Proceeding Paper

The Co-Evolution of the Extended Mind and the Internet †

Harry Halpin

Center Leo Apostel, Vrije Univerteit Brussel, 1060 Brussels , Belgium; hhalpin@ibiblio.org
† Presented at the Conference on Theoretical and Foundational Problems in Information Studies, IS4SI Summit 2021, Online, 12–19 September 2021.

Abstract: In this essay, we use the provocation of the possibility of artificial intelligence involving the physical merger of humans and machines to explore cognitive extension and the co-evolution of humans and technology. Our aim is to show that far from the stereotypical vision of cyborgs requiring chips in their brain, we are already becoming cognitively extended by the internet, and this is the product of a co-evolutionary development of the internet itself.

Keywords: extended mind; evolution; philosophy; internet; artificial intelligence

1. Introduction

The interaction between biological humanity and artificial intelligence raises the status of some increasingly pressing concerns that were once relegated to speculative fiction to its proper philosophical level. What about us crucially changes when we replace our flesh-and-blood arm with a robotic arm? Or, as predicted by Google CEO Eric Schmidt, when our brain is directly interfaced with Google via a neural implant [1]? Yet the fantastic vision of seeing a human with a mechanical arm or a chip implanted directly in their brain can actually cause a type of picture-thinking that obscures a crucial philosophical question: what are the precise conditions that form a biological and technological hybrid in the first place? Despite the hype about neuro-prostheses and cyborgs, in what ways does—or does not—the more mundane example of a human using a web-enabled smartphone exhibit some degree of bio-machine hybridization, and thus the creation of a web-extended mind? With the invention and subsequent failure of glasses that directly wire into Google, the stakes associated with whether or not the costly implanting of a chip into the brain is necessary to extend one's cognitive capabilities, or whether simply using a smartphone is sufficient, are becoming less theoretical, and more of a pressing issue for the general population.

Before going further off the rails into philo-fiction, what is needed is clarity on the conditions of cognitive extension in the here and now. How do we know if a given human and a technical apparatus count as an extended intelligence? We posit an answer by appealing to the Extended Mind Hypothesis [2], which can be iterated to produce a number of fairly straightforward conditions that can determine whether or not a particular relationship between a given human and their technologies qualifies as a fully cognitively extended hybrid, with the caveat that we assume (for humans at least) any such interesting bio-machine hybrid would also have to qualify as cognitive. Interestingly enough, this argument will show that actual cognitive extension does not require the direct interfacing of flesh and machine, so that genuine cognitive extension can take place in the much more common-place scenario of the aforementioned smartphone. To get to the point, implanting a chip in a human brain in order to become a cyborg is a rather absurd way to think about extending intelligence, as well as the wider relationship between humans and techniques, given our historical trajectory through technology leading up to the increasingly ubiquitous internet.

Yet our critique still leaves in place the idea of a future convergence between the biological and technological, as predicted by various proponents and detractors of the

Singularity Hypothesis [3]. Countering this idea that some sort of artificially intelligent singularity will be a future newcomer on the world-historical scene [4], we argue that the relationship between the biological and the technological is best seen through what may be a termed a co-evolutionary lens, where, to use the slogan of Andy Clark, humans have always been "natural-born cyborgs" [5]. Technologies for cognitive extension are the result of a fairly long, if philosophically under-considered, technological exploration of "man–computer symbiosis" and human augmentation, articulated by Licklider and Engelbart in the 1950s [6], a historical trajectory that should be explored in future work in detail in order to prove our point. So, in this extended abstract, we can only lay out the argument that the winding path to cognitive integration follows a reading of evolution that follows the work of philosophers such as Daniel Dennett, who conceive of evolution as an abstract algorithm [7], and of historian of technology George Basalla, who tries to apply such an evolutionary trajectory to technical artifacts [8].

2. The Conditions of Cognitive Extension

The Extended Mind Hypothesis states that under certain conditions, parts of the outside environment, such as technology, may qualify as part of the mind in some strict sense [2]. The paradigmatic example is Otto, who, due to being damaged with some kind of neurological short-term memory loss, can only navigate his way to the Museum of Modern Art (MOMA) with the aid of a notebook. It is claimed by virtue of what Clark and Chalmers term "the parity principle" that the notebook be given some of the cognitive credit, despite being outside the boundaries of the brain and skin [2]. Despite its seemingly strange implications, the Extended Mind Hypothesis is a straightforward extension of standard philosophical accounts of functionalism, such that any function that fulfills a cognitive role should count as genuinely cognitive regardless of what substratum a particular function is implemented on, be it biological or technological.

The standard objection would be that the notebook cannot count as part of the mind, due to some characteristic or another that renders it obviously non-cognitive, such as being made of paper rather than neurons. What matters may not be "where" a function is located, but whether or not it fulfills some criteria that justify it being part of a cognitively integrated whole system. To briefly iterate the criteria for cognitive extension including (but not limited to) technology, a summary of the conditions needed for a technology to qualify as a cognitive extension is necessary, before going further than Clark would likely approve of [5].

The first criterion is that the technology that fills the functional role must be portable, so that capabilities that count as part of the extended mind are accessible when needed. This does not mean they are always accessible, but accessible at least when needed. Even our biological mind may forget incidents and facts, and under non-normal conditions we may suffer brain damage after a catastrophe such as an automobile accident.

The second criterion is that the technology must be robust, so that capabilities are present when needed, and ensuring also that the function works consistently and within the needed time-frame necessary to solve problems. Having access to a French dictionary does not count as knowing French unless one can fluently respond to a French question in the socially appropriate time-frame.

