the deep neural networks used in cities like Baltimore are scrutinized and rendered compliant with rules and yet continue to learn to recognize and misrecognize people and to infer intent, to generate rules from the contingent and arbitrary data of many past moments of associative life on the city streets, to refine and edit the code for future uses in unknown future places. I may feel that some notion of legible rights is protected, and yet the attributes generated from my data, in correlation with yours and others', continue to supply the conditions for future arbitrary actions against unknown others. I draw a distinction here between ethics as code, or what Michel Foucault describes as "the code that determines which acts are permitted or forbidden," and ethics as the inescapably political formation of the relation of oneself to oneself and to others.¹⁶ My argument is that there is a need for a certain kind of ethical practice in relation to algorithms, one that does not merely locate the permissions and prohibitions of their use. This different kind of ethical practice begins from the algorithm as always already an ethicopolitical entity by virtue of being immanently formed through the relational attributes of selves and others. My desire for a different mode of critique and ethical account is animated not by the question, How ought the algorithm be arranged for a good society?, but by the question, How are algorithmic arrangements generating ideas of goodness, transgression, and what society ought to be?

In this book I propose a different way of thinking about the ethicopolitics of algorithms. What I call a *cloud ethics* is concerned with the political formation of relations to oneself and to others that is taking place, increasingly, in and through algorithms. My use of the term *cloud* here is not confined to the redefined sovereignties and technologies of a "cloud computing era," as understood by Benjamin Bratton and others, but refers to the apparatus through which cloud data and algorithms gather in new and emergent forms.¹⁷ The *cloud* in my cloud ethics is thus closer to that envisaged by John Durham Peters, for whom clouds are media in the sense that they are "containers of possibility that anchor our existence and make what we are doing possible."¹⁸ To consider algorithms as having ethics in formation is to work with the propensities and possibilities that algorithms embody, pushing the potentials of their arrangements beyond the decisive moment of the output.

A cloud ethics acknowledges that algorithms contain, within their spatial arrangements, multiple potentials for cruelties, surprises, violences, joys, distillations of racism and prejudice, injustices, probabilities, discrimination, and chance. Indeed, many of the features that some would like to excise from

the algorithm—bias, assumptions, weights—are routes into opening up their politics. Algorithms come to act in the world precisely in and through the relations of selves to selves, and selves to others, as these relations are manifest in the clusters and attributes of data. To learn from relations of selves and others, the algorithm must already be replete with values, thresholds, assumptions, probability weightings, and bias. In a real sense, an algorithm must necessarily discriminate to have any traction in the world. The very essence of algorithms is that they afford greater degrees of recognition and value to some features of a scene than they do to others. In so doing, algorithms generate themselves as ethicopolitical beings in the world. If to have ethics is not merely to have a code prohibiting, for example, bias or assumptions, but to work on oneself via relations, then the ethicopolitics of algorithms involves investigations of how they learn to recognize and to act, how they extract assumptions from data relations, and how they learn what ought to be from relations with other humans and algorithms.

To be clear, the cloud ethics I propose here does not belong to an episteme of accountability, transparency, and legibility, but on the contrary begins with the opacity, partiality, and illegibility of all forms of giving an account, human and algorithmic. To advance a cloud ethics is to engage the ungrounded politics of all forms of ethical relations. The significant new ethical challenges that algorithms seem to present to society actually manifest novel features of some profoundly old problems of the grounds for ethical action. As Judith Butler explains in her Spinoza lectures, the demand to give an account of oneself will always fall short, for “I cannot give an account of myself without accounting for the conditions under which I emerge.”¹⁹ If one assumes that the determination of an unequivocal *I* who acts is a necessary precondition of ethics, as Butler cautions, then this identifiable self is “dispossessed” by the condition of its emergence in relation to others. For Butler, this persistent failure to give a clear-sighted account does not mark the limit point of ethics. On the contrary, the opaque and unknowable nature of making all kinds of acting subjects is the condition of possibility of having an ethicopolitical life.²⁰

In short, and in contrast to the equation of ethics with transparency and disclosure, ethical responsibility is sustained by conditions of partiality and opacity. My notion of a cloud ethics extends the opacity of the human subject, envisaging a plurality of venues for ethical responsibility in which all selves—human and algorithmic—proceed from their illegibility. The apparent opacity and illegibility of the algorithm should not pose an entirely new problem for human ethics, for the difficulty of locating clear-sighted action was already present. The *I* who forms the ethical relation was always in question and

is now, with algorithms, in question in new ways. Though the mathematical propositions of algorithms cannot be made fully legible, or rendered accountable, they can be called to give accounts of the conditions of their emergence. These conditions include some relations that are identifiably between humans and algorithms—such as the selection and labeling of training data, the setting of target outputs, or the editing of code “in the wild,” for example—but others still are relations of algorithms to another algorithm, such as a classifier supplying the training data from which a neural network will learn. In all such instances of iterative learning, the significant point is that the conditions of an algorithm’s emergence—a composite of human-algorithm relations—are venues for ethicopolitics.

