Technology Ethics and Practical Deliberation

by Ken Archer
archerovi@hotmail.com

1. Introduction

This paper is motivated by two decades of work in technology firms, observing the wide gulf that exists between the normative appeals of technology ethicists and the daily decisions made by technology practitioners. At conferences on technology ethics one hears the normative claims that engineers should do certain things – they should make algorithms transparent; they should nudge users toward more ethical ways of life through product design. When asked how to account for the fact that engineers do not do these things, even after learning that they should, the response is generally that this is a question for psychology, or for business ethics, or for some other field. AI ethics, no less than previous domains of technology ethics such as bioethics, struggles to provide a philosophical account of the relevant decisions that practitioners, faced with daily decisions to be made, find applicable.

Motivated to answer this question, to account for decisions that engineers make day-to-day, I have found such decisions to be largely mischaracterized in philosophy of technology. Most projects in technology ethics seek to teach a set of principles that practitioners and designers are expected to apply. The message, sometimes explicit, of such projects is that, whereas artifacts may not be morally neutral, those who build them *are* applying instrumental, calculative reasoning in their work, and must be taught ethics. This paper argues that technical rationality is never inherently calculative but is always embedded in practices with various internal and external goods. The critical role of technology ethics, it seems to me, is to unpack and account for the full ethical content of the decisions faced by the designer and developer, rather than instrumentalize such decisions. Such an account would restore a sense of ethical agency that the notion of technology as applied science espoused by engineering educators conceals from engineers.

This paper attempts to develop this account of technical decision-making as embedded in practices, particularly with reference to the most technical of activities, the development of statistical models within AI and other fields.

2. Phenomenology of Technical Decision-Making

Technological artifacts – tools, machines – provide structure to practices. Sometimes the material novelty of an artifact can conceal the activity which the artifact is intended to enable or to enforce. The prominence of the artifact can create the appearance that new practices are co-invented along with artifacts, which blurs the distinction that exists in reality.[1]

[1] Gadamer's argument that art gives structure to a way of life, then, is true of artifacts as well. "I call this change, in which human play comes to its true consummation in being art, *transformation into structure*." Hans-Georg Gadamer, *Truth and Method*, (New York: Continuum International Publishing, 2004), p. 110. For an architect, "His plan is determined by the fact that the building has to serve a particular way of life". ibid, p. 156. Art, of course, was originally a techne. It's precisely when art was split into fine art of genius vs mere craftsmanship that techne lost its grounding in bringing ways of life into being and was reduced to instrumental reasoning.

A simple example that is paradigmatic of modern technology is the assembly line. A misconception about the rise of factory work from home-based craft work is that this transition was driven fundamentally by technical innovations, such as the division of labor enabled by the assembly line. In fact, the assembly line doesn't enable the division of labor, it enforces the pace of work in each task. The division of tasks – conducting one task across several units and then conducting a subsequent task across those same units – for the sake of efficiency emerged at an earlier time within artisan practices such as textile weaving.[2] Once workers were organized into a single place of work, the new breed of owners realized the opportunity to increase output by dividing tasks between workers, and by enforcing the pace of work with assembly lines. The tasks didn't change, they were abstracted as processes from the practice within which they emerged, codified as technical knowledge, and then structured through artifacts.[3]

It is not just industrial mechanization that provides structure to processes that emerge first within a practice. Word processing software structures the practice of writing, home appliances structure homemaking practices such as cooking. In all cases the practice comes first, and the artifact structures or restructures a series of tasks, a process, that emerged within the practice and from which it derives meaning in the world.

Some examples of technology that arise from scientific research removed from day-to-day practice – for example, genetic testing - may seem to resist this distinction. These are the more difficult questions that this paper addresses. For now, we can simply ask why genetic testing is meaningful, why it is done at all?

