



Against the iDoctor: why artificial intelligence should not replace physician judgment

Kyle E. Karches¹ 

Published online: 10 July 2018
© Springer Nature B.V. 2018

Abstract

Experts in medical informatics have argued for the incorporation of ever more machine-learning algorithms into medical care. As artificial intelligence (AI) research advances, such technologies raise the possibility of an “iDoctor,” a machine theoretically capable of replacing the judgment of primary care physicians. In this article, I draw on Martin Heidegger’s critique of technology to show how an algorithmic approach to medicine distorts the physician–patient relationship. Among other problems, AI cannot adapt guidelines according to the individual patient’s needs. In response to the objection that AI could develop this capacity, I use Hubert Dreyfus’s analysis of AI to argue that attention to the needs of each patient requires the physician to attune his or her perception to the patient’s history and physical exam, an ability that seems uniquely human. Human physician judgment will remain better suited to the practice of primary care despite anticipated advances in AI technology.

Keywords Technology · Artificial intelligence · Electronic health records · Physician judgment

Introduction

Ours is an age of ambivalence about technology. On the one hand, we find ourselves enamored of its promise to extend our abilities and make our lives easier. Consider a recent viewpoint article by Ravi Parikh et al. in the *Journal of the American Medical Association* arguing for the incorporation into clinical care of “predictive analytics,” algorithms that use historical information to predict future outcomes [1]. The authors cite examples such as Amazon’s product recommendation system for online shopping, and they claim that such “sophisticated machine learning algorithms”

✉ Kyle E. Karches
kyle.karches@health.slu.edu

¹ Department of Internal Medicine, Saint Louis University Hospital, 3635 Vista Ave, Saint Louis, MO 63110, USA

could analyze “big data” from electronic health records (EHRs) to predict patient risk and direct resources accordingly [1, p. 651]. For example, Parkland Memorial Hospital in Dallas already uses such an algorithm to identify patients at high risk for hospital readmission. Although the authors note concerns that these systems could threaten patient privacy and diminish the role of the physician’s judgment, they claim that “algorithms routinely outperform practitioners’ clinical intuition” and may help reduce the cost of care in the United States [1, p. 652]. If they are indeed correct about the power of algorithms, one might wonder whether or not advances in artificial intelligence could eventually replace physicians altogether, just as corporations such as Google promise to produce driver-less cars.

On the other hand, there is a nagging sense that the advance of technology effaces something important. A recent article for the *New York Review of Books* surveys the emerging literature about how smartphones are changing the way young people interact with each other and with the world [2]. For example, a psychologist concludes from a series of hundreds of interviews at schools that smartphones corrode the capacities for empathy and conversation [3], creating “students who don’t make eye contact or respond to body language, who have trouble listening and talking to teachers, and can’t see things from another’s point of view, recognize when they’ve hurt someone, or form friendships based on trust” [2]. She cites the example of a family that holds arguments on Gchat because the “value proposition” of face-to-face conflict is low [3, p. 127]. The online chat tool simply makes them more “productive.” In fact, designers of smartphone applications intend this colonization of daily life. Many of these designers studied how to manipulate human psychology at Stanford’s Persuasive Technology Lab, learning ways to produce behavioral loops that keep consumers returning to applications time and again. Yet the author of the review resolves that the proposed solutions to these concerns, such as “more thoughtful apps” or exhortations to “reclaim conversation,” are “wildly inadequate” [2]. The problems with technology seem hopelessly intractable.

In this paper, I critique the notion that primary care physicians are able to be replaced by an artificially intelligent “iDoctor.”¹ I develop two related lines of argument. First, I explain the critique of technology developed by Martin Heidegger and Albert Borgmann and use it to bring into focus the potential deleterious consequences of replacing physicians with artificial intelligence (AI) machines. I then draw on Hubert Dreyfus’s analysis of AI [4], which in turn relies on Heidegger’s epistemology and the work of Michael Polanyi [5, 6], to argue that the judgment of human physicians is better suited to the practice of medicine than is AI. As an example of what Borgmann calls a focal practice [7], primary care would inevitably be distorted by the loss of the human physician.

¹ For ease of reference, I use the term *primary care physician* throughout this paper. However, in so doing I do not mean to exclude other providers involved in primary care, such as nurse practitioners and physician assistants.

Heidegger's and Borgmann's critique of technology

Heidegger's critique of technology derives from his more general critique of metaphysics as a succession of ontotheologies [8–10]. Heidegger claims that, starting with Plato, Western metaphysics divided the question of Being into two parts. The first is ontological and asks what makes an entity an entity, an inquiry into the essence of entities and into what all entities share in common. The second is what Heidegger calls theological and asks what the source or ground of being is, an inquiry into the existence of entities and the question of which entity is the highest being. Far from a mere intellectual exercise, metaphysics creates an ontotheology that shapes the way individuals understand Being itself—determining their basic presuppositions about everything in existence, including those about themselves and, as regards physicians, their patients. Heidegger thus builds on the Kantian notion that people participate in making their worlds intelligible [10, p. 53]. The way in which Being appears to them depends upon their metaphysical assumptions, the ontotheology with which they engage the world.

Heidegger interprets the history of Western thought as a succession of ontotheologies, ending with Nietzsche. Whereas Nietzsche thought of himself as anti-metaphysical, Heidegger perceives an ontology in Nietzsche's will-to-power—which Nietzsche believed to be fundamental to all entities—and a corresponding theology in the “eternal return” of the will-to-power fully actualized, the pinnacle that any entity can achieve [10, p. 21]. Heidegger argues that the ontotheology of our late modern epoch is fundamentally Nietzschean. Being itself appears to us in the form of power relationships, as the constant interaction between competing forces with no inherent nature or purpose other than self-perpetuation through the actualization of their will-to-power. We perceive things in the world as intrinsically meaningless, awaiting the operation of human will-to-power to fashion their meaning.

According to Heidegger, modern technology is part of this Nietzschean ontotheology. He argued that “the essence of technology is by no means anything technological” [5, p. 311], but rather a way of engaging the world called *Gestell*, or “enframing” [5, p. 322]. The Nietzschean notion that things have no inherent meaning allows technology to conceive of everything in the world as *Bestand* or “standing-reserve,” that is, as mere raw materials or resources awaiting the imposition of order by the human will [5, p. 322]. Thus, when we moderns look at a tree, we see it in terms of the amount of lumber that it yields rather than on its own terms. Technology is a background assumption about the world that shapes the way the world appears to us.