In a discussion on the impossibility of the transparent algorithm, the brilliant and generous scholar of black studies and machine learning Ramon Amaro once said, “Well what *would* it be if we even *could* open it? It’s just math.”²¹ Of course, he intended the comment as a provocation, but mathematics is never only “just math,” as Amaro’s work vividly shows. To reflect on the conditions of an algorithm’s emergence is also to consider how, as mathematical knowledge forms, algorithms have achieved the status of objective certainty and definiteness in an uncertain world.²² Ludwig Wittgenstein observed mathematical propositions to be “given the stamp of incontestability,” a mark of the “incontrovertible” and an “exemption from doubt” that other propositions, such as “I am called,” are not afforded.²³ For Wittgenstein, mathematics as practice—like all other language games—is concerned with particular uses of propositions, where “what a proposition is, is in one sense determined by the rules of sentence formation, and in another sense by the use of the sign in the language game.”²⁴ His concern is that the mathematical proposition has achieved a particular status of certainty in an otherwise uncertain world, so that it becomes “a hinge on which your dispute can turn.”²⁵ For Wittgenstein, the mathematical proposition should be regarded as no less doubtful or uncertain than the “empirical propositions” otherwise made about the world. Indeed, Wittgenstein’s point is to address mathematical propositions *as empirical actions* that are “in no way different from the actions of the rest of our lives, and are in the same degree liable to forgetfulness, oversight and illusion.”²⁶ Following Wittgenstein, the use of mathematical propositions is profoundly social and, in my reading, ethicopolitical. An algorithm is formulated through a series of truth claims about its match to the world, and yet, in its use in the world it is as prone to forgetfulness, oversight, misrecognition, and illusion as any other language game.

Algorithms such as those used to detect latent social unrest in the city

may appear in the world as Wittgenstein's "hinge on which your protest can turn" in the most direct sense that the hinge delimits and circumscribes the arc of the politics of protest. But the algorithm as hinge does not merely mark the limit point of resistance; rather, it presents something as a singular optimal output, when it is actually generated through multiple and contingent relations. My cloud ethics considers the algorithmic hinge to be akin to Karen Barad's scientific apparatus, which decides what matters in the world, what or who can be recognized, what can be protested, and which claims can be brought.²⁷ Understood in this way, the algorithm is not the hinge as an incontrovertible axis, exempted from doubt, on which all social, political, and economic life turns. "The hinge point," as Foucault differently identifies, can also be the point of "ethical concerns and political struggle," as well as the point of "critical thought against abusive techniques of government."²⁸

So, a principal ethicopolitical problem lies in the algorithm's promise to render all agonistic political difficulty as tractable and resolvable. Where politics expresses the fallibility of the world and the irresolvability of all claims, the algorithm expresses optimized outcomes and the resolvability of the claim in the reduction to a single output. In the following pages, I specify how this book approaches what an algorithm is and what it does. Among the many problems of studying algorithms is the matter of specifying which type of algorithm one is addressing. Though my primary focus in this book is on machine learning algorithms, and predominantly deep neural networks, in most of the instances I discuss, these algorithms are used in conjunction with some much older and less fashionable rules-based and decision tree algorithms. The form of the algorithm is not delimited by its name but by its coming into being, its use in the wild. As Nick Seaver has argued, "Rather than offering a 'correct' definition," a critical study of algorithms could begin from "their empirical profusion and practical existence in the wild."²⁹ For example, an advanced deep neural network for object recognition is intimately connected to some much older classifiers that are used in the preparation of the data, and it meaningfully comes into being as it is modified through its deployment in the world. It is not possible to identify a secure boundary around one specific named algorithm because, as a calculative device, it is a composite creature. I propose three routes into understanding what algorithms are in the world, each with its own distinctive implications for what is at stake ethicopolitically: algorithms as *arrangements of propositions*; algorithms as *aperture instruments*; algorithms as *giving accounts of themselves*.

“Generate All the Rules and Debate Which Is Interesting”: Algorithms as Arrangements of Propositions

The spatiality of the logic of algorithms is most commonly figured as a series of programmable steps, a sequence, or a “recipe,” governed by “precise instructions” within a “finite procedure.”³⁰ The calculative logic is thus most widely represented as a logic of the series, where each step in a calculative procedure is defined by its position in a finite series.³¹ This notion of a recipe, or a series of steps, contributes to limiting the imagination of what the ethicopolitics of algorithms could be. The spatial imaginary of the algorithm as a series of steps nurtures a particular set of ideas about how to intervene in the series to change the outcomes.³² So, for example, if the negative outcome of a credit-checking algorithm was found to include a racial or gendered bias in one of the steps, then the removal of this element in the recipe could be considered significant in notions of accountable and responsible algorithmic decisions. Similarly, to envisage a “kill switch” in the algorithms of autonomous weapons systems, for example, is to imagine a sequence of steps in which a human in the loop could meaningfully intervene and prevent a lethal decision.³³ Where algorithms are represented as a sequence, or what Manuel DeLanda calls “mechanical recipes specified step by step,” the addition or deletion of a step destroys the outcome, halts the decision in its tracks.³⁴ In this sense the spatial imagination of the algorithm as series appears to make possible all kinds of human oversight of otherwise automated decisions.

