

Do Artifacts Have Morality?

Introduction

How do we come to understand the moral dimension of technology?¹ Now that we have seen that technologies have moral relevance, and that ethics needs to expand its “humanist focus” to take this into account, the question rises how to conceptualize the morality of technology. What could it imply to say that technologies have a moral dimension? Do the examples that we have seen so far—ultrasound, speed bumps, cell phones—urge us to consider technologies to be moral entities, even moral agents? Or are there other ways to conceptualize the morality of technological artifacts?

Approaching things in moral terms is not a self-evident enterprise. It goes against the grain of the most basic assumptions in ethical theory. After all, it would be foolish to blame a technology when something immoral happens. It does not make sense to condemn the behaviour of a gun when somebody has been shot; not the gun but the person who fired it needs to be blamed. Tsjalling Swierstra is a good representative of such hesitations regarding “moralizing things.” He discusses how the moral community has been expanded many times since classical antiquity. “Women, slaves, and strangers were largely or entirely devoid of moral rights,” but “over time all these groups have been admitted” (Swierstra 1999, 317).² The current inclination to also grant things access to the moral community, however, goes too far, he argues from the two predominant ethical positions: deontology and consequentialism.

Consequentialist ethics evaluates actions in terms of the value of their outcomes. When the positive consequences outweigh the negative ones, an action can be called morally correct. From this perspective, Swierstra says, things can indeed be part of a moral practice, since they can incite human beings to behave morally—and from a consequentialist perspective it is only the result that counts. But things can do this only when human beings use

them for this purpose. Things themselves are not able to balance the positive and negative aspects of their influence on human actions against each other. They can only serve as instruments, not as fully fledged moral agents that are able to render account for their actions.

Deontological ethics is directed not at the consequences of actions but at the moral value of the actions themselves. From a Kantian perspective, for instance, the morality of an action depends on whether the agent has intended to act in accord with rationally insightful criteria. Artifacts, of course, are not capable of taking up such considerations. Moreover, if they incite human beings to act in ways that are morally right from a deontological point of view, these actions are not results of a rationally insightful moral obligation but simply a form of steered behavior.

This means that both from a deontological and a consequentialist perspective, artifacts can only be *causally* responsible for a given action, not *morally*. Artifacts do not possess intentions, and therefore they cannot be held responsible for what they “do.” In Swierstra’s words: “Compelling artifacts, therefore, are not moral actors themselves, nor can they make humans act truly morally. Therefore . . . there is no reason to grant artifacts access to the moral community.” (Swierstra 1999).

I share Swierstra’s hesitations regarding a too radically symmetrical approach to humans and things (cf. Verbeek 2005b, 214–17). Yet the argument that things do not possess intentionality and cannot be held responsible for their “actions” does not justify the conclusion that things cannot be part of the moral community. For even though they don’t do this intentionally, things do mediate the moral actions and decisions of human beings, and as such they provide “material answers” to the moral question of how to act. Excluding things from the moral community would require ignoring their role in answering moral questions—however different the medium and origins of their answers may be from those provided by human beings. The fact that we cannot call technologies to account for the answers they help us to give does not alter the fact that they do play an actively moral role. Take technology away from our moral actions and decisions and the situation changes dramatically. Things can be seen as part of the moral community in the sense that they help to shape morality.

But how to account for this moral role of technology in ethical theory? As stated in chapter 1, to qualify as a moral agent in mainstream ethical theory requires at least the possession of *intentionality* and some degree of *freedom*. Both requirements seem problematic with respect to artifacts—at least at first sight. Artifacts do not seem to be able to form intentions, and neither do

they possess any form of autonomy. Yet both requirements for moral agency deserve further analysis. From the amodern approach set out in chapter 2, the concept of agency—including its aspects of intentionality and freedom—can be reinterpreted in a direction that makes it possible to investigate the moral relevance of technological artifacts in ethical theory. This will be the main objective of this chapter. First, I will discuss the most prominent existing accounts of the moral significance of technological artifacts. After that, I will develop a new account in which I expand the concept of moral agency in such a way that it can do justice to the active role of technologies in moral actions and decisions.

The Moral Significance of Technological Artifacts

The question of the moral significance of technological artifacts has popped up every now and then during recent decades. Several accounts have been developed, all of which approach the morality of technology in different ways. I will discuss the most prominent positions as a starting point for developing a philosophical account of the morality of technological artifacts.

LANDON WINNER: THE POLITICS OF ARTIFACTS

In 1980 Langdon Winner published his influential article “Do Artifacts Have Politics?” In this text, which was later reprinted in his 1986 book *The Whale and the Reactor*, Winner analyzed a number of “politically charged” technologies. The most well-known example he elaborated concerns a number of “racist” overpasses in New York, over the parkways to Jones Beach on Long Island. These overpasses, designed by architect Robert Moses, were deliberately built so low that only cars could pass beneath them, not buses. This prevented the African American population, at that time largely unable to afford cars, from accessing Jones Beach. Moses apparently had found a material way to bring forth his political convictions. His bridges are *political* entities. The technical arrangements involved *preceded* the use of the bridges. Prior to functioning as instruments to allow cars to cross the parkways, these bridges already “encompass[ed] purposes far beyond their immediate use” (Winner 1986).

Winner’s analysis obtained the status of a “classic” in philosophy of technology and in science and technology studies—even though it became the focus of controversy in 1999, when Bernward Joerges published the article “Do Politics Have Artefacts?” (Joerges 1999). In this article he showed that

Jones Beach can also be reached via alternative routes and that Moses was not necessarily more racist than most of his contemporaries. The controversy, however, did not take away the force of Winner's argument. Even as a thought experiment, the example shows how material artifacts can have a political impact—and in this case, a political impact with a clearly moral dimension (see Woolgar and Cooper 1999; Joerges 1999).

The low-hanging overpasses are not the only example Winner elaborated. For Winner, the political dimension of artifacts reaches further than examples like this, in which technologies actually embody human intentions in a material way. Technologies can also have political impact without having been designed them to do so. Many physically handicapped people can testify to this—unintentionally the material world quite often challenges their ability to move about and to participate fully in society.