It is impossible to answer without reference to a practice from which it derives its meaning and its being. The practice of medicine is meaningful because of the potential to restore another to health, and physicians have developed evidence-based practices for centuries that are meaningful as revealing how to restore another to health. The concern of medical research to unearth new forms of evidence – radiological, chemical, genetic – is meaningful only within this practical context. However, genetic technology can make use of genetic evidence in ways to enhance the practice of medicine, or that abstract from medical practice a diagnostic process that is then automated. Such technical decisions are never just technical. The latent content of such decisions is disclosed in the analysis below.

What then is the relationship of those engaged in technical work to practices? Is the practice that affords meaning to the use of a technical artifact a sort of distant backdrop that is of only passing relevance? Or is it present in day-to-day technical decisions, and just unrecognized most of the time? A closer look at practices themselves is helpful.

---

[2] Stephen A. Marglin, "What do Bosses do?: The Origins and Functions of Hierarchy in Capitalist Production", Review of Radical Political Economics, 6 (1974): 64

[3] The "chief advantage of the industrial assembly line is the control it affords over the pace of labor, and as such it is supremely useful to owners and managers whose interests are at loggerheads with those of their workers." Harry Braverman, *Labor and Monopoly Capital*, (New York: Monthly Review Press, 1998), p. 146.

Practices – An Overview

Practices emerge from the concern that humans have for themselves. A human is always reflexively a matter of concern for oneself and is thus always concernfully engaged in various practices that acquire their meaning and purpose from this self-concern.[4]

For example, one is concerned for one's health, and so one engages in the practice of medicine. Some concerns are more existential to one's being, such as one's health or one's concern to have a sense of purpose in one's life. Other concerns are more peripheral, such as one's concern to dwell or inhabit a space, alone or with others, which requires the practice of architecture, or one's concern to be clothed, which requires the practice of textiles. Practices are pervasive throughout and constitutive of human existence due to our concern for ourselves and the practical orientation towards our daily existence that results.

Our perception of artifacts and of other people is derivative of this practical concern, such that the meaning of an artifact or a person is always by reference to that for which we deal with the artifact or person. Artifacts and other persons are not understood primarily by observation, but by a use of things and soliciting of others in order to realize opportunities for ourselves.

A space of meaning is opened up by use and by soliciting, in which worldly objects and people become of concern to us *as* embedded in a network of worldly purposes. A hammer acquires its meaning originally not by observation of its properties, but by use *as* that which it is *for*, *as* useful *for* a type of work, *for* the sake of another's self-concern, in turn *for* the sake of oneself. A patient acquires her meaning for a physician *as* one concerned for their health, in turn *as* enabling a physician to realize their concern to practice medicine.

Furthermore, these meanings are not opened up from scratch in every new situation. One's past engagements with similar things and persons provides the network of references within which similar things and persons are recognized as meaningful to a certain end. Our routine, everyday interpretation of things and of other people is a hermeneutic interpretation in which they acquire meaning *as* useful in a pre-given network of references, *for* some type of work, *for the sake of* another's concern for themselves and, in turn, *for the sake of* one's self-concern.

Practices and Deliberation

Two features of the practical grounding of the meaning of things and persons are important to note.

[4] "Dasein is an entity which does not just occur among other entities. Rather it is ontically distinguished by the fact that, in its very Being, that Being is an *issue* for it." Martin Heidegger, *Being and Time*, (New York: Harper & Row, 1962), p. 32.

First, "as" perceptions are not simply abstracted through observation, independent of our self-concern. Our self-concern and the "as structure" of meaning are not two separate things, but are two moments of the same self-concern, viewed from the perspective of the self and of the world.

Second, the "as structure" of things and persons frequently breaks down in practice. The pre-given "as structure" of referential meaning works through similitude, recognizing the things and persons in a new situation *as* similar in the respects that are meaningful to the "as structure" of a practice.

