Philos. Technol.
DOI 10.1007/s13347-013-0130-y

SPECIAL ISSUE

Bio-machine Hybrid Technology: A Theoretical Assessment and Some Suggestions for Improved Future Design

Tom Froese

Received: 8 March 2013 / Accepted: 12 August 2013

Abstract In sociology, there has been a controversy about whether there is any essential difference between a human being and a tool, or if the tool–user relationship can be defined by co-actor symmetry. This issue becomes more complex when we consider examples of AI and robots, and even more so following progress in the development of various bio-machine hybrid technologies, such as robots that include organic parts, human brain implants, and adaptive prosthetics. It is argued that a concept of autonomous agency based on organismic embodiment helps to clarify the situation. On this view, agency consists of an asymmetrical relationship between an organism and its environment, because the continuous metabolic and regulatory activity of the organism gives rise to its own existence, and hence its specific behavioral domain. Accordingly, most (if not all) of current technologies are excluded from the class of autonomous agents. Instead, they are better conceptualized as interfaces that mediate our interactions with the world. This has important implications for design: Rather than trying to help humans to achieve their goals by duplicating their agency in artificial systems, it would be better to empower humans directly by enhancing their existing agency and lived experience with technological interfaces that can be incorporated into their embodiment. This incorporation might be especially facilitated by bio-machine hybrid technology that is designed according the principles of biological autonomy and multi-agent coordination dynamics.

Keywords Embodiment · Consciousness · Human–computer interface · Agency

T. Froese (✉)
Departamento de Ciencias de la Computación Instituto de Investigaciones en Matemáticas Aplicadas y en Sistemas, Universidad Nacional Autónoma de México, Mexico City, Mexico
e-mail: t.froese@gmail.com
URL: http://froese.wordpress.com

T. Froese
Centro de Ciencias de la Complejidad, Universidad Nacional Autónoma de México, Mexico City, Mexico

Published online: 30 August 2013

1 Introduction—Background Motivation

Technology has long been in the blindspot of philosophy, going back to Ancient Greek philosophers who privileged reason over praxis (Stiegler 1998). However, in modern times, it has been increasingly recognized that the phenomenon of humanity cannot be understood apart from the phenomenon of technology (Leroi-Gourhan [1964–1965] 1993). In cognitive science, this view has recently been popularized by the extended mind hypothesis, which argues that cognition can be distributed over the whole tool–user relationship (Clark 2003, 2008; Clark and Chalmers 1998). One influential idea of this hypothesis, which has become known as the "parity principle," is that, if using a tool in the world has the same function as if it were accomplished by some cognitive processes inside the mind/brain, then we should accept that such tool-use is an external part of cognition. However, this specific view of cognitive extension is not without problems, some of which stem from its functionalist and internalist premises (Di Paolo 2009a; Froese et al. 2013). Going beyond this functionalist focus on cognition, the enactive approach to cognitive science grounds lived experience in organismic embodiment (Thompson 2007), whereby mind is seen from the start as a relational phenomenon that is actively brought about by the precarious self-production of an organism in relation to its environment. We will discuss this enactive approach in more detail later; here, it is worth highlighting that this approach has also recently been inflected by the claim that technology is anthropologically constitutive (Stewart et al. 2010). In other words, there is a convergence in the social and mind sciences that technology is an inseparable part of what makes us human.

Given this technological dimension of being human, it quickly becomes evident that the normative evaluation of prospective new kinds of technologies, including various bio-machine hybrid technologies, is a philosophical task that is as profoundly important as it is beset by difficulties (Jonas 1979). The creation of new technologies is inseparable from the creation of new ways of being human, but these ways of being are difficult to predict in advance. The philosopher of technology Ihde has analyzed this dilemma in terms of a "prognostic antinomy":

> If philosophers are to take any normative role concerning new technologies, they will find from both within the structure of technologies as such, and compounded historically by unexpected uses and unintended consequences, that technologies virtually always exceed or veer away from 'intended' design. (Ihde 1999, p. 45)

Ihde builds on Heidegger's ([1927] 1962) famous phenomenological analysis of the hammer, which showed that the essence of a tool is defined by what it is used for, and that this always occurs within a context of practical significance. Since the usage of a tool is not intrinsic to its intended design, Ihde argues that all technologies display ambiguous, multistable possibilities of usage. All of us are familiar with the malleability of technology, so much so that we mostly take it for granted in our everyday practical dealings with the world. The indeterminate usage of technology has also been made explicit in the popular imagination. We can recall, for example, that, in the South African film *The Gods Must Be Crazy* (1980), a group of San people, living isolated in the Kalahari Desert, come across a discarded Coca-Cola bottle for the first time and quickly notice that it affords them many (to us rather unusual and funny) uses, such as making music and efficiently pounding vegetables. But this new "technology" also starts to bring about undesirable behaviors with detrimental effects on social relations, such as envy and violence, until it is decided that the technology should be banished from society. The unpredictability of the uses of new technology is of importance because it intimately related to new ways of being human.

The inherent fluidity of the uses of technology and the technological constitution of humanity should make us doubly cautious. In particular, because the possible uses of technology cannot be limited to its intended uses, an exhaustive risk–benefit assessment and ethical evaluation in advance of its actual dissemination is impossible in principle. For example, it might have once seemed like a good idea to create cocaine out of coca leaves—at the time, it was hailed as the new wonder drug—but this same technology eventually led to the creation of one of the world's biggest health problems today, with the lives of countless people being determined by their addiction. This may be an extreme example, but who can really tell what the proliferation of new technologies entails? And yet, despite the dangers that follow from the profound uncertainties about future uses (and users), a stream of new technologies for various uses will continue to arrive in our world, for better or for worse. Once we accept that technogenesis is constitutive of anthropogenesis, this is simply a fact of human nature (Leroi-Gourhan [1964–1965] 1993). Given humanity's essential dependence on this technically constituted lifeworld, there is no return to the idealized vision of a tool-free Garden of Eden; no such thing ever existed. The origins of humans and tools is one of co-dependent emergence and joint development (Stiegler 1998). However, judging by the profound global crises facing humanity today, spanning sociological, economical, and ecological factors, there is clearly a need for us to be better able to coax technological evolution into more favorable directions, i.e., toward those that tend to minimize inflicting damage to life, mind, and society, while increasing the likelihood of benefits.[1]

Typically, this notion of regulating the evolution of technology is discussed in terms of the need for ever more resources, energy, control, and information. Consider a "low-tech" example: In some areas, the use of wells for supplying water for drinking and irrigation causes the subsoil water level to drop, which contributes to the dryness of the land and thus increases the need for irrigation, which in turn necessitates the construction of deeper wells, and so forth. Similar problem cycles apply to "high-tech" scenarios. For example, the fact that there has been a spread of monoculture in agriculture, which means that more and more plants are genetically very similar, has increased the risks that crops are affected by epidemics of pests and diseases, whose control requires the use of more pesticide, which in turn has adverse effects on bees and other pollinators, which, if bee populations really were to collapse as some scientists are beginning to fear, would then require the design of some large-

[1] Because capitalism is founded on the untenable ideal of constant economic growth in a world of limited resources, it might be argued that these global crises have more to do with the economic system in which most technology is nowadays embedded, rather than with the technology per se. But this is not entirely accurate. The development of new technology already grew exponentially long before the modern rise of capitalism (Ambrose 2001), and the demise of many pre-capitalist ancient civilizations is closely linked to the technologically enabled overexploitation of resources (Wood 2005). Nevertheless, capitalism has surely contributed to the scale of the crises.

scale artificial pollination technology to sustain global levels of food production, but this technology might have its own unintended side-effects, and so forth. Similarly, the spread of fossil fuel technology is a leading cause of climate change, which is a popular argument to support nuclear power technology, but which brings its own set of seemingly insoluble problems, as the recent disaster in Japan has once again made painfully evident. Relatedly, as temperatures rise in today's megacities, increasingly more energy is used to cool down areas inside cars, shops, and offices, an inefficient process that further adds to the external heat.