Technological enframing leads us to extract things from their context within the world so that, for example, we view natural entities as sources for abstract gains like energy, to be harvested and stored. Indeed, the natural object is eclipsed from our view, as the energy becomes more real than its source. We reduce the qualitative to the quantitative, to mere information or data, until it is only the quantifiable that matters at all. Most perniciously, according to Heidegger, in late modernity we have turned the frame on ourselves, such that we become “human resources” awaiting optimization for maximally flexible use [10, p. 60]. This technological

understanding of human beings makes them, like everything else in nature, available for manipulation by external forces, such as the market economy.

Within the technological enframing, the primary criterion for evaluation of anything is efficiency. In the debate over clean energy, for example, one side claims that wind power is better than coal because it more efficiently meets our desire for minimally toxic energy, whereas the other side claims superiority for coal because it more efficiently meets our desire for cheap energy. Yet both sides look at objects in the world according to *Gestell*, not as things with their own nature but rather as resources for us to subject to our own desires. Heidegger calls this stance toward things in nature a *challenging-forth*, a demand that they conform to our wishes, the efficient fulfillment of which determines their moral worth [5, p. 320]. We even subject ourselves to this analysis, explaining why we so readily accept the replacement of human labor by machines, despite the human consequences: if a machine can accomplish the task more efficiently, then so be it. In a world ordered by the will-to-power, human beings enjoy no unique status but rather become one intrinsically meaningless force among others [11, p. 184].

For Heidegger, the primary problem with *Gestell* is that, like all ontotheologies, it reveals certain aspects of Being but obscures others. The smartphone applications cited in the introduction may make us more efficient, for example, but in using them we miss other important aspects of human interaction. As technology advances, we even lose the ability to perceive some of nature's properties, particularly its qualitative properties. Heidegger contrasts the challenging-forth of technology with the *bringing-forth* of traditional arts and crafts, which are ontological "openings" that allow things to reveal what they truly are in their fullness and to order our knowledge of them, prior to the will-to-power's operation upon them [11, p. 184]. For Heidegger, the solution to the problem of technology is ultimately this way of "letting beings be" [12], rather than imposing an ontotheology upon it at the outset. The only way to gain a free relation to technology is to understand it for what it is, an enframing that discloses some aspects of the world in certain ways and conceals other aspects of it in other ways. If we apprehend technology along these lines, perhaps we can step outside *Gestell* so that it no longer dominates us, and thus our epoch's technological ontotheology would give way to another.

Borgmann extends this Heideggerian critique, suggesting a concrete way in which to overcome the technological enframing. He describes the pattern of modern technology in an analogous way, as the *device paradigm*, which reduces the full significance of some part of human life to its essential function and then realizes that function as efficiently as possible [7, p. 40–48]. Within this paradigm, the only relevant criteria for moral evaluation of technology are instantaneity, ubiquity, safety, and ease of use. Devices make commodities easily available for consumption while typically concealing their inner workings, thus disburdening human beings of the labor that used to be required to realize the goods in question. Although they make human life more convenient in some ways, they also degrade the importance of skilled human labor, compressing human excellence into the small group of elite engineers responsible for technological design. The resultant attenuation of skill in the population then necessitates the development of ever more technology to compensate. Devices also disengage human beings from the real world and from each

other, rendering unnecessary the bodily expertise and caring attentiveness that characterize pre-technological practices.

To illustrate the device paradigm, Borgmann contrasts the pre-modern hearth with the heating systems of modern houses. In addition to providing warmth, the hearth was the center of work and leisure for a family's home. It required bodily engagement and the development of skills to procure wood and build the fire, necessitating the division of labor among members of the household. Thus, the hearth was inseparable from the practices that shaped families' experiences of the world and developed their capacities. Modern heating systems, by contrast, dissociate the good of warmth from this rich context, commoditizing it for easy consumption at the flip of a switch. Such systems conceal their machinery, making them inaccessible and incomprehensible to their users, so that a specialist is needed when the machinery fails. Borgmann claims that this disengagement from the material world occasioned by technology deprives human beings of opportunities to pursue excellence within a community dedicated to the achievement of certain goods. Ironically, although technology promises to augment our freedom and our abilities, it has instead undermined these human capacities.

Borgmann argues for a reform of technology centered upon what he calls "focal things" and "focal practices" [7, p. 196–219]. *Focal things*, such as the wilderness and the family meal, have dignity and greatness in their own right and therefore ennoble human life [7, p. 217–220]. They engage human capacities fully, in part because they cannot be possessed or controlled. *Focal practices*, then, are the socially established activities dedicated to focal things. They build up habits of engagement with the real world, fostering discipline and excellence of mind and body. According to Borgmann, focal things and practices are fragile, ever in danger of being undermined by technology, which tends to make the sustained effort required for focal practices unnecessary or unappealing. Yet they can also provide the orienting force for a reform of technology, because they give people reasons to use technology more selectively as a means of promoting the kind of human excellence that can be achieved only through focal practices, thereby restraining the device paradigm. For example, instead of using a microwave to heat up frozen food for a family dinner, thus bypassing human effort entirely, one might decide to use only kitchen gadgets that help to extend one's cooking skills, thus preparing a truly better meal. Borgmann calls for a renewal of the political discourse that attests to the importance of focal things in human life, revealing them in their greatness, much like poetry. Such discourse may help to identify the points at which technologies threaten the achievement of authentic human excellence within focal practices.

Technology in primary care

With this conceptual framework in place, one can analyze the way in which primary care physicians currently engage two different technologies: the stethoscope and the electronic health record (EHR). Although the stethoscope is certainly a form of technology, it also fits Heidegger's description of what he calls a simple tool. In a phenomenological account of the use of such tools, Heidegger points out that they

recede from our attention as they are used. When one uses a hammer, for example, one focuses not on the hammer but rather on the nail [13, p. 7]. Similarly, the physician using a stethoscope directs his or her attention to the sounds of the patient's bodily functions. The withdrawal of the tool's own presence allows the human operator to assimilate it and inhabit it, such that it becomes an extension of his or her body. Thus, as opposed to other technologies, simple tools promote direct human contact with the world. The stethoscope brings the physician closer to the patient's authentic body and helps the physician attune his or her senses to it.

The stethoscope therefore exemplifies Borgmann's ideal of technology, casting its use as a means of promoting the goods attainable within focal practices. At its best, primary care is a focal practice dedicated to the care of the patient, a focal thing of ultimate concern. Success in this endeavor requires the physician to attune his or her senses and intuitions to the patient's body and its needs. Because it acts as an extension of the physician's own human senses, the stethoscope facilitates this process. It directs the physician's attention to the patient's body as it really is, allowing the physician to conform his or her judgment to this bodily reality. Thus, the stethoscope helps the physician "bring forth" health from the patient's body, in Heideggerian terms.