To elaborate the nonintentional political dimensions of technological artifacts, Winner discusses the example of mechanical tomato harvesters. These machines have had an important impact on tomato-growing practices. Because of their high cost, they require a concentrated form of tomato growing, which means that once they are in use small farms have to close down. Moreover, new varieties of tomatoes need to be bred that are less tasty but can cope with the rough treatment the machines give them. There was never an explicit intention to make tomatoes less tasty and to cause small farms to shut down—but still these were the political consequences of the mechanical tomato harvester.

The example of Moses's bridges shows that technologies can have an impact that can be morally evaluated—the first kind of moral relevance of technologies. Moreover, the example of the tomato harvester shows that such impacts can occur without human beings explicitly intending them—they are in a sense “emergent,” which suggests a form of “autonomy” of technology, albeit without a form of consciousness or intentionality behind it. Technologies, according to Winner, are “ways of building order in our world.” Some technologies bring about this order at the intentional initiative of human beings, serving as “moral instruments” like Moses's bridges, and other technologies give rise to unexpected political impacts.

Winner's account is highly illuminating, yet in the context of this study his analysis leaves many knots untied. Showing that technologies can have a *politically* relevant impact on society, even when this impact was not intended by their designers, does not yet reveal how technologies can also have a *moral* impact. Moreover, we are still in the dark about the *ways in which* this impact comes, and an understanding of this is needed if we are to link me-

diation theory to ethical theory. Winner paved the way, but we need a more detailed account of the roles of technologies in moral actions and decisions if we are to grasp their moral significance.

BRUNO LATOUR:
THE MISSING MASSES OF MORALITY

A second prominent voice in the discussion about the moral significance of technological artifacts is the French philosopher and anthropologist Bruno Latour. In 1992 he published an influential article titled “Where Are the Missing Masses? The Sociology of a Few Mundane Artifacts.” In this text he elaborates the idea that morality should not be considered a solely human affair. Everyone complaining about the alleged loss of morality in our culture should open their eyes. Rather than looking only among people, they should direct their attention toward material things too. The moral decision about how fast one drives, for example, is often delegated to speed bumps in the road, which tell us to slow down. In some cars, blinking lights and irritating sounds remind us to fasten our seat belts. Automatic door closers help us to politely shut the door after entering a building. The “missing masses” of morality are not to be found among people but in things.

By attributing morality to material artifacts, Latour deliberately crosses the boundary between human and nonhuman reality. For Latour, this boundary is a misleading product of the Enlightenment. The radical separation of subject and object that is one of the cornerstones of Enlightenment thinking prevents us from seeing how human and nonhuman entities are always intertwined. Latour understands reality in terms of networks of agents that interact in manifold ways, continually translating each other. These agents can be both human and nonhuman. Nonhumans can act too; they can form “scripts” that prescribe that their users act in specific ways, just as the script of a movie tells the actors what to do and say at what place and time. Neither the intentions of the driver nor the script of the speed bumps in the road exclusively determines the speed at which we drive near a school. It is the network of agents in which a driver is involved which determines his or her speed.

Ten years later, Latour augmented his analysis in an article titled “Morality and Technology: The End of the Means.” Here he shows how inadequate it is to think that technologies belong in the realm of means while human beings inhabit the realm of ends. This view results in what he calls an “archaic split between moralists in charge of the ends and technologists controlling

the means” (Latour 2002). Latour proposes instead to understand technology in terms of the notion of *fold*. In technical action, time, space, and the type of “actants” are folded together. Technologies cross space and time. A hammer, for instance, “keeps folded heterogeneous temporalities, one of which has the antiquity of the planet, because of the mineral from which it has been moulded, while another has the age of the oak which provided the handle, while still another has the age of the 10 years since it came out of the German factory which produced it for the market” (ibid., 249). The same holds true for *space* here: “the humble hammer holds in place . . . the forests of the Ardennes, the mines of the Ruhr, the German factory, the tool van which offers discounts on every Wednesday on Bourbonnais streets,” et cetera.

By “the type of actants,” the third element that is folded into technical action, Latour means that both human and nonhuman agents are involved and help to shape each other. Technologies should not be understood merely in terms of functionality, for this would limit us to seeing only how human intentions can be realized with the help of nonhuman functionalities serving only as means of extension. Technologies are not simply *used* by humans—they help to *constitute* humans. A hammer “provides for my fist a force, a direction and a disposition that a clumsy arm did not know it had” (ibid., 249). In the same way, speed bumps are not simply neutral instruments that fulfill the function of slowing down drivers. “What they exactly do, what they suggest, no one knows, and that is why their introduction in the countryside or in towns, initiated for the innocent sake of function, always ends up inaugurating a complicated history, overflowing with disputes, to the point of ending up either at the State Council or at the hospital” (ibid., 250). Technologies are not intermediaries, helping human intentions to be realized in the material world; they are *mediators* that actively help to shape realities. Technologies do not merely provide means but also help to form new ends; they do not provide functions but make possible detours. “Without technologies, humans would be contemporaneous with their actions, limited solely to proximal interactions. . . . Without technological detours, the properly human cannot exist” (ibid., 252).

The moral significance of technologies, for Latour, is part of this phenomenon of folding. Morality is a “regime of mediation” as well (ibid., 254). We usually recognize morality in the form of *obligation*, but this is not the only form it can take, since it “derives just as much from contract, from religious events, . . . from chains of references, from the law,” et cetera (ibid., 254). Rather than being a merely human affair, morality is to be found in nonhuman entities as well. “Of course, the moral law is in our hearts, but it is also in our apparatuses. To the *super-ego* of tradition we may well add the

under-ego of technologies in order to account for the correctness, the trustworthiness, the continuity of our actions” (ibid., 253–54). This “under-ego” is present in the speed bumps that tell us how fast to drive, or the coin locks on supermarket carts, demanding that we put the carts back in their rack rather than leaving them beside our parking place.

This does not imply, to be sure, that we need to understand technologies as moral agents *in themselves*. “In themselves” entities are quite meaningless anyway—they are given a character in the relations in which they function. In Latour’s words, “Nothing, not even the human, is for itself or by itself, but always *by other things* and *for other things*” (ibid., 256; emphasis in original). Both morality and technology are “ontological categories” for Latour: “the human comes out of these modes, it is not at their origin” (ibid., 256). Technologies help to constitute humans in specific configurations—including the moral character of our actions and decisions.