And at the same time, more research is always needed to understand and address each crisis, because of the various complexities and interdependencies that tend to be involved. More data are always accumulated, and new technologies must be created to manage the situation, which includes techniques to deal with the ever-increasing amounts of digital information—so-called "Big Data"—that poses new problems of large-scale data management, potential abuses, and privacy.

What seems to be happening is that the increasing extent of negative side effects of modern technological power and globalized functionality often gives rise to the creation of even more powerful technologies to meet those challenges, thus locking humanity's future into an ever-escalating spiral of problem-creation and problem-solving. Even worse, this self-sustaining and accelerating dynamic of technological evolution is currently set in opposition to the environmental and biological underpinnings of our lives. The attainment of unprecedented powers to exploit the natural world has come at the price of a widespread inability to control these forces in order to prevent overexploitation and contamination. This collision course between technology and biology seems to be inevitable, with technology appearing to have the upper hand. Accordingly, some scientists are discussing, in all seriousness, the eventuality of having to abandon earth and are considering terraforming a nearby planet in order to evacuate the human population (or, more likely, just its paying elite). Yet others have resigned themselves to their impending fate, having already accepted the inevitable end of humanity's biological existence, and are fantasizing about their near-future prospects of digital immortality by downloading their personal consciousness into computerized selves.

Leaving these fatalistic endgame scenarios and techno-fantasies aside, one thing is clear: Far from being the masters of the technological world we have helped to create, the existence of humankind is precariously dependent on it for its very survival. And yet, at the same time, we are powerless to fully understand, let alone properly control, the various kinds of forces (environmental, demographic, technological, financial, political, etc.) that are hurtling us at breakneck speed toward an ecological catastrophe of global proportions. How do we explain the paradox that the more we create advanced tools to empower us, the less we are in control of our own lives and humanity's fate?

From the start, I want to set aside the myth that, if we had more information, we would be able to solve the problems faced by our society. Information helps but is not sufficient. To take another topical example: Systematic overfishing is well documented for many parts of the world, and it does not even require advanced science to realize that fish catches are not what they used to be. But, instead of allowing fish populations to regrow, the typical response is to invest even more effort in catching the few remaining individuals, thereby further exacerbating the problem, sometimes until the point of no return when all are gone—extinction. The situation is a little bit like that of an addiction: Most smokers know very well that smoking is unhealthy and can lead to fatal diseases, but this information does not automatically compel them to quit smoking. In general, what is needed is not simply more information but a change in one's way of being in the world. And it is in precisely this existential aspect of the problem that potential future technology, if we take better advantage of its anthropologically constitutive nature, could play a significant positive role.

It is within this context that I want to evaluate the arrival of new technologies that increasingly blur the distinction between our organic flesh and artificial material, i.e., bio-machine hybrid technologies. After comparing these new technologies with existing tools, I will end on an optimistic note, but only on the condition that we shift our understanding of the nature of technology. I will take my cue for future design from Heidegger's ([1954] 2008) analysis of the question of technology. Instead of thinking of technology as people's thoughts and ideas that are congealing into external objects, and thereby enframing and concealing the world behind our devices and systems, technology should be a medium for better revealing the world via our mediated interactions. I conclude by offering some tentative suggestions of how this could be achieved in practice.

2 Technology: Sociological, Phenomenological, and Biological Perspectives

In order to better understand the essence of technology, it is helpful to consider what constitutes a tool.[2] There is a consensus in various areas of the science and philosophy of technology that the ontological statuses of a potential "tool" and a potential "user," i.e., when each is taken as an independent entity, differ significantly from their statuses when they join to become a "tool-user" (Latour 1994; Pickering 1995; Clark 2003; Ihde 2003; Stewart 2010; De Preester 2012). This point is easy enough to understand. It is simply the notion, also familiar in artificial life and cognitive science more generally (Beer 1995; Clark 1997), that properties of a brain–body–world system as a complex whole cannot be reduced to the properties of its isolated components. The nonlinear interactions between the components (i.e., a tool and a user) give rise to an emergent interaction process at a higher level of description (i.e., a tool-user), the existence of which, in return, conditions and modifies the components. But precisely how the tool and user are transformed within the shared context of the tool-user is still a matter of ongoing debate.

The sociologist Latour (1994) has put forward an influential theory that tool and user play wholly symmetrical roles in their interaction within social institutions. He claims that this symmetry forces us to accept that, at least within the context of the tool-user, both tool and user equally deserve the ontological status of being agents or,

[2] It is an interesting open question to consider what the essential differences are between tools and technologies more generally. As one reviewer asked, can one extend the concept of tool to all techniques: house, books, robots, cities, factories, or machines? We may want to differentiate the tool itself from the totality of equipment to which it belongs and also the kinds of relationships humans can have with tools when compared with large-scale techniques (De Preester 2012). In any case, there are good reasons to believe that all techniques mediate our relationship with the world to some extent (Khatchatourov et al. 2007), and that is what is most important here.

to use the term Latour prefers, *actors*. Radically, this co-actor principle is intended to apply to all technical devices, even to objects as inert and lifeless as roads and speed bumps. The following paragraph illustrates how Latour conceives of humans and tools as identical, at least with regard to their status as goal-driven agents, and therefore as symmetrically interchangeable for most theoretical and practical purposes:

> Agents can be human or (like the gun) nonhuman, and each can have goals (or functions, as engineers prefer to say). Since the word *agent* in the case of nonhumans is uncommon, a better term is *actant*, a borrowing from semiotics that describes any entity that acts in a plot until the attribution of a figurative or nonfigurative role ("citizen," "weapon"). […] These examples of actor-actant symmetry force us to abandon the subject-object dichotomy, a distinction that prevents understanding of techniques and even of societies. It is neither people nor guns that kill. Responsibility for action must be shared among the various actants. (Latour 1994, p. 33–34)

There is much to like about this proposal, especially the idea of overcoming the traditional subject–object dichotomy when talking about tool-mediated actions. But there are problems with Latour's attempt to realize this proposal by claiming that there is a functional identity between a tool and its user. Similar to Clark's *cognitive* "parity principle" (Clark and Chalmers 1998), we may think of this as Latour's social or *agential* parity principle. Both parity principles suffer from the problem of overextension. Provocatively, Latour argues that, at least from the point of view of sociology, there are no essential differences between people and things, objects and subjects, and nonhuman actants and human actors. Thus, at some sufficiently abstract sociological level of description, there is no functional difference between the agency of a speed bump and a policeman enforcing traffic regulations (Latour 1994, p. 38).

The intended shock-value of this position may be less noticeable for members in the fields of AI and artificial life, which have long been used to the outrageous claims made by some of its members. As a point in case, consider how Franklin (1995, p. 233) defines "an autonomous agent as a creature that senses its environment and acts on it so as to further its own agenda," which is reasonable enough, but he then goes on to claim that "any such agent, be it a *human or a thermostat*, has a single, overriding *concern*—what to do next" (emphasis added). Again, it is claimed that there is no essential difference between the action of a human and the functioning of a tool as far as their agency is concerned. Here, we have a modern version of that age-old cybernetic fallacy, according to which there is no essential difference between a purposeful agent and a mere negative feedback circuit (Rosenblueth et al. 1943), and which can be rightfully criticized on biological and phenomenological grounds (Jonas [1966] 2001).

Thus, although Franklin and Latour arrive at the co-actor symmetry principle for different reasons, they are in agreement that both the user and the tool can share the same status as autonomous agents, who act according to goals in order to further their own agenda. As I have argued elsewhere at length (Froese and Ziemke 2009), this sort of discourse trivializes the notion of agency. At the very least, the notion of autonomous agency must allow for the agent in question to be a self-constituting individual that adaptively regulates its own conditions of operation (Jonas 1968). Living beings achieve this by continually creating their bodily identity metabolically, which is epitomized by a single cell producing its own membrane boundary, and by regulating their environmental exchanges so as to remain within conditions of viability (Barandiaran et al. 2009). This dependence on self-production and self-regulation makes the existence of living beings precarious, and this is precisely the root cause of their concerned life (Di Paolo 2009a); their continued existence at each moment cannot be guaranteed by mere inertness or external conditions alone but essentially depends on their own intrinsic metabolic, regulatory, and interactive activities. On this view, the very being of an autonomous agent is essentially its own doing (Fig. 1), and it is precisely this necessarily precarious co-dependence between being and doing that constitutes for the organism a meaningful perspective, i.e., as an embodied agent in the world (Weber and Varela 2002). In this manner, classic mind–body dualism can be overcome by recognizing that embodied life itself is the link that integrates the living body (biology) with the lived body (phenomenology); we exist primarily as living-experiencing bodies (Thompson 2007).