The EHR, by contrast, tends to distance the physician from the patient's living body. Often the physician has access to the patient's EHR before he or she meets the patient. Before the physician walks into the exam room, the patient's vital signs appear in the record and the *problem list*, a list of the patient's symptoms and medical conditions, is generated. Thus, the EHR draws the physician's attention away from the patient's actual body toward a collection of facts about the patient's body, to the point that physicians now spend more time at the computer than at the bedside [14]. The lack of attention to the physical examination of patients has raised concern among medical educators that the new generation of physicians will no longer be able to perform an adequate exam. Reliance on technology leaves physicians less equipped to perceive aspects of patient care that cannot be captured technologically.

Further developments in EHR technology, such as so-called best practice advisories (BPAs), tend to treat these facts about the body as things of ultimate concern. In the Epic EHR system that I use in my practice, for example, every patient encounter triggers BPAs that remind the physician to meet all of the federally mandated quality measures appropriate for the patient's age and sex, such as cancer and cholesterol screening. Just as technological enframing leads one to view a tree in terms of the amount of lumber it can yield, so does the EHR encourage physicians to view patients' bodies in terms of their most basic characteristics. It especially focuses physicians' attention on those quantitative data such as blood pressure and cholesterol that can be manipulated with medications. As a set of algorithms, it can neither account for the individual patient's preferences and circumstances nor leave room for the physician to interpret the guidelines. Instead of drawing the physician closer to the patient's body, in all its uniqueness, the EHR treats it abstractly, challenging it forth from its context and constituting it as a set of facts over which the physician can exert control.

The EHR also exposes patient care to manipulation by forces external to the doctor–patient relationship, such as the state or insurance companies. Because the

record is digital, unlike older paper charts, third parties have easy access to patient information. Insurers, including the federal government, have begun to extract data from EHRs to provide primary care physicians with reports of the rates at which their patients meet certain quality measures, such as blood pressure control or colonoscopy at age fifty. In the near future, they will tie physician reimbursement to these benchmarks, encouraging physicians to complete them for each patient as soon as possible. Thus, the EHR draws the physician's attention toward an even higher level of abstraction, away from the health of the individual patient's body and toward the health of the body politic. It turns information about patients' bodies into the means by which forces such as the state can exert power over them, achieving the good of health with maximal efficiency.

Advocates of EHRs insist that algorithms like BPAs improve preventative care, pointing out that they increase rates of completion of screening measures [15]. On this model of primary care, the best physician is the one who performs all screening measures on each patient at every visit. Yet such a physician may neglect to explain the purpose of these tests or to ask patients whether or not they want the tests done. He or she may also focus less attention on problems that patients consider more important. In other words, this model of primary care leaves no space for an alternative vision of the good primary care physician as one who applies guidelines to each patient's unique circumstances, taking the patient's own preferences into account. Such a physician may even choose temporarily to ignore the guidelines altogether in order to address a patient's main concerns. Advocates of EHRs may argue that physicians using EHRs can still maintain this standard of excellence, but when EHRs tie compensation to completion of screening measures, the busy primary care physician is unlikely to resist the temptation to focus on them at the expense of other matters.

The EHR and its economic incentives thus threaten to turn primary care physicians themselves into resources to be optimized for efficiency. Physicians in the era of high throughput find themselves under administrative pressure to see more patients in less time using fewer resources and to devote less time to uncompensated activities such as teaching students. The BPAs are intended to facilitate higher throughput by simplifying patient care, providing physicians with a checklist to accomplish at each visit. Yet physicians subjected to such expectations of efficiency have begun to feel as though they provide "care on a production line" [16]. Due to these constraints, they cannot offer the individualized attention and care that each patient expects. It is perhaps little wonder, then, that physician burnout has become a "public health crisis" [17].

As this analysis indicates, medical technologies such as EHRs are not merely neutral tools, to be used as a means to whatever ends their human operators designate. Rather, as Peter-Paul Verbeek has argued, medical technology mediates our experience of the human body [13]. It constitutes the way in which the patient's body appears to the physician. Whereas the stethoscope draws the physician closer to the patient's actual physical body, the EHR constitutes the patient as a body of facts that can be made into inputs for algorithmic reasoning. Currently, this logic remains confined in certain ways. My practice's Epic EHR system has BPAs only for certain preventative care measures and chronic conditions, and even though the EHR mediates the doctor-patient relationship in powerful ways, physician and patient still

encounter each other face-to-face in the office. Yet the replacement of primary care physicians with artificial intelligence would extend the EHR's technological frame to cover all aspects of patient care. As I show in the following sections, such an extension would have far-reaching effects on the doctor–patient relationship.

AI and the doctor–patient relationship

The EHR, with all its algorithms and dictates for patient care, opens up the possibility of the iDoctor, an artificially intelligent machine designed to deliver primary care. This introduction of AI to replace physician judgment would subject the already strained doctor–patient relationship even further to technological enframing. On the patient side, an AI computer would require a mechanical view of human beings as entities comprehensible according to laws that can be programmed into the machine. It would thus isolate the human being from its rich context in the world and deconstruct the human body into its component systems. This tendency toward deconstruction is already latent in modern medicine—as evident in the systems-based reporting of the physical exam and the EHR's collation of a problem list that presents the patient as a concatenation of his or her medical conditions. The context of the patient's life is included in a so-called social history, but usually only insofar as it informs diagnosis and treatment. While many primary care physicians pride themselves on maintaining a focus on the whole patient, AI would extend and systematize the trend toward deconstruction, making its own rationalized representation of human beings more real, in a way, than actual patients. As Heidegger argues, the iDoctor's reduction of human beings to a set of medical concepts necessarily fails to understand life as it truly is, in its fullness.

On the physician side, the aforementioned *Journal of the American Medical Association* piece by Parikh et al. [1] offers an example of the morality that accompanies the technological enframing analyzed by Heidegger. Its highest good is efficient diagnosis and therapy, and its means include cost–benefit analyses and the creation of protocols for care. If AI can execute these directives better than human physicians can, then so much the worse for the humans. Concerns about patient safety and privacy are apparently outweighed by the prospect of ever better technology. As Parikh et al. note, detractors of EHRs initially expressed such apprehension, but advances in technology neutralized their criticisms [1]. Every technological problem seems to have a technological answer.