ALBERT BORGMANN:

TECHNOLOGY AND THE GOOD LIFE

North American philosopher of technology Albert Borgmann has proposed a third position to describe the moral significance of technology. He has developed a neo-Heideggerian theory of the social and cultural role of technology. In this theory, he elaborates how our culture is ruled by what he calls the “device paradigm.” According to Borgmann, the technological devices that we use call for a quite different way of taking up with reality than did pretechnological “things.” While “things”—like water wells, fireplaces, musical instruments—evoke practices in which human beings are engaged with reality and with other people, devices primarily evoke disengaged consumption.

Borgmann understands devices as material *machineries* that deliver consumable *commodities*—for example, the boiler and radiators of a heating installation form a machinery that delivers warmth as a commodity. Devices ask for as little involvement as possible; they create the *availability* of commodities by keeping their machinery in the background as much as they can and putting their commodities in the foreground. Against this, “things” do not separate machinery from commodity. Rather, they *engage* people. Using a fireplace, for instance, requires people to collect and chop wood, to clean the hearth regularly, to gather around the fireplace to enjoy the warmth it gives, and so on.

In his article “The Moral Significance of Material Culture,” Borgmann explains how his theory of the device paradigm makes visible a moral dimension in material objects. He focuses on the role of material culture in human practice, and shows how “material culture constrains and details practice

decisively” (Borgmann 1995, 85). In line with his device paradigm, he makes a distinction between two kinds of reality: one commanding, the other disposable. While a traditional musical instrument is a commanding thing, one that requires a lot of effort and skill and needs to be “conquered,” a stereo literally puts music at our disposal. The quality of sound can be even better than that of a live performance, but the music’s presence is not commanding like that of the music performed live by a musician. According to Borgmann, the device paradigm increasingly replaces commanding reality with disposable reality. In his book *Real American Ethics*, Borgmann elaborates the concept of *moral commodification* to analyze this phenomenon: “*a thing or practice gets morally commodified when it is detached from its context of engagement with a time, a place, and a community and it becomes a free-floating object*” (Borgmann 2006, 152; italics in original).

We find the moral significance of the material culture in its role in shaping human practices. While commanding reality “calls forth a life of engagement that is oriented within the physical and social world,” disposable reality “induces a life of distraction that is isolated from the environment and from other people” (ibid., 92). Human practices take place not in an empty space but in a material environment—and this environment helps to shape the quality of these practices. “If we let virtue ethics with its various traditional and feminist variants stand in for practical ethics, we must recognize that virtue, thought of as a kind of skilled practice, cannot be neutral regarding its real setting. Just as the skill of reading animal tracks will not flourish in a metropolitan setting, so calls for the virtues of courage and care will remain inconsequential in a material culture designed to produce a comfortable and individualist life” (Borgmann 1995, 92).

Even if we do not entirely follow Borgmann in his rather gloomy approach to technology—I think there is engaging technology as well: see Verbeek 2005b—his position highlights a significant form of the moral relevance of technology. Material objects, to summarize his position, help to shape human practices. And because the quality of these practices is ultimately a moral affair, material objects have direct moral relevance. Technological devices and nontechnological “things” help to shape the ways we live our lives—and the question of “the good life” is one of the central questions in ethics. Human actions and human life do not take place in a vacuum but in a real world of people and things that help to shape our actions and the ways we live our lives. And therefore, the good life is not formed only on the basis of human intentions and ideas but also on the basis of material artifacts and arrangements. Technologies provide a setting for the good life.

LUCIANO FLORIDI AND J. W. SANDERS:
ARTIFICIAL MORAL AGENCY

A radically different but equally interesting approach was elaborated in 2004 by Luciano Floridi and J. W. Sanders in their influential publication “On the Morality of Artificial Agents.” Their article deals with the question to what extent artificial agents can be moral agents. Rather than focusing on the moral significance of technologies in general, they focus on *intelligent* technologies that could actually qualify as “agents.” Examples of such artificial agents are expert systems that assist people in making decisions, driving assistants that help people to drive their cars, and automatic thermostats in houses.

The approach Floridi and Sanders develop is so interesting because they give an account of artificial moral agency in which moral agents do not necessarily possess free will or moral responsibility. This way, they take away the obvious objection that technologies, lacking consciousness, can never be moral agents as human beings are. It is crucial to Floridi and Sanders’s analysis that they explicitly choose an adequate “level of abstraction” at which it becomes possible and meaningful to attribute morality to artificial agents—such an abstraction is needed in order to avoid the obvious objection that artifacts cannot have agency as humans do. As criteria for *agenthood*, therefore, Floridi and Sanders use “interactivity (response to stimulus by change of state), autonomy (ability to change state without stimulus) and adaptability (ability to change the ‘transition rules’ by which state is changed).” This implies that a system that interacts with its environment but is also able to act without responding to a stimulus and has the ability to learn how to “behave” in different environments could qualify as an agent. They use the ability to cause good or evil as the criterion for *morality*: “An action is said to be morally qualifiable if and only if it can cause moral good or evil. An agent is said to be a moral agent if and only if it is capable of morally qualifiable action” (Floridi and Sanders 2004, 12).

Their approach reveals what Floridi and Sanders call “aresponsible morality” (*ibid.*, 13). They consider intentions—“intentional states,” in the vocabulary of the analytic tradition from which they work—as a “nice but unnecessary condition” for moral agency. The only thing that matters for them is whether the agent’s actions are “morally qualifiable”—that is, whether they can cause moral good or evil. However, Floridi and Sanders do not aim to declare the concept of responsibility obsolete. Rather, they separate it from moral agency as such, which opens for them the space needed to clarify the role responsibility actually plays in morality (*ibid.*, 20).

