The Epistemic and Performative Dynamics of Machine Learning Praxis

Paul Kockelman, *Yale University, USA*

ABSTRACT

This article is about machine learning as it relates to classic concerns in anthropology and the social sciences, regarding meaning, value, and culture, as well as agency, power and performativity. It focuses on the role of machine learning, with its peculiar manner of modeling phenomena, in mediating: (i) the sensibilities and assumptions agents have (qua interpretive grounds and algorithmic models) insofar as these mediate their actions, inferences, and affects; and (ii) the actions, inferences, and affects of agents (qua computational processes and interpretive practices) insofar as these drive their sensibilities and assumptions. More generally, it offers a model of the process of modeling per se, so far as this process unfolds in contexts of machine learning and beyond. In this respect, the metamodel offered is meant to capture some of the key dynamics of the tense and mutually transformative relations linking objects (of analysis), data (drawn from those objects), models (of such objects, as informed by such data), and actions (grounded in such models, and often transformative of such objects). It foregrounds the wily, epistemic, performative, and often violent dynamics of such processes when the objects being modeled are themselves agents capable of modeling.

Algorithmic dystopia is all too easy to imagine. This is not just because it has long been the theme of so many movies and books but mainly because it's (almost) already here. Recording devices constantly capture all the happenings on a city's streets. Their sequenced images and sounds constitute continuous input to an ensemble of algorithms. As the output of such algorithms, entities and events, people and actions, utterances and gestures, affects

Contact Paul Kockelman at Department of Anthropology, 10 Sachem Street, New Haven, CT 06511-3707, USA (paul.kockelman@yale.edu).

Signs and Society, vol. 8, no. 2 (Spring 2020). © 2020 by Semiosis Research Center at Hankuk University of Foreign Studies. All rights reserved. 2326-4489/2020/0802-0006$10.00

https://doi.org/10.1086/708249 Published online by Cambridge University Press

and relations are identified, connected, contextualized, and interpreted. A real-time history, or always-on ethnography, of the city is thereby constructed, itself accessed through a variety of search engines, one of which will surely be called Mal(ware)inowski. After typing in a few key words (Alice Shepard, 10/25/22, 3–4 am), and entering the correct security code (or paying a modest subscription fee), a user—who may be just another algorithmic device—will receive an output like: *On Saturday, October 25th, on the corner of 23rd and Ivy, at 3:33 a.m, Alice Shepard nervously received a package from Dave Riggens, the brother of her fiance, and disingenuously told him she would return again shortly*.

That is, not only will individuals be identified (Alice Shepard), but so will their instruments, actions, and roles. And not only will their public behavior be classified and labeled, but their private motivations and hidden affects will be inferred: *she gave him the package because she felt X, and so that he would think Y*. Such algorithms will not just classify and label, reason and infer; they will also predict and retrodict and thereby add new entries into the same history (albeit under different modalities: not the "observed" and "inferred," or even the "subjunctive" and "conditional," but the "expected" and the "suspected"). Finally, using various ethic-metrics, they will evaluate such actions, intentions, and identities as good or bad and then act: locking and unlocking doors, opening and closing accounts, adding additional cameras and algorithms to predicted routes, or simply calling in more human-like authorities to act in their stead (however inhumanely).

To be sure, humans have long been subject to similar modes of scrutiny. As long as there have been languages and minds, eyes and memories, people have watched the behavior of others, characterized the actions being undertaken, and thereby ascribed meanings to such actions, and motivations to such actors. Moreover, insofar as each of us is an actor (to be observed) as well as an observer (of other actors), we have long engaged in countermeasures to thwart the full disclosure of the depths—and shallows—of our selves: feign and mask, occlude and wait, internalize and counteract. That is, we can understand—at least to some degree—the *interpretive grounds* (or "algorithmic models") that others use to understand our behavior, for we use similar grounds to do the same to them. So far as we can internalize their grounds in this way, we can exploit—or even hack—their interpretations. And insofar as we know they are reflexive agents like ourselves, who share our interpretive grounds to some degree, we expect them to do the same to us; and therefore we deploy safeguards to thwart such exploits (however unsuccessfully).

This article is about various algorithms used in machine learning: how they come to identify, characterize, and reason about our actions. It analyzes the

https://doi.org/10.1086/708249 Published online by Cambridge University Press

tense (and soon to be traumatic) coupling between the interpretive grounds of humans and the algorithmic models of such "machines"—not just the ways that each kind of agent mirrors and modulates the other but also the ways each kind of agent mutates in the face of the other.[1]

More carefully, this article offers a model (or interpretation) of the process of modeling per se, so far as this process unfolds in contexts of machine learning and beyond. It focuses on the role of machine learning in mediating (i) the sensibilities and assumptions agents hold (qua interpretive grounds and algorithmic models) insofar as these mediate their actions, inferences, and affects; (ii) the actions, inferences, and affects of agents (qua interpretive practices and computational processes) insofar as these drive their sensibilities and assumptions. In this respect, the metamodel offered is meant to capture some of the key dynamics of the tense and mutually transformative relations linking objects (of analysis), data (drawn from those objects), models (of such objects, as informed by such data), and actions (grounded in such models, and often transformative of such objects). It foregrounds the wily, epistemic, performative, and often violent dynamics of such processes when the objects being modeled are themselves agents capable of modeling.[2]

1. It should be emphasized that I use the term *machine* in no special sense: it can refer to a piece of software, algorithms used by that software (or functions called by such algorithms), various configurations of hardware (running such software), assemblages of hardware, networks of such assemblages, populations of such networks, and beyond. Such agents can thus be radically distributed (Enfield and Kockelman 2017).

2. For more on algorithmic violence and its precursors, see Galtung (1969), Anglin (1998), and Onuoha (2018). For more on semiotic grounds, see Peirce (*EP* 1:1–10), Parmentier (1994), and Kockelman (2012, 2015, 2016a, 2016b). For complementary takes on models, mediation, or performativity, see Goffman (1959), Mertz and Parmentier (1985), Silverstein and Urban (1996), Lee (1997), Knorr-Cetina (1999), Agha (2007), and Enfield and Sidnell (2017). For related work on performativity from the standpoint of science and technology studies, and opposed to linguistic anthropology, see Hacking (1995, 2002), Moll (2002), and the collection of essays edited by MacKenzie et al. (2007). Especially relevant in this rich collection are the essays by Callon (2007), Lépinay (2007), and Mackenzie (2007). For related work on algorithms from an anthropological perspective, see Kockelman (2013, 2017), Gillespie (2016), Knox and Walford (2016), and Lowrie (2017, 2018). Two superb essays by Seaver (2017, 2018) were particularly bracing. For great introductions to neural networks and deep learning, see Goodfellow et al. (2016) and Nielson (2018). For related work on numbers, see Klausner (2018), and the collected essays in Lippert and Verran (2018). On the history of cybernetics, ecology, French theory, and surveillance, see Geoghegan (2011, 2019). For particularly important work on deep learning per se, its history, and its relation to core concerns of anthropology and critical theory, see Castelle (2018a, 2018b). For more on the notion of coupling, as it is used here, see Kockelman (2017). For more on control, in a Deleuzean tradition, see Goffry (2015). For algorithms and machine learning practices as agents of racism, sexism, and oppression, see Keyes (2018) and Noble (2018), inter alia. For more on gradients, as of utmost importance to social science, outside of their role in machine learning per se, see Kockelman (2016a). For the role of gradients and grading in channeling and policing intensity, see Carruthers (2017, 2019). On the importance of contact and coupling, in relation to counterperformativity, see Elyachar (2010) and Edwards (2018). On technological innovators, their "attentuated understanding of the social," and the flattening of politics generated and entailed by such imaginaries, see Nelms et al. (2018). On *thresholds*, as the term is used here, see Carruthers (2019) and Kockelman (2019). On inalienable possessions, and the quasi-personal fringe, see Veblen (1898) and Kockelman (2007). William Stafford, Kamala Russell, Stéphane Gros, and Liam Taylor gave me very helpful feedback. Finally, many thanks to Asif Agha and an anonymous reviewer for their extremely helpful comments.

https://doi.org/10.1086/708249 Published online by Cambridge University Press

1. Agents, Practices, and Grounds

To understand the relation between the interpreting practices of humans and the algorithmic processes of machines, it is helpful to: (i) sketch the relation between semiotic practices, agents, and grounds; (ii) show how such distinctions apply to humans as well as to machines; (iii) demonstrate how both kinds of agents, along with their respective grounds, are radically coupled in contexts of machine learning; and (iv) see the transformative effects all this has on social relations and cultural values. The rest of this section will introduce readers to these core themes.

As shown in figure 1, *semiotic practices* involve three interrelated components: a sign (whatever stands for something else); an object (whatever is stood for by a sign); and an interpretant (whatever a sign gives rise to, or creates, insofar as it is taken to stand for an object). Interaction is a semiotic process: you hail me (sign), indicating your desire to ask a question (object), and I turn to look at you (interpretant). Similarly, identification is a (meta)semiotic process: having witnessed the foregoing interaction (sign), another person infers that you are the law and I am a criminal (object), and so directs their question (deference, fear, or aim) to me as opposed to you (interpretant).

Semiotic practices only unfold insofar as there are *semiotic agents*: entities that can take signs to stand for objects, and act in relation to those objects, insofar as such objects relate to their own interests. Such semiotic agents, such as the "I" and "you" in the above interaction, link signs and objects for the purpose of, and through the practice of, signification and interpretation. They do so only in reference to *semiotic grounds*: the sensibilities and assumptions they (and others) have regarding possible sign-object relations, as evinced in their signifying and interpreting practices (and/or encoded in their metasemiotic practices).

Such grounds are legion. From a semiotic stance, there is the realm of qualities that agents can sense and compare, such that they may be noticed in different events, such that iconic correspondences might be drawn between those events. For example, insofar as *this* and *that* were both sweet to some agent, *this* may direct the agent's attention (affect and action) to *that*. There is the realm of causal processes that agents may notice, internalize, or theorize, such that indexical contiguities may be tracked and instigated. For example, if an agent believes that sun bathing leads to skin cancer, they may predict skin cancer having witnessed sun bathing, or suspect sun bathing having noticed skin cancer. There is the realm of codes and conventions, such as group-specific systems of typical sign-object, or behavior-circumstance pairings. For example, agents use the natural languages

