AI & SOCIETY (2025) 40:6135–6148
https://doi.org/10.1007/s00146-025-02430-7

RESEARCH

The practices and politics of machine learning: a field guide for analyzing artificial intelligence

Francis Lee[1,2]

Received: 22 October 2024 / Accepted: 6 June 2025 / Published online: 25 June 2025

Abstract
This article develops an analytical and methodological field guide for studying the mundane practices that constitute machine learning systems. Drawing on science and technology studies (STS), I move beyond the opacity/transparency dichotomy that has dominated critical algorithm studies to examine how machine learning is assembled through everyday work. Rather than treating algorithms as black boxes or magical entities, I focus on four empirical moments of translation—feature extraction, vectorization, clustering, and data drift—where technical work becomes political choice. By ethnographically attending to practitioners' tinkering, negotiations, and valuation practices in these moments, we can trace how classification systems are constructed and stabilized. This approach allows us to ask: How are particular features of the world selected as relevant for prediction? Through what practices are people and phenomena translated into mathematical vector spaces? How are temporal assumptions encoded in data? By studying these mundane processes of construction, we can understand how machine learning systems enact particular ways of seeing, classifying, and predicting the world. This field guide thus contributes methodological tools for analyzing how the politics of machine learning is assembled in practice, opening analytical space for critical engagement beyond calls for transparency or fairness.

Keywords Machine learning ethnography · Data practices · Algorithmic assemblages · Moments of translation · Critical AI studies · Science and technology studies

1 Making AI mundane: general analytical problems

> *"Change the instruments, and you will change the entire social theory that goes with them" (*Latour 2010*, 153).*

In the social sciences today, there is an ongoing debate about how to approach algorithms, artificial intelligence, and machine learning as analytical objects. [1]Sometimes, algorithms are proclaimed to be black boxes whose inscribed logics are impossible to scrutinize and understand (Pasquale 2016), or that "[t]here may be something in the end impenetrable about algorithms" (Gillespie 2014). Their opacity of machine learning often being argued to be even more severe than other algorithms (Burrell 2016, 10).[2]

However, the worry about the impenetrability of algorithms and machine learning—understanding algorithms as a black box—has also been argued to "prevent research more than encouraging it" (Bucher 2016, 84). The magical discourse surrounding machine learning shielding their creators from social scrutiny (Campolo and Crawford 2020).

✉ Francis Lee
francis@francislee.org

1 Södertörn University, Huddinge, Sweden

2 Chalmers University of Technology, Gothenburg, Sweden

[1] This title is of course a wink to (Martin and Lynch 2009).

[2] The black box problematic has, of course, led the field of machine learning to tackle this challenge by constituting a field of its own, FAccT, see https://facctconference.org/.

Many have asked how we can productively move away from algorithms understood as neatly bounded technical objects—black boxes—that can be "opened" (cf. Ananny and Crawford 2018; Geiger 2017; Lee 2021).

This tension around the black boxed algorithm is based on a tendency to reify the objects of computer scientists, so they seem to be stable objects that just exist—having an "uncontroversial thingness" (Suchman 2023)—even though the objects *algorithm* or *machine learning* is just as negotiated as any other object (Suchman 2023; Lee 2021; Muniesa 2019)—much like bush pumps (de Laet and Mol 2000) or bicycles (W. B. Bijker 1995).[3]

This article proposes to move away from this analytical deadlock by *outlining some ways to make the opaque magic of machine learning mundane through focusing on the practices of their construction.* The article asks: how can we make sense of the construction of machine learners in practice? Which moments, situations, or practices make the politics and practices of machine learning visible? Where do actors make choices about how machine learners should predict and classify the world?

These types of questions are not new in the social study of science and technology. By learning from a deep well of studies in the field of science, technology, and society (STS), the article outlines some moments where the politics of machine learning is possible to study in practice, and offers some strategies for analyzing and understanding the practices and politics of machine learning. The aim is to help open up a methodological and analytical space for studying the practices and politics of machine learning classification (cf. Bowker and Star 1999; Bechmann and Bowker 2019; cf. also Martin and Lynch 2009).

To achieve this aim, the article suggests some situations where the classification work of machine learners can be made mundane, where tinkering happens in practice, where the opaque magic of machine learning turns into mundane engineering work. By approaching machine learning as mundane practice—work, tinkering, and engineering—the article aims to provide a methodological and analytical strategy that bypasses some of the hype around machine learning (cf. Ziewitz 2016; Neyland 2018).

1.1 The ethnography of machine learning is non-trivial

> The activity of creating black boxes, of rendering items of knowledge distinct from the circumstances of their creation, is precisely what occupies scientists the majority of the time. The way in which black boxing is done in science is thus an important focus for sociological investigation.
> (Latour and Woolgar 1986, 259)

The ethnography of algorithms and machine learners is non-trivial methodologically (cf. Star 1999; Latour and Woolgar 1986). Understanding algorithms and machine learning in practice is complicated by a myriad factors which have been cataloged and lamented extensively, they are secret, they are difficult to understand, sometimes our informants do not understand what is happening, they may escape human comprehension, and so on (Burrell 2016; Bucher 2016; MacKenzie 2019; Lange et al. 2019; Seaver 2022). There is a growing number of studies outlining tactics, strategies, and sensibilities to study ethnographically these types of systems. For instance, Neyland (2018), Seaver (2022), and Jaton (2021b) have done long-term situated ethnographic studies of algorithmic work and grappled with how to engage machine learning algorithms in practice.

Still, just like in other ethnographic work, it is a challenge to understand and analyze our interlocutors' practices. However, this difficulty of analyzing practices is not a unique feature of black boxed algorithms or machine learning, but a difficulty in dealing with all activity in the world. Anthropology for instance has long reflected on how to approach a modern, global, and interconnected world through multi-sited ethnography (Marcus 1995), polymorphous engagement (Gusterson 1997), global flows (Appadurai 2003), or networks (Burrell 2009). In a telling reflection on the intelligibility of a "crowded, noisy city street scene, where different languages, different cultures, diverse social microworlds, and discordant frames of meaning are all thrown together in the normal course of things" Ferguson (1999, 208) writes:

> Here, there is much to be understood, but none of the participants in the scene can claim to understand it all or even to take it all in. Everyone is a little confused (some more than others, to be sure), and everyone finds some things that seem clear and others that are unintelligible or only partially intelligible. (Ferguson 1999, 208)

Perhaps, much like the modern city scene that Ferguson describes, the study algorithms and machine learning is one of these areas—where it is thorny for an ethnographer to grasp the object at hand. But maybe not only because the algorithms are impenetrable black boxes, but also because the work of observing cultures and societies is an exercise in the humility of making situated knowledge (Haraway 1988).

[3] Avoiding this reification, some have proposed that we need to treat algorithms as culture—not as "technical rocks in a cultural stream" but just as more eddies in that cultural brook (Seaver 2017, 5). That we should approach them as quasi-objects (Lange et al. 2019). As heterogeneous assemblages—as collectives of human and non-human actors (Ananny and Crawford 2018).

1.2 Classification and machine learning as practical work and tinkering

"Once an item of apparatus or a set of gestures is established in the laboratory, it becomes very difficult to effect the retransformation into a sociological object." (Latour and Woolgar 1986, 259)

Much of the critique against the black boxed and hardened conceptualization of algorithms and artificial intelligence has been in dialog with actor-network theory (ANT) (Gillespie 2014; Bucher 2016; Seaver 2017, 2022; Ananny and Crawford 2018; Muniesa 2019; Lee 2021). Importantly, in ANT, the notion of the black box has never been meant to criticize an object which is inscrutable for the analyst, but an object that our interlocutors—not the analyst—treat as a set of inputs and outputs (Latour 1987; 1993).