It is an important contribution to understanding the moral significance of technology to reveal how normative action is possible even when there is no moral responsibility involved. The approach of Floridi and Sanders offers an answer to an obvious objection against attributing morality to technologies, that technologies do not have consciousness and therefore cannot “act” morally. If moral agency can be adequately understood in terms of showing “morally qualifiable” action, greater justice can be done to the moral relevance of technological artifacts than mainstream ethical theory allows. The problem remains, however, how to deal with forms of artifact morality that cannot be considered results of artificial agency. How should we deal with ultrasound imaging, for instance, in terms of this framework? And with Winner’s example of Moses’s bridges? These examples do not meet Floridi and Sanders’s criteria for agency—but they do actively contribute to moral actions and have impacts that can be assessed in moral terms. However illuminating Floridi and Sanders’s position is, we need more if we are to understand the moral relevance of technology. Artificial moral agency constitutes only a part of the moral relevance of technology; we need a broader understanding of “artifactual morality.”

Moral Mediation

In the positions discussed above, various approaches to the morality of technology played a role. All authors agree that technologies are morally significant because they have a morally relevant impact in society. Technologies help to shape actions, inform decisions, and even make their own decisions, as some information technologies do; in all cases, they have an impact that can be assessed in moral terms. Yet there appear to be many ways to understand this “morally relevant impact.”

MORAL INSTRUMENTALISM

The first and minimum option is to approach technologies as moral instruments. Winner’s bridges, Latour’s speed bumps and door closers, and Hans Achterhuis’s turnstiles are examples of technologies that bring about a moral effect that humans seek to achieve through them. From the approach of technological instrumentalism, artifacts like these provide human beings with means to realize their moral ends: racial segregation, safety on the road, neatly closed doors, paying passengers on trains.

However, this approach is far too shallow to do justice to the complex moral roles of technologies. And to be sure, none of these authors actually

think that technologies are merely neutral means to realize human moral intentions. Winner's example of the tomato harvester, for instance, shows that technologies can have *unintended* consequences. Latour would readily acknowledge that speed bumps can invite local skaters to engage in behavior that actually diminishes rather than enhances traffic safety, and that automatic door closers might also embody forms of impoliteness by slamming doors in people's faces and making it difficult for elderly people to open them. Even though technologies can certainly function as moral instruments that enable human beings to generate specific moral effects, they always do more than this.

The behavior of technologies is never fully predictable—a thought that is vividly illustrated in Edward Tenner's book *Why Things Bite Back* (1996). Moral instrumentalism is too poor a position to account for the moral relevance of technology. Technologies inevitably enter into unforeseeable relations with human beings in which they can develop unexpected morally relevant impacts. Obstetric ultrasound is a good example again: this technology was not designed to organize new moral practices, and yet it plays an active role in raising moral questions and setting the framework for answering them.

TECHNOLOGIES AS MORAL AGENTS

Does this imply that we should take the opposite direction and approach technologies as moral agents? Should we simply start to acknowledge the fact that technologies can act morally? This is the position Floridi and Sanders defend. From the level of abstraction they elaborate, an entity is a moral agent when it is able to cause good or evil. This approach allows them to conclude that artificial agents can qualify as moral agents because they can "do" evil or good by producing effects that can be assessed morally. This approach is highly interesting and relevant, but unfortunately it applies to only a limited set of technologies. Not all morally significant technologies could qualify as agents based on Floridi and Sanders's criteria of interactivity, autonomy, and adaptability. Ultrasound imaging, for instance, would fail the criterion of autonomy, yet it has a moral impact beyond what human beings designed into it.

The position of Bruno Latour also attributes agency to technologies, but in a radically different way. While Floridi and Sanders focus on *artificial* agency, one could say that Latour focuses more broadly on *artifactual* agency. In his symmetrical approach, both humans and nonhumans can be agents, and nonhuman agents can also embody morality by helping to shape moral action. Yet, as indicated above, from a Latourian point of view it would not

be adequate to attribute moral agency to technologies “themselves”—as if “agency” were some intrinsic property of technology. Latour’s claims that nonhumans can be agents as well and that there is morality in technology need to be read in the context of his actor-network theory, in which all entities are understood relationally. From this perspective, technologies do not have moral agency *in themselves*; rather, when humans use technologies, the resulting moral agency is not exclusively human but incorporates nonhuman elements as well. Contrary to the position of Floridi and Sanders, for Latour technologies only “have” agency and morality in the context of their relations with other agents.

MORAL MEDIATION

Actually, Latour’s approach occupies a third position with respect to the moral relevance of technology. Rather than moral instruments or moral agents, Latour’s work makes it possible to see technologies as moral *mediators*. This position does justice to the active moral role of technologies in moral actions and decisions, without reducing this role entirely to human intentions. At the same time, it avoids characterizing morality as an intrinsic property of the technologies themselves. By mediating human experiences and practices—as elaborated in chapter 1—technologies mediate moral decisions and actions. Technologies help us to phrase moral questions and find answers to them, and they guide our actions in certain directions.

The notion of “mediator” expresses both the active moral role of technologies and the relational character of this moral role: they mediate, rather than being some kind of neutral “intermediary,” but mediators can function only in the context of an environment for and in which they mediate. The moral significance of Latour’s speed bumps and Winner’s overpasses can be understood best in terms of moral mediation. Understanding them as moral instruments for realizing the racist intentions or safety ambitions of city planners appeared to fall short, because this does not recognize the *unintended* roles these artifacts can play. Understanding them as moral agents would go too far, at least in the sense of being moral agents “in themselves,” capable of moral action. Only in the context of the practices in which they function do their moral roles emerge. Sometimes these roles coincide with the intentions of their designers, sometimes they don’t. In all cases, the moral roles of technologies come about in the context of their relations with their users and the environment in which they function.

Borgmann’s approach to the moral significance of technology is an inter-

esting supplement to the notion of moral mediation. He broadens the discussion from action-oriented ethics to the classical ethical question of the good life by focusing on technologies as providing a material setting for the good life. In Borgmann's approach, the moral role of technologies is not to be found in the ways technologies help to shape human actions but in how they help to answer the classical question of "how to live." Borgmann's example of the difference between a stereo set and a musical instrument does not revolve around the different actions involved in operating the two but around their roles in shaping a way of life.