https://doi.org/10.1086/708249 Published online by Cambridge University Press

https://doi.org/10.1086/708249 Published online by Cambridge University Press

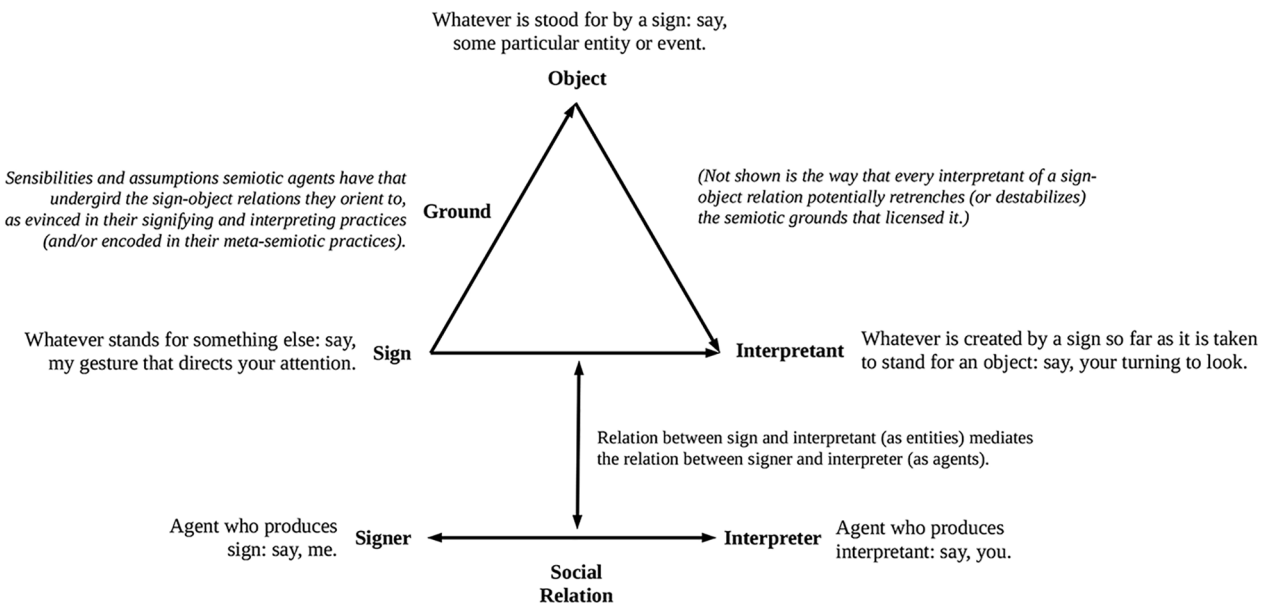


Figure 1. Practices, grounds, and agents

and symbolic codes they have in common to share their thoughts and coordinate their actions (not to mention lie, cheat, pray, and steal).

Such semiotic grounds have many faces, and go by many names. Here are just a few (in no particular order): grammars and lexicons, diagnostics and causal logics, cosmologies and astrologies, taxonomies and partonomies, concepts and schema, registers and repertoires, conventions and codes, paradigms and models, theories of minds and beliefs about signs, understandings about societies and psyches (from Smith's *Wealth of Nations* to Nietzsche's *Genealogy of Morals*), ontologies and epistemes, stereotypes and prejudices, felicity conditions and deictic grounds, superstition and common sense, knowledge and ideology, sensoria of qualia and their mediated extensions (from microscopes to ultrasound), criminal profiles and affective propensities, tropes and tensors, situated contexts and medical diagnostics, hermeneutics and horoscopes, metaphors and theories, and far far beyond.

In effect, a semiotic ground is anything that potentially relates one entity or event to another, such that a mind, or agent more generally, may "move" from the former (qua sign) to the later (qua object). They are called grounds, because they seem to stand beneath, or remain in the background of, semiotic practices per se (which tend to be more striking figures). That said, they can easily be turned into figures through metasemiotic practices that point to them, characterize them, and reason about them, as well as practices that teach them, make them, or critique them. Finally, as should be clear, such semiotic grounds, when they are understood to be relatively shared, and self-reflexively so, by members of some agentive collectivity (in contrast to other collectivities, who share other grounds), as condition for, and consequent of, their semiotic practices, constitute a large part of what is often called "culture."[3]

Machines, along with the algorithms they run, are semiotic agents (however derivative their interests, or deterministic their interpretants). As will be shown in the next section, they engage in semiotic practices (such as calculating the output of a function given an input) based on semiotic grounds (such as the specification of a mediating function, and/or the rules of an algorithm). And they may update such grounds as a result of such practices (say, by changing the parameters used in a particular function, and/or the functions called by a

3. As anthropologists, we observe the semiotic processes of others as signs (their utterances, actions and affects, their rituals, artworks and instruments), and we use our own semiotic grounds to draw inferences about their semiotic grounds, so far as such grounds constitute the roots and fruits of such processes. Contra Geertzian anthropology, culture is not public because meaning is; cultural patterns and meaningful processes are partly observable and partly inferrable, just like everything else. They lie in the movement between, and mediation of, such figures and grounds, our own and those of others.

https://doi.org/10.1086/708249 Published online by Cambridge University Press

particular algorithm). From such a stance, machine learning is a mode of *machine semiosis*.

As semiotic agents, computers are bound to us through social relations: we relate to them as signers to interpreters, as mediated through our signs and their interpretants (and vice-versa). To some degree, their semiotic grounds not only derive from our semiotic grounds, they also drive them (and vice-versa). Finally, with new modes of machine learning (such as so-called deep learning and the like), the grounds of computers have gotten sophisticated enough to undertake human-like (as opposed to super-human) semiotic processes. Optimistically, they can recommend enjoyable songs and compose captivating narratives (and not just factor prime numbers and calculate complicated trajectories). Pessimistically, as highlighted in the introduction, they may come to write real-time histories of our actions and intentions, and offer always-on ethnographies of our relations and affects. For this reason, along with newly available (and often massively-large) data sets, as well as recent developments in computer science and engineering, *our* respective semiotic grounds (i.e., those of both humans and machines), and calculating/interpreting practices, to an unprecedented degree, and on an enormous scale, have become intrinsically *coupled*.

2. Algorithmic Model as Interpretive Ground

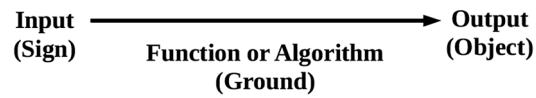
Having just seen the role that semiotic grounds play in the interpretive practices of humans, we may now look at the role they play in the algorithmic processes of machines.[4] Just as a semiotic ground relates signs to objects, a (parameterized) function relates inputs to outputs. That is, when given an input, and a set of parameters, a computer calculates the function, and/or runs an algorithm, and thereby returns an output. Schematically, this might be represented as follows: $O = F(I, P)$.

To start with a very simple example (that nonetheless scales), we may take as our ground a linear function, such as $y = m \cdot x + b$. In this case, x is the input, y is the output, and m and b are parameters (to be determined) that correspond to the slope and y-intercept of the function in question. (Where all such variables have real numbers as values.) Such a function, when properly parameterized, might return the weight of a person (y) when given their height (x), as shown in figure 2.

To move to a more interesting example, we may take as our ground a more complicated function: $\bar{y} = \sigma(M \cdot \bar{x} + \bar{b})$. In this case, $\bar{y}$ and $\bar{x}$, as the output and input of this function, are vectors (and so lists of numbers $[x_1, x_2, x_3, \ldots]$, rather

4. And hence, at one degree of remove, in the interpretive practice of humans.

https://doi.org/10.1086/708249 Published online by Cambridge University Press



Output = Function(Input, Parameters)

$$y = m \cdot x + b$$

(Where y is output, x is input, and m and b are parameters that correspond to the slope and y-intercept of a linear function.)

Such a function (when properly parameterized) might give the weight of a person when given their height.

Figure 2. Algorithmic model as interpretive ground

than numbers per se). The parameters of the function, M and $\bar{b}$, are a matrix and a vector (and hence a two-dimensional array of numbers and a list of numbers, respectively). The operation $M \cdot \bar{x}$ represents the multiplication of a vector by a matrix. This results in a vector that is then added to the vector $\bar{b}$. Finally, this operation (of vector addition) results in a vector that then becomes the input of $\sigma()$, the softmax function (or something similar), which reshapes the values in this vector so that they can serve as a probability distribution. (Obscured by the compactness of this second example is the fact that we have moved from two parameters [the slope m and y-intercept b of a line] to possibly thousands of parameters: all the values contained in the matrix M and the vector $\bar{b}$.)

Framed another way, itself evocative of a completely different imaginary, every x_i in $\bar{x}$ constitutes an input to a layer of "neurons"; each y_i in $\bar{y}$ is the output of one of the neurons in this layer; and M and $\bar{b}$ are the "weights" and "biases" of the neurons in the layer.

In such an imaginary, note how terms like *weights* and *biases*, which we might rechristen "parameterized prejudices," simply and starkly capture the relation between algorithms, values, standards, and violence. Note as well the metaphorical—and actual—embodiment of such violence in the "nervous system" of the machine (however distributed). As intimated by such a metaphor, and as will be further developed in what follows, machine learning is, in no uncertain terms, the highly successful, lucrative, and brutely physical training of a kind of a body to be biased to certain kinds.

Figure 3a fills in some of the details of such an imaginary, diagramming the connections in the network as well as writing out the entries in the matrix. Figure 3b

https://doi.org/10.1086/708249 Published online by Cambridge University Press

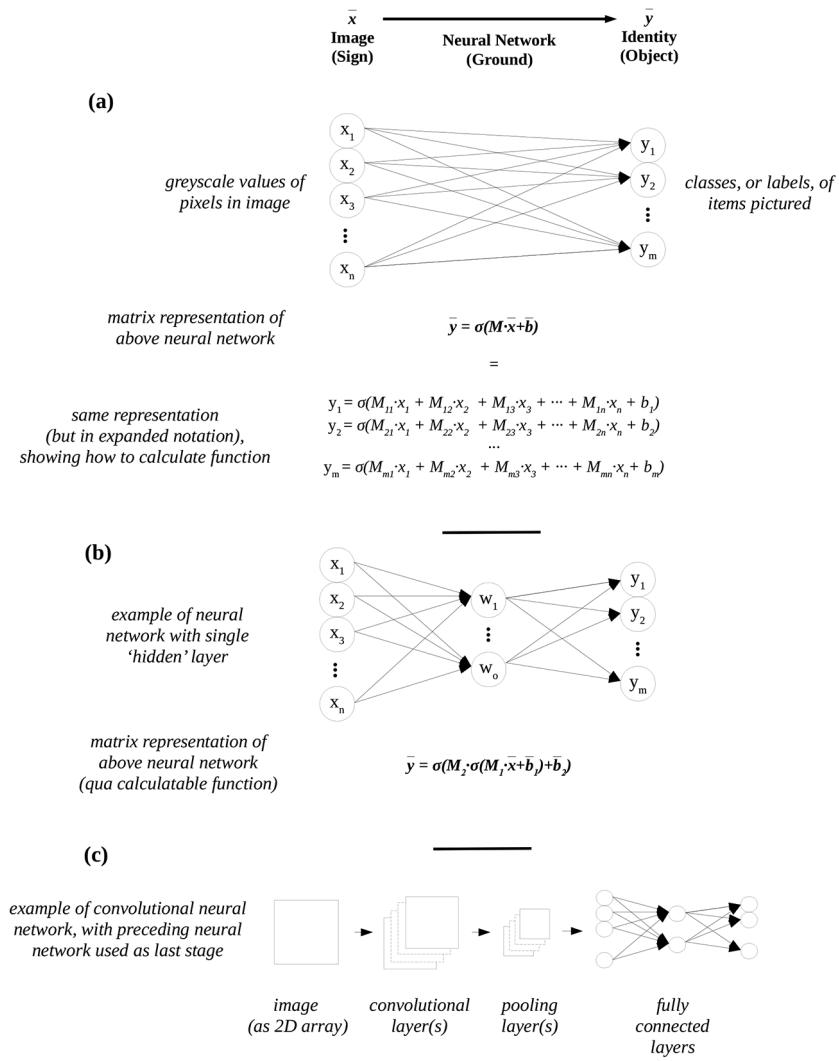


Figure 3. Neural network as interpretive ground

demonstrates how any such layer of neurons may be networked with other layers of neurons, in more and more complicated arrangements or "architectures." Finally, figure 3c schematizes one such arrangement: a convolutional network, in which each square represents, in a highly condensed fashion, a two-dimensional array of neurons (along with their respective weights and biases, qua parameters), themselves selectively connected to other neurons in adjacent layers. Note that, however complicated the architecture, all such arrangements ultimately

https://doi.org/10.1086/708249 Published online by Cambridge University Press

constitute computable functions: assuming their parameters have been established, they will return an output (qua object) when given an input (qua sign). Such a calculation is usually described as "forward propagation."