The representation of algorithms as a logical series, however, seriously overlooks the extent to which algorithms modify themselves in and through their nonlinear iterative relations to input data. The machine learning algorithms that are so categorically redefining our lives are characterized less by the series of steps in a calculation than by the relations among functions. Within computer science these relations are understood to be recursive functions, whereby the output of one calculation becomes the defining input for another, and so on, with each function nested within others like an infinite nesting of Russian dolls. These recursive functions, as Paulo Totaro and Domenico Ninno have argued, are having specific and durable effects on contemporary society.³⁵ Significantly, the removal or deletion of one function does not destroy the overall arrangement. Indeed, intrinsic to the logic of machine learning algorithms is their capacity to learn which outputs from which of their layers to pay greater attention to and which to bypass or discard.³⁶ This matters greatly, because the removal of a step one assumed to contain sensitive data on race, for example, would not remove or delete the process of learning via proxies how to recognize by means of racialized attributes. Consider, for

example, the machine learning algorithms that are being used in sentencing decisions by the courts to anticipate the optimal outcome of a prison sentence versus noncustodial measures.³⁷ Such technologies do not deploy a sequential logic that could be amenable to oversight and intervention in the steps. On the contrary, the optimal future outcome is defined solely in relation to an array of recursive functions, a “different mode of knowing,” as Adrian Mackenzie explains, in which the input data on past convictions and sentencing outcomes of hundreds of thousands of unknown others supply the contingent probabilities to all the layers within the algorithm.³⁸

Algorithms are not merely finite series of procedures of computation but are also generative agents conditioned by their exposure to the features of data inputs. As Luciana Parisi has argued, this is “a new kind of model,” which “derives its rules from contingencies and open-ended solutions.”³⁹ When algorithms learn by inductively generating outputs that are contingent on their input data, they are engaging experimentally with the world. As computer scientist Rakesh Agrawal explains the shift from rules-based to generative learning algorithms, past forms “used a statistical notion of what was interesting” so that “the prevailing mode of decision making was that somebody would make a hypothesis, test if it was correct, and repeat the process.” With machine learning algorithms, such as recursive neural networks, Agrawal notes that “the decision process changed,” and algorithms would “generate all rules, and then debate which of them was interesting.”⁴⁰ In their contemporary form, algorithms generate output signals that open onto uncertainty as to what is interesting, useful, or optimal. These output signals are not mere mathematical abstractions but are actionable propositions, such as “this person poses high risk of overstaying their visa,” or “this object is threatening the security of the street.” Often, when algorithm designers describe how they work with clients on a particular application, they suggest that they “tune” the algorithm as part of a discussion with the client of what is useful or optimal. This experimental tuning enacts the process Agrawal describes as debating which of the outputs is interesting, where the observation of the output of the model modifies and adjusts the weightings and thresholds of the algorithm. A kind of science of emergent properties, as I describe in my book *The Politics of Possibility*, such techniques significantly transcend and undercut traditional statistical notions of what matters, what is interesting, and what is optimal.⁴¹

In this book I understand the spatial logic of algorithms to be an *arrangement of propositions* that significantly generates what matters in the world. In contrast to the spatiality of the series or recipe, the arrangement of propositions articulates the algorithm’s capacity to engage experimentally with the

world, to dwell comfortably with contingent events and uncertainties, and yet always to be able to propose, or output, an optimal action. Practically, in the research for this book, I have studied the algorithm not as a finite series of programmable steps but as perennially adjustable and modifiable in relation to a target output. What does it mean to understand algorithms as arrangements of propositions? In Alan Turing's famous paper "Systems of Logic Based on Ordinals," he reflected on what he called "the exercise of ingenuity in mathematics." For Turing, ingenuity was important for mathematical reasoning because it provided "arrangements of propositions," which meant that the intuitive mathematical steps could not "seriously be doubted."⁴² The arrangement of propositions was thus a kind of mathematical architecture that supported the intuitive and the inferential faculties. As I use the notion, an arrangement of propositions extends to the experimental and iterative capacities of algorithms to propose things in and about the world. This is most likely not an interpretation Turing would approve of. Indeed, when he attended Wittgenstein's 1939 Cambridge lectures, their manifest disagreements in the lecture theater concerned precisely the matter of whether mathematical propositions could have normative effects. When Turing asserted that "from the mathematical theory one can make predictions," Wittgenstein replied that "pure mathematics makes no predictions." For Wittgenstein, " $30 \times 30 = 900$ is not a proposition about 30" but rather a proposition that finds its expression only in the "grammar" of its arrangement.⁴³ While for Turing, the numeric value 30 itself has a kind of agency, for Wittgenstein this is afforded only by its arrangement in a wider grammar through which it comes into use.