Accordingly, the proliferation of black boxes in society has been used as a *starting point* for further investigation in ANT—not a lamentable endpoint (cf. Callon and Law 1995). One of the core tenets of actor-network theory has been to see black boxes as an opportunity to scrutinize the work, actors, and processes that make these seemingly self-evident objects appear as a stable and static entities in practice (Latour 1987). For an ethnographer of AI drawing on ANT sensibilities, the act of proclaiming that we live in a black box society should be a call to arms to study these black boxes.

How do we then approach machine learning as an analytical thing in an advanced digital and socio-technical society? Machine learning is made to work through mundane engineering work, and more specifically a combination of work on data, mathematics, statistics, and programming work. At its simplest, a machine learner attempts to use a set of measurable quantities—data—to predict how they affect other measurable quantities—other data. It is sometimes described as finding approximations of functions (Mackenzie 2017, 81–83). That is, a set of mathematical, computational, and statistical techniques that are used to find a (sometimes very complex!) function that can describe a dataset. In essence, machine learners are trained on big sets of data, to find mathematical functions that can describe the datasets, and produce predictions from the sets of data they were given.

In practice, this process is not magical. It is not solely data-driven, solely statistical, solely mathematical, or solely computational—on the contrary, it is very messy and practical and includes all these elements (cf. Latour 2011). This means that there is an abundance of mundane work and tinkering involved in making a machine learning model work. Teaching a system to classify or train it to perform "better than a human" means tinkering with what desired output means (Bechmann and Bowker 2019; Henriksen and Bechmann 2020). To teach a machine about the world the world, it is sometimes bounded to a "ground truth" dataset (Neyland 2018; Jaton 2017; 2021a). And to teach a machine learner how to classify and recommend music demands that an engineer listens to the machine as well as music to tweak it to our human tastes (Seaver 2022).

That is, there is a constant interplay between the machine learner, and its creators: As Seaver describes it: "perception is 'in the loop,' so to speak, and we could argue that 'the neural network' does not end at the computational nodes but extends to include [its creators'] own mind" (Seaver 2022, 114). Indeed, to build an algorithmic system is to constantly tinker to reach some imagined—and sometimes elusive—stabilization. To understand why a machine learning system does what it does, we need to—just as with laboratory practices to construct facts—understand how these systems are constructed in practice (cf. Latour and Woolgar 1986; Latour 1987).

2 Four moments of translation

Below, I bring to the fore four situations—inspired by Callon (1984), one might say four moments of translation, four moments of work, tinkering, and transformation—where actors work to construct and stabilize the machine learner. The four moments focus on what our interlocutors call "feature extraction," "vectorization," "clustering," and "data drift." As the themes are derived from practice, the actual moments of translation might in practice be infinitely different, but these are suggestions for empirical moments that might provide purchase in understanding how the magic of machine learners is put together.[4]

Importantly, I define these four moments in an emic manner: they are based on how the actors talk about, write about, organize, and work with the construction of machine learners. They are not clear-cut temporally, conceptually, or analytically, but are meant to provide inroads to engaging ethnographically and analytically with the construction of machine learners for prediction and classification. As these moments are defined by the actors, they provide both a language to communicate with our interlocutors as well as a direction for ethnographic attention.[5]

[4] These themes focus on the work of constructing machine learners in practice. Other possible moments of translation would be for instance attending to model evaluation (Jaton 2023). There are also whole fields of research that focus on the critical engagement with datasets (cf. Ciston 2023), which are a crucial part of understanding the politics that precedes the machine learner. In this field guide, I focus on the tinkering to make the machine learner—but the datasets are of course crucial for making the machine learning magic happen.

[5] However, they also risk becoming analytical jails that reify our analyses and limit our understanding of the work that is done in practice.

2.1 Feature extraction: translating data into objects with features

The first moment of classification deals with what our interlocutors call *feature engineering*, *feature extraction*, or *feature selection*. This moment of translation might be thought of as the moment when the actors start relating to what they call data that need to be transformed—tinkered with, sliced and diced—for a particular machine learning application, with particular objectives and goals. During this moment, actors work to transform data into sets of objects that have what they call "features." In practice, working with feature selection is often an iterative process oscillating between selecting features and testing the predictive qualities of the model. One thus has to select among the features, which ones are worth looking for and which are not.

This process of tinkering with data to create features involves transforming what the actors sometimes define as an "observable phenomenon's measurable property" into some type of symbolic system that can be digitized and computed ("What Is Feature Vector? Role in ML & Applications," 2024). For example, the height of a person can become a feature in the form of a number representing the height in centimeters, "178," while the gender of a person might be turned into a feature in the form of the word "female."

An object in machine learning can in theory have an infinity of features—often there are thousands of features that can be relevant. For instance, the features of a person-object in machine learning could include various medical records, medical test results, their DNA, their grades, diplomas, and certifications, their online behavior, their tax information, their travel history, address, email, their payment history, loans, what they buy, what kind of car they drive, their insurance coverage, etc. The list can be made infinite. Accordingly, there are potentially as many different techniques for transforming data into objects with features as there are facets of the world.

2.1.1 The curse of dimensionality and the cost of features

In machine learning practice, it is understood that each additional feature of an object adds a dimension to the dataset. Thus, an object with three features is said to have three dimensions. Datasets that have objects with a multitude of features introduces various statistical problems for the engineers of machine learning, and is sometimes referred to as the "curse of dimensionality." The curse of dimensionality is identified as a central challenge in machine learning and statistics more generally. It means that the possible combinations and permutations of correlations are exponentially multiplied in high-dimensional spaces.

The curse of dimensionality is compounded in a world with almost infinite amounts of data, where the possible dimensions seem endless. Some describe it as an "affliction" suffered by predictive systems, resulting from too many features being added to the observed entities ("Curse of Dimensionality," 2024). Objects often become high dimensional in a world of overabundant data. This means that there is an overabundance of features that are ascribed to the objects that are under scrutiny.

The potential overabundance of features about different objects is why the reduction of dimensionality—the tinkering with and choice of what are the essential features of an object—is a crucial practice in statistics and machine learning. First, each dimension adds a cost to the computation. The training of the model becomes slower and more costly as the dimensions increase. The reduction or compression of dimensions hinges on the reduction of the "feature space"—the removal of features and thereby dimensions—to simplify a complex dataset. Second, each dimension might add or subtract to the predictive power of the resulting machine learning model.

Feature selection can involve expert knowledge about which features are understood as being most important for prediction, as well as computational techniques for evaluating which features turn out to be most predictive in a machine learning model. For example, when predicting the optimal move for a chess piece it would perhaps seem to be unnecessary to look for players' blood pressure, while ignoring this information or classifying it as noise or normal in public health monitoring might be a severe mistake (cf. Chandola et al. 2009). There are several ways to choose which features are important for prediction. For instance, by constructing decision trees to rank features.

2.1.2 Local taxonomies for the twenty-first century

How can a social scientist then understand and analyze the practices and politics of feature selection? One way in which we might de-mystify this process as a social object is by thinking through historical practices of classification. For instance, the practices and politics of feature selection can be said to be related to the practices of classification that Foucault (2007) and Desrosières (1998) has written about when analyzing classification in natural history. Desrosières has observed the contrast between Linné's classification system which was based on particular fixed criteria that formed the basis for his classification system, and Buffon's system of classification which was based on comparison and a flexible approach to which traits were relevant for distinguishing species.