Regardless of the framing, such a function (when properly parameterized) and/or such "neurons" (when properly networked with others of their kind, in more or less complicated architectures) might give the identity of a person when given their image. That is, when given a relatively long list of numbers (corresponding to the grayscale values of particular pixels in an image), such a function returns a relatively short list of numbers (corresponding to the probabilities that the person pictured belongs to various possible identities: Dave, Alice, Susan, or Joe; diseased or healthy; happy, sad, angry, or indifferent; knife, fork, or spoon); or, to return to that violent, embodied imaginary, and algorithmic dystopias more generally: male or female; white or black; straight or queer; rioting or queuing; citizen or alien; saluting or plotting; ethnic majority or ethnic minority; healthy or sick; critter or vermin.

3. Learning Machines and Gradient Descent

Everything just described is typically the easiest part of machine learning: when given an input, return an output, by using an already parameterized function. Much more difficult is determining good parameters for such a function, or determining a good function to be parameterized. (And by "good" we don't of course mean non-evil but, rather, adequate to the kinds of patterns our model is meant to capture or enclose.) Indeed, it's really only this second step—determining good parameters—that seems to resemble "learning" per se. To return to the notion of interpretive grounds, this is not the case of inferring fire having perceived smoke; this is the case of associating smoke and fire in the first place. So how does a "machine" do it?

The most straightforward way a machine may be made to learn, or "trained," is as follows. First, we need some training data: a large set of inputs whose outputs are already known. For example, the actual ages and weights (or images and identities) of various people in a given population. We can represent such training data as a collection of already paired input-output values: $[y_i, x_i]$, where i is an index that ranges over all the people in the collection (say, from 1 to N).[5]

5. Do not confuse this use of i (which indexes a data point) with the i used in the last section (which indexed a component of a vector). Such mathematical uses of the term index (qua numerical subscript on a variable which allows one to refer to a particular element in the list or collection of elements that the variable refers to) are different from, yet closely related to, the semiotic sense of an index introduced by Peirce: an indexical sign is one that stands for its object via a relationship of causality and/or contiguity.

https://doi.org/10.1086/708249 Published online by Cambridge University Press

While such collected, categorized, and curated training data sounds innocuous enough, it is not just the source of most of the "patterns" to be modeled but of many of the prejudices as well.

Next, we need a function to be parameterized. For example, having examined the training data (and perhaps even plotted the points), we might assume that a linear function ($y = m \cdot x + b$), with two parameters (m and b), is complicated enough to capture the correlation between age and weight. The general strategy will be to initial these parameters to some unmarked values, and then slowly tweak them until the model (= function + parameters) adequately captures the patterns in the training data and thereby arrives at a kind of epistemic closure.

To do this, we need a good measure of the discrepancy between our model and the training data. That is, we need a way to quantify the magnitude of the difference between the predicted outputs of our model and the actual outputs of our training data. While there are many such measures (collectively referred to as "loss" functions), here is a particularly simple example:

$$L = \frac{1}{2N}\sum_{i=1}^{N}(y_i - (m \cdot x_i + b))^2.$$

As may be seen, this measure returns the average value of the square of the difference between the actual values (y_i) and the predicted values ($m \cdot x_i + b$), taking into account all members of the "population" ($i = 1, 2, 3, \ldots, N$). It thereby folds together the training data, and the model (= function + parameters). The smaller the difference between our model and the data, the smaller the value (L) returned by this measure. In short, such a formula aggregates all the differences between what the model predicts and what the training data shows into a single number; and that number thereby constitutes a sign of how well the model seems to be working.

Finally, we need some way to minimize this loss (L), qua measured discrepancy between our model and the training data. To do this, we need to find values for our parameters, m and b, that make L as small as possible; and we need to find these values in a relatively automated way. Readers familiar with calculus will already know that we can take the (partial) derivative of L with respect to our parameters m and b. In particular, the derivative of L with respect to m ($\partial L/\partial m$) tells us how L changes as m changes. If this derivative is positive (> 0), it means that an increase in m leads to an increase in L, whereas a decrease in m leads to a decrease in L. If this derivative is negative (< 0), it means that an increase in m leads to a decrease in L, whereas a decrease in m leads to an increase in L. And if this derivative is neither negative nor positive ($= 0$), it means

https://doi.org/10.1086/708249 Published online by Cambridge University Press

that a small change of m, in either direction, will not affect the value of L. Similar considerations hold for the derivative of L with respect to b ($\partial L/\partial b$) and any other parameter our model might employ.

All this being the case, it means that if we know the partial derivatives of L with respect to its various parameters, we know how to change the values of those parameters such that L gets smaller and smaller, such that our model (= function + parameters) offers a better and better representation of the data. Such partial derivatives are called the *gradient* of L ($\nabla L(m, b) = \partial L/\partial m,\ \partial L/\partial b$), and such a practice of way finding in a space of possible parameter values, such that the value of L gets lower and lower, is called "gradient descent" (see fig. 4). In particular, we calculate the gradient of L (given our training data, the function, and our current parameter values), adjust the values of our parameters so that they displace us a small amount (α) in the desired direction, and then repeat the process over and over again until we find ourselves in the most desirable location. Simply stated, such a procedure iteratively calculates which direction of parameter adjustment will minimize the loss function, and then adjusts those parameters in that direction. This process is usually described as "backward propagation."

Note, then, that in forward propagation, as was depicted in figure 2 and figure 3, the parameters of a function are presupposed, and we use the parameterized function as a ground that returns an output when given an input: $O = F(I, P)$. In backward propagation, in contrast, various input-output patterns are given as inputs (the training data), and we use a metafunction (or algorithm) to set the

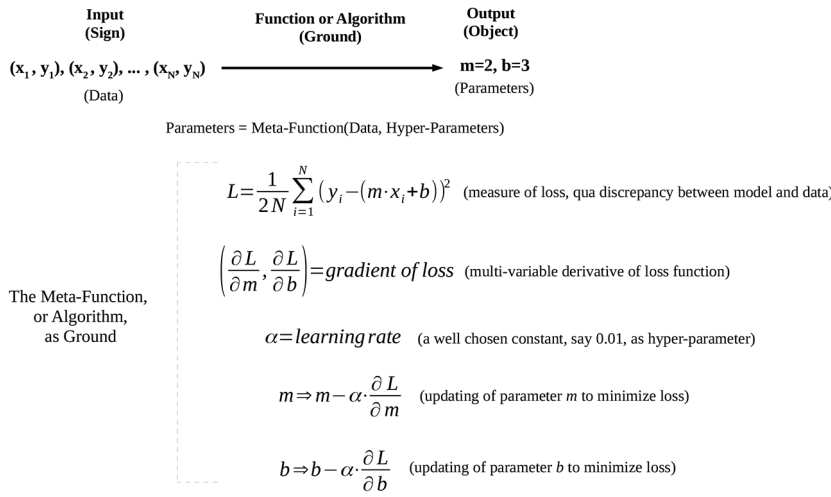


Figure 4. Gradient descent

https://doi.org/10.1086/708249 Published online by Cambridge University Press

parameters of the first function as outputs: $P = MF([I, O], HP)$. Crucially, in determining such parameters, we usually need to have set various hyper-parameters, or "options," of this metafunction. One such hyper-parameter was described above: alpha (α), or the learning rate. But there are many others, with names like batch-size, number of epochs, stride length, and so forth. Unlike parameter values, which are determined by training the model on data, hyper-parameter values are typically set before the learning process begins, and are often adjusted via trial and error, feedback and tuning. That is, one must select not just a good architecture (qua calculable function) for solving a problem but also good hyper-parameter values for efficiently finding parameter values (for that function).

Such an algorithm (gradient descent in the service of back propagation) is the workhorse of machine learning. In broad strokes, it generalizes for a wide variety of functions (such as the convolution networks used in image recognition) and is thus capable of minimizing a loss function that has thousands, or perhaps even millions, of parameters. To be sure, as the networks become "deeper" (with more and more layers of "neurons," qua recursively applied functions), the math gets hairier, the algorithms more complex, the data more massive, and the calculations more resource intensive. To be sure, there is an art to it (if not a magic) as much as a science. And, to be sure, the proceeding discussion barely scratches the surface of a rapidly changing field. Nonetheless, this preliminary sketch should be enough to build on as we make our way into some of these complications, themselves as ethical as they are technical, focusing on the ways such algorithmic processes interact with interpretive practices.

4. Objects, Data, and Models

Having looked at some of the details of machine learning, it is useful to step back for a moment and examine the relation between objects, data, and models. The meaning of these terms, and the multiple ways their referents mutually mediate each other, may be unpacked with the following example. Imagine we are trying to understand some perceived and/or putative correlation: the weight of children as a function of their age; the price of apartments as a function of their square footage; the emotions of people as a function of their facial expressions; the identities of individuals as a function of their images; the sexuality of people from their speaking habits; and so forth. For the moment, we are not going to worry about the reason for, or reality of, the correlation.

In such a context, we might have some *data*, understood as pairs of potentially correlated variables, such as the age (x) in months and weight (y) in

https://doi.org/10.1086/708249 Published online by Cambridge University Press

pounds of various children (in some population of people, at some point in time). For this data, the *object* is the actual population of children whose ages and weights were measured, including perhaps all the biological, environmental, economic, and social processes that may have given rise to the weight-age relation being modeled. Needless to say, such objects are usually infinitely deep: there is no way to exhaust, predetermine, or fix the varieties and specificities of data they could possible generate. That is, infinities of other correlations may be found in infinities of other data drawn from one and the same population. Finally, the *model* is some parameterized function (such as $y = m \cdot x + b$), and set of parameter values (such as $m = 7.6$ and $b = 9.2$), that best captures the age-height relation evinced in the data (given some process of modeling, as described in the past two sections).[6] Figure 3 showed some relatively complicated models.

As may be seen from figure 5, object, data, and model relate to each other—and thereby mediate each other—in a variety of ways. Moving clockwise around the diagram, the object causes the data, the data informs the model, and the model represents the object. Moving counterclockwise around the diagram, the object constrains the model, the model conditions the data, and the data indexes the object. In the terms of Section 1, data relates to object as sign, and model relates to data as interpretant. Note, then, the proliferation of semiotic grounds, mediating not only object-sign relations, but also sign-interpretant and interpretant-object relations, in multiple ways. (We will see still other modes of mediation in later sections.) Each of these three elements is inseparably coupled to the others and so comes into being through them, while not necessarily being determined by them.

More specifically, the object-data relation has two directions of mediation. In one direction, the object causes, or gives (rise to), the data (as fire causes smoke, say, or overeating leads to indigestion).[7] In the other direction, the data indexes the object and thus functions as a sign (symptom, index) of it, to an agent that understands the causal relation.