In a sense, the disagreement between Turing and Wittgenstein is undercut by a twenty-first-century world in which algorithms arguably generate their grammars and propositions through their exposure to numbers in the form of input data. As historian of mathematics Keith Devlin reminds us, mathematics involves not only numeric values based on a "count" but also, crucially, transformations based on "processes you perform."⁴⁴ The arrangements of propositions I envisage are not only numeric but transformative and performative. They contain within them multiple combinatorial possibilities and connections. The multiplicity of the algorithm—its plural possible combinations, pathways, weights, and connections—matters greatly to my tracing of the empirical processes by which algorithms learn and reach decisions, and to my desire to shift the ethicopolitical terrain on which this is understood to take place. The arrangement of propositions means that an apparently optimal output emerges from the differential weighting of alternative pathways through the layers of an algorithm. In this way, the output of the algorithm is

never simply either true or false but is more precisely an effect of the partial relations among entities. As Isabelle Stengers has noted of the proposition, “It is crucial to emphasize that the proposition in itself cannot be said to be true or false” because in itself “it is indeterminate with regard to the way it will be entertained.”⁴⁵ As a proposition, the algorithm can similarly not be said to be true or false—it cannot be held to account for its relation to truth in this sense. A pattern of false positives from a biometric algorithm, for example, can never be simply false because the threshold is immanently adjustable. Understood in Stengers’s sense of the proposition, algorithms are “indeterminate with regard to the way [they] will be entertained.”

Consider, for example, the multiple arrangements of propositions as they are manifest in the Asimov Institute’s Neural Network Zoo. Described as an “almost complete chart of neural networks,” the neural net zoo displays a spatial mapping of the arrangements of propositions of machine learning algorithms. This is not a zoo that categorizes its flora and fauna by the characteristics of their genus and species. Instead, it displays the algorithms’ architectures as arrangements of proximities, distances, intensities, and associations. The Asimov researchers, in their depiction of the arrangement of a convolutional neural network (CNN)—commonly used for facial recognition, feature extraction, and image classification—explain how each node concerns itself only with its close neighboring cells. The nature of the function performed within the node is decided by the close communion of weighted probabilities in the neighboring cells. As an arrangement of propositions, one could not meaningfully open or scrutinize the 60 million probability weightings that make it possible for a CNN algorithm to recognize the attributes of a face in a crowd, declaring them to be true or false. Indeed, following Stengers’s formulation of the proposition, the output of a facial recognition algorithm is never either “true” or “false” but instead is a useful proposition that can be infinitely recombined. Unlike a series of steps or a recipe, one could never have oversight of the infinite combinatorial possibilities of the algorithm as proposition. Nor would the deletion of a step render the whole unworkable. Once the algorithm is understood as an arrangement of propositions, the mode of ethics must work with the partiality and illegibility of the relations among entities.

“Reduced to That Which Interests You”: Algorithms as Aperture Instruments

Critical accounts of the rise of algorithms have placed great emphasis on the power of algorithms to visualize, to reprogram vision, or indeed even to “see” that which is not otherwise available to human regimes of visibility. Similar to the spatial arrangement, this primacy of the visual register has also annexed what could count as the ethics and politics of algorithms. There are two curiously twinned accounts of contemporary algorithms in relation to regimes of sight and vision. The first is that algorithms operate on a plane in excess of human visibility and at scales that are inscrutable to the human. The second is that algorithms themselves have an enhanced capacity to visualize the invisible, to see, scan, and search volumes and varieties of data heretofore unavailable to human senses. Indeed, this intersection of machinic and human vision comes to the fore in the espoused ethics of public inquiries into the state’s deployment of automated algorithms for the government of the population. For example, in the UK parliamentary inquiry following the Edward Snowden disclosures of widespread automated data analysis, a peculiar kind of virtue was found in the notion that, in machine learning intelligence, “only a tiny fraction of those collected are ever seen by human eyes.”⁴⁶ Similarly in the United States, the former director of national intelligence James Clapper likened the NSA’s algorithmic analysis of citizens’ data to a form of library in which few books are “actually read” and where the output of the system supplies “the books that we need to open up and actually read.”⁴⁷ There is an acute problem, then, with the widespread appeal to ethical codes that regulate what algorithms or humans do or do not see. It is a problem, I suggest, with its roots in the privileging of sight and vision over other forms of making things perceptible. “Vision cannot be taken,” writes Orit Halpern in her wonderful book *Beautiful Data*, “as an isolated form of perception” but must be considered “inseparable from other senses.”⁴⁸ To act and to be responsible for action, an algorithm need not “see” or “read” but need only make something or someone perceptible and available to the senses.

In this book I situate the ethics and politics of algorithms within a genealogy of technologies of perception. Contemporary algorithms are changing the processes by which people and things are rendered perceptible and brought to attention. This is definitively not merely a matter of making things amenable to vision and indeed is frequently a matter of sustaining something beneath the visual register and yet perceptible. As art historian Jonathan Crary writes, “Ideas about perception and attention were transformed” alongside the historical “emergence of new technological forms of spectacle, display, pro-

jection, attention, and recording.”⁴⁹ Understood in this way, the transformation of perception involves changes in how the perceiving subject thinks about what could be brought to attention, changes in the horizon of possibility of human action. As with the advent of the technologies of printing press, camera, or cinema, so the advent of the machine learning algorithm implies a reworking of what it means to perceive and mediate things in the world.⁵⁰ This is not a process that is effectively captured by the idea that automated systems are undermining or superseding human forms of perception and action. To foreground instruments of perception, or what Henri Bergson terms “organs of perception,” is to breach conventional distinctions between humans and machines and acknowledge the entangled nature of all forms of perception.⁵¹ Bergson insists on the shared limits of perception across science and ordinary everyday experience, so that “ordinary knowledge is forced, like scientific knowledge,” to divide up time into perceptible slices, to “take things in a time broken up into particles.”⁵² Whether the organ of perception is microscope, telescope, eye, camera, or algorithm, perception is attuned to action, to the dividing up of movement into points on a trajectory so that they can be acted on. “What you have to explain,” he writes, is not “how perception arises, but how it is limited, since it should be the image of the whole, and is in fact reduced to the image of that which interests you.”⁵³ Following Bergson’s insight on how an organ of perception seizes the object of interest from its environment, to consider algorithms as instruments of perception is to appreciate the processes of feature extraction, reduction, and condensation through which algorithms generate what is of interest in the data environment.