Yet many primary care physicians express frustration at the introduction of technologies like EHRs precisely because they undermine the relationship whereby physician and patient determine together, in conversation, how best to proceed. The sort of algorithms required for AI presume that medical judgment can be made abstractly, without regard for each patient's particular circumstances, for any such consideration would come at the cost of reduced efficiency. The substitution of such techniques for physician judgment would thus effect the ultimate transformation of primary care into a production line. As manifest in the introduction of best practice advisories into the EHR, these algorithms conceive of patients as fungible. For example, they cast every person who turns fifty years old as a patient-in-need-of-a-colonoscopy,

regardless of the patient's circumstances and preferences. The purpose of these algorithms is to convert every such individual into a patient-who-has-had-a-colonoscopy, much like the worker on an assembly line who performs the same action on the same product over and over again.

Whereas best practice advisories currently impose machine algorithms on preventative care, the replacement of physicians with AI would extend the algorithmic model to cover all aspects of primary care. All patients with similar symptoms would be classified together for analysis, and the diagnosis and therapy considered best for such patients would be provided. For example, AI would tend to treat all patients with knee osteoarthritis the same way, according to guidelines, heedless of the differences between particular patients that, in everyday clinical medicine, often determine the best course of treatment. This sort of medicine might maximize efficiency, but the best word to describe it is *mindless*: it delegates no role to the human mind's capacities for creativity and adaptability in accounting for patients' unique circumstances and preferences. Artificially intelligent medicine would replace this type of flexible knowledge with the uniform logic of the assembly line.

The introduction of AI into medicine would thus commodify patient care according to Borgmann's device paradigm. Much like modern heating systems, AI would replace an organic context of skillful human interaction with a machine, whose inner workings remain obscure to anyone without the requisite technological knowledge. Instead of interacting with a fellow human being, the patient would encounter a device created and maintained by experts who need not be present in the medical context at all. Such a "relationship" would be purely functional, focused solely on the provision of diagnosis and therapy as commodities for patient consumption. As with all commodities, the primary moral imperatives for this sort of care would be safety, instantaneity, and ease.

This disburdening effect of AI undermines physicians' capacity for true excellence. The reliance on machine algorithms would absolve physicians of the responsibility to develop their own medical judgment, a form of practical wisdom. This virtue arises in part from individual experience and personal habits of excellence, such as reading the medical literature and learning new skills. But it also arises through collaborative enterprise: it cannot be sustained apart from a community of practitioners mutually dedicated to achieving the goods of medicine or an educational system designed to cultivate good judgment. In other words, as Borgmann shows, excellence in human endeavors like primary care requires a practice. Alasdair MacIntyre has argued that human practices not only promote excellence in the achievement of certain goods but also foster moral virtues such as justice, courage, and honesty [18]. The use of AI in primary care may eventually prove to be safer and more efficient than the use of human physicians, but such use would render unnecessary the formation of communities of good practice that give rise to intellectual and moral virtues. It would also deprive patients of their own role as teachers. Because primary care requires practical knowledge, physicians can develop skills and virtues only in relationship with the patients for whom they take responsibility.

One might object here that AI's potential benefits to medicine outweigh its risk to the doctor-patient relationship. After all, many patients likely see physicians primarily to obtain diagnosis and treatment as efficiently as possible.

As the technology develops, it may turn out that AI machines are simply better at this task than human physicians are. Yet, according to Heidegger's analysis, this objection rests upon a notion of the goods of primary care that is already enframed. For example, the argument that BPAs improve preventative care is quantitative, based on statistics. The BPAs make physicians order colonoscopies earlier and more often than they would otherwise, and the increased rate of such orders is held to be good. As Heidegger points out, however, this enframing reveals certain aspects of Being while concealing others. It is true that, according to evidence from observational studies of the population of all patients aged fifty or older, screening colonoscopy may prevent some deaths from colorectal cancer [19]. Yet a physician who focuses on the provision of such screening measures at the population level interprets his or her individual patients' bodies according to such facts in order to facilitate efficient provision of services. He or she sees the patient as a member of a cohort of individuals that share only their age in common and treats that status as quasi-pathological, requiring a medical procedure as soon as possible. Such a physician might attend to the particular patient's circumstances, but he or she would feel pressure to use this information to convince the patient to undergo the procedure. After all, the physician must ensure that the rate of colonoscopy screening among his or her patients remains as high as possible, so as to maximize the diagnosis of early colon cancer in the population.

Heidegger's and Borgmann's critiques of technology suggest the possibility of an alternative conception of primary care as a focal practice. Instead of constituting the patient's body as a set of facts available for efficient manipulation, primary care physicians may let beings be, so to speak. The physician may attune his or her attention to the patient's body, allowing the body, in all its uniqueness, to call forth the proper response. The focus is primarily on the patient's specific wants and needs, which may not always be amenable to efficient management but rather require a long-term relationship between physician and patient. In most cases, the physician provides care consistent with generalized guidelines, but he or she remains willing to set the guidelines aside when appropriate. This sort of primary care requires the ability to attend to each patient as a unique individual rather than as part of a cohort whose members can all be managed in kind. It resists commodification because it cannot be easily separated from its context within the doctor–patient relationship.

Proponents of AI would likely contend that, eventually, AI machines will become capable of the thought processes underlying this type of primary care practice. Programmers simply need to develop more sophisticated and adaptable algorithms to match the capacities of the human mind. After all, engineers have devised technological solutions to many other problems that once seemed insurmountable. Yet just as Heidegger's critique of technology reveals the iDoctor's potential to dehumanize medicine, so does his epistemology provide a response to this objection. As I show in the next part of this paper, Dreyfus's Heideggerian analysis of AI gives reasons to believe that AI may never be able to attune to human beings in the way that good primary care requires. It may be the case that human thought is not simply more complex than machine thought, but rather wholly different from the type of intelligence that machines can achieve.

Dreyfus's critique of AI

In *What Computers Still Can't Do*, Dreyfus calls attention to the often unspoken rationalist theory of mind underlying AI [4]. Rationalists like René Descartes think all mental knowledge consists of representations containing the fixed, context-free, abstract features of a domain. On this view, such representations can be expressed in propositional form, and indeed the ability to represent knowledge of a domain discursively explains that domain's intelligibility to the human mind. As Dreyfus puts it, this type of representational rationalism "assumes that underlying everyday understanding is a system of implicit beliefs" [4, p. xvii]. Even common-sense know-how is really a type of knowing-that. Here, *know-how* refers to the type of practical knowledge that enables day-to-day interactions with the world and allows one to acquire skills, whereas *knowing-that* refers to the theoretical knowledge that can be adequately contained in the form of propositions.