> Of all the features available, Linné chose certain among them, and created his classification on the basis of those criteria, excluding the other traits. [...] For Buffon, on the other hand, it seemed implausible that

> the pertinent criteria would always be the same. It was therefore necessary to consider all the available distinctive traits a priori. (Desrosières 1998, 241)

The local classification and comparison of exemplars that Buffon proposed—which compared species based on a locally applicable and flexible approach to classification—is reminiscent of the "feature selection" practices in machine learning. Just as Foucault describes the systematic observation of natural history and taxonomy, machine learning engineers—through the selection of features—produce particular and local boundaries and discontinuities between objects: "To observe, then, is to be content with seeing—with seeing a few things systematically." (Foucault 2007, 146).

Just as Linné and Buffon chose particular visual and quantifiable "features" as the basis for his classification project, machine learning experts choose particular ways of quantifying the world. *The selection of features and dimensions constitute a complex politics of selection—where the selection of features of objects are decided through statistical techniques and particular choices. These techniques and choices define the conditions of possibility for classification, prediction, and decision making—and therefore creates the conditions of possibility for objects that are classified and constructed in machine learning.*

As ethnographers of machine learning, we can observe, and ask our interlocutors about what becomes chosen as "relevant" features in the practices of designing machine learning systems. We can ask about what makes up a "charismatic" feature, the features that are included, because they "should" be predictive. We can also observe which features are seen as "naturally" predictive of particular phenomena (Bowker 2000). Observing practices of feature selection holds the promise of bringing to light some politics of classification and prediction for the machine learning age.

Building on these insights, the social analyst of machine learning can ask several questions about the politics of "features." For instance, we can ask about how features are chosen, engineered, and deemed relevant for prediction. On what grounds are they chosen as relevant? Conversely, we can ask how features are excluded and deemed irrelevant. By paying attention to the practices and politics of feature selection, we might gain insight into how the "art and science" of feature selection fuses preconceived notions of what is seen as characteristic of a phenomenon with the objectivizing mathematics of machine learning prediction.

2.2 Vectorization: from "features" to "vectors" and "tensors"

Machine learning is a mathematical endeavor. Transforming the world into mathematical, calculable, and quantified entities is therefore another of the central activities of making machine learning. One of the most important techniques for constructing machine learners is vector mathematics, which is one of the dominant ways in which machine learners are constructed today. Vector mathematics is understood as an easy and intuitive way to represent data for programmers. Vectors are also easy to apply basic mathematical operations on. It is possible to add, subtract, divide, and multiply vectors. In addition, there exists specialized and cost-effective hardware—Graphical Processing Units—that can make vector calculations quickly and efficiently.

The application of vector-based machine-learning techniques is based on making several transformations, from "raw data" into "features" and then from "features" into mathematical objects with names, such as "scalars," "vectors," and "tensors." This means that an object's "features" needs to be transformed into numerical form (numbers are often called scalars in machine learning contexts) to be useful for the machine learning model. "Whatever data you need to process—sound, image or text you must turn them into integers [i.e., whole numbers] […] which is called data vectorization" (Patidar 2019). Vector-based machine learning thus depends on quantifying the world—making any object in the world into an array of numbers that can be treated with mathematical, statistical, and computational techniques.

When using a vector-based machine learning model, an object's "features" are "extracted" from "raw data" and transformed into scalars. Several scalars can then be combined into mathematical objects with several dimensions that are called "vectors" or "tensors."[6] A vector is a computational and mathematical concept that is often used in machine learning to describe how an object can have several features. For instance, a person-object might be described with several vectors containing weight, height, and age.[7] Several person-objects would form a matrix with each person-object contained in a column with vectors, and each new column describing a new person-object with specific vectors. Additional dimensions can be added to the vector space forming "tensors" that can have several dimensions.

2.2.1 Seeing the world as a vector space

The transformation of the world into numbers involves choices of what to quantify—which "features" to quantify—as well as *how* to quantify them—how to transform the world into numbers, scalars (real numbers), vectors, matrixes, and tensors. These transformations of the world into scalars,

[6] A tensor is a concept that describes the relationships between different scalars or vectors.

[7] Vectors and tensors are mathematical objects that combine several scalars for which there exists efficient (computationally inexpensive) mathematical techniques that—by way of their efficiency—are good for machine learning purposes.

vectors, and vector spaces also come to embody and *spread specific ways of seeing and understanding the world*.

Through the transformation of "features" into "scalars," the properties of features become translated into a mathematical world of continuous numbers that can be translated to coordinates in a multidimensional Cartesian coordinate system. The multidimensionality of this coordinate system—each quantified feature adding another dimension to the coordinate system. Through vectorization, the world becomes translated into vector spaces. Thus, one of the things that machine learning does is propagate particular ways of handling, dissecting, and quantifying the world.

2.2.2 Seeing the world as a multidimensional Cartesian space

When using a vector-based model, songs, language, DNA, and medical diagnoses all need to be transformed into coordinates in a space (cf. Seaver 2021). For machine-learning modelers, vectorization entails a particular way of thinking about the world. For instance, in a manual about applied text analysis, the authors propose a shift in how one thinks about language:

> For this reason, we must now make a critical shift in how we think about language—from a sequence of words to points that occupy a high-dimensional semantic space. Points in space can be close together or far apart, tightly clustered or evenly distributed. Semantic space is therefore mapped in such a way where documents with similar meanings are closer together and those that are different are farther apart. By encoding similarity as distance, we can begin to derive the primary components of documents and draw decision boundaries in our semantic space. (Bengfort et al. 2018, chap. 4)

The practice of vectorization thus not only propagates a practice of quantifying features to make them amenable for computation, but also spreads particular geometrical and spatial ways of thinking about the objects that the machine learners are set to analyze.

The practices and politics of vectorization are possible to study ethnographically. They entail making technical and epistemological choices in how to translate the world into scalars and vectors. There are numerous techniques for transforming the world into vectors, and they involve choices in how data become translated into suitable formats for machine learners. These practices of vectorization also involve a politics of quantification—how is the world quantified, by whom, based on what criteria?

2.2.3 Data that do not fit a vector space: the challenge of nominal data

For instance, one of the recurring challenges in machine learning is to transform categorical and nominal data—data that do not have an ordered or hierarchical relationship like numbers do—into numbers that do have an ordinal relationship—data where objects can naturally be ordered. "There is no mapping from categorical to numerical values that is semantically meaningful" (Andritsos and Tsaparas 2010, 154). Much things in our world are nominal and do not have an inherent or apparent order, such as medical diagnostic codes (ICD codes), types of rocks, ethnicity, or names of people.

Representing nominal data in terms of vectors means translating for instance ethnicity or gender—which has no inherent ordinal relationship—into a scalar—where the numbers have an inherent ordinal relationship. Representing nominal values in the form of a vector creates ordered relationships between objects that might not exist. What is the ordinal relationship between different ethnicities or genders? Translating nominal data into vectors therefore introduces particular mathematical relationships between entities.

A constant struggle in the application of machine learning is how images, text, DNA, or medical codes can be transformed into numbers that machine learners can handle. For instance, the International Classification of Diseases (ICD-10) has around 70,000 codes for different ailments, and is applied by human coders that apply their own standards for classification when they apply these codes (Kaur et al. 2021; cf. also Lee 2022a). The practices of mathematization and vectorization enact a quantified world, reminiscent of Descartes and Leibniz philosophies, where all phenomena are possible to translate into numbers.