Not everyone will agree with this enactive approach to agency. But, as I hope to show, it allows us to better grasp the differences between users and artifacts without giving up the crucial insight that tool-use transforms the user's agency. Moreover, we can retain the claim that actions are attributes of extended brain–body–world systems, but without committing to the agential parity principle that attributes human-like agency to all artifacts. This restriction makes theoretical and practical sense, because the implementation of basic biological agency in an artificial system is a hard problem for AI and has so far met with little success (Froese and Ziemke 2009). It is interesting to inquire to what extent recent advances in bio-machine hybrids satisfy these criteria of agency, and we will return to this question at the end. For now, it is sufficient to note that Latour's speed bump is not an agent in this specific sense. A speed bump does not actively constitute its own physical identity. The boundary between a speed bump and the rest of the street, which demarcates it as a distinct kind

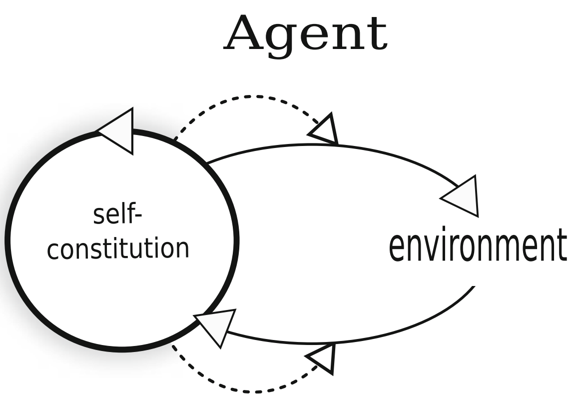

Fig. 1 Illustration of autonomous agency in a biological sense: The organism self-constitutes an identity (*circular arrow*), which is conserved during structural coupling with its environment (*full arrows*); agency requires additional regulation, which is aimed at adjusting this coupling relationship appropriately (*dotted arrows*) (Figure adapted from Froese and Di Paolo (2011).)

of entity, only exists as such for people who interact with it as a part of the system of road traffic. Its “goal” to slow down traffic only exists for the external observer’s perspective of its designers. Nor does the speed bump engage in any activity to regulate its interactions with the environment in any way. Simply put, a speed bump has neither existence nor agency of its own; it merely persists due to physical inertia.

But, you might say, Franklin’s thermostat does actively regulate its interactions in order to adjust the temperature of its environmental conditions. So does the principle of co-agent symmetry not apply here? Not quite, at least not according to the notion of agency I am defending here. The existence of the thermostat as a distinct entity is only externally defined by the activity of its designers and users, and so is the “goal” of its temperature-regulating activity. Normative evaluation of this function’s success depends solely on the criteria of an external observer; the distinction between function and malfunction of the regulatory activity does not exist for the thermostat as such. In contrast, an animal actively determines the physical boundaries of its own body as well as the goals of its adaptive activity, such as the regulation of its body temperature, according to its own internal criteria of viability, since the very existence of the animal as such depends on its successful bodily construction and regulation. In other words, the purpose of the animal’s regulation of temperature, but not the thermostat’s, is intrinsically defined and regulated by its own being and doing. It is this internal relationship between existence and activity that provides us with a justification in attributing intrinsic teleology, as well as a perspective of lived concern, to an autonomous agent in the biological sense (Weber and Varela 2002) but not a mere negative feedback system (Froese and Ziemke 2009).

At the same time, Latour does have a point that the ontological status of tools can be transformed by their context of usage, such as when the same physical device that was designed for writing (a pencil) can become appropriated in another context for stabbing (a weapon). This is akin to Ihde’s (1999) idea of multistable possibilities of usability. Similarly, Latour is right to emphasize that a tool and its user form an extended process of interaction, which will have its own emergent dynamics. This becomes readily apparent when we give up the idea that our mind is something that is contained inside a brain and instead adopt a relational concept of mind and agency (Froese et al. 2013). On this view, the tool of a tool-user serves as an interface that mediates and modulates the user’s perception–action loop with features of their environment (see Fig. 2).

I therefore also readily agree with Latour’s important claim that this interactive tool-using process has profound transformative effects on the user. Indeed, this idea has been receiving much empirical support, including on neural, behavioral, and experiential levels of description (for a review and analysis, see Froese et al. 2012b). According to the enactive approach (Khatchatourov et al. 2007), all tools and technical systems mediate our interaction with the world, as such, they can all be considered as interfaces. An empty lot of grass appears as a playing field by means of a ball, as a parking place by means of a car, as a garbage dump by means of a trash bin, or as a botanical experiment by means of a magnifying glass. Mediation is not limited to tools; it includes social institutions in which we are situated (Gallagher in press). The same empty lot of grass appears as undeveloped real estate by means of a construction company, as lost tax revenue by means of the city council, or as a nature reserve by means of Greenpeace. Even having a thermostat in our room will change

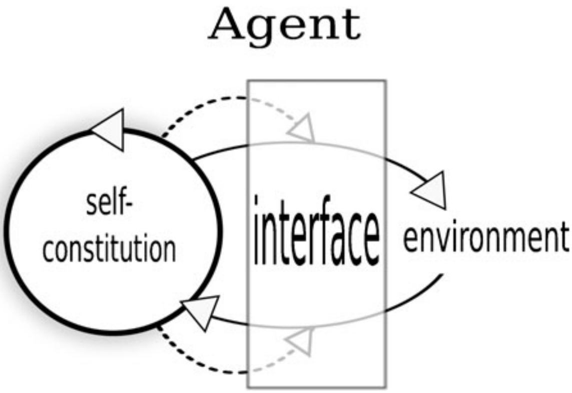

Fig. 2 The tool is conceived of as a mediating interface, which implicitly modulates and augments the user's actual target of perception and interaction, ideally without simultaneously occluding the center of attention in the world (Figure adapted from Froese et al. (2012b).)

the way we perceive temperature, if ever so slightly. For instance, our experience of a hot summer day may turn from something that must be endured to something that could be potentially adjusted by setting the desired temperature. Importantly, the tool–user relationship is still characterized by an essential asymmetry: The tool owes its very existence *as a tool* to the concerned activity of a potential user who will employ it for some purpose, such as cooling a room, and the user's autonomous agency depends on her embodiment as an organism.

3 Human–Nonhuman Co-actor Networks: The Case of Organisms

The mediated actions performed by a human using a tool are a property of the whole brain–body–world system, i.e., the tool-user. The agency of this relationship is normally centered on the biological embodiment of the human. But what if the tool includes some biological components, as is the case with bio-machine hybrid technology? Does this alter the nature of the tool-user? To consider the potential role of biological components for our understanding of these new technologies, let us first consider relations with nonhuman animals.