When this breakdown occurs, when a new situation resists a practice, this "as structure" of the practice is put in question and made more explicit. Again, there is no recourse to another form of coping with the world when such a breakdown occurs – there is only the "as structure" which constitutes the world as meaningful, and so a rehabilitation of our self-concern is only possible by making this "as structure" more explicit, more an object of our attention. When a particular situation resists the "as structure" of a practice initially, we are thrown into deliberation, deliberating between the situation and the various "as-structures" that constitute the practical knowledge of a practice.[5]

Donald Schon in *The Reflective Practitioner* observes that every practice is characterized by this type of deliberation, which he calls reflection-in-action.

> "The practitioner has built up a *repertoire* of examples, images, understandings, and actions….When a practitioner makes sense of a situation he perceives to be unique, he sees it *as* something already present in his repertoire. To see *this* [architectural] site as *that* one is not to subsume the first under a familiar category or rule. It is, rather, to see the unfamiliar, unique situation as both similar to and different from the familiar one, without at first being able to say similar or different with respect to what. The familiar situation functions as a precedent, or a metaphor, or – in Thomas Kuhn's phrase – an exemplar for the unfamiliar one."[6]

A patient doesn't respond to medication for a routine condition, which prompts her physician to ask questions about diet, stress, and unusual symptoms, trying to reframe the patient's presentation in different ways until the patient's concernful presentation of illness is revealed in a way that is serviceable through practical treatment. When front-line medical interventions fail, most subsequent interventions are not applications of double-blind research,

---

[5] "There is no deliberation about the sciences (*episteme*) that are exact and self-sufficient, as, for instance, about letters, since we are in no doubt about how to write them. Rather, we deliberate about what results through our agency, but in different ways on different occasions – about, for instance, medicine and money making. We deliberate about navigation more than about gymnastics, to the extent that it is less exactly worked out, and similarly with other [crafts]." Aristotle, *Nichomachean Ethics*, (Hackett Publishing, 1999), NE III, 3, Sec 8, 1112b1-7.

[6] Donald Schon, *The Reflective Practitioner* (United States: Basic Books, 1984), p. 138

but research-in-practice in which the physician relies upon their repertoire of case history – from their own practice and those of their colleagues.

When the "as structure" of a practice is put into question, the entire network of references is put into question. The understanding of things and persons within the space of meaning opened up by the practice, but also the meaning of the practice itself - what it means to be a good physician, a good father, a good teacher - is put into question.[7] When the meaning of the practice is better revealed, then the "as structure" that derives its meaning within this space is modified as well.

Practitioners don't just *see as* and *do as*, they evaluate each reframing of an unfamiliar situation. This evaluation happens through an experimenting in practice that prefigures the more formal experimental methods of scientific research.[8] The practitioner asks, what would happen if I treat this *as* that? This experimental reflection on a situation is the core of deliberative practice observed by Schon. An architect views a site *as* another site, and mentally moves a building in accordance with this reframing to see the result. This *move testing experiment* can occur silently, even unconsciously, or it can occur on paper or computers (what Schon calls "virtual worlds"). In either case, the exemplars, the *as* structures, are made more explicit through evaluation-in-practice.

Schon also describes two other types of experimental reflection-in-action. In *exploratory experiments*, a practitioner manipulates an unusual object without predictions as to its behavior. A carpenter probes a new type of wood to observe its response to different conditions. A software engineer makes requests to a new module to see which requests return errors, which return valid responses, and the errors and responses they return. Finally, in *hypothesis testing experiments*, a practitioner evaluates different hypotheses. "If a carpenter asks himself, What makes this structure stable? and begins to experiment to find out – trying now one device, not another – he is basically in the same business as the research scientist."[9]

This reflective conversation between exemplars from practical experience and an unfamiliar situation is meaningful within the context of the practice itself. In fact, the practitioner experiences deliberation in the face of unfamiliar situations as more meaningful, more gratifying,

---

[7] "The essence of thinking" is "an indwelling releasement [*gelassenheit*] to the worlding of the world". Martin Heidegger, *Country Path Conversations*, (Indiana: Indiana Univ Press, 2016), p. 99, GA 77: 151.