The data-model relation also has two directions of mediation. In one direction, the data informs the model: the pairs of numbers are used to determine which function to use, and what values should be assigned to its parameters. In the other direction, the model (or modeling process more generally) constrains

6. This is a particular way of rendering the object-sign-interpretant relation introduced in the last section.

7. Instead of "causes" we might say "produces" to foreground the labor that goes into the creation and curating of data; but that is another essay (as is adversarial learning, unsupervised learning, feature extraction, inter alia).

https://doi.org/10.1086/708249 Published online by Cambridge University Press

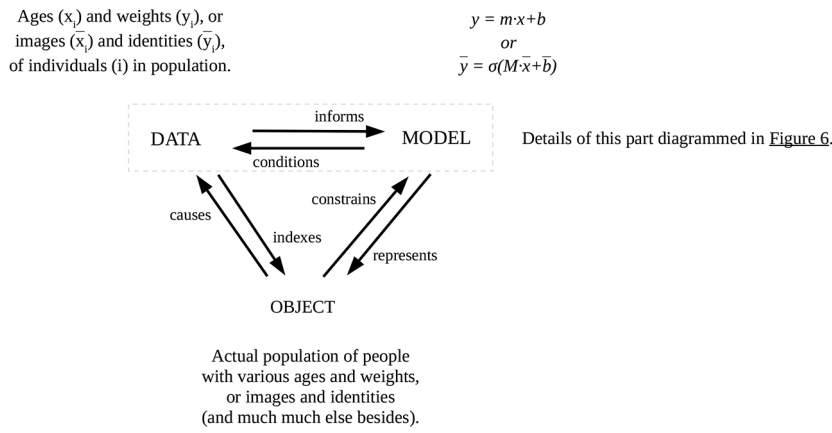


Figure 5. Modes of mediation

the data: depending on what set of models we are using, we might produce, collect, and curate different kinds of data.

Finally, the model-object relation also has two directions of mediation. In one direction, the model is a representation of the object: it is meant to stand for features of the object, and thereby provide information about that object, by drawing an agent's attention to those features. In this regard, it is not just a symptom of the object; it is an intentionally designed representation of that object, as addressed to some interpreting agent. In the other direction, the object constrains the model: it is a key means to judge whether or not the model offers a good, or at least adequate, representation. It can prove the model wrong, incomplete, or uninformative and thereby send the modeling agents back to the drawing board (to calculate better parameters for their function or to change the function per se).

Needless to say, figure 5 is quite optimistic as it stands. As will be shown in the next two sections, it leaves out the (rhizome-like) fine structure of each particular mode of mediation; and it omits a variety of other modes of mediation that make the dynamics of modeling somewhat messy (if not chaotic). That said, as simple as this initial example is, the modes of mediation in question will be shown to productively generalize and thus be easily ported to an enormous range of other objects, data, and models.

5. The Data-Model Relation

We now want to open up the data-model relation, as it was portrayed in figure 5, and detail some of its internal structure or "stages." If our function, along with

https://doi.org/10.1086/708249 Published online by Cambridge University Press

our parameter values, constitutes a good model of the data, then we can calculate an output given an input. Recall our discussion of forward propagation as summarized in figure 2. If this function is well chosen, but we don't yet know the proper values of the parameters, then we may use another strategy (typically an algorithm) to find these parameters. Recall our discussion of backward propagation as summarized in figure 4. Finally, if we cannot adequately capture the data with this function for any parameters, we might look for another function altogether (with more parameters, higher-order terms, additional variables, more exotic shapes, and so forth), and then see if this will capture the data and thereby offer some kind of epistemic closure.

We might call this last process "selection and creation" (of a model), as opposed to "propagation" of that model (be it forward or backward). To take the simplest example of such a process, perhaps the relation between x and y is not line-like ($y = m \cdot x + b$), but rather parabola-like, and so requires a higher-order term (so that $y = n \cdot x^2 + m \cdot x + b$). Or perhaps the value of y depends not just on the value of one independent variable (x) but on the values of two independent variables (so that $y = m_2 \cdot x_2 + m_1 \cdot x_1 + b$). To take a more complicated example, perhaps the relation between x and y cannot be captured by a single layer of "neurons" ($\bar{y} = \sigma(M \cdot \bar{x} + \bar{b})$ but requires an additional hidden layer as well, or a different architecture altogether (such as a multilayered, and many channeled, convolutional network). Recall figure 3. Crucially, generation requires not just selecting (or creating) a parameterizable function that is suitable to the patterns in the data (qua model); it requires selecting (or creating) a hyper-parameterized metafunction that is suitable for determining the parameters of this function (qua method).

Figure 6 is an ideal-typic representation of the complex flow of meaning (and math) involved in such a modeling process. Not so much a hermeneutic circle as a semiotic rhizome. It shows some of the details of such a possibility space (understood as a set of stages, sequentially accessed in multiple ways), represented as a directed graph. The three thick horizontal arrows, labeled with capitalized Roman numerals (I, II, III), indicate the three semiotic processes, or stages, just described: determine output (given input); determine parameters (given data); determine model and method (given problem). The three vertical ellipses, labeled with lowercase Roman numerals (i, ii, iii), indicate modes of evaluation that occur after such processes: are the outputs, parameters, or models acceptable (given some criteria). The winding arrows, labeled by number-letter combinations (1a, 2b, 3c, etc.), indicate movements from possible "destinations" (on the right-hand side) to possible "origins" (on the left-hand side), as determined

https://doi.org/10.1086/708249 Published online by Cambridge University Press

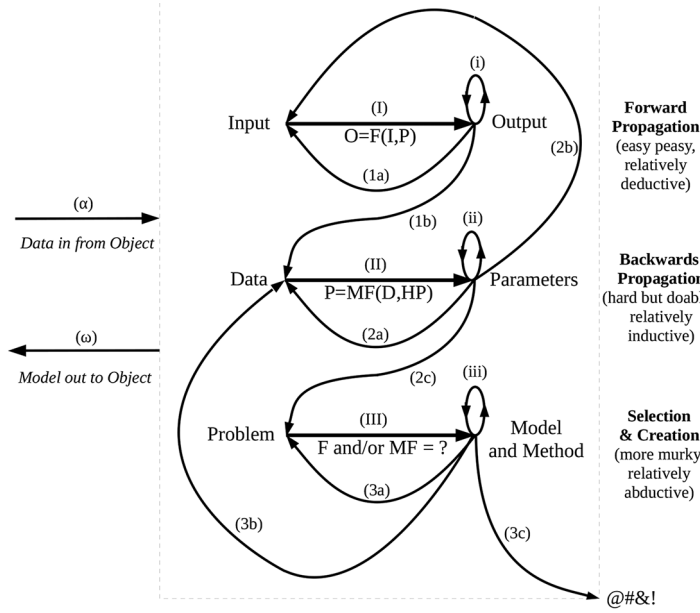


Figure 6. Interaction of model and data

by such evaluations. The large box with a dotted outline, that encloses all the foregoing paths, indicates the data-model relation. It is an expansion of a similar box shown in figure 5. Finally, the two leftmost arrows, labeled with Greek letters, indicate the object-data and model-object relation (from that same figure). To be sure, as an ideal type, there is a lot of fine structure not shown in this diagram, as well as additional loops not added. We will delve into some of this structure, and some of these loops, below.

Suppose we are at the beginning of arrow I, which is just the stage known as forward propagation. We give our parameterized function an input (such as an image), and it returns an output (such as an identity). We are now at the beginning of ellipse i, which evaluates this output. Does it seem reasonable (given some criteria, trials, or thresholds). For example, does it correctly identify images in our test data (within a desired margin). If so, we follow arrow 1a back to the beginning of arrow I, and thus prepare to engage in forward propagation again (with a new input). If not, we follow arrow 1b down to the beginning of arrow II, and thus prepare to update the parameters of our function (and thereby "train the model" some more).

Suppose we are at the beginning of arrow II, which is just the stage known as backward propagation. We give our hyper-parameterized metafunction training

https://doi.org/10.1086/708249 Published online by Cambridge University Press

data (as an input), and it returns values for the parameters of our function (as an output). We are now at the beginning of ellipse ii, which evaluates this output. Does it seem reasonable (given some criteria, or thresholds)? For example, is the value returned by the loss function small enough (to suit our purposes)? If so, follow arrow 2b back up to the beginning of arrow I and thus prepare to engage in forward propagation again (with a new set of values for the parameters of the function). Is the value returned by the loss function smaller than it was before (so it looks like our model is learning), but not yet small enough to suit our purposes? If so, follow arrow 2a back to the beginning of arrow II and thus prepare to train the model some more. If not, it may be that our model cannot capture the patterns in this data, and so follow arrow 2c down to the beginning of arrow III (and thus prepare to find a new model, or a new method for training the old model, if only by adjusting some hyper-parameter).

Suppose we are at the beginning of arrow III, which is just generation. This is a much more open-ended and much less easily automated process. In part, it consists of selecting, or creating, a good model (qua parameterizable function) to capture the patterns in some data. In part, it consists of selecting, or creating, a good method (qua hyper-parameterized metafunction) for setting the parameters of such a model. Having selected, or created, a model and/or method, we are now at the beginning of ellipse iii, which evaluates the model and/or method. Do they seem reasonable given some criteria (if only our intuitions about what patterns the model is capable of capturing, or our intuitions about how adequately the method can find parameters for the model). If so, follow arrow 3b back up to the beginning of arrow II and begin training the model (with this method). If not, and we still haven't exhausted our options (qua already existing models and methods) or ingenuity (qua capacity to create new models or methods), follow arrow 3a back to the beginning of arrow II and try again. If not, and we have exhausted our options and ingenuity (if only for the moment, R&D is probably being undertaken somewhere), follow arrow 3c. As may be seen, this last path indicates a page turn or paradigm shift that may occur when all currently existing models and methods are found wanting, and the typical techniques for creating new models or methods are radically revised or reenvisioned (given past failures of the old methods and models).

It should be emphasized that the feedback loops in figure 6 (arrows i, ii, iii) are hugely consequential for the operation and existence of such models and modeling practices, notwithstanding their delicate appearance. In part, this is because they constitute key places where human minds (and additional machinery) are looped into the process. In part, this is because they provide information

https://doi.org/10.1086/708249 Published online by Cambridge University Press

which controls large-scale patterns of flow within the modeling process. In part, this is because they implement additional trials for assessing the performance of the model and/or assisting in the debugging of the code. In part, this is because they constitute additional semiotic processes, themselves dependent on semiotic grounds, which turn on relatively evaluative acceptability ranges, or relatively context-dependent thresholds: are the values converging fast *enough* (given some expectation or deadline); is the accuracy good *enough* (given some application); is the machine consuming *too* many resources (given some budget for memory, time, energy, or cash flow); and so forth. Such trials, qua evaluative loops, tuned to particular thresholds of acceptability (Kockelman 2019) mediated by particular values and directed by certain biases, are no less subject to interpretation—and hence mediated by interpretive grounds and/or algorithmic models—than any other entity or event in this process.

We have just offered a metamodel of the modeling processing itself and hence described the innards of a relatively automatized ontology and/or episteme, one in which a model—by being repetitively forced to internalize patterns in data—comes to represent the object that generated that data (where this sentence is meant to capture some of the violence, algorithmic and otherwise, in such a process). As should be clear, the space diagrammed in figure 6 may be segmented, scaled, or framed (by algorithms, subroutines, hardware, social relations, spacetimes, commodity chains, cultural values, power dynamics, ideologies, divisions of labor, etc.) in a wide variety of ways, none of which map directly onto the paths sketched out. As should also be clear, not shown in figure 6 are all the parasites that might lie along any particular path (qua arrow), themselves capable of diverting the process onto other paths (many of which are probably orthogonal to this plane, if not downright otherworldly with respect to this imaginary).