Confronted by something of a moral panic surrounding the expansive volumes of “big data” and powers of surveillance of automated systems, my emphasis on practices of perception foregrounds precisely the opposing process: the reducing, distilling, and condensing of particles of interest from a whole. A defining ethical problem of the algorithm concerns not primarily the power to see, to collect, or to survey a vast data landscape, but the power to perceive and distill something for action. Algorithms function with something like an aperture—an opening that is simultaneously a narrowing, a closure, and an opening onto a scene. Let us consider, for example, an algorithm designer working in the UK defense sector, demonstrating the capacity of his deep neural network algorithms to recognize a mistaken civilian target amid a crowded data environment of drone images.⁵⁴ He shows a slice through time as a vehicle travels away from the center of Kandahar, Afghanistan. He explains the problem for the decision: that this could be a suspect vehicle or, crucially, a school

bus taking children home to villages outside the city. The algorithms had learned to recognize a school bus through training data that supplied images of predominantly yellow US-style buses. The designer explains that his algorithm is aggregated with many others to generate a single actionable output—target/no target—but to do this he must necessarily reduce and condense the patterns of interest from a volume of input data. The training data—and the humans who labeled it—have conditioned the CNN algorithm to carve out and value some objects and to discard others. In this part of Afghanistan, some of the school buses are indeed in the spatial form of a US-style school bus, but others still are open-back trucks repurposed for transporting scholars. My point is that a potential act of violence, such as a school bus wrongfully targeted by a drone strike, resides not primarily in the vertical surveilling and collecting of data, but in fact in the horizontal arraying of possible patterns of interest lodged within the algorithm itself.⁵⁵ As an aperture instrument, the algorithm's orientation to action has discarded much of the material to which it has been exposed. At the point of the aperture, the vast multiplicity of video data is narrowed to produce a single output on the object. Within this data material resides the capacity for the algorithm to recognize, or to fail to recognize, something or someone as a target of interest.

The ethical stakes of what Mark Hansen calls “potential perceptual reconfiguration” applied to my cloud ethics necessarily involves something like a re-opening of the process of the algorithm's reduction of a multiplicity to one.⁵⁶ What is happening in this process of condensing plural possible pathways to a single output? When an algorithm determines whether a vehicle is a military or a civilian target, or when it decides if a public protest contains latent dangerous propensities, it reduces the heterogeneity of durational time to perceive the attributes of an object and their differences of degree from other objects encountered in a past set of data. The question of what this crowd could be, what this vehicle might do, the frustrations or discomforts of the actual lived experience of waiting or gathering persist as indeterminacies in the hidden layers of the algorithm. Even within the archive of training data—sometimes just a Google ImageNet dataset of labeled images of school buses—are the residual contingencies of durational time, with all the past lived moments supplying norms and anomalies for the algorithm to learn. “My own duration, such as I live it in the impatience of waiting,” reflects Gilles Deleuze, “serves to reveal other durations that beat to other rhythms, that differ in kind from mine.”⁵⁷ To respond to the perceptual power of the algorithm and to prize open the aperture of the single output is to trace the other durations that continue to

beat in the discarded data, the multiple other potential pathways that could be mapped between fragments.

“At the Limits of What One Knows”: Algorithms Give Accounts of Themselves

Perhaps the most widespread concern in public and scholarly discourse on the operation of algorithms in society is that they are unaccountable or that they cannot meaningfully be held to account for harmful actions. Indeed, a kind of proxy form of accountability is emerging, in which the designers of algorithms are made the locus of responsibility for the onward life of their algorithms. Virginia Eubanks proposes a “Hippocratic oath for data science,” in which the designers of algorithms would be accountable to human beings and not “data points, probabilities, or patterns.”⁵⁸ Similarly, in a 2016 *Nature* editorial calling for “more accountability for big data algorithms,” the editors propose that “greater transparency” could be achieved if the designers of algorithms “made public the source of the data sets they use to train and feed them.”⁵⁹ They urge greater disclosure of the design of algorithms and an “opening up to scrutiny” of their workings. Indeed, so widespread is this notion that the accountability of algorithms can be grounded in their design or source code that technical solutions for “explainability” are being developed to trace an apparent bias back to a design problem. For example, automated systems for the assessment of creditworthiness are thought to be rendered transparent by a tool that traces the specific credit score output back to a data element, such as an unpaid bill.⁶⁰ Such techniques are thought to anchor the accountability of the algorithm precisely in an intelligible knowledge of its workings.