[8] For Aristotle, deliberation brings to bear knowledge of means to produce a *techne*'s end that, while truly a knowledge of causes in the classical model of science, is *incomplete* and thus not true for all particulars. This type of incomplete knowledge is central to all inquiry, and deliberation, according to Aristotle, "is a type of inquiry (*zetesis*)". NE VI, 9, Sec 1 1142a33 The other type of inquiry, dialectic, which addresses "the things we search for (*ta zetoumena*)", also begins with knowledge that is partially true, *endoxa*, the opinion of the many or the wise. PA II, 1, 89b23-24 Whereas "the principles of demonstrations are definitions" (PA II, 3 90b25) dialectic "reason[s] from opinions that are generally accepted". (*Topics*, T I.1 100a21) *The Complete Works of Aristotle*, ed. Jonathan Barnes (Princeton: Princeton University Press, 1984). That's they the "end of the medical art is health", according to Aristotle (NE I, 1 Sec 3 1094a8), and "the health of this human being even more, since he treats one particular patient at a time". NE I, 6, Sec 16 1097a12-15

[9] Schon, p. 147

than the routine execution of a practice in familiar situations, as the practitioner is more personally and fully involved.

When professionals speak of the internal goods of their practices – the architect's sensitivity to context, the evidence-based practice of the physician – these internal goods are manifestations of deliberation within the context of a specific practice. The internal goods of practices are thus grounded in the self-concern of the practitioner which opens up the space of meaning within which practical reasoning occurs.[10]

Practices and Technical Reasoning

Should this deliberation still fail to account for the uncertainty of a particular situation, one further externalizes the "as structure" into a predicative, propositional theorization of worldly objects. Again, there is no recourse to another form of coping with the world when deliberation fails to restore a practice in a particular situation – there is only the "as structure" which constitutes the world as meaningful, and so a rehabilitation of our self-concern is only possible by making this "as structure" even more explicit, even more an object of our attention. One takes a more scientific approach to objects.

---

[10] It is helpful to situate this account of practices relative to the well-known account of MacIntyre. Both accounts have the same key constituents. MacIntyre defines a practice as "any coherent and complex form of socially established cooperative human activity through which goods internal to that form of activity are realized in the course of trying to achieve those standards of excellence which are appropriate to, and partially definitive of, that form of activity". Alasdair MacIntyre, *After Virtue,* (Indiana: Univ of Notre Dame Press, 1984), p. 187.

The key elements of MacIntyre's definition are that (a) the activity is complex ("Bricklaying is not a practice; architecture is"), (b) the activity requires social cooperation and (c) the activity possesses goods internal to the activity. The phenomenological account of practices given here, likewise, describes practices as (a) deliberative in face of unfamiliar situations, (b) soliciting of others whose self-concern opens a space of meaning and (c) possessing goods internal to the practice than cannot be attained through other activities.

An objection to MacIntyre's account is what to do when one must choose between the internal goods of conflicting practices. For it is precisely the breakdown of a single unifying tradition that organizes practices and goods into a hierarchy that has led to the ethical conflict that animates MacIntyre's study. MacIntyre's response is that such a tradition is necessary but has been lost because the external goods of acquisitive institutions that house practices have replaced the internal goods of practices with bureaucratic, technical reasoning. He calls the present moment a "new dark ages" and suggests we build communities of practice isolated from industrial society, in order to begin the development of a new tradition. MacIntyre defends his grounding of virtues, and the practices within which they develop, in tradition, what MacIntyre calls the "social teleological account" of virtues, because the Aristotelian grounding of virtue in a teleological account of human nature is no longer a coherent option.

The phenomenological account of practices given here presents a third way of grounding goods internal to practices. The internal goods of practices simply *are* the manifestations of deliberation specific to each practice, and deliberation - as the subjective side of the "complex form" of activity identified by MacIntyre as constitutive of a practice – is the essential core of practical reasoning. A phenomenological account of practices discloses the structures of their presence and absence in more detail than is provided by MacIntyre's historically contingent account of practices.