6. Questioning Presuppositions

One way to interpret the foregoing metamodel (ideal type, semiotic ground, or imaginary) is as a hierarchy of presuppositions and questions (see fig. 7). In regards to arrow I, we presume object, data, class of models, model, and parameters; and we question output (i.e., we ask what some y is, given some x, in the context of such presuppositions). In regards to arrow II, we presume object, data, class of models, and model; and we question parameters. In regards to arrow III, we presume object, data, and class of models; and we question model (and/or methods used to set its parameters, including its hyper-parameters). In regards to arrow 3c, we presume object and data; and we question class of models (or modeling process per se). In regards to arrow α, understood as a

https://doi.org/10.1086/708249 Published online by Cambridge University Press

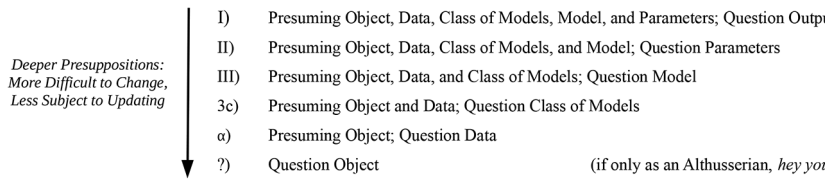


Figure 7. Hierarchy of presuppositions and questions

path that is back-trodden, we presume object; and we question our data. Finally, not shown in figure 6, we might even question the object, understood as that which the data derives from, and/or that which is represented by our model (as per the discussion in Sec. 4). While the object might at first seem to be independent of our modeling process, it is usually—and decisively—not so. Indeed, as will be seen in later sections, it is usually best understood as an agent in itself.

All that is just a rough sketch of course, and there are many caveats. Nonetheless, there is a fundamental logic behind such an organization: something is held fixed or remains invariant (the presuppositions); while something else is allowed to fluctuate or vary (the questions and, in particular, their answers). That is, the questions, when answered, gives rise to new "beliefs" (if only as newly updated values of particular variables). The presuppositions, in contrast, are older beliefs (or deeper beliefs), and so are not just "further down" (in the sense of being more difficult to change, and perhaps even to become aware of) but also "more stable" (in the sense of being subject to updating less often). This does not mean that they are hidden deeper in some subject, but rather that they are rooted more widely in some world. To be sure, it often takes something like a failure to function, or a frustration that only emerges in functioning, to cause a modeling agent to dig down deeper (or rather range out wider)—qua arrows 1b, 2c, or 3c. Relatively speaking, that which goes without question is that which grounds our questions.[8]

(Note, by the way, how such presuppositions may constitute shared values [held by some particular collectivity of agents, composed of humans and machines alike, engaged in algorithmic modeling]; and note the ways such values, in their updating or transformation, diagram a space-time of social relations within this collectivity, however fleeting: who is committed to what [output,

8. In terms of their stakes for human-being, or their role in the production of knowledge, such epistemic dynamics might be usefully compared to a range of other ideas—not just Foucauldian epistemes, Wittgenstein's language games, and Kuhnian paradigms but also, as one reviewer helpfully pointed out, issues pertaining to confirmation holism (via Quine and Duhem), proofs and refutations (via Lakatos), and various kinds of experimental and theoretical regress (via Collins and Kennefick). I have yet to work through, or sort out, such connections.

https://doi.org/10.1086/708249 Published online by Cambridge University Press

model, method, parameter, hyper-parameter, data, object, etc.], when, for how long, why, and with what repercussions?)

All that describes some of the ways the model-data relation might, through a repeated set of failures (to model the data), given all those trials (qua evaluative loops), lead to a backward-directed causal process, whereby the object being modeled, as opposed to the model of the object, is called into question. Indeed, so far as it goes "all the way," and one begins to question the object itself (as some small swatch of the world), this model-data relation might change "the world" (if only by causing "it" to turn as a function of having been stopped or caught, hailed or questioned).

Crucially, one cause of this self-undercutting questioning is *semiotic strain*: when the output of a process is out of whack with "our" expectation (given some other set of assumptions), and so functions as evidence (sign) of a parasite (object), in the extended sense of Serres and Shannon. Such a divergence between predicted and expected values functions as a symptom of a bad, or poorly functioning, model (given some higher-order or, more often, simply other ontology or model). All the elliptical paths in figure 6 are also, in some sense, attempts to grapple with more manageable varieties of semiotic strain—qua relatively imaginable varieties of relatively unexpected results. That said, there are surely more subtle, if not unimaginable, varieties of semiotic strain—the machine learning equivalents of neurosis, parapraxes, noise, and dreams. Time will tell what kinds of psycho-(social-semiotic-cybernetic) dynamics will result, and what will lead to, or forestall, their diagnosis and cure, their demonization and exorcism, their enclosure or capture.

7. The Model-Object Relation

We may now inquire into the ways that our modeling process, as summarized in figure 6, may come to perturb, if not transform, the object being modeled. We just sketched one somewhat surreptitious and possibly infrequent way this might happen: in attempting to model an object, and repetitively failing, we may question deeper and deeper presuppositions of the modeling process, until we go so far as to question not just the data (generated by the object, and constraining of the model) but also the very object itself (and/or our relation to it) and all the devices that help stabilize such relations (including our very own signs, models, machines, and modes of being). In this section we analyze a more straightforward way such object-perturbing effects may happen: our model of the object, insofar as it functions as a representation of the object, is interpreted by a range of agents; and the interpretants of those agents ultimately have some

https://doi.org/10.1086/708249 Published online by Cambridge University Press

effect on that object, making it more or less like the way it was represented by the model (or simply different from, or otherwise than, how it was before it was subject to such a representation). Or, to cast all this in the terms of Section 1, *the interpretant of a sign brings about a change in the object that is stood for by that sign*.

As will be shown, the performativity of modeling loses much of its magic and mystery when one includes the interpretant (not to mention grounds and agents) in one's model of the sign-object relationship.

Here is a simple example. Suppose we have found a good model of the correlation between two variables (say, age and weight, for a population of children): $y = m \cdot x + b$. We use this model to predict the weight of children given their age and also to inform parents (or pediatricians) how particular children compare to such a prediction—are they heavier or lighter than expected, for example. Thus, while such an agent might not learn the model per se, they might learn the predicted weight for a child's age (given the model) and hence the discrepancy between the actual and predicted values. Crucially, such predicted values are not just *modeled* values—they often become *modalized* values: not just predicted, but expected; not just average, but normative; not just inferred, but preferred. Assuming such agents have some image of an ideal range, or acceptable threshold, for a value like weight—one in which the child is, say, not too much above (or below) the average or expected value for their age—such agents may act in reference to such an ideal. For example, a parent might begin to provide the child with more or less vitamins, milkshakes, salads, or soda (given other models they have, however reasonable or ridiculous, regarding relevant cause-effect relations, such as diet-weight correlations). Assuming weight is indeed correlated with diet (even if not in the way the parent expects), the parent's dietary actions, themselves interpretants of the model in light of their own beliefs and values, may thereby have an effect on the object: the population of children being modeled and, in particular, the age-weight correlation within that population. Finally, insofar as it perturbs the object, it will change the data generated by that object and hence will call into question the accuracy of the original model which was created using data that was generated by the object prior to the intervention. Somewhat provocatively, a true sign-object relation leads to an interpretant that changes the object in such a way that the sign-object relation becomes false.

All this may be put in slightly more formal terms and thereby related to our earlier discussion of semiotic practices and mediating grounds. As may be seen in figure 8, there are two conjoined semiotic practices, each of which consists of

https://doi.org/10.1086/708249 Published online by Cambridge University Press

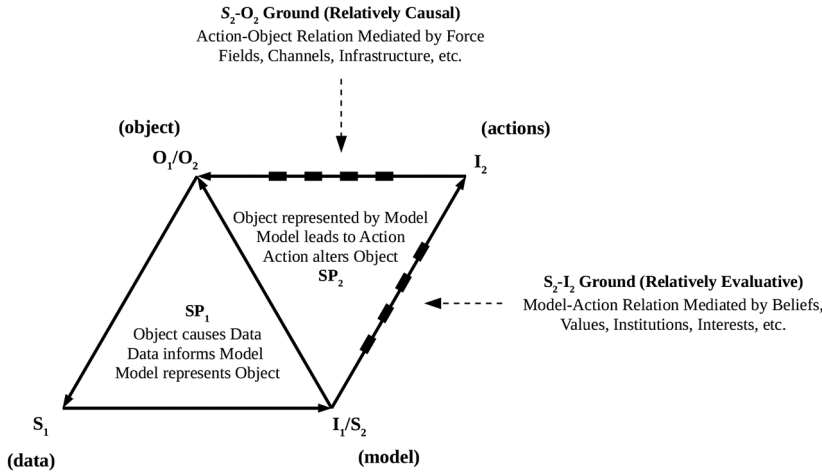


Figure 8. Object transformed by model by means of actions as mediated by grounds

an object-sign-interpretant relation. In the first semiotic practice (SP_1), the object (O_1) causes the data (S_1), which informs the model (I_1), which represents the object (O_1). These relations were analyzed in prior sections, and might best be labeled "epistemic dynamics." In the second semiotic practice (SP_2), the model (S_2) represents the object (O_1), and thereby gives rise to an action (I_2), which ultimately transforms the object (O_2). These relations are the focus of the present section and might best be labeled "performative dynamics."

Note, then, that the model is framed as an interpretant from the standpoint of the first semiotic practice, whereas it is framed as a sign from the standpoint of the second semiotic practice. The second semiotic practice thereby builds on, and partially overlaps with, the first semiotic practice. Note that the second object is just the first object post-actions as opposed to pre-actions. It may be more or less the same as the first object, depending on the magnitude of the transformation brought about by the actions: from tiny perturbation to radical alteration. Such a transformation, whatever its magnitude or nature, may make the object more or less like the model. Indeed, perhaps most likely, the object may have its qualities, and correlations therein, transformed in a way that is not captured by the model (so far as the model was designed with other qualities and correlations "in mind"). For example, perhaps the children become anemic, listless, depressed, or obese as a function of their caregivers' actions.

Such performative dynamics are, of course, not at all specific to the modeling practices analyzed here. What is somewhat unique in regards to machine learning practices, rather, is the ways in which we may use the mathematical

https://doi.org/10.1086/708249 Published online by Cambridge University Press

machinery of machine learning practices to both specify and quantify the performative effects of the model in situ. To continue with our simple example of modeling a weight-age correlation, researchers can sample data from the same population before and after a model is deployed (by pediatricians and parents, say, over some period of time) and pose mathematically tractable questions such as: to what degree, and in what way, does the object (qua population of people) come to generate data in alignment with the model that was deployed during that period. One can then measure, for example, changes in averages and standard deviations over time, for a particular age group, in relation to the expected weight given by the model: does the average weight of seven-year-olds come into alignment, or go out of alignment, with the model's predicted weight for that age; does the standard deviation of such weights become wider or narrower; and so forth. In what ways are the temporal trajectories of such changes in averages and standard deviations coupled to practices involved in deploying—and even generating—the model? Moreover, one can even use the "loss" function, as described in Section 3, to measure the goodness-of-fit of the model to the object, for all ages, through the new data. That is, one and the same measure of loss (or really fidelity of model to data) can be used to measure the performative effect of a model's deployment as much as the success of an epistemic undertaking. In this way, the "loss" can lessen over time, not by more efforts to accommodate the model to the data by updating its parameters, but simply by letting the model, with its parameter values now fixed and presupposed, run wild in the world.