Within these demands for algorithmic accountability lies a specific form of giving an account. The locus of a truthful account is in the apparent “source” of the algorithm, in its origins, whether in the source code or in the algorithm designer as an author. This locus of original account, I suggest, is profoundly limiting the capacity to demand that algorithms give accounts of themselves. It imagines a secret source or origin to which all potential future harms could be traced. As Foucault writes on the disappearance of the author, “the task of criticism is not to bring out the work’s relationship with the author” but rather to “analyse the work through its structure, its architecture, its intrinsic form, and the play of its internal relationships.”⁶¹ For our contemporary times, the call for accountability of algorithms has precisely targeted the work’s relationship with the author, seeking to render transparent the intent and the workings of the design. At one level there are clear limits to identifying the source of the algorithm, not least that each apparent “one” contains multiple

elements from multiple sources, with much of this aggregation concealed even from the designer. My point, however, is that the problem of anchoring accountability in a source or origin is not confined to algorithms but is a persistent and irresolvable ethicopolitical problem.

The problem of the unidentifiable origin of the algorithm extends to all notions of an authoring subject *I* who can give a clear-sighted account of herself. If one's account of oneself can never be fully secured, then the full disclosure of one's grounds for action is impossible. To give an account is always to give an uncertain narrative that risks falling short or failing to be recognized. Let us begin not from a search for secure grounds for accountability, then, but from the very ungroundedness of all forms of giving an account. "Ethics and politics only come into being," Thomas Keenan writes, "because we have no grounds, no reliable standpoints" from which to forge foundations.⁶² To be responsible for something—an errant output, a fatal decision, a wrong judgment—is less a matter of securing the grounds for the action than a matter of responding even when knowledge is uncertain and the path is unclear. "What could responsibility mean," asks Keenan, if it is "nothing but the application of a rule or decision."⁶³ There is no great origin or source of responsibility without uncertainty and undecidability.

In this book I propose that algorithms are not unaccountable as such. At least, they should not be understood within a frame of ethical codes of accountability in which the source of the problem could be secured. Algorithms, I propose instead, are giving accounts of themselves all the time. These accounts are partial, contingent, oblique, incomplete, and ungrounded, but, as N. Katherine Hayles vividly documents, this is not a condition unique to the cognitive complexities of algorithms.⁶⁴ Far from it. The condition of giving an account that is never transparent or clear sighted is already the ethicopolitical condition with which we must live. To attend to the accounts that algorithms give is to "stay with the trouble" of "unexpected collaborations and combinations," as Donna Haraway has captured the "method of tracing, of following a thread in the dark."⁶⁵ Given the conditions of following collaborative threads in the dark, perhaps one should not create a special category of opaque and illegible agency for the identification of algorithms. If, as Judith Butler suggests, "my account of myself is partial, haunted by that for which I have no definitive story," then might it be that "the question of ethics emerges at the limits of our schemes of intelligibility, where one is at the limits of what one knows and still under the demand to offer and receive recognition?"⁶⁶ This is the ethics at the limits of schemes of intelligibility, of following threads in the dark, that I envisage for my cloud ethics. Refuting the many demands for an impartiality

of the algorithm, excised of bias and prejudice, I wish to be alert to the always already partial accounts being given by algorithms.

As feminist scholars of technoscience have long reminded us, the partial account is not an account devoid of insight or real purchase on the world. “There is no unmediated photograph or passive camera obscura in scientific accounts of bodies and machines,” writes Donna Haraway. “There are only highly specific possibilities, each with a wonderfully detailed, active, partial way of organizing worlds.”⁶⁷ To attend to algorithms as generating active, partial ways of organizing worlds is to substantially challenge notions of their neutral, impartial objectivity. To foreground partiality is also to acknowledge the novel forms of distributed authorship that newly entangles the *I* who speaks in composite collaborations of human and algorithm. If one element of my past presence on a London street for a “Stop the War” campaign march enters a training dataset for a multinational software company’s neural net, which, one day in the future, intervenes to detain some other person in a distant city, how is some part of my action lodged within this vast and distributed authorship? What is the possibility of my ethicopolitical responsibility for the dark thread to a future intervention made partially on the attributes of my past data?

To be attentive to the accounts algorithms are giving of themselves, then, is to begin with the intractably partial and ungrounded accounts of humans *and* algorithms. As is so manifestly present with the surgical robotics algorithms I discuss in chapter 2, the embodied accounts human surgeons give of what they can do, how they decide on boundaries between diseased and healthy tissue, how they can reach decisions, are not meaningfully separable from the surgical robots with whom they share a cognitive workload. The robot can only act to suture a wound because its algorithms have learned from the past data of many thousands of instances of human surgeons suturing. The surgeon can only reach a difficult kidney tumor because the robot’s data yield precise coordinates from an MRI that make it recognizable amid occlusions. When an error is made, the identification of a unified and identifiable source of the error is not possible. The figure who would be required to give a clear-sighted account is an impossible figure.