What is notable here is that the scientific approach to practical reasoning is not alien to practices. A common theme in many critiques of modernity is that practical reasoning has been undermined by instrumental, scientific thinking that runs roughshod over the particularities of practice, and that has its roots in philosophical errors made by Descartes or Plato or Parmenides.[11] The weakness of such accounts is their inability to account for much of the evidence in question, namely their inability to account for the widespread *appeal* of modernity.

The methods of scientific research, then, are not alien to practical deliberation. Not only is the propositional presentation of objects an externalization of the practical "as structure" from which they acquire their meaning, but the formal methods for evaluating hypotheses expressed propositionally are themselves externalizations of the evaluation-in-practice that characterizes practical deliberation.

Exploratory experiments, move testing experiments and hypothesis testing experiments *are* the methods that constitute scientific research, though now considered directly. By formalizing these approaches to experimental evaluation, science is able to control for bias in the more deliberative evaluation of hypotheses. A practitioner, even a school of practitioners sharing their case histories, is subject to the prejudice of their particular experience and their limited memory of their experience. A scientific approach to evaluation can eliminate much of this bias and thus restore practices in many situations in which they break down.

The scientific approach to practices that occurs when we raise the "as structure" and evaluation-in-practice of deliberation to the level of formal research, however, is problematic. The scientific framing of worldly objects, when viewed as a corrective to the bias of deliberative practice, suggests the possibility of an entirely new comportment towards worldly objects. By grounding one's comportment towards beings on the externalized predications of beings, an epistemic comportment seems to better satisfy the striving for universality inherent in the "as structure".

While a *deliberative comportment* towards beings emerges in practices, and constitutes the internal goods of practices, an *epistemic comportment* towards beings is grounded only in the proper use of method and rejects deliberation as subject to prejudice and bias.[12] Within this epistemic comportment, meaning is grounded in proper use of method, and method is severed

[11] This critique was noted in the discussion of MacIntyre above but is also given by the historicist metaphysics of the later Heidegger.

[12] This distinction echoes Heidegger's distinction between two types of solicitude: leaping in for others and leaping ahead of others. "With regard to its positive modes, solicitude has two extreme possibilities. It can, as it were, take aware 'care' from the Other and put itself in his position in concern: it can *leap in* for him. This kind of solicitude takes over for the Other that with which he is to concern himself....In contrast to this, there is also the possibility of a kind of solicitude which does not so much leap in for the Other as *leap ahead* of him in his existentiell potentiality-for-Being, not in order to take away his 'care' but rather to give it back to him authentically for the first time. This kind of solicitude pertains essentially to authentic care." *Being and Time,* pp. 158-159.

from its origins and conceived as an application of pure science.[13] It is this epistemic comportment that the later Heidegger critiques as enframing beings, as blocking the deliberation, the thinking, through which we allow unfamiliar situations, and the meaning of the practices that encounter them, to reveal themselves.[14]

In the epistemic comportment towards a practice, the internal goods of the practice, grounded in the "as structure" of deliberation, are externalized into the external goods of universal validity and avoidance of prejudice. Again, this contrasts with critiques of technology and technical reasoning that see the external goods of capital and power as undermining the internal goods of practices.[15] The appeal of technical reasoning is anchored in practice itself.

Practices and Technical Reasoning with Statistics

Both possibilities of technology – the automation of practices and the support for the deliberative core of practices by appeal to their internal goods - are present for AI. The ethical issues raised by AI are thus not new in the history of technology. AI, even at its most abstract level as the development of statistical models, is embedded in practices, and works within statistical paradigms that themselves arose within practices.