By conceptualizing technologies as moral mediators, we can bring the postphenomenological approach to technological mediation into the realm of ethics. As we saw in the example of obstetric ultrasound, technologies-in-use establish a relation between their users and their world. Ultrasound imaging organizes a specific form of contact between expectant parents and unborn child, in which the parents and the child are constituted in specific ways with specific moral roles, responsibilities, and relevance. Along the same lines, larger-scale technologies mediate moral actions and decisions; energy production systems, for instance, help to organize a way of living in which it becomes ever more normal and necessary to use large quantities of energy, and in doing so they help to shape moral decisions regarding how we deal with environmental issues.

To be sure, approaching technologies as moral mediators does not imply that we need to reject Latour's ideas about nonhuman agency. Indeed the notion of moral mediation implies a form of technological agency. Moral mediation always involves an intricate relation between humans and nonhumans, and the "mediated agency" that results from this relation therefore always has a hybrid rather than a "purely nonhuman" character. When technologies are used, moral decisions are not made autonomously by human beings, nor are persons forced by technologies to make specific decisions. Rather, moral agency is distributed among humans and nonhumans; moral actions and decisions are the products of human-technology associations.

The way I use the notion of moral mediation is different from the way Lorenzo Magnani uses it in his book *Morality in a Technological World* (2007). Magnani lays out an approach to morality and technology that is congenial to the approach set out in this book but reaches different conclusions. Because his approach departs from the perspective of cognitive science rather than phenomenology, it cannot take into account the hermeneutic and pragmatic dimensions of technological mediation that are so central to the account developed here. For Magnani, moral mediators mediate moral *ideas*.

In his definition, “moral mediators . . . are living and nonliving entities and processes—already endowed with intrinsic moral value—that ascribe new value to human beings, nonhuman things, and even to ‘non-things’ like future people and animals.” Even though he discusses Latour’s work assentingly (ibid., 25–26), he does not acknowledge that Latour’s actor-network theory radically differs from his cognitive approach. Magnani’s strong focus on *knowledge* as the primordial variable in ethics and in moral mediation is rather remote from Latour’s focuses on practices, interactions, and materiality.

For Latour, and for the postphenomenological approach that uses his work, the cognitive approach makes too sharp a distinction between (subjective) minds that have knowledge and the (objective) world that this knowledge is about. In the approach I follow in this book, morality should not be understood in terms of cognitive “templates of moral doing” (Magnani 2007, 187–93) but in terms of ways of being-in-the-world which have both cognitive and noncognitive aspects and which are technologically mediated in more-than-cognitive ways. In my postphenomenological approach, technological mediation concerns action and perception rather than cognition; and moral mediation is not only about the mediated character of moral ideas but mostly about the technological mediation of actions, and of perceptions and interpretations on the basis of which we make moral decisions.

The concept of moral mediation has important implications for understanding the status of objects in ethical theory. As indicated in my introduction, in mainstream ethical theory “objects” have no place apart from being mute and neutral instruments that facilitate human action. Now that we have seen that technologies actively help to shape moral actions and decisions, we need to expand this overly simplistic approach. The mediating role of technologies can be seen as a form of moral agency—or better, as an element of the distributed character of moral agency.

I will rethink the status of the object in ethical theory in two ways. First, I will offer a “nonhumanist” analysis of two criteria that are usually seen as conditions *sine qua non* for moral agency. An entity can be called a moral agent if it can be morally responsible for its actions, and to be morally responsible, it needs at least (1) *intentionality*—the ability to form intentions—and (2) the *freedom* to realize its intentions. I will show that these two criteria can be reinterpreted along postphenomenological lines in such a way that they also pertain to nonhuman entities. Second, I will investigate the possibility that the predominant ethical approaches propose to take seriously the moral dimension of technologies. By elaborating what role objects could play in

deontological ethics, consequentialism, and virtue ethics, I will create the space needed to take the moral significance of technologies seriously.

Technological Intentionality

The first criterion for moral agency—the possession of intentionality—directly raises a serious problem for anyone who intends to defend some form of moral agency for technology. While agency is not thinkable without intentionality, it also seems absurd to claim that artifacts can have intentions. Yet a closer inspection of what the concept of intentionality can mean in relation to what artifacts actually “do” makes it possible to articulate a form of “technological intentionality.”

The concept of intentionality actually has a double meaning in philosophy. In ethical theory, it primarily expresses the ability to form intentions. In phenomenology, though, the concept of intentionality indicates the directedness of human beings toward reality. Intentionality is the core concept in the phenomenological tradition for understanding the relation between humans and their world. Rather than separating humans and world, the concept makes visible the inextricable connections between them. Because of the intentional structure of human experience, human beings can never be understood in isolation from the reality in which they live. They cannot simply “think” but always think *something*; they cannot simply “see” but always see *something*; they cannot simply “feel” but always feel *something*. As experiencing beings, humans cannot but be directed at the entities that constitute their world. Conversely, it does not make much sense to speak of “the world in itself.” Just as human beings can be understood only through their relation with reality, reality can be understood only through the relation human beings have with it. “The world in itself” is inaccessible by definition, since every attempt to grasp it makes it a “world for us,” as disclosed in terms of our particular ways of understanding and encountering it.

In the context of this discussion of the possibility of “artifactual moral agency,” these two meanings of the concept of intentionality augment each other. The ability to form intentions to act in a specific way, after all, cannot exist without being directed at reality and interpreting it in order to act in it. Actually, the two meanings of intentionality have a relation to each other similar to the relation between the two dimensions of technological mediation that I discerned in chapter 1. The “praxical” dimension, concerning human actions and practices, cannot exist without the “hermeneutical” dimension, concerning human perceptions and interpretations—and vice

versa. Forming intentions for action requires having experiences and interpretations of the world in which one acts.

From the perspective of technological mediation, both forms of intentionality are not as alien to technological artifacts as at first they might seem. As for the phenomenological interpretation of the concept: the work of Ihde shows that the human-world relations that are central in the phenomenological tradition often have a technological character. Many of the relations we have with the world take place “through” technologies or have technologies as a background—ranging from looking through a pair of glasses to reading temperature on a thermometer, from driving a car to having a telephone conversation, from hearing the sound of the air conditioner to having an MRI scan made. Ihde shows that intentionality can work *through* technological artifacts, it can be directed *at* artifacts, and it can even take place against the *background* of them.