Finally, as may be seen by the series of bumps along the model-action and action-object relations in the second semiotic practice, both of these relations are heavily mediated by a variety of semiotic grounds, and so involve a lot of fine structure that is otherwise not shown in figure 8. For example, the model can only give rise to an action in the context of mental states, moods, social statuses, and so forth; or the intentionalities, affects, and identities of a variety of interrelated agents. That is, to act on the new information offered by the model regarding the object, an agent must relate that information to many other beliefs, values, obligations, institutions, and interests. Similarly, the action only transforms the object in the context of intermingled force fields, channels, assemblages, and so forth. That is, for the action to affect the object, a complicated ensemble of cause-effect and sign-interpretant relations must unfold. (And, of course, all this must typically happen on the scale of a population [many different agents engaged in many different actions], or else the object won't be transformed enough to make a difference.) Both such grounds and their mediation of such relations are complicated enough and probably singular enough so as to

https://doi.org/10.1086/708249 Published online by Cambridge University Press

require an ecological outlook (think Darwin), genealogical stance (think Nietzsche), and processual imaginary (think Whitehead) to fully understand. For this reason, it is probably unlikely that a single theory, or metamodel, can be offered to capture the probably infinite variety of possibly transformative processes.

8. When Object Modeled Is Ground of Agent

Here is another example of a subtle, yet pervasive, transformation: when the object being modeled is itself the interpretive ground of a population of agents. Suppose we have training data, culled from people's actual judgments and behavior, regarding the appropriate type for a token or the appropriate sign for an object. For example, we take photos of various handwritten digits and have people classify them as particular letters of the alphabet: that is an *A*, that is a *D*, that is an *R*, and so forth (itself a classic benchmark for judging machine learning practices). Or we take photos of different tools and have people label them with particular words: that is a blender, that is a toaster, that is a cutting board, and so forth. Such token-type relations, or sign-object pairs, as data will be used to set the parameters of some algorithmic model. Simultaneously, such judgments were themselves generated by the grounds of the people in question: in part, the people who wrote the letters, or designed the appliances; in part, the people who classified the letters, or labeled the appliances. That is, the data itself was caused by a very particular object: the (relatively shared) grounds of the agents who made up the collectivity engaged in such behaviors—their lexicons and orthographies, their dictionaries and scripts, their ontologies and habits, their techniques of the body and ways of seeing.

Now suppose there is some other agent, or collectivity of agents, who is using some algorithm to model this training data. And suppose that, through some of the techniques discussed in earlier sections, this "modeling agent" has settled upon a set of parameters for their model that captures the patterns in this data (to the degree desired and/or within some acceptable threshold). Such an agent may then interpret this model, in the sense of using it as a means for some particular end, in a variety of ways (qua I_2 in fig. 8). To focus on the relatively simple case of typing character tokens, for example, if the modeling agent belongs to some branch of the post office, they might use such a model to automate the "reading" of addresses written on letters so as to know where to send them. In the context of such an interpretive regime, for a letter to "arrive at its destination," most of the characters in the address of that letter must be unambiguously typed by such an algorithmic model. In particular, those letters whose

https://doi.org/10.1086/708249 Published online by Cambridge University Press

addresses are read incorrectly might be sent to the wrong destination, and those letters whose addresses cannot be read per se might not be sent anywhere at all.

In either case, a modal logic, or deontic ground, may emerge as the result of a sorting practice, itself licensed by a modeling process: one *must* write letters in accordance with the ground of the algorithmic model or else they *may* not get where one wanted them to go. Note, then, that a model built to *represent* a semiotic ground (of some agentive collectivity) may thereby come to *regiment* the semiotic grounds of that collectivity (if only a day, or decade, later).

Figure 9 shows all these steps as a potentially iterated (and ideal typic) sequence. As discussed, the object is, in part, the ground of the collectivity prior to its being modeled (O_1). Such an object, insofar as it consists as an ensemble of sign-object (or type-token) relations (S_i–O_i), is itself a kind of metaobject. The data (S_1) is this ground as "externalized" by the collectivity (through their judgments and practices) and collected by the modeling agent. The model (I_1/S_2) is this ground as "internalized" by the modeling agent (in the sense of embodied in the algorithms, functions, and parameters of their model). The actions (I_2) of the modeling agent in light of its model is the ground as "externalized" by the agent (so far as such actions come to regiment the collectivity). Finally, the object (O_2) is, in part, the ground of the collectivity subsequent to, and/or as perturbed by,

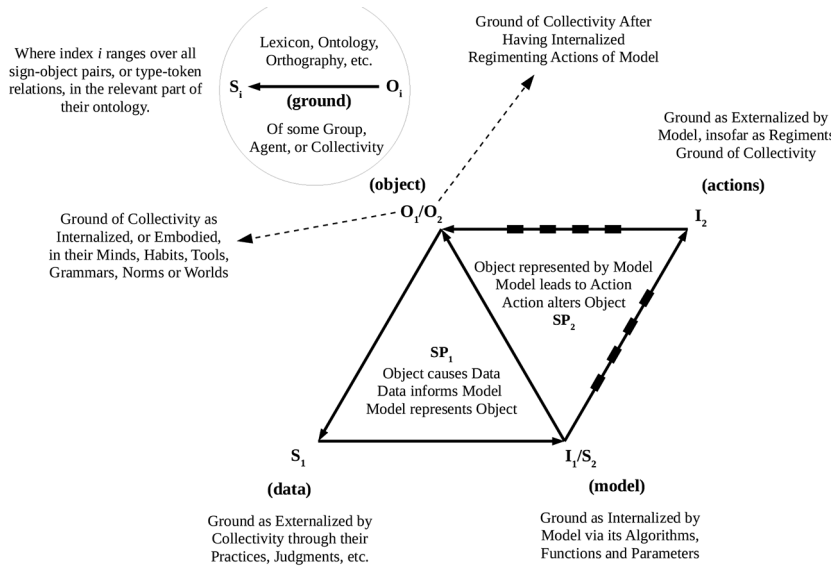


Figure 9. When object modeled is ground of some agent

https://doi.org/10.1086/708249 Published online by Cambridge University Press

its being modeled (so far as such a model leads to regimenting actions that alter the ground of the collectivity in question, qua ensemble of S_i–O_i relations).[9]

Note that, even if there is no external form of regimentation per se—qua heavy fines or jail time—agents might self-regiment in the face of such sieving practices: fear of the dead letter, so to speak, leading them to keep all their tokens more or less in line with their imaginaries of the "ideal types" their tokens will be sieved by. That is, the agents whose semiotic ground is being modeled may come to internalize the modeling agent's ground (so far as it is externalized in various actions), which was itself the modeling agent's internalization of the agent's semiotic ground (so far as this last ground was originally externalized in various judgments and practices qua data).

Such strange and loopy processes—whereby the modeling agent internalizes and externalizes the ground of a modeled agent, who then internalizes and externalizes the ground of the modeling agent, each agent all the while transforming the other agent's ground as well as its own, and so on indefinitely, on various time scales, along various dimensions, to various degrees, in relation to still other agents with still other grounds—may become a mainstay of "our" existence, the pivot on which machine-human social relations twist and turn. If so, modeling these internalization-externalization processes (and "our" metamodels of them, so far as these affect our interventions in them), as opposed to modeling the grounds being processed, will be central to understanding the rich and wily dynamics of machine-human interactions, of humanomachinic affects and action. Indeed, unlike Ian Hacking's (1995) take on related processes, these are not at all specific to human kinds, nor even human general: not only do they require machinic kinds to happen at all, but they happen to such machinic kinds as well.

To conclude, let me add one important caveat. While the processes described in these last three sections (7, 8, 9) could be called "performative," that is hardly illuminating and probably not quite right anyway, given the way this term was originally defined by John Austin. Moreover, it should be pointed out that, given our original definition of semiotic practices (signs stand for objects, and give rise to interpretants) *all* semiotic practices—insofar as they involve interpretants—lead to changes in the world. In particular, object-sign relations lead to, or create, interpretants, which not only reflect but also transform the grounds of semiotic agents: their beliefs and values (however derivative these might be). Such interpretants may be embodied in decisions, actions, and affects (to cast them in

9. Note that there is nothing "internal" about being internalized: the parameters are not "private" as opposed to "public"; rather, there is simply a different set of practices that would open them up to analysis.

https://doi.org/10.1086/708249 Published online by Cambridge University Press

human-like terms), or they may be embedded in variable updates, function calls, and rule changes (to cast them in machine-like terms). Indeed, most of the processes discussed in Section 5 (and diagrammed as arrows in fig. 6) were precisely movements, creations, or transformations of this sort: we learn a new output; we calculate a new parameter; we choose a new model. (And we move somewhere else in the diagram, qua directed graph.) So far as such arrows involve changes to the world (if only as the outputting of a value, the updating of a parameter, or the deployment of a model), they are transformative of the world in a relatively straightforward way. That said, such model-internal transformations—even though they directly involve "our" beliefs and values (if only as the updating of parameters in our models), and so are quite interesting in themselves—were not the focus of these last three sections, which were focused rather on the way our model may affect the object per se, however indirectly, slightly, or surreptitiously.

9. Recursive Warfare

Having just analyzed the case where the object being modeled by some algorithmic ground is itself an interpretive ground, we now turn to situations in which the object being modeled (by some modeling agent) is itself an agent capable of reacting to, perturbing, and even thwarting the modeling agent's model (of itself as an object). Figure 10, a slightly modified version of figure 8, shows the modeling agent (A) at the bottom right. As may be seen, such an agent "senses signs" (in this case, by utilizing data drawn from the object as the input for its model) and "instigates interpretants" (in this case, by undertakings actions based on the output of its model, actions that may more or less directly affect the object).[10] Such interpretants make sense in relation to such signs, given the way this relation mediates the relation between the interests of such an agent and the features of such an object.