I agree with Ihde (2003) that Latour's theory of a symmetrical network of human and nonhuman actors is most applicable to human relationships with nonhuman animals, as exemplified by companion animals. For example, Haraway (2003, 2008) has argued at length that dogs and humans mutually invented each other and that wild wolves domesticated our ancestral hominids at the same time as those hominids domesticated the wolves. I propose that this form of social co-determination between two co-evolving species is the most paradigmatic case of a symmetrical human–nonhuman co-actor network. Indeed, experimental studies have demonstrated that, in some situations, human infants and dogs are more alike in their response to human social cues than dogs and wolves (Topál et al. 2009). To be sure, this special relationship does not turn dogs into humans:

> Dogs' special social–cognitive skills are not "normal" in that they do not gesture for or teach humans reciprocally, and they do not use their comprehension abilities

with other dogs. They have evolved specialized skills for dealing with their unique situation in which they benefit by taking orders from humans. (Tomasello and Kaminski 2009, p. 1214)

And yet dogs often train their owners almost as much as the owners train their dogs. This should not be surprising, since self-other co-determination is a generic property of all kinds of social interaction, especially if they involve some kind of empathy (Thompson 2001). Particularly empathic people may bring about such social human–nonhuman co-actor networks even with non-domesticated animals in the wild (Smuts 2001). But this mutual recognition of agency is not necessary for co-actor symmetry. According to the enactive approach to agency, which is grounded in the embodiment of an organism, all relations between biological individuals form multi-agent systems (Froese and Di Paolo 2011).

The widespread existence of self-other co-determination is not a mysterious process. In dynamical terms, it can be shown to spontaneously follow from the mutual entrainment of two agents within an interaction process, whereby the dynamics of interaction form an extended body (Froese and Fuchs 2012). These effects can happen outside of explicit awareness of each other. For example, the relative stability of an interaction process can lead to the spontaneous self-organization of mutual mimicry of behaviors (Froese et al. 2012a). In this basic dynamical sense, co-determination based on co-actor symmetry is a property of all networks involving two or more biological individuals (Fig. 3).

Of course, more often than not, it is the case that human–nonhuman relationships are heavily biased. One only needs to think of the modern meat industry or of the mass extinctions that tend to follow in the wake of civilization and “progress” and wherever humanity spreads, a worrying pattern possibly dating back to the extinctions of large mammals at the end of the last Ice Age. Unfortunately, even human–human relationships are heavily biased as well. Indeed, the effect of many social institutions is not so much an empowering of inanimate objects to the status of autonomous agents, but the degrading of human beings to nothing but their effective roles and functions, as Marx has observed. For example, to a large extent, the essence of labor is being paid to partly give up one’s personal autonomy. In extreme institutional situations, such as warfare, forced labor, prison camps, and mental institutions (e.g., Foucault [1961] 2001), humans are systematically reduced toward

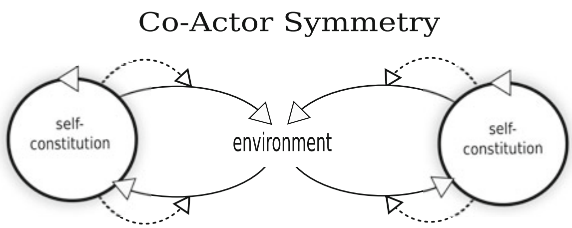

Fig. 3 Illustration of the relationship between two agents sharing the same environment: The manner in which one agent’s movements affect the environment can result in changes to sensory stimulation for the other agent, and vice versa, leading to multi-agent recursive interaction. Through this interaction, the two agents form an extended system characterized by co-actor symmetry (Figure adapted from Froese and Di Paolo (2011).) This idea of co-actor symmetry is stricter than the one proposed by Latour (1994) because it excludes extended systems formed by tool-users alone (e.g., Figs. 2 and 4)

the status of nonhuman objects. In other words, while I accept that basic co-actor symmetry holds for all relations between living beings, in accordance with the biological definition of agency, I emphasize that the nature of this relationship can vary in other important respects.

One means of variation that is directly relevant for understanding bio-machine hybrid technology are machine or tool-based interfaces that mediate agent–agent interactions. This practice started in a systematic manner with the domestication of plants and animals. For example, the use of yeast bacteria in the process of baking bread and brewing beer is an example of biology-based engineering that dates back to the times of Ancient Egypt and Mesopotamia. Given that we are dealing with a process involving several human agents and countless nonhuman agents, we have another case of human–nonhuman co-actor symmetry in a strict biological sense. The role of the tools for baking and brewing is to mediate this relationship and to make it possible in the first place, although this mediation falls short of enabling a genuine social interaction. The only thing that matters from the point of view of the human brewer is that the bacteria's metabolism can be coordinated with the rest of the brewing process as a mechanism of chemical transformation. On that view, the bacteria are only components of a technical system. Their situation can be compared with that of our bodily cells when viewed from the perspective of the whole organism, except that only the components of our body, but not the cells of the brewing system, actively contribute to the coordination of the larger system. Our body is a chemically mediated meshwork of living beings who are coordinating to realize our intentions (Varela 1991).

4 The Case of Nonhuman Bio-machine Hybrid Agents

Having concluded our discussion of mediation by tools, nonhuman animals, and their integration into larger systems, we can now better understand more recent developments in bio-machine hybrid technology.

To begin with, researchers have begun to explore the possibility of artificially creating new autonomous agents in the strong biological sense of the term. For example, there are robots that incorporate neuronal cells to regulate their sensorimotor loops (Warwick et al. 2010) and robots that include microbial cells to generate energy (Ieropoulos et al. 2010). However, although these kinds of bio-machine hybrid robots include more co-dependencies between their biological components and the rest of their technical system, compared with the brewery, for example, the overall system still does not count as an agent in its own right. The hybrid system is unlike the body of a multi-cellular organism because the co-dependencies between the biological and technical components are not of the right kind. Specifically, the existence of the machine processes would have to depend on the existence of the biological processes, and vice versa, in order to co-constitute one operationally closed network as a whole. In this sense, the enactive approach provides novel design principles for an AI based on the principles of biological embodiment (Froese and Ziemke 2009).

What are the prospects for this kind of bio-machine hybrid technology? So far, this kind of living system has not been produced in practice, although it should be possible to do so in principle given the right materials and conditions. It has been argued that so-called "living technology" will provide a number of benefits that derive from its biological organization:

> We deem technology to be *living* if it is powerful and useful precisely because it has the core properties of living systems, including such properties as the ability to maintain and repair itself, to autonomously act in its own interests, to reproduce, and to evolve adaptively on its own. (Bedau et al. 2010, p. 91)

Again, we find the aim of engineering a system that is able to act according to its own goals so as to further its own interests. However, the status of such living technology, if it were to be fully realized, would be qualitatively distinct from all previous technology. In contrast to Franklin's (1995) thermostat, these technical products would have to be interpreted as being *autonomous agents* in the strong biological sense of the term. Accordingly, here we would have a case where the traditional asymmetrical tool–user relationship is replaced by a new version of co-actor symmetry. Ideally, this symmetrical relationship would still be useful for humans, perhaps because the hybrid agent is embedded in a larger technical system (like yeast), bodily system (like an organ), or has shared interests (like a companion species). But this utility cannot be guaranteed from the outset. The main problem we have to consider is the displaced locus of agency, which is explicitly moved away from the user and into the tool itself, and which thereby upsets the asymmetry of the typical tool–user relationship.

To be sure, already, Franklin's thermostat shows superficial similarities with goal-directed behavior, and some tools can appear even more life-like when we consider how they condition the social processes in which they are embedded (McGregor and Virgo 2011). And we know that all technologies are designer-independent to some extent, because their potential range of uses cannot be predetermined (Ihde 1999), and their developmental "lifecycle" follows its own intrinsic affordances and laws (Stiegler 1998). However, the proponents of living technology are trying to push this relative independence of technology to its ultimate conclusion. If they achieve their stated aims, they would have created the conditions for the emergence of a new class of bio-artifacts: The operation of these hybrid agents would be defined by their own intrinsic values; they would adapt their behavior according to their own goals, and they would progress by means of their own open-ended evolution (Bedau et al. 2010).