In most of these cases—with an exception for human relations that are directed *at* artifacts—human intentionality is *mediated* by technological devices. Humans do not experience the world directly here but via a mediating technology that helps to shape a relation between humans and world. Binoculars, thermometers, and air conditioners help to shape new experiences, either by procuring new ways of accessing reality or by creating new contexts for experience. These mediated experiences are not entirely “human.” Human beings simply could not have such experiences without these mediating devices. This implies that a form of intentionality is at work here—one in which both humans and technologies have a share. And this, in turn, implies that in the context of such “hybrid” forms of intentionality, technologies do indeed “have” intentionality—intentionality is “distributed” among human and nonhuman entities, and technologies “have” the nonhuman part. In such “hybrid intentionalities,” the technologies involved and the human beings who use the technologies share equally in intentionality.

The ethical implications of the second meaning of the concept of intentionality are closely related to those of the first. Intentions to act in a certain way, after all, are always informed by the relations between an agent and reality. These relations, again, have two directions; one pragmatic, the other hermeneutic. Technologies help to shape actions because their scripts evoke given behaviors and because they contribute to perceptions and interpretations of reality that form the basis for decisions to act. In the Netherlands, to give an example in the pragmatic direction, experiments are done with crossings that deliberately include no major road. The script of such crossings contributes to the intention of drivers to navigate extra carefully in order to be able to give priority to traffic from the right (Fryslân Province, 2005).

Genetic diagnostic tests for hereditary breast cancer, as mentioned in chapter 1, are a good example in the hermeneutic direction. Such tests, which can predict the probability that people will develop this form of cancer, transform healthy people into potential patients and translate a congenital defect into a preventable defect: by choosing to have a double mastectomy now, you can prevent breast cancer from developing in the future. Here technologies help to interpret the human body; it organizes a situation of choice and also suggests ways to deal with this choice.

In all of these examples, technologies are morally *active*. They help to shape human actions, interpretations, and decisions that would have been different without these technologies. To be sure, artifacts do not have intentions as human beings do, because they cannot *deliberately* do something. But their lack of consciousness does not take away the fact that artifacts can “have” intentionality in the literal sense of the Latin word *intendere*, which means “to direct,” “to direct one’s course,” “to direct one’s mind.” The intentionality of artifacts is to be found in their directing role in the actions and experiences of human beings. Technological mediation therefore can be seen as a distinctive, material form of intentionality.

There is another element that is usually associated with intentionality, though, and it is one that technologies seem to miss: the ability to form intentions that can be considered *original* or *spontaneous*, in the literal sense of “springing from” or “being originated by” the agent possessing intentionality. Yet the argument above can be applied here as well. For even though because of their lack of consciousness artifacts evidently cannot form intentions entirely on their own, their mediating roles cannot be entirely reduced to the intentions of their designers and users. If they could be, the intentionalities of artifacts would merely be a variant of what John Searle called “derived intentionality” (Searle 1983), entirely reducible to human intentionalities. Quite often, though, as pointed out already, technologies mediate human actions and experiences in ways that were never foreseen or desired by human beings.

Some technologies are used in different ways from those their designers envisaged. The first cars, which could go only 15 km/h, were used primarily for sport and for medical purposes; driving at a speed of 15 km/h was thought to create an environment of “thin air,” which was supposed to be healthy for people with lung diseases. Only after cars were interpreted as a means of long-distance transport did the car come to play its current role in the division between labor and leisure (Baudet 1986). In this case, unexpected mediations come about in specific use contexts. Unforeseen mediations can also emerge when technologies *are* used as intended. The introduction of mobile phones

has led to a different way of dealing with appointments, especially for young people—making plans far in advance for a night out does not make much sense when everyone can call each other anytime to make an ad hoc plan. This change in behavior was not intended by the designers of the cell phone, even though the phone is being used in precisely the context the designers had envisaged. And nobody foresaw that the introduction of the energy-saving lightbulb would actually cause people to use more rather than less energy. Apparently such bulbs are often used in places previously left unlit, such as in a garden or on the front of a building, thereby canceling out their economizing effect (Steg 1999; Weegink 1996).

It seems plausible, then, to attribute a form of intentionality to artifacts—albeit a form that is radically different from human intentionality. The intentional “dimension” of artifacts cannot exist without human intentionalities supporting it; only within the relations between human beings and reality can artifacts play the mediating roles in which their “intending” activities are to be found. For example, when expectant parents face a decision about abortion on the basis of technologically mediated knowledge of the chances that the child will suffer from a serious disease, this decision is not “purely” human, but neither is it entirely induced by technology. The very situation of having to make this decision and the very ways in which the decision is made are coshaped by technological artifacts. Without these technologies, either there would not be a situation of choice or the decision would be made on the basis of a different relation to the situation. Yet the technologies involved do not *determine* human decisions. Moral decision making is a joint effort of human beings and technological artifacts. Technological intentionalities are one component of the eventually resulting intentionality of the “composite agent,” a hybrid of human and technological elements.

Strictly speaking, then, there is no such thing as “technological intentionality”; intentionality is always a hybrid affair involving both human and nonhuman intentions, or, better, “composite intentions” with intentionality distributed among the human and the nonhuman elements in human-technology-world relationships. Rather than being “derived” from human agents, this intentionality comes about in associations between humans and nonhumans. For that reason it could best be called “hybrid intentionality” or “composite intentionality.”

Technology and Freedom

A second requirement that is often connected to moral agency is the possession of freedom. If moral agency entails that an agent can be held mor-

ally responsible for his or her actions, this requires not only that the agent needs to have the intention to act in a particular way but also that he or she has the freedom to realize this intention. Now that we have concluded that artifacts may have some form of intentionality, can we also say that they have *freedom*?

The answer obviously seems to be no. Again, freedom requires the possession of a mind, which artifacts do not have. Technologies cannot be free agents as human beings are. The only degree of freedom that could be ascribed to them is their “ability” to have unintended and unexpected effects, like the increase in energy use brought on by the energy-saving lightbulb. But this is not freedom, of course, in the sense of the ability to choose and to have a relation to oneself and one’s inclinations, needs, and desires. Still, there are good arguments not to exclude artifacts entirely from the realm of freedom.