Chapter Outlines

In sum, my cloud ethics has a principal concern with the algorithm’s double political foreclosure: the condensing of multiple potentials to a single output that appears as a resolution of political duress; and the actual preemptive closure of political claims based on data attributes that seek recognizability in

advance. Confronted with this double foreclosure, each chapter of this book elaborates a speculative strategy for reinstating the partial, contingent, and incomplete character of all algorithmic forms of calculation. In their partial and incomplete way of generating worlds, we can locate their ethicopolitics. Part 1, “Condensation,” comprises two chapters that detail how algorithms condense and reduce the teeming multiplicity of the world to a precise output. Chapter 1, “The Cloud Chambers,” examines how algorithms are acting through cloud data architectures to produce a new paradigmatic alliance between sovereign authority and scientific knowledge. Developing an analogy with the scientific experiments of the cloud chamber of early twentieth-century particle physics, where the chamber made it possible to perceive otherwise invisible subatomic particles, I explore the capacities of the cloud in cloud computing. How does the cloud apparatus render things perceptible in the world? I address the character of cloud architectures across two distinct paradigms. The first, Cloud I, or a spatiality of *cloud forms*, is concerned with the territorial identification and location of data centers where the cloud is thought to materialize. Here the cloud is understood within a particular history of observation, one where the apparently abstract and obscure world can be brought into vision and rendered intelligible.

This notion of cloud forms, I propose, has led to a distinct ethicopolitical emphasis on rendering algorithms explainable or legible. Cloud I is founded on a misunderstanding of the nature of algorithmic reason within the cloud, so that the cloud is thought to obscure or obfuscate what is “really” going on in the world. On the contrary, I propose, the cloud is not an obfuscation at all but is a means of arranging the models for otherwise incalculable processes for a condensed decision. In my second variant—Cloud II, or the spatiality of a *cloud analytic*—the cloud is a bundle of experimental algorithmic techniques acting on the threshold of perceptibility. Like the cloud chamber of the twentieth century, contemporary cloud computing is concerned with condensing that which cannot be seen, rendering things perceptible and actionable.

Chapter 2, “The Learning Machines,” turns to the question of machine learning algorithms and their complex and intimate relationships with what we think of as human practices. Through a study of the deep neural network algorithms animating surgical robots, I show the entangled composite bodies of surgeons, robots, images of organs, and cloud-based medical data of past surgeries—all of which learn together how to recognize and to act amid uncertainty. Where there has been public concern and moral panic around machine learning, what is most commonly thought to be at stake is the degree of autonomy afforded to machines versus humans as a locus of decision. I propose that

the principal problem resides not with machines breaching the limit but in sustaining a limit point of the autonomous human subject—the oft-cited “human in the loop”—who is the locus of decision, agency, control, and ethics. Such an autonomous human disavows the *we* implicated in the entangled learning of humans with algorithms. Much of this entangled learning takes place in a space of play between a target output and an actual output of the algorithm. The indeterminate combinations of weights, parameters, and layers in the algorithm are the traces of rejected pathways and alternative correlations.

Part 2, “Attribution,” elaborates the practices of attribution through which algorithms write themselves into being in the world. Taking seriously the attribute as it is understood in computer science, the chapters connect this attribution back to genealogies of writing and ethics. In chapter 3, “The Uncertain Author,” I am concerned with how the search for ethical codes to govern algorithms has located an author function in the “source code” of algorithms. When an algorithm appears to have precipitated a crisis or to have caused a harm, often the reflex response is to seek out its origin: Who designed the model? Who wrote the code? Who or what labeled the training data? Was it biased? Who is the author of the algorithm? In this chapter I explain the limits of locating an ethicopolitical response in the authorship of source code. Focusing on the algorithmic techniques for natural language processing—where a corpus of literary texts are used to train an algorithm to recognize style and sensibility—I suggest that the algorithm’s ways of being in the world are not all present in the source code and, indeed, substantially exceed the design of an authoring subject. The authorship of the algorithm is multiple, continually edited, modified, and rewritten through the algorithm’s engagement with the world. Juxtaposing novelists’ accounts of their own uncertain authorship with computer scientists’ reflections on how the fragments of an algorithm come together in acts of writing, I suggest the impossibility of identifying a definitive author.

To invoke the call for attributing authorship is not only insufficient as critique, but it also risks amplifying the ability of the algorithm to bind together a unity of incompatible, fraught elements as though the difficulties and differences could be resolved. In place of the call for securing responsibility via authorship, I propose that there is a possibility of ethicopolitics in the act of writing, in the bringing together of scattered elements, and in the opening of a space of uncertainty. A cloud ethics must begin to reopen the act of writing as a site of ethical significance. The profound uncertainty that is brought to the writing of an algorithm—“I do not know how adjusting this weight is changing the output”; “I cannot be sure how my training data has produced

these clusters”; “I am curious whether shifting this threshold will reduce the false positives”—invites an iterative process of writing that is never completed. Here is an ethicopolitical tension that is worth holding on to—the algorithm promises to complete everything, to condense to a single optimized output and action, and yet it enacts a process of writing that opens on to an indeterminate future.