Much of the early work in statistics arose from the demand to control industrial processes. In the 1920s the formalization of probability to manage problems of risk and uncertainty in industry, in particular the risk of production errors, was in full force. The Bell Labs inspection department, which accounted for 12 percent of the workforce,[16] historically transitioned inspection from full lot inspection to random sampling between 1923 and 1927.[17]

---

[13] For example, early industrial boosters such as Andrew Ure maintained that "The principle of the factory system then is, to substitute mechanical science for hand skill." While these boosters contrasted the skills of the artisan with the "applied science" that legitimized the new industrial order, numerous historians have made it clear that the applied science of the Industrial period was largely an appropriation of artisanal knowledge. Andrew Ure, *The Philosophy of Manufacturers; or, an Exposition of the Scientific, Moral, and Commercial Economy of the Factory System of Great Britain* (London: Charles Knight, 1835; rpt., New York; Kelley, 1967), pp. 20, 23. Steven Shapin, "The Invisible Technician," American Scientist, 1989, 77: 554-563; Pamela H. Smith, "Art, Science, and Visual Culture in Early Modern Europe," Isis, 2006, 97:83-100; Charles F. Sabel and Jonathan Zeitlin, eds., *World of Possibilities: Flexibility and Mass Production in Western Industrialization* (Cambridge: Cambridge Univ. Press, 1997); Maxine Berg, *The Machinery Question and the Making of Political Economy* (Cambridge: Cambridge Univ. Press, 1980), pp. 154,250; and Berg, "The Genesis of 'Useful Knowledge'" History of Science, 2007, 45: 123-133.

[14] This paper thus follows Crowell who distinguishes the phenomenological observations of Heidegger from this historicism of the later Heidegger. When read as works in transcendental phenomenology, the later Heidegger can be embedded within the early Heidegger. Steven Crowell, "The Challenge of Heidegger's Approach to Technology: A Phenomenological Reading," in *Heidegger on Technology*, ed. Aaron James Wendland, Christopher Merwin and Christos Hadjioannou, (New York: Routledge, 2019), pp. 74-95.

[15] Examples of such a critique include MacIntyre as well as Marxist historians of deskilling.

[16] Juran, *Architect of Quality,* p. 75

[17] Before the rise of the factory, the autonomous workman inspected his own product, and this practice continued in the early factories. Scientific management changed this, creating independent inspection departments in line with the sharp separation of management functions and labor activities (or, head and hand activities). "The Taylor System and Quality Control", Dr. Joseph M. Juran, p. 5. One of these independent inspection departments, at Bell Laboratories, was faced with the massive engineering challenge of building

There were two distinct motivations for sampling-based inspection that led statistics in two divergent directions. In the context of the antagonistic relationship between labor and management, the most common approach was the use of acceptance sampling *tests*, which applied statistical significance testing to reject work from vendors (in the case of inputs) or workers (in the case of outputs).

This desire for acceptance sampling *tests* from industry drove much of the development of statistical significance tests by statisticians such as Ronald Fisher and, subsequently, Egon Pearson and Jerzy Neyman. Significance tests such as p-values, f-values and R^2 values use arbitrary thresholds of statistical significance to infer the presence of an effect in any object of study or to select from competing statistical models of an effect, and have been applied across dozens of fields from agriculture and psychology to medicine and industrial acceptance sampling.

This application of statistics has been criticized as having impeded as much as it advanced technical progress, by moving substantive domain expertise "out-of-the-loop" of work in favor of tests of statistical significance. These arbitrary tests of an effect replace substantive significance, which is concerned with the magnitude of an effect, with statistical significance, which sets magnitude aside and makes binary assertions of the presence or absence of effects based on thresholds of statistical significance.[18]

At Bell Labs, Walter Shewhart and his younger colleague Edwards Deming, were early critics of the displacement of workers' domain expertise with the use of arbitrary significance tests to judge quality of work. For them, variations in industrial output quality help workers better understand the processes themselves and inform new hypotheses into how to improve these processes.[19]

---

and inspecting a universally accessible telecommunications network. Inspection, both before and after the rise of industry, was traditionally full-lot inspection, which imposed significant costs as well as delays for large-scale industrial operations based on extreme specialization. Paul J. Miranti, "Corporate Learning and Quality Control at the Bell System, 1877-1929", pp. 55-57.