2.2.4 The politics and practices of vectorization

However, as sociologists of machine learning, there are numerous places where choices are possible to study—to open the black box of vectorization. As ethnographers of machine learning, we can ask several questions about vectorization building on work in sociology and STS: What dilemmas and difficulties of transforming the world into numbers and spatial thinking do our interlocutors face? What tools and algorithms do they use? How do they transform non-numerical things—such as nominal data—into vectors in a space? What challenges arise? How do actors' work to make features and vectors commensurate (Espeland and Stevens 1998)? What features and objects are included and excluded when the world is made into vectors (Lee 2022b)? Are there objects or features that are more difficult to vectorize, and which therefore are excluded from the machine learning model by the actors? What metadata is included in the vector space (Edwards et al. 2011)?

In essence, vectorization can be approached through the lens of actors' practices to translate the world into "raw data," into "features," and then into "vectors" in high-dimensional spatial representations. Like all translation, there is not complete fidelity. After all, to translate is to betray (Law 1997). Attending to actors work to transform the world into vectors promises to open up their choices, valuations, dilemmas of vectorizing, in short to open up the micro-politics of vectorization.

2.3 Clustering: algorithms for boundary making

A third way in which we can approach ethnographically the practices of constructing machine learners is by attending to actors' work to cluster their data in high-dimensional vector spaces. Cluster analysis is a common statistical technique that groups data points that are understood as being more similar to each other. Clustering has become important in machine learning as it allows an analysis of unlabeled data—data where humans do not tell the machine the characteristics or classification of the data. In a world where the promised data deluge seems to have become a reality the need for unsupervised machine learning—where data are analyzed without human intervention—is becoming more and more common.

Just as in your run-of-the-mill statistical analysis, the process of using cluster analysis in machine learning aims to group data in clusters. A cluster analysis can, for instance, be used to classify data in a binary manner (e.g., "yes" allow the loan, or "no" deny the loan), divide data into several groups, or to find anomalies in the data. For instance, it could be used to identify outliers in a dataset, or to group data into types.

Just as other machine learning techniques, clustering involves an iterative and practical process of choosing the clustering techniques that give actors' their desired results. This iterative and practical process can be followed ethnographically. In this section, we attend to the practices, tools, and techniques of clustering datasets—and in particular, we attend to actors work with algorithms for clustering.

2.3.1 Dividing the world: working with clustering algorithms

Each clustering algorithm divides data in different manners, with different effects on the resulting clusters of objects. There is no inherent or optimal way of clustering data—each dataset and application involved choices about what to cluster on and how. What works in one application, might not work in another setting. Furthermore, different clustering algorithms have different characteristics and assumptions built in.

For instance, clustering algorithms produce clusters that have different characteristics based on their ways of making the boundaries between clusters, some clustering algorithms demand that the number of clusters are specified in advance, and some algorithms are more sensitive to noise and outliers. Importantly, for the ethnographer, using certain clustering algorithms, the actors always make assumptions and choices about what is important in drawing boundaries between data points in a space.

Consequently, the actors' iterative work to cluster data involves a politics of classification which can be studied in practice. As there is no natural or optimal algorithm that works for every setting, the actors need to negotiate what a good clustering algorithm should entail. As Wikipedia describes the process, deciding on a clustering algorithm is fraught with difficulty:

> Evaluation (or "validation") of clustering results is as difficult as the clustering itself. Popular approaches involve "*internal*" evaluation, where the clustering is summarized to a single quality score, "*external*" evaluation, where the clustering is compared to an existing "ground truth" classification, "*manual*" evaluation by a human expert, and "*indirect*" evaluation by evaluating the utility of the clustering in its intended application. ("Cluster Analysis" 2024)

In sum, actors must work to choose, adapt, tinker with clustering. This means valuing where good boundaries between classes should be drawn, as well as to negotiate what the desired outcomes of the cluster analysis are. We can study these practices and politics of clustering ethnographically.

2.3.2 Classification in a world with machine learning

Actors work to classify people, objects, and things through information infrastructures is a classic topic for the social study of classification. The work to construct systems of classification of various things, for instance, species, diseases, or race have been the focus of intense scrutiny. In this moment, I want to highlight not only learning to see "features" but also actors work in practice of constructing systems of classification with machine learners. Each construction of a cluster becomes the basis for classifying the world.

However, there are differences in how the systems of classification are constructed with the help of machine learners compared to the bespoke and manually constructed databases (Bowker 2000), standards (Edwards et al. 2011), and systems of classification (Lee 2022a; Bowker and Star 1999). In making a system of classification with machine learners, the actors' work, tinkering, and practice seem to take on a different tenor. This is a difference in degree, not a difference in kind.

In constructing machine learners the actors' work is not focused on dealing with the fate of individual people, objects, or boundaries, nor is it focused on delving into the correct design of the classification system (Bowker and Star 1999). Rather the actors' focus and work with constructing classifiers with machine learning seem to become focused on:

First, the constant valuation of what a successful clustering algorithm should be, what it should be able to do, and how to measure its failure or success rate (cf. Lee and Helgesson 2020). The valuation of what a successful algorithm is can vary within the same situation. These different evaluations can, for example, take the form of competitions between teams, or in a computer to select most successful algorithm based on some sort of measurement, it can take the form of dividing up the dataset you are working on, so it is successful both on the training data, and data that you set aside for testing, and it can take the form of identifying a "ground truth" dataset that you test against (Jaton 2017).

Second, actors work to handle the problematic cases of classification. Partly by removing problematic things from the data—i.e., data cleaning—and partly by tinkering with the selection and parameters of different clustering algorithms to handle edge cases that are difficult to classify. The crucial point, again, is that the construction of a machine learner is not a black boxed process that happens magically without intervention of domain experts, programmers, or data stewards. The selection of a clustering algorithm is a situated process of tinkering and work by actors.

2.3.3 Tinkering with the k-means clustering algorithm

As I have pointed out above, each clustering algorithm brings with it specific challenges. One of the most well-known and well-used clustering algorithms is k-means clustering. This clustering algorithm works by trying to minimizing the total distance between any number of clusters to the datapoints. It works to "partition the observations into K clusters such that the total within-cluster variation, summed over all K clusters, is as small as possible" ("RPubs—K-Means Clustering Tutorial," 2023). There are three central choices that the actors must make to make k-means clustering work: first, the number of clusters that they want to divide the dataset into, second, where the search for cluster centers should start, and third how cluster variation should be defined mathematically—common choice being squared Euclidian distance ("RPubs—K-Means Clustering Tutorial," 2023).

K-means clustering is known to have several challenges. For instance, an online course on machine learning identifies five challenges in using the k-means algorithm for machine learning ("K-Means Advantages and Disadvantages" 2022): the first challenge identified is that the actors must initially choose how many clusters are present in the dataset. This might be easy for certain datasets where it is easy to pre-identify clusters. In large and complex datasets with a lot of different "features", it might be very difficult, if not impossible, to identify how many clusters the dataset has. A second challenge is that k-means clustering is dependent on the initial positions given to the "centroids" of each cluster, which then spawns new technologies and techniques for "initializing" the algorithm (Celebi et al. 2013). Furthermore, "k-means has trouble clustering data where clusters are of varying sizes and density." All these challenges point to areas where actors need to make choices, tinker, and work with the algorithm.

However, the two last challenges identified in the course are perhaps most relevant in relation to this field guide, as they connect to classic questions and themes in science and technology studies.