First of all, even though freedom is obviously required if one is to be accountable for one’s actions, the thoroughly technologically mediated character of our daily lives makes it difficult to make freedom an absolute criterion for moral agency. This criterion might exist in a radical version of Kantian ethical theory, where freedom is understood in terms of autonomy and where the moral subject needs to be kept pure of polluting external influences. But many other ethical theories take into account the situated and mediated character of moral agency. People do not make moral decisions in a vacuum, after all, but in a real world, which inevitably influences them and helps to make them the persons they are. The phenomenon of technological mediation is part of this. Technologies play an important role in virtually every moral decision we make. The decision how fast to drive and therefore how much risk to run of harming other people is always mediated by such things the layout of the road, the power of the car’s engine, the presence or absence of speed bumps and speed cameras. The decision to have surgery or not is most often mediated by all kinds of imaging technologies and blood tests, which help to constitute the body in specific ways and organize specific situations of choice.

Moral agency, therefore, does not require complete autonomy. *Some* degree of freedom can be enough for one to be held morally accountable for an action. And not all freedom is taken away by technological mediations, as the examples of abortion and driving speed make clear. In these examples, human behavior is not determined by technology but rather coshaped by it, with humans still being able to reflect on their behavior and make decisions about it. Nevertheless, we can in no way escape these mediations in our moral decision making. The moral dilemmas of whether to have an abortion and of how fast to drive would not exist in the same way without the technologies

involved in these practices. Such dilemmas are rather *shaped* by technologies. Technologies cannot be defined away from our daily lives. In this respect, technologically mediated moral decisions are never completely “free.” The concept of freedom presupposes a form of sovereignty with respect to technology that human beings simply do not possess.

This conclusion can be read in two distinct ways. The first is that mediation has nothing to do with morality at all. If moral agency requires freedom and technological mediation limits or even annihilates human freedom, only non-technologically mediated situations leave room for morality. Technology-induced human behavior then has a nonmoral character. Actions that are not products of our free will but induced by technology cannot be described as “moral.” This position does not help us much further, though. Denying that technologically mediated decisions can have a moral character throws out the baby with the bathwater, for it prevents us from conceptualizing the undeniably moral dimension of making decisions about unborn life on the basis of ultrasound imaging.

Therefore, an alternative solution to the apparent tension between technological mediation and ethics is needed. Rather than taking freedom from (technological) influences as a *prerequisite* for moral agency, we need to reinterpret freedom as an agent’s ability to *relate* to what determines him or her. Human actions always take place in a stubborn reality, and for this reason, absolute freedom can be attained only if we ignore reality and thus give up the ability to act at all. Freedom is not a lack of forces and constraints; rather, it is the existential space human beings have within which they can realize their existence. Humans have a relation to their own existence and to the ways it is coshaped by the material culture in which it takes place. The materially situated character of human existence *creates* forms of freedom rather than impeding them. Freedom exists in the possibilities that are opened up for human beings so that they might have a relationship with the environment in which they live and to which they are bound.

This redefinition of freedom, to be sure, does not imply that we need to actually attribute freedom to technological artifacts. Yet it does make it possible to take artifacts back into the realm of freedom, rather than excluding them from it altogether. Just as intentionality appeared to be distributed among the human and nonhuman elements in human-technology associations, so is freedom. Technologies “in themselves” cannot be free, but neither can human beings. Freedom is a characteristic of human-technology associations. On the one hand, technologies help to *constitute* freedom by providing the material environment in which human existence takes place and takes its form. And on the other hand, technologies can form associations with hu-

man beings, which become the places where freedom is to be located. Technological mediations create the space for moral decision making. Just like intentionality, freedom is a hybrid affair, most often located in associations of humans and artifacts. In chapter 4, which deals with the role of the technologically mediated subject in ethical theory, I will give a more extensive reinterpretation of the concept of freedom in relation to moral agency and technological mediation.

Material Morality and Ethical Theory

By rethinking the concepts of intentionality and freedom in view of the morally mediating roles of technology, I have dispatched the major obstacles to including technological artifacts in the domain of moral agency. But how does this redefined notion of moral agency relate to mainstream ethical theory? Can it withstand the obvious deontological and consequentialist objections presented by Swierstra? And how does it relate to virtue-ethical approaches?

Let me start by discussing the deontological approach. The deontological argument against attributing moral agency to nonhumans revolves around the fact that objects lack rationality. Applying Kant's categorical imperative—the most prominent icon of deontological ethics—to things immediately makes this clear: “Act only in accordance with that maxim through which you can at the same time will that it become a universal law” (Kant [1785] 2002, 37). Technologies are obviously not able to follow this imperative—unless maybe they embody an advanced form of artificial intelligence. Yet that does not necessarily imply that there is no room for nonhuman moral agency in deontological ethics at all. It implies only that technologies cannot have moral agency *in themselves*. The position I have laid out in this chapter is based on the idea that the moral significance of technology is to be found not in some form of independent agency but in the technological *mediation* of moral actions and decisions—which needs to be seen as a form of agency itself.

Technologically mediated moral agency is not at odds with the categorical imperative at all. After all, technological mediation does not take away the rational character of mediated actions and decisions. A moral decision about abortion after having had an ultrasound scan can still be based on the rational application of moral norms and principles—and even on the Kantian question whether the maxim used could become a universal law. However, the rational considerations that play a role in the decision may be thoroughly technologically mediated. As we saw, the ways in which ultrasound constitutes the fetus and its parents help to shape the moral questions that are relevant and also the answers to those questions. The moral decision to have an

abortion or not is still made by a rational agent—but it cannot be seen as an *autonomous* decision. Human beings cannot alter the fact that they have to make moral decisions in interaction with their material environment.