[18] "If you yourself deal in medicine or psychiatry or experimental psychology, …we would recommend that you focus on clinical significance. If you deal in complete life forms, environmental or ecological significance. If you deal in autopsies or crime or drugs, forensic or psychopharmacological significance. And so forth…An arbitrary and Fisherian notion of "statistical" significance should never occupy the center of scientific judgment." *The Cult of Statistical Significance: How the Standard Error Costs us Jobs, Justice and Lives*, by Deirdre McCloskey and Stephen Ziliak (Univ of Michigan, 2008), p. 20

[19] Deming consistently critiqued the application of statistical significance to management, countering that "Statistical theory shows how mathematics, judgment, and substantive knowledge work together to the best advantage. *Sample Design in Business Research*, by Deming, W. Edwards, (John Wiley & Sons, 1990), p. v. Deming had this to say in his *Out of the Crisis*.

> There are many other books on so-called quality control. Each book has something good in it, and nearly every author is a friend and colleague of mine. Most of the books nevertheless contain bear traps, such as reject limits,…areas under the normal curve, acceptance sampling….The student should also avoid passages in books that treat confidence intervals

This school of thought has generally been overshadowed by the more dominant school of Fisher, whose followers have argued for the use of tests of statistical significance as a more objective approach to model selection that eliminates the subjectivity inherent in human model building and decision making. Statisticians such as Shewhart and Deming, however, reject this elimination of subjective judgment. For them, technical jobs are essentially deliberative, model building activities, the ongoing refinement of mental models, before statistics enters the scene. This view of technology, as a mental modeling activity that is natural to man's practical life, is precisely the deliberation in response to unfamiliar and uncertain situations that characterizes all practical activity. The value of statistics, on this view, is the measurement of variation to inform model building.

This tension that runs through the history of statistics is present at its founding. Prior to statistical models, there was a prejudice against the variation of particulars in favor of mental models. In all fields, researchers "took simple averages of nearly perfectly replicated determinations of the same quantity; but the idea that accuracy could be increased by combining measurements made under different conditions was slow to come. They feared that errors in one observation would contaminate others, that errors would multiply, not compensate".[20]

The achievement of statistics was to measure variation, thus learning from particulars by introducing models that diverged from mental models. The Shewhart & Deming school of thought envisioned statistical and mental models checking and modifying each other, such that statistical models would be tools to be used to enhance our understanding of processes used in production.[21]

However, this school was not to win the day. The Fisher school of thought was an overreaction that denied the value of mental models, analogous to the previous denial of the value of particulars, which has plagued statistics from the beginning. The Fisher school is rooted in a positivist view of science which seeks to remove any trace of subjectivism from objective science, a bias which when applied to the workplace seeks to remove the need for a workers' practical knowledge in favor of objective processes specified by management.

This epistemic comportment of Fisherian tests of statistical significance has been central to the subsequent use of such tests in medical research as the only valid basis for medical practice, and to the use of predictive models in AI as autonomous from human interference. In medical research, initiatives exist to integrate the evidence of case-based practice with the evidence of controlled trials. Similarly, Explainable AI seeks to inform practices and integrate

---

and tests of significance, as such calculations have no application in analytics problems in science and industry.

Out of the Crisis, by Deming, W. Edwards (MIT, 1982), p. 369

[20] *The History of Statistics: The Measurement of Uncertainty before 1900*, by Stigler, Stephen (Belknap Press, 1986), p. 4

[21] This same school of thought is found in present-day Bayesians such as Gelman and Kruschke.

practical and model-based knowledge. In both cases, the technical work required to develop statistical models proceeds in divergent directions, depending on the appeal of the internal goods of the practice in question.