One is that outliers in the data are problematic in k-means clustering, as "Centroids can be dragged by outliers, or outliers might get their own cluster instead of being ignored" ("K-Means Advantages and Disadvantages" 2022). The proposed solution in a google machine learning course is to remove outliers before the clustering algorithm is applied. That is, the solution is for actors to remove the things that do not fit into the clusters. The actors' handling of that which does not fit into the classification systems is of course an important and classic topic in the study of the politics of classification systems: how are the things that do not fit into the grids of classification handled? Are they assigned to the "Other" category (cf. Bowker and Star 1999)? Are they classified as abnormal or pathological (Canguilhem 1966; Foucault 2007)? Here, we as ethnographers of machine learning can find practices and work that can potentially have huge political consequences for how the world is put together with machine learning. What people, things, or phenomena become treated "outliers" and removed from the datasets? How are the decisions made to remove outliers made? What metrics (and which features) are used to decide if something is an outlier? Is there a discussion of the consequences? What are the possible consequences of not being part of that dataset?

Another is that the k-means algorithm has a difficult time in handling an increasing number of dimensions of a dataset, as "As the number of dimensions increases, a distance-based similarity measure converges to a constant value between any given examples" ("K-Means Advantages and Disadvantages" 2022; Aggarwal et al. 2001). Thus, the problem is that the number of "features" of your dataset makes it more difficult to discern differences between "objects." This points us again toward the actors' practices of feature selection, and the choice of what is relevant and important aspects of the world. Which of course involves more choices, tinkering, algorithms, and techniques for reducing dimensionality for the algorithm to output the desired results.

As ethnographers of machine learning, we can ask question such as: *How do actors decide what clustering algorithms are desirable, and based on what criteria? On what criteria do the actors assess what a successful clustering is? Which features are possible to cluster with this specific technique?* Machine learning is one of the new tools for classifying and ordering the world, and clustering is one of many techniques where a politics of classification can be studied.

2.4 Data drift/concept drift: data and temporalities

The last moment that I want to highlight here relates to what our actors call "data drift" or "concept drift." That is, the challenge that is posed to machine learners by the fact that the world is constantly changing. The world during one period might look very different to what it looks like during another period. For the actors that design machine learners, the constantly changing state of the world poses design challenges. Machine learners are often trained on a snapshot of data—and training a machine learning model is often the most expensive part of constructing a machine learner. Frequently particular snapshot of the world as it is encoded in a model entails as much data as the actors can possible get their hands on within their projects' constraints. However, if there is no sense of the temporality of the world, and the data that are collected, the machine learner might miss important patterns in an evolving dataset. This points our ethnographic sensibilities toward the practices of handling "data drift" and "concept drift."

Machine learners exist in a changing world, where classification is in constant flux. This poses another challenge for machine learning. Another challenge in handling the rhythms and dynamics of data. This challenge is called "data drift" or "concept drift" in the parlance of machine learning, and points to the fact that phenomena in the world change over time. The changing nature of the world leads to challenges in designing machine learners and machine learning processes that can adapt temporally to the changing face of anomalies, abnormalities, and normalities. In one article, the problem is described in the following manner:

> In an operational setting we face the additional difficulty that human behaviors are dynamic and what is considered "normal" behavior is likely to change over time. This problem is known in the machine learning community as concept drift, and requires a system that can dynamically update the user profile and classification parameters (Lane and Brodley 1999)

The machine learner needs continuous work and updating to stay relevant. Thus, the machine learning system needs to be in constant flux and re-learning to handle how the world and its data changes. A concrete example is "fingerprint drift" where our fingerprints subtly change over time, so that the system needs to continuously re-learn what our fingerprints look like.

2.4.1 Shifting boundaries of abnormality

Some facets of classification in machine learning deal with the temporalities of anomalies. That is, in some machine learning applications, temporal changes in data are the basis for identifying something as anomalous. That is, "normal" and "abnormal" data change over time for instance in nature there are constant cyclical changes due to seasons, while on another time scale, changes in temperature might point to climate change. Here, the actors have assumptions about what time cycles and scales are relevant, and they need to work with and attempt to identify temporal dynamics in the data to detect anomalies.

For instance, if there is a sudden change in the characteristics of a data stream, this could be treated as a signal that there is something anomalous happening. In other applications, it is acknowledged that the temporalities of data can also include "concept drift." That is, the boundaries between the normal shift over time in a particular data stream. Rather than treating the world of data as something static, it is treated as a stream of ever changing rhythms and dynamics that need to be handled by the machine learner.

For instance, in a security application that attempts to identify anomalous network traffic, changes in the character of the data stream over time are the basis for drawing boundaries between anomalous and normal data. The definition hinges on defining the anomalous as "abrupt changes" between data sampling intervals. In the quote below, the boundaries between abrupt change hinges on the sampling interval in the data stream. Thus, the anomalous becomes performed as abrupt temporal change in the analyzed data stream:

> In statistical analysis, a network anomaly is modeled as correlated abrupt changes in network data. An abrupt change is defined as any change in the parameters of a time series that occurs on the order of the sampling period of the measurement. For example, when the sampling period is 15 s, an abrupt change is defined as a change that occurs in the period of approximately 15 s. (Thottan and Ji 2003)

In this example, anomalous data are identified through the rhythms and periods of the machine learner sampling intervals and the dynamic changes of the data stream. The boundary between the anomalous and the normal being dependent on the temporal configuration of the machine learner. Temporal changes in dynamics come to define the anomalous and normal. In this example, the power and politics of machine learning depend on configuring and erecting temporal boundaries through sampling intervals.

2.4.2 The temporalities of data

The temporalities of data are crucial for knowledge production. There are many temporal complexities in the construction of datasets, and there are tensions between different temporalities and ways of thinking about how time moves when a dataset is constructed. How temporality is understood is encoded in databases and metadata, and it is reflected in how phenomena are analyzed and understood. For instance, in analyzing the presence or absence of species in habitats, it is sometimes assumed that data are cyclical, reflecting the yearly seasonality of ecologies (Lee 2021).

As Bowker (2005, 190-) has pointed out, thinking about how time moves can involve thinking about time as *contingent* on grand historical events, such as meteor strikes or floods, it can involve thinking about time as *cyclical* commonly in climatology, or it can involve *secular* thinking, as change that occur over very long periods of time. It can also involve thinking about *stasis* and *equilibrium* versus a constant state of change. In a dataset—in this case a biodiversity dataset—Bowker argues that:

> As far as databases go is that there is no uniform way of separating off the data objects (which themselves enfold complex histories) from their spatial and temporal packaging […]. As you nest cycles one inside the other, you find secular change irrupting into the story; as you nest secular narratives, cycles emerge. […] The work of flattening out all the narrative sciences into a single narrative timeline is a productive effort that articulates data formats with relative power relationships between disciplines—through the mediation of classification systems and data standards. The manipulable second nature created within the computer is structured by an organizationally, politically, and morally inflected set of temporalities. (Bowker 2005, 192–93)

The multiple temporalities that Bowker points out about biodiversity data have many resonances with machine learning actors' struggles to practically handle the challenge that data drift and concept drift pose for their construction of prediction and classification. Just as the biodiversity databases that Bowker deals with, the datasets that machine learners are trained on are treated by the actors as encoding certain temporalities. For instance, facial recognition systems assume stability and stasis in time—while variation of ethnicity is hotly debated (Buolamwini and Gebru 2018)—and fingerprint systems are commonly constructed to handle "drift."