Similarly, proponents of AI assume that all knowledge can, in principle, be represented as propositions in a vast database, fed to a computer as inputs. Furthermore, they assume that the human mind's manipulation of this knowledge can be reduced to a set of formal rules, which can also be programmed into the computer. The challenge for AI is to generalize background knowledge, such as common sense, away from its context, in much the same way that ontology seeks to describe context-free entities as a foundation for philosophy. Yet as Dreyfus shows in his extensive review of the history of AI, attempts to overcome this challenge have repeatedly failed. Versions of AI produced to date have had difficulty answering certain simple questions rendered in ordinary English, for example. Dreyfus suggests that these failures stem from an inability to capture our *background knowledge*, the "understanding we normally take for granted," in propositional terms [4, p. xix].

Dreyfus turns to Heidegger's epistemology for an alternative to the rationalist theory of mind. As a phenomenological account of skill acquisition shows, human know-how cannot simply be represented as knowing-that. A novice learning a skill such as chess must indeed begin by "learning and applying rules for manipulating context-free elements," but as one gains experience, one learns to find meaning in "context-dependent characteristics such as unbalanced pawn structure" and to see these elements in terms of one's own goals [4, p. xii]. Over time, one leaves the rules behind altogether, having begun to see immediately what to do. The chess master's past experience shapes even her perception of the present moment in the game, such that a response comes to mind without needing to consult a set of general rules. Indeed, the best players are those who creatively transcend the rules. If one does not experience the mastery of a skill as the application of the rules learned through its acquisition, Dreyfus reasons [4], there is no basis for assuming that these rules still play any role at all. The master may sometimes, but not always, be able to explain post hoc the process by which she arrives at a correct response; but, even so, the rule derived from this explanation is not the correct response's cause, as AI proponents assume.

Furthermore, on a rationalist model, for a computer to respond appropriately to a specific situation, it needs to categorize the situation, use rules to search its database for additional rules that could be relevant, and then deduce a conclusion about how

to proceed. The more information about new situations that the computer receives, the more difficult such a process becomes. Dreyfus points out that this very problem has confounded attempts to increase AI's scale, for as the number of inputs increases, the system takes longer to retrieve information [4, p. xxi]. By contrast, as the example of the chess master shows, when a human being gains experience, he or she responds more easily, retrieving the relevant information even faster. This observation indicates that humans rely upon a very different type of information storage than that presumed in the AI rationalist model. When one gains enough experience to become an expert, one's knowledge comes to structure one's very perception such that one "directly experiences which events and things are relevant and how they are relevant" [4, p. xxviii]. As Heidegger would say, objects appear to the expert "not in isolation and with context-free properties but as things that solicit responses by their significance" [4, p. xxviii].

Know-how, then, seems to be built not on rules but on a type of pattern recognition that can be extended to new situations, as well as an immediate, intuitive perception of proper response. Inspired by Polanyi's account of tacit knowledge [6], Dreyfus draws attention to several features of background understanding, which underlie human know-how and would be difficult for AI to replicate [4, pp. xxvii–xxix]. First, as the example of chess indicates, this type of knowledge seems to require vast experience of typical cases. The reason why children "find it fascinating to play with blocks and water day after day for years," Dreyfus suggests, is that they are "learning to discriminate the sorts of typical situations they will have to cope with in their everyday activities" [4, p. xxvii]. Any attempt to capture this knowledge in terms of propositions for a computer program would be so incredibly complex and subject to so many exceptions that its success seems nearly impossible. Yet human children possess and act upon such knowledge intuitively.

Second, much of human background knowledge requires socialization. As Heidegger and Maurice Merleau-Ponty point out, in everyday know-how, what appears as a relevant fact depends on social skills. To demonstrate this point, Dreyfus draws on Bourdieu's description of a gift-giving culture [20]. Members of this culture do not appeal to rules to deduce what to do or how to act but rather simply react in appropriate circumstances with an appropriate gift. The rules need be explained only to an outsider who lacks the requisite social skill. Hence knowledge of gift-giving "is not a bit of factual knowledge, separate from the skill or know-how for giving one" [4, p. xxiii]. Another example is reasoning by analogy or metaphor. One cannot understand a metaphor such as "Sally is a block of ice" representationally by listing the qualities of Sally and ice and then comparing them [21]. Rather, as Heidegger and Merleau-Ponty contend, the social and cultural context in which such devices are deployed allows the relevant association to make itself known.

Third, and most important, know-how often requires the type of knowledge acquired by sensation, emotion, and imagination. One example of the sort of simple English sentence that can stymie AI is: "Mary saw a dog in the window; she wanted it." Computer programs have had difficulty determining whether "it" refers to the window or the dog [22, p. 200]. Dreyfus suggests that interpreting the contextual antecedent of the pronoun in this sentence appeals to "our ability to imagine how we would feel in the situation, rather than requiring us to consult *facts* about dogs and

windows” [4, p. xix]. It also relies on our experience of dogs as objects of desire, an experience gained through socialization. If the latter part of the sentence were instead “she pressed her nose up against it” [22, p. 200], then interpreting the contextual antecedent of “it” would involve appealing to the sensory experience of contact with a window. Such examples require imagination, and recollection of certain experiences, to organize the knowledge necessary for understanding the sentence, and thus they demonstrate the important role of the body in such understanding.

Dreyfus concludes from these observations that our experience and knowledge of the world take their structure from pre-conceptual background knowledge. Human beings approach the world with certain expectations, and meaning arises from the complex equilibrium between these expectations and information from the world. When the expectations that frame one’s perceptions are wrong, the incoming data do not make sense, and a new hypothesis is required. When these expectations do succeed, however, one finds oneself able to cope with the world well enough—a situation that Merleau-Ponty calls *maximum grasp* [23], which “varies with the goal of the agent and the resources of the situation” [4, p. 250]. Such a process underlies the pattern recognition that leads to human know-how and allows us to adapt to different situations in which we may be more or less certain of our knowledge about the world.

Because attempts at AI proceed from rationalist assumptions, computers have become adept at exactly the kind of abstract reasoning that Descartes believed to be characteristic of all knowledge. Thus, for example, computers can perform complex mathematical calculations much more rapidly and easily than can most human beings. They also succeed in situations such as games, which are able to provide clear reinforcement about whether the rules applied are correct or incorrect. The example of chess is relevant here as well. The triumph of computers over human chess grandmasters, as exhibited in the victories of IBM’s Deep Blue over Garry Kasparov, has led some proponents of AI to believe that machine learning will surpass human intuition elsewhere. However, in chess the feedback from the world is binary: win or loss. Such clear reinforcement is rarely involved in the type of situated know-how at which humans characteristically excel. As Dreyfus argues, although some utilitarian moral philosophers have proposed to interpret all moral reasoning as a utility-maximizing mathematical function—the kind of calculation at which a computer would excel—human beings are rarely able to foresee all the consequences of a moral decision in a way that would allow for this sort of accounting [4, p. xlv]. Human practices in which conclusions are less certain seem to rely more upon the embodied, situated knowledge that Dreyfus describes.