Latour made an attempt to expand Kant's moral framework to the realm of nonhumans by providing a "symmetrical" complement to the categorical imperative. In *Groundwork for the Metaphysics of Morals* Kant actually gave several formulations of his categorical imperative. While the formulation given above is the so-called first formulation, Latour focused on the second, which reads "Act so that you use humanity, as much in your own person as in the person of every other, always at the same time as end and never merely as means" (Kant [1785] 2002, 46–47). In his book *Politics of Nature* Latour augmented this formulation with the imperative to act in such a way that you use *nonhumans* always at the same time as ends and never merely as means (Latour 2004, 155–56). In this way he tried to make room for ecological issues in ethical thinking; such issues by definition require us to bring nonhuman reality into the heart of ethical reflection.

This reformulation of the categorical imperative, though, approaches nonhumans primarily as moral *patients*, while the approach I develop here is primarily interested in nonhumans as moral *agents*—or, better, as active moral mediators. But Latour's reformulation leaves room for this other interpretation as well. "Using nonhumans at the same time as means and as ends," after all, can imply that using a technological artifact brings in not only means but also "ends"—the ends that are implied in the means of technology. Because of their mediating capacities, after all, technologies belong not only to the realm of means but also to the realm of ends (cf. Latour 1992b). And this makes possible a paraphrase of yet another formulation of the categorical imperative. Kant's third formulation reads "Every rational being must act as if it were through its maxims always a legislative member in a universal realm of ends"—but the approach of technological mediation makes clear that not only "rational beings" but technologies as well are "members in the universal realm of ends."

With regard to *consequentialist* ethics, the same line of argument applies. Utilitarianism, as the predominant variant of consequentialism, seeks to assess the moral value of actions in terms of their utility. This utility can be located in various things: the promotion of happiness (Jeremy Bentham's "greatest happiness for the greatest number of people"), the promotion of a plurality of intrinsically valuable things, or the fulfillment of as many preferences as possible. Obviously, technological artifacts are generally not able to perform an assessment like this—with the possible exception of artificially

intelligent devices. Yet such assessments are not products of autonomous human beings either. In our technological culture, the experience of happiness, the nature of intrinsically valuable things (like love, friendship, and wisdom), and the specific preferences people have are all technologically mediated.

Making a utilitarian decision about abortion, to return again to this example, clearly illustrates this. A hedonistic-utilitarian argument in terms of happiness, for instance, inevitably incorporates a thoroughly technologically mediated account of happiness. The medical norms in terms of which the fetus is represented, and the fact that ultrasound makes expectant parents responsible for the health of the unborn child, changes how abortion is connected to the happiness of the people involved here. Similarly, a preference-utilitarian argument will rest upon preferences that are highly informed by the technology involved. Preferences to have a healthy child, to avoid feelings of guilt if a child is born with a serious disease, or to prevent a seriously ill child from threatening the happiness of other children in the family—to mention just a few preferences that are likely to play a role in this case—could not exist without the whole technological infrastructure of antenatal diagnosis and abortion clinics.

From a *virtue-ethical* position it is much easier to incorporate the moral roles of technologies. As Gerard de Vries has noted (de Vries 1999), this pre-modern form of ethics does not focus on the question of “how should I act” but on the question of “how to live.” It does not take as its point of departure a subject that asks itself how to behave in the outside world of objects and other subjects. It rather focuses on “life”—human existence, which inevitably plays itself out in a material world. From this point of view, it is only a small step to recognize with de Vries that in our technological culture, not only ethicists and theologians answer this question of the good life but also all kinds of technological devices tell us “how to live” (ibid.). The next chapter, in which I will discuss the technologically mediated moral subject, will give a more extensive elaboration of the importance of classical virtue-ethical conceptions for understanding the moral significance of technologies.

Conclusion: Materiality and Moral Agency

Technologies appear to be thoroughly moral entities—yet it is very counterintuitive to attribute morality to inanimate objects. In this chapter I have developed a way to conceptualize the moral significance of technological artifacts which aims to do justice to both of these observations by developing the concept of moral mediation in the context of ethical theory. This concept

makes it possible to address the moral significance of technologies without reverting to a form of animism that would treat them as full-blown moral agents. The example of the gun, used at the beginning of this chapter, can also serve as a conclusion. Now we can come to a more nuanced picture of the moral significance of a gun. Rather than simply stating that it would be ridiculous to blame a gun for a shooting and using this as an argument against the moral agency of technology, we can find our way to a more sophisticated understanding via the concept of moral mediation. After all, it would not be satisfactory either to completely deny the role of the gun in a shooting.

This related to an example explored by Latour: the debate between the National Rifle Association in the United States and its opponents. In this debate, those opposing the virtually unlimited availability of guns use the slogan “Guns Kill People,” while the NRA replies with the slogan “Guns don’t kill people; people kill people” (Latour 1999, 176). The NRA position seems to be most in line with mainstream thinking about ethics: if someone is shot, nobody would ever think of holding the gun responsible. Yet the antigun position also has a point: in a society without guns, fewer fights would result in murder. The problem in this discussion, however, is the separation of guns and people—of humans and nonhumans. Only on the basis of such a modernist approach does the question “can technologies have moral agency?” become a meaningful problem. From an amodern perspective, as I suggested in chapter 2, this question leads us astray. It seeks to find agency in technology *itself*, isolated from its relations with other entities, human and nonhuman.

A gun is not a mere instrument, a medium for the free will of human beings; it helps to define situations and agents because it offers specific possibilities for action. A gun constitutes the person holding the gun as a potential killer and his or her adversary as a potential lethal victim. Without denying the importance of human responsibility in any way, we can conclude that when a person is shot, agency should not be located exclusively in either the gun or the person shooting, but in the assembly of both. The English language even has a specific “amodern” word for this example: *gunman*, as a hybrid of human and nonhuman elements. The gun and the man form a new entity, and this entity does the shooting.

The example illustrates the main point of this chapter: in order to understand the moral significance of technology, we need to develop a new account of moral agency. The example does not suggest that artifacts can “have” intentionality and freedom, just as humans are supposed to have. Rather, it shows that (1) intentionality is hardly ever a purely human affair—most often it is a matter of human-technology associations; and (2) freedom should

not be understood as the absence of “external” influences on agents but as a practice of *dealing* with such influences or mediations. Chapter 4 will further explore this new understanding of moral agency—not from the perspective of the object but from the point of view of the technologically mediated *subject*.