There is one important difference between this open-ended adaptability of living technology and the multistable adaptability of a normal tool. In contrast to the non-living tool, whose ambiguous and multistable functional properties are still constrained by having to be functional in a human-defined context of usage, a genuinely living artifact would have its own autonomous agenda, whether it is useful for humans or not. It might develop and evolve in ways that are beneficial for us, like a companion species (like dogs), but it might also find its niche as a useless nuisance (like cockroaches), or as a direct competitor for resources (like locusts). In other words, the stated ideal of a truly autonomous and self-adapting living technology would present us with an extreme version of Ihde's (1999) prognostic antinomy. One way of understanding the aim of living technology is therefore as a decoupling of Ihde's notion of multistable uses from their human context. New uses of non-living technology may exploit affordances that already reside within the tool, yet they ultimately depend on humans for their realization. Similarly, living artifacts will also exhibit surprising new contexts of functionality, and yet, because they are themselves autonomous agents evolving according to their own criteria, new niches of existence might be independent of, or even in direct competition with, human usage.

I therefore share Di Paolo's (2010) skepticism when he questions: What is the practical benefit of making a tool that is genuinely autonomous such that it can decide that it will no longer do what we want to use it for? In engineering terms, this seems to be counterproductive to say the least, and there are already plenty of issues related to the relative independence of existing technical systems. I have already mentioned the problems caused by the self-reinforcing cycle of the globalized technological–industrial system, and similar worries can be raised about harmful effects caused by the emergence of autonomous fluctuations in social systems and financial systems. If societies really were like organisms and people their cells, as has occasionally been proposed by theorists, our personal freedoms would be drastically curtailed (Di Paolo 2009b). Thus, in spite of the powers promised by genuinely living technologies, or perhaps rather because of them, it is not immediately clear that such hybrid agents can be directly "useful" to humans in the sense of using a tool. In order to direct their agency in useful ways, engineers would have to ensure that their existence is strongly constrained by being dependent on functioning in a human-defined context of usage. And this context may include the inner milieu of the human body itself, which places strong constraints on the operation of its parts. On this view, a better approach for future design of bio-machine hybrid technology would be creating new kinds of partial hybrid *organs*, not whole hybrid *organisms*. This would be a return to the original Greek meaning of organon (ὄργανον), which is interchangeably used to refer to a sense organ, instrument, or tool.

5 From Living Technology to Lived Technology

There is another benefit of shifting our design perspective from creating hybrid organisms to hybrid organs. I believe there may be a chance that bio-machine hybrid interfaces may equip us with the tools we need to better manage the various social and ecological systems that constitute our world. The technology-based nature of human existence means that the presence of a self-sustaining cycle of problem-solving and problem-creation cannot be avoided altogether. However, it is possible to mitigate its worst excesses by gaining insight into its complexities and then modulating its underlying conditions accordingly. We do need more data, but more importantly we also need better insight and understanding. This is certainly true for the people working in all kinds of high-level institutional positions whose decisions have immense impact, but it is also true of billions of other people whose individual decisions sum up to global proportions.

In order to facilitate the process of improved insight creation in our everyday lives, we should take into account how usage of technology in all its forms implicitly mediates our ways of acting, thinking, and perceiving. The key insight is that *all technologies modulate our capacities for interaction and therefore alter the way in which the world is disclosed to us in our lived experience*. On this view, the philosophy of technology is essentially a philosophy of interfaces and human–world mediation. And the key hope is that eventually we will be able to design more beneficial forms of technological mediation that provide direct perceptual access to the complex problems underlying our global crises, perhaps similar to the way in which complex neural systems are starting to be studied by mixed virtual reality systems (Betella et al. 2012).[3] But our world's problems cannot just be solved at the large-scale level alone, they also require changes in the attitudes of billions of individuals who, for example, are concerned more about today than tomorrow, about themselves than others, about comfort than sustainability, and so forth. That means most of us. A good starting point would therefore be to design interfaces that more directly reveal to us our dependencies on others and on the natural world so as to effect existential adjustments.

In fact, our attitudes toward others and the world are already being implicitly modulated everywhere and all of the time. At the start of this article, we observed that all technologies mediate how we experience the world to some extent. Nevertheless, this modulation of our relationship with the world is especially profound in the case of tools that are specifically designed to replace or extend our natural perceptual faculties, such as spectacles, hearing aids, as well as scientific instruments (Ihde 2011a). A classic example is the blind man's walking stick. As Merleau-Ponty has observed, during usage, the cane becomes a part of the man's embodied interaction with the environment through which his immediate world is perceptually revealed.

> The blind man's stick has ceased to be an object for him, and is no longer perceived for itself; its point has become an area of sensitivity, extending the scope and active radius of touch, and providing a parallel to sight. […] The position of things is immediately given through the extent of the reach which carries him to it, which comprises besides the arm's own reach the stick's range of action. (Merleau-Ponty [1945] 2002, pp. 165–166)

Once we accept that the conscious mind is embodied in the world, we realize that it is possible to systematically vary user experiences by systematically varying their effective manner of being in the world via their usage of different user-interfaces (Froese et al. 2012c). Much research in human–computer interfaces, including haptics, tangible, and enactive user-interfaces, prosthetic and assistive technology, and so forth, has started to study how to best make use of our propensity for technological extension. Interfaces typically need to satisfy two requirements during usage: (1) the interface itself must become as experientially transparent as possible (Thompson and Stapleton 2009) and (2) the interface should extend our natural sensorimotor capacities so as to give rise to novel ways of experiencing the world. This idea of *augmented sense-making* (Froese et al. 2012b) is more strongly exemplified by Merleau-Ponty's cane than by Heidegger's hammer. Nevertheless, both tools share the property of serving as a pre-reflective medium for engaging with a tool-inflected world, rather than being themselves the objects of attention. I therefore

[3] This work is part of a European Commission FP7 funded project on the Collective Experience of Empathic Data Systems (CEEDs), which aims to develop novel, integrated technologies to support human experience, analysis, and understanding of very large datasets. I share their idea of placing our experience at the center of technical solutions, which I try to motivate as being applicable to all kinds of complex problems, including our understanding of major crises.

propose to refer to all kinds of user-interfaces as "lived technology" to emphasize that, during normal usage, these tools become an implicit part of *how* the world as such is experienced, rather than merely an explicit part of *what* is experienced in the world among other objects.[4] We should aim to modulate the structure of experience itself.

To some extent, this shift to lived technology has already taken place several decades ago. When AI moved from the scientific arena into the practical domain of engineering, its biggest successes were not in the creation of autonomous robots (except for especially designed industrial contexts). On the contrary, it is by means of "intelligent" human–computer interfaces, of which "smart" phones are only the latest development, that the information technology revolution has pervaded all aspects of our modern lives. However, current interface designs based on the symbolic principles of AI can also be problematic, in particular, if we want to go beyond mere functionality and take effects on first-person experience into account. The relative neglect of user experience has surely contributed to fact that the unprecedented computerization of all aspects of our lifeworld has also been a source of considerable apathy and alienation. And where the user's experience has been a key concern, the goal has typically been to replace the real world with a fantasy. In modern lives, the real world has largely been replaced by a more captivating virtual world (Baudrillard [1981] 1995). This technologically mediated removal from our concrete reality further complicates current efforts to mobilize the necessary motivation and resources to address the ecological and social challenges we are facing. There are two aspects of information technology that contribute to this distraction.

First of all, there is the widespread rationalist's fallacy that we could solve the world's problems if we just had more information. But this simply confuses raw data with interpreted information, and interpreted information with practical understanding. Most importantly, there is a fact-value gap. As Hume noted long ago, that there is no logical way of proceeding from an "is" to an "ought" (Froese 2009). Reason (and, by technological extension, symbolic AI) is able to provide us with information about relations of ideas and matters of fact, but it cannot in itself decide why we *ought* to choose one course of action over another. Such a decision requires making a value judgment based on a normative understanding of the situation. Yet we currently have no systematic method of warranting the validity of this judgment. Science is no help either because, according to current standards, it ought to be subject-less and value-free. In other words, our senses are swamped by a relentless massive data flow, and yet we have little means of making sense of all the data except by amassing even more data, and so forth. Here we are faced by another version of the self-sustaining problem-solving and problem-creation cycle, which is specific to the case of information technology.