In an attempt to overcome the limitations of the rationalist model of AI, some computer programmers have designed AI machines based on a *neural net*, a form of AI designed to approximate the way in which humans learn the kind of pattern recognition that underlies know-how. Instead of giving the computer pre-formed propositional knowledge, neural net programmers feed the computer raw data and allow it to make its own connections spontaneously, in a manner analogous to the way in which human beings seem to learn. Yet such programs have nevertheless encountered problems for the reasons that Dreyfus anticipates. One problem is that the rules guiding these spontaneous connections must still be written out propositionally for

the computer. On a deeper level, however, because AI machines necessarily lack a human body with human perception, they seem incapable of determining relevance in the same way that humans do. In other words, they make spontaneous connections between inputs that no human would ever make.

The results of Google's DeepDream AI project provide a vivid example of the differences between human and AI perception [24]. DeepDream is an image-analysis program based on a neural net. Its programmers intended for it to learn to determine what an image "looks like" by picking out patterns in images and using a form of probabilistic logic to reinforce them, in much the same way that a human might think that a cloud looks like a boat or a turtle. Yet when asked to process simple images, DeepDream found all sorts of random objects where none was apparent to the human eye. For example, when given a picture of President Obama, DeepDream produced an image with eyeballs, birds, and psychedelic whorls of color scattered across the president's face and the background. It found within the image groups of pixels resembling these objects and then simply interpreted them as such. These results, while rather dazzling to behold, show how different the computer's probabilistic grasp of similarity is from that of the human mind, and how difficult it is for a computer to "think" the way humans do without human faculties of perception. Lacking a human body, AI machines may be destined to produce knowledge that, to humans, proves more odd than helpful.

Since human beings cannot step outside their own pre-conceptual background knowledge so as to describe it propositionally, they cannot impart the features of this knowledge to a computer in terms of explicit rules. Without such rules, a computer will remain without the capacity to distinguish relevant features of the world like humans can. An AI computer may be able to identify connections between its inputs that humans cannot see, and some proponents of AI have argued that, in this way, AI can contribute to our knowledge about the world. More likely, however, is that, lacking a human pre-conceptual understanding of relevant knowledge, it would produce connections that are uninteresting at best and unintelligible at worst.

Dreyfus's analysis draws our attention to the importance of know-how, the type of human practical knowledge that underlies both our common-sense grasp of the world and our ability to acquire skills. Computer engineers have struggled to reproduce this type of knowledge in AI machines not only because it is difficult to represent propositionally, but also because it relies upon various forms of pre-conceptual background knowledge, such as sensory perception and socialization, that seem uniquely human. As shown below, humans thus remain particularly well suited for practices, such as primary care, in which know-how plays a central role.

Human know-how and primary care

Dreyfus's critique casts doubt on the notion that AI machines can replace primary care physicians, since machines lack access to the sources of background knowledge underlying the practice of medicine. Much of clinical knowledge consists of know-how, the sort of pattern recognition that Dreyfus describes. The structure of medical education in fact presumes the importance of experiential learning. The popularity

of problem-based learning in the preclinical years derives from a general recognition that bedside clinical reasoning is context-dependent, requiring a wealth of specific prior experiences to draw on. Physicians typically require abstract conceptual knowledge only secondarily, as a resource to consult when they encounter a situation that does not fit the pattern [25]. The purpose of the clinical years and residency training is to expose trainees to enough cases that they begin to perceive patterns in the diagnosis and treatment of disease. Such experience is especially crucial for primary care physicians, who often evaluate patients with new or nonspecific complaints. Dreyfus's analysis indicates that this type of clinical know-how cannot be adequately captured in propositional terms and will thus be difficult, if not impossible, to transmit to a computer.

One example of such know-how is the ability to take a "history of present illness" from a patient. As medical educators know well, medical students often fail to elicit clinically relevant information from patients despite spending more time than anyone else actually taking the history. By contrast, an experienced clinician can obtain a complete history in a fraction of the time. Such physicians would likely have difficulty explaining this ability discursively, for it rests on their capacity to perceive relevance. As the patient begins to tell his or her story, the physician immediately begins to think of other symptoms to ask about and physical exam maneuvers that might lead to a particular diagnosis. When I take a history, at no point do I resort to a rule such as "always consider asthma and ask about cough when a patient complains of difficulty breathing." Rather, I recognize what parts of the patient's history are most significant, often by recalling similarities between this patient's complaints and those of many other cases I have seen. True medical expertise consists in precisely this type of experiential knowledge, not simply in the mastery of context-free abstract concepts of disease and treatment.

Dreyfus's critique also reinforces the importance of socialization and somatic experience for clinical know-how. Social skills, bodily experience, and imagination shape our pre-conceptual expectations and thus the way we perceive the world—which for primary care physicians consists of patients and their complaints. In order to apply its algorithms, an iDoctor would necessarily render this complex perceptual environment as knowing-that, thereby overlooking important information, such as the veracity of a patient's account. Any experienced clinician knows not to rely solely on patients' verbal reports of their illnesses. Some patients understate the burden of their symptoms, for example, and others lie outright about the purpose of their visit. Over time, a human physician might learn to attend to the subtle nonverbal cues, such as averted glance or speech disfluency, that can signal the withholding of information. In such cases, the physician might ask more questions to discover the truth or decide on a safer treatment plan. Yet an AI machine, without the resources of human experience or interpretive skills, must take the patient's history at face value, as an accurate transmission of propositional knowledge, potentially leading to mistakes in diagnosis and treatment.

Furthermore, even when patients are truthful and forthcoming, their narratives of symptoms are not just reports containing factual information; they are appeals to the physician as an embodied and social human being. They invite the physician to consult his or her own bodily experience and to imagine what such symptoms would

feel like and thus call forth a certain response. A physician can only truly understand the complaint of a patient with a headache if he or she has also suffered from headaches, or at least possesses the ability to draw on other experiences to imagine what such pain might be like. A patient describing a symptom or the impact of illness on his or her life might also make use of metaphor, relying on a set of shared social assumptions in order to convey meaning. Additionally, the appeal to a human physician invites the physician to respond with appropriate conversational expressions of empathy, such as facial gestures or consoling touches on the hand. These expressions, however minor, can often augment patients' healing response to treatment [26]. Even if an AI machine could perform these actions, it would likely have difficulty knowing how and when to use such gestural tokens appropriately, and still one wonders whether patients would have the same response to the touch of the iDoctor.