Chapter 4, “The Madness of Algorithms,” addresses the moments when it appears that an algorithm has acted in a state of frenzy or has departed from its otherwise rational logic. From the lethal accidents of autonomous vehicles during tests, to the racist hate and misogyny of the Twitter chatbot Tay, an ethical frame often seems required to somehow rein in the worst excesses or to restore reasonableness to autonomous actions. Yet, philosophy has long grappled with the problem of madness and, specifically, how the identification of madness shores up and sustains the domain of reason. Reflecting on the conjoined histories of ideas of reason and madness, I propose that one cannot speak of the madness of the algorithm except in relation to the very form of reason the algorithm embodies. While the contemporary moral panic at each moment of the madness of algorithms urges us to police ever more vigilantly the line between reasonable and unreasonable actions, understood as a threshold, this line is precisely the condition of possibility of algorithmic rationality. Algorithms cannot be controlled via a limit point at the threshold of madness because the essence of their logic is to generate that threshold, to adapt and to modulate it over time. In short, my argument is that when algorithms appear to cross a threshold into madness, they do, in fact, exhibit significant qualities of their mode of reasoning. Understood in this way, the appearance of a moment of madness is a valuable instant for an ethicopolitics of algorithms because this is a moment when algorithms give accounts of themselves.

Contra the notion that transparency and the opening of the black box secure the good behavior of algorithms, the opacity and clouded action exhibited in the excesses and frenzies of algorithms have a different kind of fidelity to the account. Throughout this book I argue that, when viewed from the specific propositional arrangements of the algorithm, particular actions that might appear as errors or aberrations are in fact integral to the algorithm’s form of being and intrinsic to its experimental and generative capacities. I am advocating that we think of algorithms as capable of generating unspeakable things precisely because they are geared to profit from uncertainty, or to output something that had not been spoken or anticipated. Of course, this is not a less troubling situation than the one in which some controls are sought on the worst excesses of the algorithm. On the contrary, it is all the

more political, and all the more difficult, because that which could never be controlled—change, wagers, impulses, inference, intuition—becomes integral to the mode of reasoning.

Part 3, “Ethics,” develops a set of tactical routes for a cloud ethics to follow, each of these insisting on a different kind of weightiness from the calculative weights of adjustable probabilities. The apparent lightness of an optimized single output is afforded the full weight of undecidability and the difficulty of decision. In chapter 5, “The Doubtful Algorithm,” I suggest that doubt, and more precisely the idea of doubtfulness, can be a tactical point of intervention for a cloud ethics. The overwhelming logic of algorithmic systems in our society is that they can optimize decisions made where there is profound doubt and uncertainty. Indeed, the algorithm is offered as a means of deploying doubt productively so that, for example, the doubt-ridden voter or doubtful consumer can be clustered as “having a propensity to be influenced” and can be targeted with personalized media. A specific form of truth telling, established in the “ground truth” of the data environment from which the algorithm learns, has asserted its dominance in the governing of societies with algorithms.

By contrast, and cutting against the grain of the dominance of definiteness as algorithms act on doubt, I seek to reinstate doubtfulness as what N. Katherine Hayles calls “embodied actuality” within the calculative architecture of the algorithm.⁶⁸ Though at the point of optimized output, the algorithm places action beyond doubt, there are multiple branching points, weights, and parameters in the arrangements of decision trees and random forest algorithms, branching points at which doubt flourishes and proliferates. A cloud ethics reopens the contingencies of this multiplicity, giving life to the fallible things that the algorithm has learned about the world, rendering the output indelibly incomplete.

In each of the chapters, I maintain a faithfulness to the specificities of how particular algorithms learn via, and generate worlds through, their relations with data. I am wary of “algorithm talk” when it is asserted generally and without specificity, for different algorithms are as varied in their logics and grammars as languages are, and these differences, as Adrian Mackenzie argues, should be made to matter.⁶⁹ Thus, for example, in discussions of facial recognition systems, I would want to know what kinds of CNNs are being used as the basis for recognizing a face from data inputs. Though each of the chapters draws out the arrangements of propositions of particular algorithms and how they recognize, misrecognize, and target through their relations with other algorithms, data, and humans, I have also foregrounded the fallibility of the algorithm, its incompleteness and contingency. In short, a specific algo-

rithm will always exceed its name, its type, its genus, for it is immanently modifying itself through the world. In chapter 6, “The Unattributable,” I address more directly those critics who would ask what use my cloud ethics would be to their campaigns or movements against the injustices of specific algorithms. Following a lecture on machine learning that I gave in Copenhagen, a lawyer in the audience asked what my cloud ethics could ever look like in law. I had similar questions from lawyers at the Turing Institute in London. My response, in short and in those moments, was to say that it would be a crowded court in the sense that my approach multiplies the possible sites of intervention and responsibility. This notion of the crowded space, or forum, has stayed with me throughout the research and the writing of this book. What kind of political claim could be brought in the name of a cloud ethics? How does one forge a form of responsibility for the future onward life of something like an attribute as it becomes attached to others? What is the power of the unattributable as the set of qualities that cannot be attributed to a subject?

Chapter 6 maps out what it means to be together ethicopolitically—to be associated as a society, a community, a movement, a gathering of protesters—as an association of partial associations or attributes. The chapter concludes the book along three lines of argument for a cloud ethics: apertures, opacity, and the unattributable. When machine learning algorithms segment a social scene, generating clusters of data with similar propensities, everything must be attributed. Yet, that which is unattributable does remain within the scene, exceeding the algorithm’s capacity to show and tell, as well as opening onto a different kind of community and a different mode of being together, being ethicopolitical.