Second, information technology typically forces the user to focus on the symbolic output of the tool itself. For instance, engineers have long been aiming for the

[4] Interestingly, bio-machine hybrid robots could also be considered examples of lived technology in this specific sense—except that in this case it is the perceptual interactions of nonhuman organisms that are primarily being modulated. This change in perspective could be useful for future design of hybrid robots because the goal of creating an augmenting interface for one or more existing biological individuals, rather than creating another genuine hybrid agent as such, places significantly less demands on the engineering process (Froese and Ziemke 2009).

creation of a user-interface characterized by co-actor symmetry, for example, by duplicating our own heuristic problem-solving skills in the computer (Simon and Newell 1958). This gives symbolic computer interfaces a special status, because practical tool-use is turned into virtual "communication" with a "semiotic artifact" (Khatchatourov et al. 2007), whereby communication is limited to the conditions defined by the interface alone. A "critical hermeneutic" (Ihde 2011b) therefore becomes a necessary element of information technology. Moreover, the need to reflectively interpret the symbolic output turns perception away from the world, thereby acting more like a form of perceptual occlusion than mediation (Fig. 4). This is exemplified by all kinds of problems caused by people who rely more on their navigation systems than on their own senses, such as truck drivers stuck in impassable streets and ships run aground on visible reefs. This situation is in stark contrast to more traditional kinds of technology, as exemplified by the hammer and the cane, which mediate our sensorimotor relationship with the world in a more direct and semi-transparent manner because their rigid physical coupling to the body ensures an extended unity of action and perception.

In sum, information technology has been designed to change *what* is perceived in the world, i.e., symbolic representations of data, while giving little consideration to *how* it is perceived. This design stance is understandable from the traditional philosophical perspective of mind–body dualism. A mind that is considered as strictly isolated from its embodiment and the world can only be technologically supported by providing it with external processing and data. However, from the enactive perspective, we know that, like all technologies, usage of such human–computer interfaces also affects *how* the world is perceived. In traditional cases, it does so mainly negatively by forcing us to put the operation of the symbolic interface itself at the center of our lifeworld.

However, despite these misgivings, I acknowledge that a well-designed computer interface is capable of transforming our lived experience in positive ways. Since the beginnings of the information technology revolution, there has been a steady

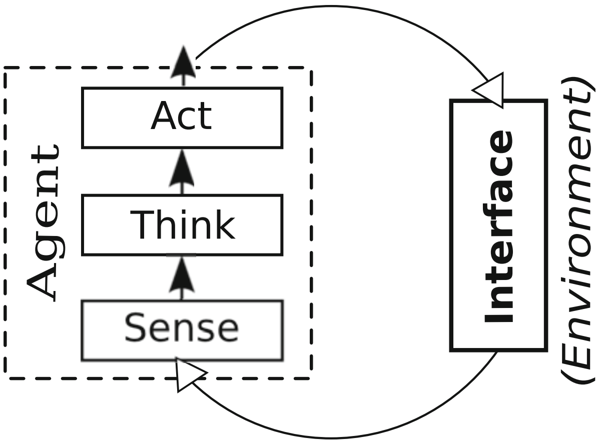

Fig. 4 An illustration of the design stance exemplified by information technology making use of symbol-based interfaces. The artifact is conceived of as an external object in the world, which is the agent's target of explicit reflection while occluding perception of the world (Figure adapted from Froese et al. (2012b).)

progression away from mere functional devices toward ever more embodied and immersive interfaces. Although this trend has largely been driven by consumer demands for enhanced user experiences, I expect this trend to further accelerate given that the mediating role of technology has recently come to attention in mainstream cognitive science (Clark 2003). But many of these efforts still remain limited by the symbolic approach to cognitive science, which neglects our living and lived embodiment (Thompson 2007). But once we recognize that experience is structured by our being in the world, including by the way in which our bodies interface with technologies (Froese et al. 2012c), a new horizon of possibilities becomes available. As already Merleau-Ponty observed:

> Scientific thinking, a thinking which looks on from above, and things of the object-in-general, must return to the "there is" which underlies it; to the site, the soil of the sensible and opened world such as it is in our life and for our body—not that possible body which we may legitimately think of as an information machine but that actual body I call mine, this sentinel standing quietly at the command of my words and acts. (Merleau-Ponty [1961] 1964, pp. 160–161)

I suggest that current symbolic computer interfaces are a source of alienation because their underlying principles are inherently alien to those of embodied life and mind. The world is normally revealed through our lived body immediately, without any necessity for reflective symbol processing. And we should take care not to be misled by Merleau-Ponty's poetic style: We normally do not reflectively "command" our body, either. Take speech, for example. When I am talking, I do not know the ultimate origin of my words; all of the various micro-movements and gestures of each body part, from the vibrations of my vocal tract to my gesticulating hands, spontaneously self-organize into a whole-body concert that realizes the meaning I am trying to express. I do not have reflective control over this process, nor do I have any need for it. And the same is true of actions more generally. Each body part is a relatively autonomous and yet reliable partner in a collective dance, each one playing its own little role while co-adapting smoothly to changes of overall context, such as intonation, subtleties of meaning, and so forth (Kelso 1995). There are currently no technological interfaces available that can autonomously and reliably realize our capacities for bodily expression in this kind of manner. I suggest that it is precisely this shortcoming that stands in the way of genuine bodily incorporation, whereby a tool becomes an integral part of our embodiment rather than just a temporary extension (De Preester 2011). And this lack of incorporation limits the full potential for augmented sense-making.

6 The Potential of Bio-machine Hybrid Interfaces

Bio-machine hybrid technology shares some of the biological principles of our living body, the basis of our embodied mind, and might therefore be in a better position to address the problems that have beset symbolic user-interfaces. It is quite likely that interfaces that are designed according to the core operational principles of life, such as self-organization, robustness, and adaptability, can be better integrated into our biological embodiment and will be more effective at enhancing our existing sense-making capacities. Admittedly, at the moment, the validity of this conjecture is a matter of empirical investigation. And even though I accept it as plausible, I currently do not have any concrete suggestions of how we should proceed in practice. We are dealing with uncharted territory, so it is hard to know where to begin. But if we are indeed serious about investigating the potential of lived technology based on bio-machine hybrid interfaces, then some additional conceptual changes are in order.

Consider the fact that during normal perception our embodiment conceals itself at the same time as it reveals the world. Moreover, this revealing is not a neutral observation of a mere set of facts; perception is a concrete understanding that is already imbued with valence and shaped by the possibilities for future action (Merleau-Ponty [1945] 2002). Thus, if we want a new kind of technology that mitigates the problems associated with symbolic computer interfaces, such as information overflow, the fact-value gap, and perceptual opacity, then its devices should ideally fit into and enhance this natural perceptual situation. This brings us back to the suggestion that, in general, future design based on bio-machine hybrid technology could benefit by shifting the design perspective from hybrid robots to hybrid interfaces. For example, it has been demonstrated that a well-integrated neuroprosthesis is able to repair and even enhance an agent's perceptual interaction with the world (Hampson et al. 2012). And it is worth considering that these internal interfaces do not necessarily have to consist of silicon and/or organic material but could also be temporally integrated in chemical form (Sessa 2008).

Several philosophical issues are raised by this proposal. For example, organ-like bio-machine interfaces raise intriguing questions with respect to the theoretical distinction between bodily extension versus bodily incorporation (De Preester 2011), as well as regarding the role of the potential separability of tools, which has been one of their defining features (Lenay 2012). In contrast to traditional tools, it may not be so easy to put down a bio-machine hybrid device once it has become incorporated. And yet it is precisely this tighter coupling with the body that promises to improve functionality, for example, by giving a feeling of touch to arm prosthetics (Kwok 2013). As suggested previously, it may even turn out that the most effective deployment of hybrid technology as a mediating interface happens to be *within* the confines of our biological body (Fig. 5).