Just as importantly, a human physician can account for the patient's unique non-medical circumstances when devising a plan of care. If an elderly patient tells a physician that anti-hypertensive medication makes him so dizzy that he can no longer drive to visit his family, the physician can understand that information as relevant, having had a similar experience of the value of visiting with loved ones. Although the "best practice" guidelines for management of hypertension would recommend treatment, the physician might devise a more permissive blood pressure goal for this particular patient so as to better meet his needs. An AI machine, by contrast, does not possess similar memories of non-medical experiences to which patients might appeal. Dreyfus implies that an AI machine, lacking socialization within a human culture and the experience of embodiment, would necessarily fail to comprehend the nuances of these aspects of patient care.

Unlike games and other constrained situations at which computers excel, medicine rarely admits of enough certainty to allow for clear feedback. As Dreyfus points out, computers "learn" best when they are able to receive unambiguous reinforcement from a programmer or from the world with regard to the rules that they apply: right or wrong, win or loss. Yet outcomes in medicine are rarely so obvious or binary. Primary care physicians in office practice diagnose and treat patients without knowing for certain if they will return, and when they do return, their response to therapy is rarely a clear "better" or "worse." Many patients experience an ambiguous mix of benefits and side effects, leaving physician and patient to determine together the best way forward in a specific situation. Often the "right" answer, such as a statin medication for high cholesterol, is not the best answer for a particular patient's concerns. Primary care thus requires flexibility and adaptability, as well as a willingness to accept as sufficient the coping with the world that Merleau-Ponty calls maximum grasp [23]. Dreyfus points out that computers tend to have more difficulty applying this type of reasoning when they are faced with an increasing variety of circumstances, whereas humans tend to improve on their application. Any designer of an AI machine that purports to replace physician judgment would need to overcome this significant obstacle.

The ongoing attempt to produce driver-less cars provides some evidence of the problems that occur when AI technology is introduced into a human social milieu. A recent article aimed at the lay public notes that private and public investment in this type of AI has been massive, with GM alone spending \$500 million and the

U.S. Transportation Secretary proposing to invest \$4 billion [27]. Yet, to date, no version of this technology has been able to approximate human *social intelligence*, the ability to guess effortlessly and correctly what other people on the road will do. Whereas research shows that “even very young children have exquisitely tuned senses for the intentions and goals of other people”—as the “core of uniquely human intelligence”—AI cars frequently fail to anticipate the actions of human drivers [27]. Furthermore, although both AI machines and humans make mistakes on the road, the machines make mistakes that no human would: “They will mistake a garbage bag for a running pedestrian. They will mistake a cloud for a truck” [27]. In other words, AI machines have difficulty with the type of know-how that comes naturally to human beings. They cannot always identify the humanly relevant information from among their many inputs, especially in situations that call for knowledge gained by socialization, precisely the problems that Dreyfus predicts. A communal human practice such as driving, and by extension medical practice, may be possible only because human beings share the same background knowledge.

It is conceivable that an AI machine might eventually be capable of providing certain aspects of primary care, such as preventative services, more efficiently than humans do. No doubt some patients go to physicians solely to obtain such appropriate treatment as efficiently as possible. Yet many others have sought not only medical therapy but also compassionate care from their relationships with primary care physicians. As Dreyfus’s analysis shows, the patient’s complaint of chest pain appeals to a fellow human being capable of a similar experience. As the physician responds, the patient can perceive not only the physician’s effort to diagnose and treat the problem but also the bodily signals, both verbal and nonverbal, of the physician’s concern for the patient. Thus, the relationship between patient and physician opens the possibility for true compassion, understood etymologically as “suffering-with.” This affective response then motivates the physician to use the tools at his or her disposal, including medical knowledge as well as empathic listening and presence, for the good of the patient.

An AI machine might be vulnerable to various malfunctions that would require “treatment” by an expert. However, because a machine cannot become incarnate, and therefore cannot be vulnerable in the same way as human beings, it cannot possibly suffer with its patients. As MacIntyre has argued, this acknowledgment of the vulnerability of the human body gives rise to certain virtues that contribute to human flourishing [28]. These virtues, such as compassion, are essential for a doctor–patient relationship that not only provides efficient diagnosis and treatment but also assures the patient that the physician cares for him or her. Thus, close attention to the human body is crucial both for clinical knowledge and for the moral practice of medicine. Dreyfus’s account suggests that AI may remain ever incapable of attuning to the human body in a way that comes naturally to other human beings.

Conclusion: Primary care as focal practice

As noted above, Heidegger and Borgmann’s critique of technology shows how technology tends to distort focal things. It thus suggests an alternative model for primary

care as a focal practice centered on a focal thing: the care of the patient. Both parts of this focal thing, the care and the patient, are of ultimate concern. Within such a practice, care must be understood holistically, as not only diagnosis and treatment but also compassion and fellowship in the face of illness—benefits a human physician may be equipped to offer. Indeed, because the physician is human, there is some degree of contingency in the doctor–patient relationship, for its outcome is not determined algorithmically in advance. This contingency seems anachronistic, if not alarming, from within the technological enframing that tends toward standardization and control, but it also preserves an opportunity for grace, as doctor and patient may experience healing as a gift [29]. Similarly, it allows physicians the opportunity to develop their knowledge and skill, achieving excellence in a distinctively human way. This notion of primary care, then, challenges physicians to design and use technology such as EHRs to extend rather than attenuate their ability to care for patients. It also holds medical educators responsible for their duty to cultivate such intellectual and moral virtues in their trainees.

This model of primary care as a focal practice upholds the importance of the individual patient. Instead of imposing a technological frame of law-like generalizations and best practices on the patient, physicians in such a practice let patients be themselves, in all their particularity. This attempt to see patients anew, outside the technological enframing, fulfills the Heideggerian hope for medicine as an ontological opening that brings forth health from the individual patient's life story, rather than isolating those aspects of the patient's life that can be therapeutically manipulated in order to deliver health as a commodity. Such a doctor–patient relationship seems possible only between two human beings. As Dreyfus points out, our shared humanity allows us to recognize the most relevant parts of our experience and thus to arrive together at the truth. In other words, as human beings, physician and patient are capable of caring about the same things, and this mutual understanding makes possible an appropriate response [30]. In some circumstances, efficient diagnosis and treatment are enough. However, in others, particularly those of chronic and terminal illness, more is required. After all, disease and death remind us of our frailty and expose as a fiction the technological conceit that we do or can control nature. Thus, those of us trapped in the technological enframing are apt to perceive them not as parts of our natural human life story but as existential threats [31]. In the face of such suffering, the physician might offer not merely more technological artifice but rather wisdom and compassion derived from his or her own experience of being human.