The temporalities of machine learners are also part of the politics of machine learners. In actors' work and tinkering with machine learners, we can expect there to exist a struggle with handling a variety of times and periodizations. And that there could exist struggles to handle these various temporalities of data. The actors' work to handle the temporalities of the world and the datasets that they have collected about the world constitutes choices about what can be discerned by the machine learner. The assumptions about the temporality of datasets—contingent, cyclical, secular, in flux or in stasis—constitutes another facet of the conditions of possibility for the predictive politics of learning machines.

In approaching machine learners ethnographically, we can start to unpack the practices and politics of temporality in machine learning systems. We can ask questions about what temporalities are already encoded in the datasets that are used to train the learning machines. We can ask questions about how the actors struggle to mangle data into various temporalities to produce the predictions that are desired. We can ask about how the actors reason about temporality, and about how they value the effect of different temporalities in relation to their assumptions about what the predictive capacities of the machine learner should be handling.

By paying attention to the practices and politics of temporality—that the data are for instance stable or in stasis—we might gain insight into how the actors perceive the world—what is seen as the natural temporality of a phenomenon, and what politics does that temporality bring with it?

3 Conclusion

*Most important of all, what values and ethical principles do we inscribe in the inner depths of the built information environment (*Star 1999*, 379).*

Today, learning machines are becoming ubiquitous. Machine learning models are implemented an increasing number of computer systems, software, and digital devices. From facial recognition and loan applications to network security and chatbots, machine learning is used to classify, predict, and make decisions. The power of learning machines in society seems to be steadily on the rise. However, contrary to the feeling of wonder and magic that often surrounds machine learning, the politics and practices of learning machines are mundane. There is a huge amount of work and tinkering by the people that make machine learners.

In this state of being, it is becoming increasingly crucial to understand these machines beyond the transparency/opacity dichotomy, beyond the auditing of fairness and bias of algorithms, and beyond legal debates about algorithmic accountability. These perspectives are hugely valuable in highlighting the visible politics of the machine learning society. However, they tend to simplify the politics of algorithmic assemblages into questions of how to change the algorithm. The politics of algorithms risks becoming reduced to striving for unbiased and transparent computation (cf. Lee et al. 2023). In these perspectives—highlighting the fairness of the algorithm—agency seems to become punctualized to the machine learner and thus risks hiding the actors work and tinkering (cf. Callon and Law 1995). This in turn risks reproducing the machine learner as a Deus ex machina—a magical artifact where the politics of choice resides.

However, by pushing our ethnographic sensibilities to attend to the construction of machine learners in practice, we might get a glimpse of a more complex politics of machine learners, that goes beyond auditing, computational ethics, and fixing the algorithmic bias. By attending to actors' mundane practices of constructing machine learners, we can start delineating and understanding how the politics of these heterogeneous assemblages of machine learning is put together. As ethnographers, we can attend to actors' work to classify and value the world through sweat, negotiations, tinkering, failures, and successes. We can start to understand how the actors' decide in practice what a good machine learner should do (cf. Ziewitz 2011). Thus, the magic of machine learning evaporates in a mist of mundaneness, to highlight choice, valuations, and a lot of tinkering and work.

The crucial methodological and analytical move here is to move beyond the dichotomy where the machine learner is treated separate from the work of actors. This constructivist perspective on technology draws on a decades long history. Just because our tools are shiny and new—and seem wondrous and magical—does not mean that they are magic. The insight that all technology is sociotechnical—inseparable from society—is by now decades old (Bijker et al. 1987; Bijker and Law 1992).

A second analytical move—that this ethnographic sensibility enables—is to attend to computation, calculation, and valuation as a practice, rather than treating them as separate from actors work. Just because our computers are no longer human nor huge mechanical machines, and just because our calculative and statistical practices have been combined into new assemblages of computation does not mean that they are not possible to pick apart with the analytical sensibilities of science and technology studies. The insight that calculation and quantification are political and closely intertwined with power are insights that have been developed in the social study of statistics and calculation (Gigerenzer et al. 1989; Hacking 1990; Porter 1995; Desrosières 1998).[8]

In studying the practices and politics of machine learners, we need to learn to reach forward toward a sometimes mystified technological present, while we embrace our analytical past where we have for decades turned our analysis toward the construction of facts, infrastructures, and technologies. We must re-learn to pay ethnographic attention to the magical tools of technology, to make them mundane again. We must train our analytical attention on the actors' work and tinkering. Through actors' practices, we will be able to understand some facets of the politics of machine learning.

The intent with this article has been to construct a methodological and analytical field guide—a starting point—for ethnographers studying the construction of machine learning. This field guide is meant to direct our analytical gaze toward moments where practices and politics come to the fore, both in an empirical sense but also in an strategically analytical sense—with a concern for the politics of machine learners in society. Doing this, I wanted to draw out questions and issues that are emic concerns, concerns that our interlocutors deal with, but also etic concerns, concerns that we as sociologists of machine learning have about a society of machine learning. I aimed to make the opacity and magic of machine learning mundane, to lay the ground for asking questions about the practical politics of machine learning.

I wanted to open the door to ask questions about actors' assembling of the politics of machine learning: how does the actors' work, tinkering, and negotiations affect the construction of machine learners? How do actors include things, people, objects, and phenomena in machine learners? How do the actors construct datapoints, dimensions, and features to be included and excluded in their making of machine learners? How do they translate the world into mathematics, and with what effects? How do the actors make decisions about how to group and classify the world into different categories? How do they think and enact temporalities? How do the actors divide the world into different timescales, and how do they think about and handle the nature of stasis and change?

By attending to these questions, we can start grappling with the politics of machine learning in earnest. How does the work of the actors shape how machine learners come to

[8] A third methodological and analytical possibility—which I have not highlighted in this particular account, as I have wished to demystify the construction of machine learners rather than data and data work—is of course to attend ethnographically to data practices. Data studies also have a decades long history in science and technology studies, attending to the mundane work of constructing databases and datasets (Hine 1995; Bowker 2000; Hine 2006; Fujimura and Fortun 2006). During the past decades, critical data studies exploded as more and more knowledge production is premised on huge datasets and databases (Leonelli 2009; Edwards et al. 2011; Boellstorff and Maurer 2015; Douglas-Jones et al. 2021; Beaulieu and Leonelli 2022; Thylstrup et al. 2022; Ciston 2023).

predict, classify, and enact the world? What becomes enacted as problematic? What is included? What is excluded? How are things, people, objects, and phenomena translated into the world of machine learners? And with what effects?

How can we then study the politics and practices of machine learning? Above I have suggested that we need to attend to the practices of constructing machine learners. I have also suggested that there are many insights, questions, and research problems in the study of science and technology that can be applied to the ethnographic study of machine learning in practice. I have suggested that we can pay attention to four moments of translation where machine learners are constructed.

There are of course a myriad of other possible moments of translation highlight the politics of constructing machine learners. In developing this particular set of analytical sensibilities, the ethnographer of machine learning can both learn from classic questions in science and technology studies, but of course from the work in critical studies of algorithms or data.

A critical lesson for the ethnographer of machine learning is to make the magic of machine learning mundane by attending to work, tinkering, negotiations, and valuations. By attending to the magic of machine learning as a practical concern, we can de-elevate the politics of machine learning for social analysis. The social facts of today are not made by magic. They are made in particular situations, by particular actors, with particular ideas about the world. Let us go forth and de-mystify the magic of machine learning, to bring political agency back where it belongs.

Author contributions The article was solely researched and written by the author.