In addition, we need to keep Ihde's prognostic antinomy in mind. The design of bio-machine hybrid technology is not without its dangers. For example, do we really want the development of neural implants to the extent of allowing remote control of behavior (e.g., Talwar et al. 2002)? Apart from ethical considerations, there may also be significant practical drawbacks compared with more traditional methods of training behaviors. Because this externally induced neural activity is not an expression of the organism's own autonomous agency (Nasuto and Bishop 2013), the controlled behavior lacks the inherent flexibility and adaptability of intentional action. Similarly, complete autonomy of an implant or prosthesis is not a desirable engineering goal, either. For example, what would be the use of creating a prosthetic hand with anarchic hand syndrome? To be sure, the various parts (e.g., cells and organs) of our living body also enjoy a relatively autonomous existence (Varela 1991), but they still join together in harmony to express our intended goals. If they do not, we are in serious health trouble. Accordingly, the principle that traditionally ensured a unity of action

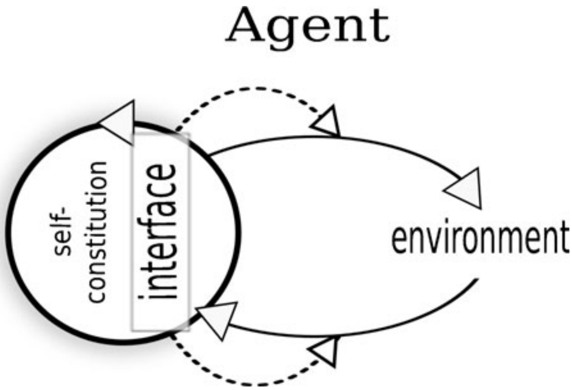

Fig. 5 An illustration of a design stance exemplified by lived technology that may become possible on the basis of bio-machine hybrid technology. The hybrid artifact is conceived of as a mediating interface within the user's biological embodiment, which implicitly modulates and thereby augments the user's potential for perception and interaction in the world

and perception during tool-use, i.e., rigid physical coupling between body and non-autonomous tool, must be replaced by a new principle of design: ensuring flexible coordination between relatively autonomous organic and hybrid body parts. The interaction dynamics of the whole body need to effectively constrain the autonomy of its parts.

As we have seen, the same problem of too much autonomy applies if our society insists on going ahead with the creation of genuinely autonomous agents by using bio-machine hybrid technology. How do we curtail their agency in useful ways? How could we ensure that, for example, we end up creating a companion species rather than something useless or even worse? Thus, while there is indeed a pressing need to improve upon the symbolic information interfaces that currently dominate our technological lives, we also must make sure that we do not overshoot our target. A successful and justifiable development of living and bio-machine hybrid technology must therefore go hand in hand with a more systematic study of the conditions leading to more effective mutual coordination and harmonious cooperation both *between* and *within* living beings.

7 Conclusion

We are faced by various crises of global proportions, which partially derive from our technologically facilitated capacities for the exploitation and destruction of our natural and social environment. The social systems we have created so as to govern our lives, now based on huge networks of information technology, are too complex to be properly understood and effectively regulated. At the same time, there is an ongoing population explosion, sustained by advances in agricultural and medical technologies, and yet we are impotent at shaping the behaviors of these billions of people toward more sustainable attitudes. There is no end in sight: given the open-endedness of human being, and the fact that technology is anthropologically constitutive, continuous change is the default condition. It is impossible to curtail the relentless march of "progress" as such. Prohibition is not an option. Instead, the pragmatic aim should be to mitigate technology's negative side effects as best as possible. I have proposed that a crucial condition for this task is the design of awareness raising technology that will help us to better understand and deal with these negative consequences.

Accordingly, while acknowledging some serious concerns, I tentatively support the development of new bio-machine and living technologies on the basis of the following conjecture: *Bio-machine and living technology will be a gateway toward improved lived technology*. I base this conjecture on the theory of the embodied mind. If our experience is not locked up inside the head but is rather a relational phenomenon of the brain–body–environment system as a whole, then this opens up new perspectives for future design. We are no longer limited to trying (and failing) to duplicate our mind in external devices; rather, we can aim to create interfaces that become directly incorporated into our embodied mind. My hope is that the development of bio-machine hybrid and living technology will more readily allow for this incorporation and thereby entail a powerful enhancement of our natural ability for perceptual insight and practical understanding. On this view, albeit one limited by prognostic antinomy, the benefits are worth the risks.

Acknowledgements Tom Froese is grateful to Mark Bishop and Yasemin Erden for their support and encouragement during the writing of this article. He also thanks three anonymous reviewers for their helpful and constructive comments.

References

Ambrose, S. H. (2001). Paleolithic technology and human evolution. *Science, 291*(5509), 1748–1753.

Barandiaran, X., Di Paolo, E. A., & Rohde, M. (2009). Defining agency: Individuality, normativity, asymmetry, and spatio-temporality in action. *Adaptive Behavior, 17*(5), 367–386.

Baudrillard, J. ([1981] 1995). *Simulacra and simulation*. Ann Arbor, MI: University of Michigan Press.

Bedau, M. A., McCaskill, J. S., Packard, N. H., & Rasmussen, S. (2010). Living technology: Exploiting life's principles in technology. *Artificial Life, 16*, 89–97.

Beer, R. D. (1995). A dynamical systems perspective on agent–environment interaction. *Artificial Intelligence, 72*, 173–215.

Betella, A., Carvalho, R., Sanchez-Palencia, J., Bernadet, U., & Verschure, P. F. M. J. (2012). Embodied interaction with complex neuronal data in mixed-reality. In S. Richir (Ed.), *Proceedings of the 2012 Virtual Reality International Conference (VRIC'12)*. New York, NY: ACM.

Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge: The MIT Press.

Clark, A. (2003). *Natural-born cyborgs: Minds, technologies, and the future of human intelligence*. New York: Oxford University Press.

Clark, A. (2008). *Supersizing the mind: Embodiment, action, and cognitive extension*. New York: Oxford University Press.

Clark, A., & Chalmers, D. J. (1998). The extended mind. *Analysis, 58*(1), 7–19.

De Preester, H. (2011). Technology and the body: The (Im) possibilities of re-embodiment. *Foundations of Science, 16*, 119–137.

De Preester, H. (2012). Equipment and existential spatiality: Heidegger, cognitive science and the prosthetic subject. In J. Kiverstein & M. Wheeler (Eds.), *Heidegger and cognitive science* (pp. 276–308). Basingstoke: Palgrave Macmillan.

Di Paolo, E. A. (2009a). Extended life. *Topoi, 28*(1), 9–21.

Di Paolo, E. A. (2009b). Overcoming autopoiesis: An enactive detour on the way from life to society. In R. Magalhaes, & R. Sanchez (Eds.), *Autopoiesis in organizations and information systems* (pp. 43–68). Elsevier.