I have argued in this paper that AI should not, and perhaps cannot, replace the physician's role in providing primary care. Yet it must be conceded that these arguments could be overcome by actual technological developments. The science of AI may well progress to the point that it achieves the capacity to replace the judgment and intuition of primary care physicians, at which point patients might welcome such a change. Heidegger was famously pessimistic on this point, declaring in an interview that “only a god can save us” from the encroachment of *Gestell* [32]. If technology can undermine even human communication and empathy, as demonstrated in the introduction, then it seems likely to exert greater influence over other human practices, including medicine. If nothing else, then,

perhaps this paper serves to show what goods are at stake if, or when, we walk into our primary care provider's office only to find the iDoctor there waiting.

References

1. Parikh, Ravi B., Meetal Kakad, and David W. Bates. 2016. Integrating predictive analytics into high-value care: The dawn of precision delivery. *Journal of the American Medical Association* 315: 651–652.
2. Weisberg, Jacob. 2016. We are hopelessly hooked. *New York Review of Books*. <http://www.nybooks.com/articles/2016/02/25/we-are-hopelessly-hooked>.
3. Turkle, Sherry. 2015. *Reclaiming conversation: The power of talk in a digital age*. New York: Penguin.
4. Dreyfus, Hubert L. 1992. *What computers still can't do: A critique of artificial reason*. Cambridge: MIT Press.
5. Heidegger, Martin. 1993. The question concerning technology. In *Basic writings*, rev. ed, ed. David Farrell Krell, 307–341. San Francisco: HarperCollins.
6. Polanyi, Michael. 1958. *Personal knowledge: Towards a post-critical philosophy*. Chicago: University of Chicago Press.
7. Borgmann, Albert. 1984. *Technology and the character of contemporary life: A philosophical inquiry*. Chicago: University of Chicago Press.
8. Heidegger, Martin. 1991. *Nietzsche, Vol. 1: The will to power as art; Vol. 2: The eternal recurrence of the same*. Trans. David Farrell Krell. San Francisco: HarperCollins.
9. Heidegger, Martin. 1991. *Nietzsche, Vol. 3: The will to power as knowledge and as metaphysics; Vol. 4: Nihilism*. Trans. David Farrell Krell. San Francisco: HarperCollins.
10. Thomson, Iain D. 2005. *Heidegger on ontotheology: Technology and the politics of education*. Cambridge: Cambridge University Press.
11. Feenberg, Andrew. 1999. *Questioning technology*. New York: Routledge.
12. Heidegger, Martin. 1993. On the essence of truth. Trans. John Sallis. In *Basic writings*, rev. ed, ed. David Farrell Krell, 111–138. San Francisco: HarperCollins.
13. Verbeek, Peter-Paul. 2011. *Moralizing technology: Understanding and designing the morality of things*. Chicago: University of Chicago Press.
14. Chi, Jeffrey, Maja Artandi, John Kugler, Errol Ozdalga, Poonam Hosamani, Elizabeth Koehler, Lars Osterberg, et al. 2016. The five-minute moment. *American Journal of Medicine* 129: 792–795.
15. Dexheimer, Judith W., Thomas R. Talbot, David L. Sanders, S. Trent Rosenbloom, and Dominik Aronsky. 2008. Prompting clinicians about preventive care measures: A systematic review of randomized controlled trials. *Journal of the American Medical Informatics Association* 15: 311–320.
16. Goitein, Lara. 2015. Training young doctors: The current crisis. *New York Review of Books*. <http://www.nybooks.com/articles/2015/06/04/training-young-doctors-current-crisis>.
17. Caplan, Arthur L. 2016. Physician burnout is a public health crisis, ethicist says. *Medscape*. <http://www.medscape.com/viewarticle/859300>.
18. MacIntyre, Alasdair. 1984. *After virtue: A study in moral theology*, 2nd ed. Notre Dame: University of Notre Dame Press.
19. Brenner, Hermann, Christian Stock, and Michael Hoffmeister. 2014. Effect of screening sigmoidoscopy and screening colonoscopy on colorectal cancer incidence and mortality: Systematic review and meta-analysis of randomised controlled trials and observational studies. *BMJ* 348: g2467.
20. Bourdieu, Pierre. 1977. The objective limits of objectivism. In *Outline of a theory of practice*, 1–71. Trans. Richard Nice. Cambridge: Cambridge University Press.
21. Searle, John R. 1979. Metaphor. In *Expression and meaning: Studies in the theory of speech acts*, 76–116. Cambridge: Cambridge University Press.
22. Lenat, Douglas B., and Edward A. Feigenbaum. 1991. On the thresholds of knowledge. *Artificial Intelligence* 47: 250–285.
23. Merleau-Ponty, Maurice. 1962. *Phenomenology of perception*. Trans. Colin Smith. London: Routledge and Kegan Paul.

24. Auerbach, David. 2015. Do androids dream of electric bananas? *Slate*. http://www.slate.com/articles/technology/bitwise/2015/07/google_deepdream_it_s_dazzling_creepy_and_tells_us_a_lot_about_the_future.html.
25. Norman, G.R. 1988. Problem-solving skills, solving problems and problem-based learning. *Medical Education* 22: 279–286.
26. Brody, Howard. 2012. On talking and touching in medicine. *Journal of Pain and Palliative Care Pharmacotherapy* 26: 165–166.
27. Anthony, Samuel English. 2016. The trollable self-driving car. *Slate*. http://www.slate.com/articles/technology/future_tense/2016/03/google_self_driving_cars_lack_a_human_s_intuition_for_what_other_drivers.html.
28. MacIntyre, Alasdair. 1999. *Dependent rational animals: Why human beings need the virtues*. Chicago: Open Court.
29. Lake, Christina Bieber. 2013. *Prophets of the posthuman: American fiction, biotechnology, and the ethics of personhood*. Notre Dame: University of Notre Dame Press.
30. Reich, Warren Thomas. 1996. A new era for bioethics: The search for meaning in moral experience. In *Religion and medical ethics: Looking back, looking forward*, ed. Allen Verhey, 96–115. Grand Rapids: Eerdmans.
31. Hauerwas, Stanley. 1990. *God, medicine, and suffering*. Grand Rapids: Eerdmans.
32. Heidegger, Martin. 1976. Only a god can save us: *Der Spiegel's* interview with Martin Heidegger. Trans. Maria P. Alter and John D. Caputo. *Philosophy Today* 20: 267–28