The object (O_1/O_2) is, at first pass, that entity that produces the data (S_1) sensed by the modeling agent and that entity affected by the actions (I_2) instigated by the modeling agent. Crucially, all this can happen not just because

10. Whereas fig. 1 was designed to show how the relation between a signifying agent and an interpreting agent is mediated by the relation between a sign and an interpretant (e.g., the relation between a teacher and a student is, in part, mediated by the relation between the questions/answers of the former and the answers/questions of the latter), fig. 10 is designed to show how a single agent (here one that happens to be radically distributed) both senses signs (that emanate from some object or agent) and instigates interpretants (that are directed toward that same object or agent). A teacher can also relate to a student through the frame of fig. 10 (e.g., in hearing a question, qua sensed sign, they offer an answer, qua instigated interpretant), just as a state can relate to its citizens though something like fig. 1. Depending on the interests of the analyst, as well as the details of the interaction, either of these semiotic frames may prove useful or illuminating.

https://doi.org/10.1086/708249 Published online by Cambridge University Press

https://doi.org/10.1086/708249 Published online by Cambridge University Press

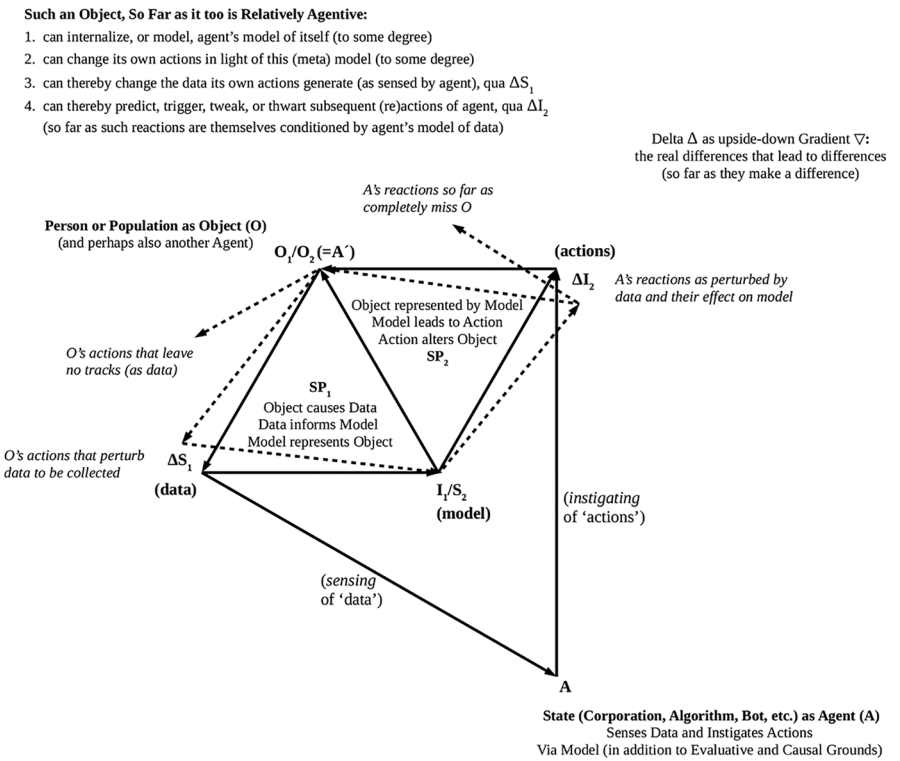


Figure 10. Object's internalization of agent's model of object

the agent has a model (of the object, through the data) but also because the agent has certain evaluative and causal grounds (as discussed in Sec. 7). While such a modeling agent could be an institution or an individual, we are focused on particularly powerful and potentially insidious agents—states, corporations, armies, police, and so forth.

The key point of figure 10, however, is that the object (be it an individual or a population) is simultaneously framed as a potential agent (A′). In particular, we are assuming that such an object, insofar as it too is relatively agentive, has a set of capacities, and so can engage in a set of practices that allow it to counter the modeling practices of the first agent (A).

In what follows, we will refer to this second agent (A′) using the pronoun *they*. For example, we are assuming that they can—to some degree—internalize and/or model the agent's model of themselves (as well as the agent's evaluative and causal grounds per se). That is, they can come to understand, to some degree, (i) the way A uses their features and practices as data; (ii) the model A constructs of them from that data; and (iii) the actions A will undertake, as grounded in that model, so far as such actions may affect them.

Moreover, we assume that they can change their own practices in light of their metamodel of A's model. For example, they can behave differently than they did before—say by feigning or masking particular actions (traits, affects, utterances, etc.) insofar as they are aware that the modeling agent may use these actions for data, and/or insofar as they are aware that such actions will be reacted to by the modeling agent in particular ways. (Some of these practices are indicated with dotted lines in fig. 10.)

Note, then, that it is precisely because they (A′) are an object (O_1/O_2) for that first agent (A) that they can so agentively trigger, perturb, thwart, or subvert its perceptions, models, and actions. To some degree, their capture (or really their being aware of the process through which they are captured) is a condition of possibility for their escape. Not in the trivial sense: there is nothing to escape from if you haven't been caught. Rather, by knowing something about the agent's models and by allowing themselves to be captured or enclosed (by the agent's models) in their own terms, they make sure such models are given erroneous, or at least insignificant, data and thereby capture a distorted image of their actual or ideal selves. One key question for future agents will be how to leverage such tactics for the greatest effect: to engage in a delta (Δ = change, perturbation, exploit) that overturns a gradient (∇ = model, algorithm, utility, trap, or enclosure).

https://doi.org/10.1086/708249 Published online by Cambridge University Press

To be sure, the first agent probably has its own meta-meta-models of their meta-models of its model; and so on, and so forth. The internalization-externalization war doesn't just occur on many fronts; it also occurs on many levels. Warfare, no less than malware, no less than *die Ware*, is recursive.

Argonauts (Δ) in a sea of juggernauts (∇).

10. The Coupling of Epistemic and Performative Dynamics

Before concluding, it is helpful to review a few key steps of the argument, especially as they concern relations among objects, data, and models, as laid out in Section 4. Recall that this essay was mainly focused on the boxed part of figure 5, qua *data-model* relation (Secs. 3, 4, 6, 7). While the innards of this box, as represented in figure 6, were quite complicated, they were nonetheless relatively easy to diagram (as a directed graph). While this fact was, in part, an effect of the ideal typic framing, it is mainly due to the relatively constrained nature of the modes of mediation themselves (notwithstanding those evaluative loops). There is a kind of finitude underlying the grounds of machine learning that does not apply to semiotic processes, or cultural grounds, more generally—at least for the moment.

The *object-data* relation, while no less complicated, was simply not the focus of this essay. For readers interested in the object-data relation, the semiotechnics of rendering the real more generally, and the relation between this kind of analysis and the kind offered by scholars like Michel Serres, Karen Barad, Annemarie Mol, Lorraine Daston, and Peter Gallison, Kockelman (2016b) offers an analysis that is complementary to this one.

In contrast to the relatively constrained nature of the data-model relation, the *model-object* relation was quite complicated (Secs. 8, 9, 10). In part, this was because another semiotic process, with additional semiotic grounds, was added to the analysis, as shown in figure 8. In part, this was because, as unruly as data and models are, objects are infinitely more unruly. In part, this was because the object being modeled (by a series of semiotic grounds) was itself a semiotic ground, as shown in figure 9. In part, this was because such a ground belonged to a relatively self-reflexive, and inherently antagonistic, semiotic agent, as was shown in figure 10.

Note, then, that while the first half of this essay focused on what we called *epistemic dynamics* (accommodating a model to an object, through data), the second half of the essay focused on what we called *performative dynamics* (assimilating an object to a model, however unintentionally, through actions).

https://doi.org/10.1086/708249 Published online by Cambridge University Press

That said, notwithstanding the ways they were separated for analysis, we saw how all such moments across both such dynamics are radically coupled.

Moreover, insofar as such inherently coupled processes fall way outside the topical bounds of what was originally called performativity (not to mention the topical bounds of what is traditionally called epistemology), and insofar as the analysis offered here was radically different in its particulars, it is probably best to leave that older analytic stance and nomenclature behind, lest we project a spurious similarity onto radically distinct practices and thereby soothe ourselves into thinking they can be handled with relatively simplistic analytics ("repetition," "ideology," "illocutionary force," and the like). It is for this reason that I have opted in this essay and others to refer to various moments of such dynamically coupled processes as *modes of transformativity* and to analyze them in relatively device-specific terms (such as the details of machine learning praxis), however general the particular patterns may turn out to be.[11]

11. Objectivity in the Social Sciences

There is nothing made by the human hand, nor conceived by the human mind, that is not affected by culture—and, indeed, is not culture per se. All such modes of making and conceiving are (parts of) semiotic processes mediated by intersubjectively shared semiotic grounds (themselves mediated by semiotic processes), in an unending recursive process, whereby semiotic processes build on semiotic grounds while building up semiotic grounds (where to "build on" can mean to stand on or exploit and to "build up" can mean to dismantle or destroy; there is nothing inherently warm and fuzzy about culture). Algorithms, and machine learning assemblages and ensembles more generally, are particularly interesting from this vantage as they not only are made by the human hand and conceived by the human mind but are making and conceiving agents in their own right—semiotic agents engaged in semiotic processes that are themselves both mediated by, and mediating of, semiotic grounds. We relate to such ensembles and assemblages not just as subjects to objects, and as rabbis to golems, but also as

11. In the context of analyzing spam filters, and Bayesean reasoning, as applicable to sieving/sorting/filtering/trapping processes more generally, Kockelman (2013) referred to various moments within such epistemic dynamics as transformativities 2, 3, and 4 and to various moments within such performative dynamics as transformativities 1 and 5. He also related such dynamics to those analyzed by other theorists, from Erving Goffman and Hannah Arendt to Mary Douglas and Ian Hacking. This early essay, then, justifies some of the presuppositions built into the current argument. It also shows the relation between such epistemic and performative dynamics and *ontological dynamics* per se (insofar as such models presume and produce individuals, kinds, indices, and worlds).

https://doi.org/10.1086/708249 Published online by Cambridge University Press

selves to others, parts to wholes, and rabbits to wolves. No less than kinship relations, body parts, shadows, exuvia, or names, they have become our inalienable possessions, our extended phenotype, our uncanny Umwelt, our friends and enemies, our ethnographers and colonists, our social network, our nervous system, our quasi-personal fringe.

That said, it should be obvious enough that algorithms (and ensembles of algorithms, machines, humans, and other agents) are modes of culture, involving semiotic processes as much as semiotic grounds. More interesting and less intuitive, perhaps, is the way that such ensembles are quickly becoming anthropologists in their own right: interpretive agents$_1$ engaged in semiotic processes that are designed to figure (out) the semiotic grounds of other agents$_2$, by attending to their$_2$ semiotic processes, all the while being both guided and led astray by the presumptions built into their$_1$ own semiotic grounds: *machine anthropologists (and historians) engaged in machine ethnography (and history)*.[12] From this vantage, the most fetishized object in all of anthropology is not the shaman, taboo, myths, *mana*, or the gift; it is not even social constructionism, relativity, culture, or the fetish per se (though those are up there); it is probably the ethnographer as a special kind of agent charged with a noble kind of task, and/or ethnography as a special mode of inquiry blessed with a singular form of insight. (But I drank the Kool-Aid long ago, and so here I am.)

Nor is performativity some mysterious process worthy of all the fetishization it has received. In particular, if one has a model of semiosis that involves grounds, agents, and interpretants, then it is a pretty obvious fact that the interpretant of a sign (qua model) can bring about a change in the object (qua world) that is stood for by that sign—making it more or less like the way it was represented by the sign or simply otherwise from how it was before it was signified and interpreted. To be sure, the fact that models can transform worlds—sometimes in ways that make them align with such models, in ways that modeling agents cannot seem to understand or do not like to acknowledge—is interesting and important. To be sure, to modeling agents prone to naturalize the patterns they find, it may come as a surprise that many aspects of the patterns they find are historically specific, culturally mediated, and/or model derived. But there is no need for some special theory of "performativity" to understand such processes; semiosis—and pragmatism more generally—had such dynamics built into it from the start, so long

12. As used here, then, *machine ethnography* is not just the ethnography (and, more generally, anthropology) of machine-learning practices; it is also the study of machines engaged in ethnographic (and, more generally, interpretive) practices, potentially through the use of, in dialog with, and/or as undertaken by the machines themselves.

https://doi.org/10.1086/708249 Published online by Cambridge University Press

as one lets agents, grounds, and interpretants do their work. My sense is that performativity gets much of its topical allure not because it is some mysterious and all important process, but simply because so many critical theorists have such rudimentary theories of semiosis. (As my mother would say, it is as if they have never been told to take out the trash, pick up their toys, do chores for their allowance, or act on what they believe.)