Di Paolo, E. A. (2010). Living technology. In M. A. Bedau, P. Guldborg Hansen, E. Parke, & S. Rasmussen (Eds.), *Living technology: 5 questions*. Birkerød: Automatic Press/VIP.
Foucault, M. ([1961] 2001). *Madness and civilization: A history of insanity in the age of reason*. London, UK: Routledge.
Franklin, S. (1995). *Artificial minds*. Cambridge: The MIT Press.
Froese, T. (2009). Hume and the enactive approach to mind. *Phenomenology and the Cognitive Sciences, 8*(1), 95–133.
Froese, T., & Di Paolo, E. A. (2011). The enactive approach: Theoretical sketches from cell to society. *Pragmatics & Cognition, 19*(1), 1–36.
Froese, T., & Fuchs, T. (2012). The extended body: A case study in the neurophenomenology of social interaction. *Phenomenology and the Cognitive Sciences, 11*(2), 205–235.
Froese, T., & Ziemke, T. (2009). Enactive artificial intelligence: Investigating the systemic organization of life and mind. *Artificial Intelligence, 173*(3–4), 366–500.
Froese, T., Lenay, C., & Ikegami, T. (2012a). Imitation by social interaction? Analysis of a minimal agent-based model of the correspondence problem. *Frontiers in Human Neuroscience, 6*(202), doi:10.3389/fnhum.2012.00202.
Froese, T., McGann, M., Bigge, W., Spiers, A., & Seth, A. K. (2012b). The enactive torch: A new tool for the science of perception. *IEEE Transactions on Haptics, 5*(4), 365–375.
Froese, T., Suzuki, K., Ogai, Y., & Ikegami, T. (2012c). Using human–computer interfaces to investigate 'mind-as-it-could-be' from the first-person perspective. *Cognitive Computation, 4*(3), 365–382.
Froese, T., Gershenson, C., & Rosenblueth, D. A. (2013). The dynamically extended mind: A minimal modeling case study. In *2013 I.E. Congress on Evolutionary Computation* (pp. 1419–1426). IEEE Press.
Gallagher, S. (in press). The socially extended mind. *Cognitive Systems Research*.
Hampson, R. E., Gerhardt, G. A., Marmarelis, V., Song, D., Opris, I., Santos, L., et al. (2012). Facilitation and restoration of cognitive function in primate prefrontal cortex by a neuroprosthesis that utilizes minicolumn-specific neural firing. *Journal of Neural Engineering, 9*(5), 056012. doi:10.1088/1741-2560/9/5/056012.
Haraway, D. (2003). *The companion species manifesto: Dogs, people, and significant otherness*. Chicago: Prickly Paradigm Press.
Haraway, D. (2008). *When species meet*. Minneapolis: The University of Minnesota Press.
Heidegger, M. ([1927] 1962). *Being and time*. New York, NY: Harper and Row.
Heidegger, M. ([1954] 2008). The question concerning technology. In D. F. Krell (Ed.), *Martin Heidegger: basic writings* (pp. 307–342). New York, NY: HarperCollins.
Ieropoulos, I., Greenman, J., Melhuish, C., & Horseld, I. (2010). EcoBot-III: A robot with guts. In H. Fellermann, M. Dörr, M. M. Hanczyc, L. L. Laursen, S. Maurer, D. Merkle, et al. (Eds.), *Artificial Life XII: Proceedings of the Twelfth International Conference on the Synthesis and Simulation of Living Systems* (pp. 733–741). Cambridge, MA: The MIT Press.
Ihde, D. (1999). Technology and prognostic predicaments. *AI & Society, 13*, 44–51.
Ihde, D. (2003). If phenomenology is an albatross, is *post-phenomenology* possible? In D. Ihde & E. Selinger (Eds.), *Chasing technoscience: Matrix for materiality* (pp. 131–144). Bloomington: Indiana University Press.
Ihde, D. (2011a). Husserl's Galileo needed a telescope! *Philosophy & Technology, 24*(1), 69–82.
Ihde, D. (2011b). Stretching the in-between: Embodiment and beyond. *Foundations of Science, 16*(2–3), 109–118.
Jonas, H. (1968). Biological foundations of individuality. *International Philosophical Quarterly, 8*(2), 231–251.
Jonas, H. (1979). Toward a philosophy of technology. *The Hastings Center Report, 9*(1), 34–43.
Jonas, H. ([1966] 2001). *The phenomenon of life: Toward a philosophical biology*. Evanston, IL: Northwestern University Press.
Kelso, J. A. S. (1995). *Dynamic patterns: The self-organization of brain and behavior*. Cambridge: The MIT Press.
Khatchatourov, A., Stewart, J., & Lenay, C. (2007). Towards an enactive epistemology of technics. In *Proceedings of ENACTIVE/07: 4th International Conference on Enactive Interfaces* (pp. 129–132). Grenoble, France: Association ACROE.
Kwok, R. (2013). Once more, with feeling: Prosthetic arms are getting ever more sophisticated. Now they just need a sense of touch. *Nature, 497*, 176–178.
Latour, B. (1994). On technical mediation—Philosophy, sociology, genealogy. *Common Knowledge, 3*(2), 29–64.
Lenay, C. (2012). Separability and technical constitution. *Foundations of Science, 17*(4), 379–384.

Leroi-Gourhan, A. ([1964–1965] 1993). *Gesture and speech*. Cambridge, MA: The MIT Press.
McGregor, S., & Virgo, N. (2011). Life and its close relatives. In G. Kampis, I. Karsai, & E. Szathmáry (Eds.), *Advances in artificial life: 10th European Conference, ECAL 2009* (pp. 230–237). Berlin: Springer.
Merleau-Ponty, M. ([1945] 2002). *Phenomenology of perception* (Trans: Smith, C.). Oxon, UK: Routledge.
Merleau-Ponty, M. ([1961] 1964). Eye and mind (Trans: Dallery, C.). In J. M. Edie (Ed.), *The primacy of perception: And other essays on phenomenological psychology, the philosophy of art, history and politics* (pp. 159–190). Evanston, IL: Northwestern University Press.
Nasuto, S. J., & Bishop, J. M. (2013). Of (Zombie) mice and animats. In V. C. Müller (Ed.), *Philosophy and theory of artificial intelligence* (pp. 85–106). Berlin: Springer.
Pickering, A. (1995). *The mangle of practice: Time, agency, & science*. London: The University of Chicago Press.
Rosenblueth, A. N., Wiener, N., & Bigelow, J. (1943). Behavior, purpose and teleology. *Philosophy of Science, 10*, 18–24.
Sessa, B. (2008). Is it time to revisit the role of psychedelic drugs in enhancing human creativity? *Journal of Psychopharmacology, 22*(8), 821–827.
Simon, H. A., & Newell, A. (1958). Heuristic problem solving: The next advance in operations research. *Operations Research, 6*(1), 1–10.
Smuts, B. (2001). Encounters with animal minds. *Journal of Consciousness Studies, 8*(5–7), 293–309.
Stewart, J. (2010). Foundational issues in enaction as a paradigm for cognitive science: From the origin of life to consciousness and writing. In J. Stewart, O. Gapenne, & E. A. Di Paolo (Eds.), *Enaction: Toward a new paradigm for cognitive science* (pp. 1–31). Cambridge: The MIT Press.
Stewart, J., Gapenne, O., & Di Paolo, E. A. (2010). Introduction. In J. Stewart, O. Gapenne, & E. A. Di Paolo (Eds.), *Enaction: Toward a new paradigm for cognitive science* (pp. vii–xvii). Cambridge: The MIT Press.
Stiegler, B. (1998). *Technics and time, 1: The fault of Epimetheus* (Trans: Beardsworth, R. & Collins, G.). Stanford, CA: Stanford University Press.
Talwar, S. K., Xu, S., Hawley, E. S., Weiss, S. A., Moxon, K. A., & Chapin, J. K. (2002). Rat navigation guided by remote control. *Nature, 417*, 37–38.
Thompson, E. (2001). Empathy and consciousness. *Journal of Consciousness Studies, 8*(5–7), 1–32.
Thompson, E. (2007). *Mind in life: Biology, phenomenology, and the sciences of mind*. Cambridge: The Belknap Press of Harvard University Press.
Thompson, E., & Stapleton, M. (2009). Making sense of sense-making: Reflections on enactive and extended mind theories. *Topoi, 28*(1), 23–30.
Tomasello, M., & Kaminski, J. (2009). Like infant, like dog. *Science, 325*, 1213–1214.
Topál, J., Gergely, G., Erdőhegyi, Á., Csibra, G., & Miklósi, Á. (2009). Differential sensitivity to human communication in dogs, wolves, and human infants. *Science, 325*, 1269–1272.
Varela, F. J. (1991). Organism: A meshwork of selfless selves. In A. I. Tauber (Ed.), *Organism and the origins of self* (pp. 79–107). Dordrecht: Kluwer Academic Publishers.
Warwick, K., Xydas, D., Nasuto, S. J., Becerra, V. M., Hammond, M. W., Downes, J. H., et al. (2010). Controlling a mobile robot with a biological brain. *Defence Science Journal, 60*(1), 5–14.
Weber, A., & Varela, F. J. (2002). Life after Kant: Natural purposes and the autopoietic foundations of biological individuality. *Phenomenology and the Cognitive Sciences, 1*, 97–125.
Wood, M. (2005). *In search of the first civilizations*. London: BBC Books.