Funding Open access funding provided by Södertörn University. This article is funded by Wallenberg AI, Autonomous Systems and Software Program—Humanities and Society (WASP-HS), 2019. 0235, Marianne och Marcus Wallenbergs Stiftelse, 2019.0235.

Data availability No datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors declare no competing interests.

References

Aggarwal CC, Hinneburg A, Keim DA (2001) On the surprising behavior of distance metrics in high dimensional space. In Database Theory — ICDT 2001, edited by Jan Van den Bussche and Victor Vianu, 1973:420–34. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg. https://doi.org/10.1007/3-540-44503-X_27

Ananny M, Crawford K (2018) Seeing without knowing: limitations of the transparency ideal and its application to algorithmic accountability. New Media Soc 20(3):973–989. https://doi.org/10.1177/1461444816676645

Andritsos P, Tsaparas P (2010) Categorical data clustering. In: Sammut C, Geoffrey I (eds) Encyclopedia of machine learning. Springer, Boston, pp 154–59

Appadurai Arjun (2003) Disjuncture and difference in the global cultural economy. In: Braziel JE, Mannur A (eds) Theorizing diaspora: a reader. Blackwell, Malden, pp 25–48

Beaulieu A, Leonelli S (2022) Data and society: a critical introduction, 1st edn. SAGE Publications Ltd., Thousand Oaks

Bechmann A, Bowker GC (2019) Unsupervised by any other name: hidden layers of knowledge production in artificial intelligence on social media. Big Data Soc 6(1):2053951718819569. https://doi.org/10.1177/2053951718819569

Bengfort B, Bilbro R, Ojeda T (2018) Applied text analysis with python: enabling language-aware data products with machine learning. 1st edition. Sebastopol, CA: O'Reilly Media. https://www.oreilly.com/library/view/applied-text-analysis/9781491963036/ch04.html. Accessed 15 Aug 2024

Bijker WB (1995) On bicycles, bakelites and bulbs. MIT Press, Cambridge, MA

Bijker WE, Law J (eds) (1992) Shaping technology/building society: studies in sociotechnical change. Inside technology. MIT Press, Cambridge, Mass

Bijker WE, Hughes TP, Pinch T (eds) (1987) The social construction of technological systems: new directions in the sociology and history of technology. MIT Press, Cambridge, Mass

Boellstorff T, Maurer B (2015) Data, now bigger and better! Prickly Paradigm Press, Chicago, IL

Bowker GC (2000) Biodiversity datadiversity. Soc Stud Sci 30(5):643–683. https://doi.org/10.1177/030631200030005001

Bowker GC (2005) Memory practices in the sciences. Inside technology. MIT Press, Cambridge

Bowker GC, Star SL (1999) Sorting things out: classification and its consequences. MIT Press, Cambridge, MA

Bucher T (2016) Neither black nor box: ways of knowing algorithms. In: Kubitschko S, Kaun A (eds) Innovative methods in media and communication research. Springer International Publishing, Cham, pp 81–98. https://doi.org/10.1007/978-3-319-40700-5_5

Buolamwini J, Gebru T (2018) Gender shades: intersectional accuracy disparities in commercial gender classification. In Conference on Fairness, Accountability and Transparency, 77–91. PMLR. http://proceedings.mlr.press/v81/buolamwini18a.html. Accessed 19 Apr 2021

Burrell J (2009) The field site as a network: a strategy for locating ethnographic research. Field Methods 21(2):181–199. https://doi.org/10.1177/1525822X08329699

Burrell J (2016) How the machine 'thinks': understanding opacity in machine learning algorithms. Big Data Soc. https://doi.org/10.1177/2053951715622512

Callon M (1984) Some elements of a sociology of translation: domestication of the scallops and the fishermen of St Brieuc Bay. Sociol Rev 32(1):196–233. https://doi.org/10.1111/j.1467-954x.1984.tb00113.x

Callon M, Law J (1995) Agency and the hybrid collectif. South Atlantic Q 94(2):481–507

Campolo A, Crawford K (2020) Enchanted determinism: power without responsibility in artificial intelligence. Engag Sci Technol Soc 6(8):1–19. https://doi.org/10.17351/ests2020.277

Canguilhem G (1966) The normal and the pathological. Zone Books, New York

Celebi ME, Kingravi HA, Vela PA (2013) A comparative study of efficient initialization methods for the K-Means clustering algorithm. Expert Syst Appl 40(1):200–210. https://doi.org/10.1016/j.eswa.2012.07.021

Chandola V, Banerjee A, Kumar V (2009) Anomaly detection: a survey. ACM Comput Surv 41(3):1–58. https://doi.org/10.1145/1541880.1541882

Ciston S (2023) A critical field guide for working with machine learning datasets. In edited by Mike Ananny and Kate Crawford. https://knowingmachines.org/critical-field-guide. Accessed 8 Feb 2023

Cluster Analysis (2024) In Wikipedia. https://en.wikipedia.org/w/index.php?title=Cluster_analysis&oldid=1238821333. Accessed 15 Aug 2024

Curse of Dimensionality (2024) Statistics.Com: Data Science, Analytics & Statistics Courses. https://www.statistics.com/glossary/curse-of-dimensionality/. Accessed 15 Aug 2024

de Laet M, Mol A (2000) The Zimbabwe bush pump: mechanics of a fluid technology. Soc Stud Sci 30(2):225–263. https://doi.org/10.1177/030631200030002002

Desrosières A (1998) The politics of large numbers: a history of statistical reasoning. Harvard University Press, Cambridge

Douglas-Jones R, Walford A, Seaver N (2021) Introduction: towards an anthropology of data. J Roy Anthropol Inst 27(S1):9–25. https://doi.org/10.1111/1467-9655.13477

Edwards PN, Mayernik MS, Batcheller AL, Bowker GC, Borgman CL (2011) Science friction: data, metadata, and collaboration. Soc Stud Sci 41(5):667–690. https://doi.org/10.1177/0306312711413314

Espeland WN, Stevens ML (1998) Commensuration as a social process. Ann Rev Sociol 24(1):313–343. https://doi.org/10.1146/annurev.soc.24.1.313

Ferguson J (1999) Expectations of modernity: myths and meanings of urban life on the Zambian Copperbelt

Foucault M (2007) The order of things: an archaeology of the human sciences. Routledge, London

Fujimura JH, Fortun M (2006) Constructing knowledge across social worlds: the case of DNA sequence databases in molecular biology. Naked Science: Anthropological Inquiry into Boundaries, Power, and Knowledge 160–73

Geiger RS (2017) Beyond opening up the black box: investigating the role of algorithmic systems in wikipedian organizational culture. Big Data Soc 4(2):1–14. https://doi.org/10.1177/2053951717730735

Gigerenzer G, Swijtink Z, Porter T, Daston L, Beatty J, Krüger L (1989) The empire of chance: how probability changed science and everyday life. Cambridge University Press, Cambridge

Gillespie T (2014) The relevance of algorithms. In: Gillespie Tarleton, Boczkowski PJ, Foot KA (eds) Media technologies: essays on communication, materiality, and society. The MIT Press, Cambridge. MA, pp 167–94. https://doi.org/10.7551/mitpress/9780262525374.003.0009

Gusterson H (1997) Studying up revisited. PoLAR Polit Legal Anthropol Rev 20(1):114–119

Hacking Ian (1990) The taming of chance. Ideas in context. Cambridge University Press, Cambridge

Haraway D (1988) Situated knowledges: the science question in feminism and the privilege of partial perspective. Fem Stud 14(3):575. https://doi.org/10.2307/3178066