Explainable AI, as an example, is one such direction. However, there is confusion about what explainability means in AI. One school of thought holds that the purpose of explainability is to build trust in a prediction. Such a purpose presumes that models deserve to be trusted. Another school, which accounts for the deliberation of decision makers, holds that Explainable AI facilitates an iterative relationship between domain experts and AI models that improves domain expertise - and improves models through iterative model debugging and model checking by domain experts.[22] This iterative vision of strong AI, however, is largely overshadowed by the vision of strong AI as autonomous from subjective, human intervention.

3. Implications for Technology Ethics

First, ethical issues raised by technology are not about technology specifically - but about our practical comportment towards the world realized in practices. Artifacts just give structure to this comportment.[23] Our comportment can be deliberative, or our comportment can be epistemic and ground itself on the externalized account of practices and of evaluation given by science. The significance of technology is that it emerges, it is designed and developed, from within one of these two worldly comportments.

Technology can either be a tool enabling better deliberative action, or it can be its own foundation for action. Because technology expresses an entire worldly comportment, considered ontologically, the upskilling or deskilling effect of technology is more broadly a moral upskilling or a moral deskilling.[24] But technology is not the comportment itself and is thus limited - one can't be a technological determinist in matters of ethics without falling into the trap of removing ethical content from technical decisions.

Second, the constructive role of philosophy in relation to practitioners within technology is to provide a philosophical account of the decision-making involved in technical practices, not to take over such decision-making.[25] Careful, phenomenological study of such decision-making

---

[22] "The first step towards improving an AI system is to understand it's weaknesses. Obviously, it's more difficult to perform such weakness analysis on black box models than on models which are interpretable." The complimentary purposes of Explainable AI are "Verification of the system", "Improvement of the system", and "Learning from the system". "Explainable Artificial Intelligence: Understanding, Visualizing and Interpreting Deep Learning Models", Samek, Wiegand and Muller, https://arxiv.org/abs/1708.08296, August 28, 2017.

[23] "The essence of technology is by no means anything technological." Martin Heidegger, "The Question Concerning Technology", in *Basic Writings*, (New York: Harper Hollins, 1993), p. 311.

[24] Shannon Vallor, "Moral Deskilling and Upskilling in a New Machine Age: Reflections on the Ambiguous Future of Character." *Philosophy of Technology* 28 (2015):107–124.

[25] This is in contrast to postphenomenology, an approach to the philosophy of technology that evaluates artifacts as mediating our experience of the world and thus as co-constituting humans. Designers and developers of artifacts struggle to find such accounts to be relevant to day-to-day technical decisions. This loss of relevance, I argue, is due to the "empirical turn" in philosophy of technology away from Heidegger,

develops and affirms a rich deliberation, as is provided by Schon. Appealing to such deliberation, we appeal to the internal goods of practices.

Third, the striving of science and technology for universality cannot be dismissed as a mistake in the history of ideas. In fact, the striving manifests our most distinctively human capacity, to generalize through the "as structure" across particulars. Instead, science must be at the service of deliberation, of thinking. Deliberation-centered science is not conservative. It accelerates the development of practical knowledge through transparent measurement of practical decisions and their measurable consequences.

---

who was deeply concerned with man's thoughtless, inauthentic daily decision-making. The "empirical turn" critiques Heidegger as essentialist based on *The Question Concerning Technology*, while overlooking the role of world as given to man and structuring man's daily orientation towards common opinion. It is the inauthenticity of daily decisions oriented towards a given world that is of concern in the early and later Heidegger, yet the topic of authenticity in Division Two of *Being and Time* is not addressed in the critiques of Ihde or Verbeek. Rather, these critiques selectively read the tool analysis in *Being and Time* as Heidegger's "positive" view of ready-to-hand artifacts and "negative" view of scientific approaches to artifacts that render than present-at-hand. Heidegger was neither "positive" nor "negative" towards these aspects of world, they simply are essential features of daily existence in a world. His concern is with an authentic orientation towards one's world. Don Ihde, *Heidegger's Technologies,* (United States: Fordham Univ Press, 2010); Peter-Paul Verbeek, *What Things Do,* (University Park, PA: Pennsylvania State Univ Press, 2005).