In addition to sketching some of our relations to such alien inalienable possessions and what they reveal about some deep-seated prejudices-qua-fetishes, this essay has also offered a model of the process of modeling through machine learning practices. Unlike the modeling processes it models, however, this metamodel was offered in the spirit of an ideal type (Weber 1955). As such, it was not meant to describe such processes in a way that may be true or false, nor even to hypothesize such processes in a way that may be tested per se. (Though, to be sure, one may find it useful to wield it that way.) Rather, it was an attempt to make clear, in a way that is both contextually portable and analytically precise, some of "the characteristic uniqueness of the reality in which we find ourselves" (if not the dystopia in which we have put ourselves), a reality that is now made, inhabited, and modeled by both human and machinic semiotic agents, working in tandem yet often at odds. Moreover, as an ideal type, each and every claim of this metamodel should usher in a score of caveats, any one of which is potentially more interesting to elucidate and investigate than the claim itself. Such is the usefulness, as opposed to the truthfulness, of the ideal type (perhaps best understood as a pragmatic typology). As such, this essay is itself a semiotic ground and so is meant to offer agents a set of assumptions and sensibilities that can enable them to signify and interpret, and thereby both represent and transform, a wide range of worlds, if only to critique, update, or overturn its own standing as a useful semiotic ground.

But all that said, we would do well to perturb Weber: the reason to model such real-world processes of modeling is not just so that we may see the discrepancies between our metamodel and such models, and thereby temper our scholarly and scientific imagination, and hence transform our models, if not our modeling practices. Rather, it is to make it easier to hack, exploit, or sabotage some of the models being made.[13] Framed as such, the point is not just to make our models correspond more and more with the world but to help the world wriggle free from many of the models being imposed on it, such that they

13. On racism in America in relation to the pragmatist tradition and Peirce's model of semiotics, see Glaude (2007) and Cummings (2018).

https://doi.org/10.1086/708249 Published online by Cambridge University Press

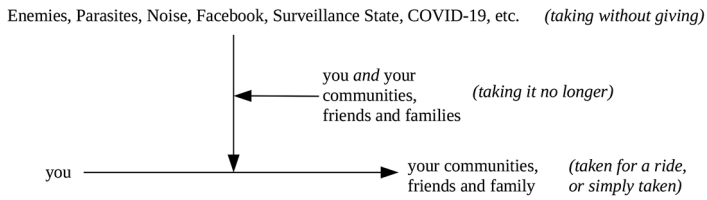


Figure 11. Take back the channel

may say "fuck you" to the agents imposing such models and regimenting such worlds (see fig. 11).

References

Agha, Asif. 2007. *Language and Social Relations*. Cambridge: Cambridge University Press.

Anglin, Mary K. 1998. "Feminist Perspectives on Structural Violence." *Identities* 5 (2): 145–51.

Callon, Michel. 2007. "What Does It Mean to Say That Economics is Performative?" In *Do Economists Make Markets? On the Performativity of Markets*, edited by Danald MacKenzie, Fabian Muniesa, and Lucia Siu, 311–57. Princeton, NJ: Princeton University Press.

Carruthers, Andrew. 2017. "Grading Qualities and (Un)settling Equivalences: Undocumented Migration, Commensuration, and Intrusive Phonosonics in the Indonesia-Malaysia Borderlands." *Journal of Linguistic Anthropology* 27 (2): 124–50.

———. 2019. "Policing Intensity." *Public Culture* 31 (3): 469–96.

Castelle, Michael. 2018a. "Deep Learning as Epistemic Ensemble." https://castelle.org/pages.

———. 2018b. "Social Theory for Generative Networks (and Vice Versa)." https://castelle.org/pages.

Edwards, Terra. 2018. "Re-channeling Language: The Mutual Restructuring of Language and Infrastructure among DeafBlind People at Gallaudet University." *Journal of Linguistic Anthropology* 28 (3): 273–92.

Elyachar, Julia. 2010. "Phatic Labor, Infrastructure, and the Question of Empowerment in Cairo." *American Ethnologist* 37 (3): 452–64.

Enfield, N. J., and Paul Kockelman, eds. 2017. *Distributed Agency*. New York: Oxford University Press.

Enfield, N.J., and Jack Sidnell. 2017. *The Concept of Action*. Cambridge: Cambridge University Press.

Erving Goffman. 1959. *The Presentation of Self in Everyday Life*. New York: Anchor Books.

Galtung, Johan. 1969. "Violence, Peace, and Peace Research." *Journal of Peace Research* 6 (3): 167–91.

Gillespie, Tarleton. 2016. "Algorithm." In *Digital Keywords*, edited by Benjamin Peters, 18–30. Princeton, NJ: Princeton University Press.

Glaude, Eddie S., Jr. 2007. *In a Shade of Blue: Pragmatism and the Politics of Black America*. Chicago: University of Chicago Press.

Geoghegan, Bernard D. 2011. "From Information Theory to French Theory: Jakobson, Lévi-Strauss, and the Cybernetic Apparatus." *Critical Inquiry* 38 (1): 96–126.

https://doi.org/10.1086/708249 Published online by Cambridge University Press

———. 2019. "An Ecology of Operations: Vigilance, Radar, and the Birth of the Computer Screen." *Representations* 147(1): 59–95.

Goffry, Andrew. 2015. "Towards a Rhizomatic Technical History of Control." *New Formations* 84/85 (16): 58–73.

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. Cambridge, MA: MIT Press.

Hacking, Ian. 1995. "The Looping Effects of Human Kinds." In *Symposia of the Fyssen Foundation. Causal Cognition: A Multidisciplinary Debate*, edited by D. Sperber, D. Premack, and A. J. Premack, 351–94. Oxford: Clarendon.

———. 2002. *Historical Ontology*. Cambridge, MA: Harvard University Press.

Keyes, Os. 2018. "The Misgendering Machines: Trans/HCI Implications of Automatic Gender Recognition." *Proceedings of the ACM on Human-Computer Interaction*, vol. 2, no. CSCW, article 88.

Knox, Hannah, and Antonia Walford. 2016. "Is There an Ontology to the Digital?" Theorizing the Contemporary, *Fieldsights*, March 24. https://culanth.org/fieldsights/is-there-an-ontology-to-the-digital.

Klausner, Martina. 2018. "Calculating Therapeutic Compliance: An Ethnographic Account of Numerical Inference and Interference in Mobile Health Care." *Science & Technology Studies* 31 (4): 30–51.

Knorr-Cetina, Karin. 1999. *Epistemic Cultures*. Cambridge, MA: Harvard University Press.

Kockelman, Paul. 2007. "Inalienable Possession and Personhood." *Language in Society* 36 (3): 343–69.

———. 2012. "The Ground, The Ground, The Ground: Why Archaeology Is So 'Hard.'" *Yearbook of Comparative Literature* 58:176–83.

———. 2013. "The Anthropology of an Equation: Sieving Spam, Algorithmic Agents, and Ontologies in Transformation." *HAU: Journal of Ethnographic Theory* 3 (3): 33–61.

———. 2015. "Four Theories of Things: Aristotle, Marx, Heidegger, and Peirce." *Signs in Society* 3 (1): 153–92.

———. 2016a. "Grading, Gradients, Degradation, and Grace: Part 1, Intensity and Causality." *HAU: Journal of Ethnographic Theory* 6 (2): 389–423.

———. 2016b. "Meeting the Universe Two-Thirds of the Way (Witchful Thinking)." *Signs in Society* 4 (2): 215–43.

———. 2017. *The Art of Interpretation in the Age of Computation*. New York: Oxford University Press.

———. 2019. "Threshold." Theorizing the Contemporary, *Fieldsights*, September 30. https://culanth.org/fieldsights/threshold.

Lee, Benjamin. 1997. *Talking Heads: Language, Metalanguage, and the Semiotics of Subjectivity*. Durham, NC: Duke University Press.

Lépinay, Vincent-Antonin. 2007. "Decoding Finance: Articulation and Liquidity around a Trading Room." In *Do Economists Make Markets? On the Performativity of Markets*, edited by Danald MacKenzie, Fabian Muniesa, and Lucia Siu, 87–127. Princeton, NJ: Princeton University Press.

https://doi.org/10.1086/708249 Published online by Cambridge University Press

Lippert, Ingmar, and Helen Verran, eds. 2018. "After Numbers? Innovation in Science and Technology Studies' Analytics of Numbers and Numbering." *Science and Technology Studies* 31 (4): 2–126.

Lowrie, Ian. 2017. "Algorithmic Rationality: Epistemology and Efficiency in the Data Sciences." *Big Data & Society* 4 (10): 1–17.

———. 2018. "Algorithms and Automation: an Introduction." *Cultural Anthropology* 33 (3): 349–59.

Mackenzie, Donald. 2007. "Is Economics Performative? Option Theory and the Construction of Derivatives Markets." In *Do Economists Make Markets? On the Performativity of Markets*, edited by Donald MacKenzie, Fabian Muniesa, and Lucia Siu, 54–86. Princeton, NJ: Princeton University Press.

Mertz, Elizabeth, and Richard J. Parmentier, eds. 1985. *Semiotic Mediation: Sociocultural and Psychological Perspectives*. Orlando, FL: Academic Press.

Moll, Annemarie. 2002. *The Body Multiple: Ontology in Medical Practice*. Durham, NC: Duke University Press.

Nelms, Taylor C., Bill Maurer, Lana Swartz, and Scott Mainwaring. 2018. "Social Payments: Innovation, Trust, Bitcoin, and the Sharing Economy." *Theory, Culture and Society* 35 (3): 13–33.

Neville, Robert Cummings. 2018. "How Racism Should Cause Pragmatism to Change." *American Journal of Theology & Philosophy* 39 (1): 53–63.

Nielson, Michael. 2018. "Neural Networks and Deep Learning." http://neuralnetworksanddeeplearning.com/.

Noble, Safiya Umoja. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press.

Onuoha, Mimi. 2018. "Algorithmic Violence." http://github.com/MimiOnuoha.

Parmentier, Richard J. 1994. *Signs in Society: Studies in Semiotic Anthropology, Advances in Semiotics*. Bloomington: Indiana University Press.

Peirce, Charles S. 1998. *The Essential Peirce: Selected Philosophical Writings*. 2 vols. Edited by the Peirce Edition Project. Bloomington: Indiana University Press [cited as *EP*].

Seaver, Nick. 2017. "Algorithms as Culture: Some Tactics for the Ethnography of Algorithmic Systems." *Big Data & Society* 4 (2): 1–12.

———. 2018. "Captivating Algorithms: Recommender Systems as Traps." *Journal of Material Culture* 24 (4): 421–36.

Silverstein Michael, and Greg Urban, eds. 1996. *Natural Histories of Discourse*. Chicago: University of Chicago Press.

Veblen, Thorstein. 1898. "The Beginnings of Ownership." *American Journal of Sociology* 4 (3): 352–65.

Weber, Max. 1955. *The Methodology of the Social Sciences*. New York: Free Press.

https://doi.org/10.1086/708249 Published online by Cambridge University Press