Henriksen A, Bechmann A (2020) Building truths in AI: making predictive algorithms doable in healthcare. Inf Commun Soc 23(6):802–816. https://doi.org/10.1080/1369118X.2020.1751866

Hine C (1995) Representations of information technology in disciplinary development: disappearing plants and invisible networks. Sci Technol Human Values 20(1):65

Hine C (2006) Databases as scientific instruments and their role in the ordering of scientific work. Soc Stud Sci 36(2):269

Jaton F (2017) We get the algorithms of our ground truths: designing referential databases in digital image processing. Soc Stud Sci 47(6):811–840. https://doi.org/10.1177/0306312717730428

Jaton F (2021a) Assessing biases, relaxing moralism: on ground-truthing practices in machine learning design and application. Big Data Soc 8(1):20539517211013570. https://doi.org/10.1177/20539517211013569

Jaton F (2021b) The constitution of algorithms: ground-truthing, programming, formulating. The MIT Press, Cambridge. https://doi.org/10.7551/mitpress/12517.001.0001

Jaton F (2023) Groundwork for AI: enforcing a benchmark for neoantigen prediction in personalized cancer immunotherapy. Soc Stud Sci 53(5):787–810. https://doi.org/10.1177/03063127231192857

Kaur R, Anupama Ginige J, Obst O (2021) A systematic literature review of automated ICD coding and classification systems using discharge summaries. arXiv. http://arxiv.org/abs/2107.10652. Accessed 1 Dec 2022

K-Means Advantages and Disadvantages (2022) Advanced Courses: Clustering. July 18, 2022. https://developers.google.com/machine-learning/clustering/algorithm/advantages-disadvantages. Accessed 20 Feb 2023

Lane T, Brodley CE (1999) Temporal sequence learning and data reduction for anomaly detection. ACM Trans Inf Syst Secur 2(3):295–331. https://doi.org/10.1145/322510.322526

Lange A-C, Lenglet M, Seyfert R (2019) On studying algorithms ethnographically: making sense of objects of ignorance. Organization 26(4):598–617. https://doi.org/10.1177/1350508418808230

Latour B (1987) Science in action. Harvard UP, Cambridge, MA

Latour B (1993) We have never been modern. Harvard University Press, Cambridge, Mass

Latour B (2010) Tarde's idea of quantification. In: Candea M (ed) The social after gabriel tarde: debates and assessments. Routledge, Milton Park, pp 145–162

Latour B (2011) Sur La Pratique Des Théoriciens. In *Savoirs Théoriques et Savoirs d'action*, edited by Jean-Marie Barbier, 131–46. Paris, France: Presses Universitaires de France. https://doi.org/10.3917/puf.barbi.2011.01.0131

Latour B, Woolgar S (1986) Laboratory life: the construction of scientific facts, 2nd edn. Princeton University Press, Princeton, NJ

Law J (1997) Traduction-Trahison: Notes on ANT. 1997. http://www.comp.lancs.ac.uk/sociology/papers/Law-Traduction-Trahison.pdf Accessed 30 Nov 2003

Lee F (2021) Enacting the pandemic: analyzing agency, opacity, and power in algorithmic assemblages. Sci Technol Stud 34(1):65–90. https://doi.org/10.23987/sts.75323

Lee F (2022a) Detecting the unknown in a sea of knowns: health surveillance, knowledge infrastructures, and the quest for classification egress. Preprint. SocArXiv. https://doi.org/10.31235/osf.io/b692n

Lee F (2022b) Ontological overflows and the politics of absence: zika, disease surveillance, and mosquitos at the ECDC. Preprint. SocArXiv. https://doi.org/10.31235/osf.io/e5ju9

Lee F, Helgesson C-F (2020) Styles of valuation: algorithms and agency in high-throughput bioscience. Sci Technol Human Values 45(4):659–685. https://doi.org/10.1177/0162243919866898

Lee F, Landström C, Palmås K (2023) Caring for the monstrous algorithm: attending to wrinkly worlds and relationalities in an algorithmic society. SocArXiv. https://doi.org/10.31235/osf.io/rg7k6

Leonelli S (2009) On the locality of data and claims about phenomena. Philos Sci. https://doi.org/10.1086/605804

Mackenzie A (2017) Machine learners: archaeology of a data practice. The MIT Press, Cambridge, MA

MacKenzie D (2019) How algorithms interact: Goffman's 'interaction order' in automated trading. Theory Cult Soc 36(2):39–59. https://doi.org/10.1177/0263276419829541

Marcus GE (1995) Ethnography in/of the world system: the emergence of multi-sited ethnography. Annu Rev Anthropol 24(1):95–117. https://doi.org/10.1146/annurev.an.24.100195.000523

Martin A, Lynch M (2009) Counting things and people: the practices and politics of counting. Soc Probl 56(2):243–266. https://doi.org/10.1525/sp.2009.56.2.243

Muniesa F (2019) Societe du comportement, information de la sociologie. Zilzel 1(5):196–207

Neyland D (2018) The everyday life of an algorithm, 1st edn. Palgrave Pivot, New York, NY

Pasquale F (2016) The Black Box Society: The Secret Algorithms That Control Money and Information. Harvard University Press, Cambridge, Massachusetts London, England

Patidar P (2019) Data preprocessing & feature engineering in machine learning. Medium. December 20, 2019. https://medium.datadriveninvestor.com/data-preprocessing-feature-engineering-in-machine-learning-530ccc38d01a. Accessed 15 Aug 2024

Porter TM (1995) Trust in Numbers: The Pursuit of Objectivity in Science and Public Life. Princeton, Princeton UP

RPubs - K-Means Clustering Tutorial (2023) https://rpubs.com/cyobero/k-means. Accessed 21 Feb 2023

Seaver N (2017) Algorithms as culture: some tactics for the ethnography of algorithmic systems. Big Data Soc 4(2):2053951717738104. https://doi.org/10.1177/2053951717738104

Seaver N (2021) Everything lies in a space: cultural data and spatial reality. J Roy Anthropol Inst 27(S1):43–61. https://doi.org/10.1111/1467-9655.13479

Seaver N (2022) Computing taste: algorithms and the makers of music recommendation. University of Chicago Press, Chicago

Star SL (1999) The ethnography of infrastructure. Am Behav Sci 43(3):377–391. https://doi.org/10.1177/00027649921955326

Suchman L (2023) The uncontroversial 'thingness' of AI. Big Data Soc 10(2):20539517231206790. https://doi.org/10.1177/20539517231206794

Thottan M, Ji C (2003) Anomaly detection in IP networks. IEEE Trans Signal Process 51(8):2191–2204. https://doi.org/10.1109/TSP.2003.814797

Thylstrup NB, Hansen KB, Flyverbom M, Amoore L (2022) Politics of data reuse in machine learning systems: theorizing reuse entanglements. Big Data Soc. https://doi.org/10.1177/20539517221139785

What Is Feature Vector? Role in ML & Applications (2024) Deepchecks. https://deepchecks.com/glossary/feature-vector/ Accessed 15 Aug 2024

Ziewitz M (2011) "How to Think about an Algorithm: Notes from a Not Quite Random Walk." Presented at the Symposium "Knowledge Machines between Freedom and Control," Kulturfabrik Hainburg, Austria, September 29

Ziewitz M (2016) Governing algorithms: myth, mess, and methods. Sci Technol Human Values 41(1):3–16. https://doi.org/10.1177/0162243915608948