The third criterion is that such a technology must be augmented by providing a genuinely new capability, significantly augmenting an existing capability, or replacing a missing capability. In Otto's case, the notebook replaces his working memory, but for many of us, notebooks (or smartphones, and perhaps one day implants) also augment our memory by allowing more precise recall of maps, phrases, and other items that escape our biological memory.

Lastly, the technology must be dove-tailed, or "evolutionary co-adapted" in the more controversial framework that we pursue in the next section [5]. This means that a new capability grows and changes in such a way that it becomes adapted to face the problems

facing the reproduction of the extended cognitive assemblage, and so is automatically trusted and endorsed, even if not "consciously" endorsed [2].

We can apply these four characteristics to a range of technologies to determine if they qualify as cognitive technologies. No technology is guaranteed to be a cognitive extension, but is only given as one in lieu of a particular well-defined problem. Given that a multitude of technologies may count as cognitive extensions under differing conditions, the key point is that there is no essence of the human, and that the mind is constantly assembled from differing configurations of biological and technical components in a given environment that can be described adequately as a cognitively integrated whole.

3. The Co-Evolution of Technology

To extend the Parity Principle to technology [2], if some technological component is essential in completing some functionally important role in the reproduction of a replicator, then we should at least consider giving that technological component its Darwinian due and classifying it as part of what Dawkins considered "the extended phenotype" of the replicator, or in other words, "all the effects of the gene upon the world" [9]. This should allow us to define in an effective manner the otherwise vague term "co-evolution" that tends to be abused by technophiles.

One of the reasons behind the sudden interest in artificial intelligence and the post-human is the dizzying technological change that threatens to eradicate the very concept of historical development that can be broadly construed as evolutionary in a Darwinian sense, as it seems that technological progress is both rapidly accelerating and veering towards some predestined end, such as the singularity. Appearances can be deceiving, and so we would like to put forward a contradictory hypothesis; namely, that the advent of cognitively integrated technologies is the product of a process of selection over often minor, incremental changes in a vast lineage of technological variations [8]. To extend the Parity Principle to technology [2], if some technological component is essential in completing some functionally important role in the reproduction of a replicator, then we should at least consider giving that technological component its Darwinian due and classifying it as part of what Dawkins considered "the extended phenotype" of the replicator, or in other words, "all the effects of the gene upon the world" [9]. This evolutionary hypothesis still holds likely true of not just the past, but also of the future of technology in a more general sense, although we will confine our analysis to the rise of the internet for the sake of using a well-known generic example. In going beyond the mundane, we state that the selection of certain technical artifacts has led to cognitive extension becoming increasingly widespread over the last half a century—it is no historical accident that we increasingly feel at one with our techniques.

The classic difference between humans and technology is that humans are the result of biological evolution, while technology is the product of human design, and thus not part of an evolutionary framework. In traditional accounts of biological evolution, phenomena such epigenetics are not entirely overthrown, and natural selection happens over genetic material that changes via mutation and sexual reproduction. Yet drawing from Dennett, evolution can be thought of functionally as an algorithmic process operating over replicators (such as DNA) in some search space [7]. Replicators are genotypes that code for particular functions in what is termed the vehicle of the replicator, where the replicator embodies these functions. The replicators are thought of as being carried by the vehicle, just as a human carries DNA. Selection is successful when vehicles, and thus the replicators, are not destroyed by their environment and get around to replicating. Evolution then selects replicators on an abstract level and can thus—controversially—operate not only over the biological vehicles that have DNA as replicators, but also over technical vehicles whose basic replicators are unknown.

Just because the replicator is not currently known does not automatically disqualify a process as being evolutionary—at least in the abstract sense but forward by Dawkins and Dennett. Darwin himself was unaware of Mendelian genetics, and thus argued in

The Origin of Species for natural selection without himself knowing about the existence of the biological replicator of DNA [7]. This is not to say that the concept of intentional technical selection would be anathema to Darwin. Darwin observed a vast number of variations in domesticated biological species caused by artificial selection under human guidance, and then argued that the environment itself could provide a process of selection without human agency, dubbed as natural selection. Darwin extended an argument about the design of biological species from domesticated species designed by humans to species that did not have a human designer. What evolutionary accounts of technology commit to is the reverse, arguing that technical artifacts that have an element of human design obey the same general principles as biological evolution. These techno-evolutionary frameworks move beyond a debate that poses natural selection as an alternative to natural selection, and become a general framework based on an abstract notion of selection that encompasses both biological and technical evolution, although biological selection is given primacy for the obvious historical reason that DNA preceded culture. Technology can be thought of as producing a vast number of variations, and a process of artificial selection operates over these varieties, of which only a small number are then replicated. Put differently, there is a vast variety of novel technological inventions, but very few are successful and so are replicated throughout their particular milieu [8].

Under this neo-Darwinian account, biological and technological evolution still have a crucial difference, as the replicators of technology are unknown: the social and economic forces that cause the processes of variation and selection in technology are complex and usually involve human intention. There are a number of theoretical options. In the work of Dawkins and other neo-Darwinians, genes are considered replicators insofar as they are considered information-bearing instruction manuals for producing organisms. While on some deep-level, such a metaphor is doubtless itself problematic, there is an oversight in that we need something equivalent to genes-as-replicators for technology. In this case, one answer is the concept of "memes" that Dawkins invented to attempt to show that genes were only one possible embodiment of evolution: "A unit of cultural inheritance, hypothesized as analogous to particular genes, as naturally selected by virtue of its phenotypic consequences in its own survival and replication in the cultural environment" [9] (p. 297). The elusive if much-lauded meme is a unit of imitation, whereby two vehicles that display the same behavior can be said to share a meme if their behavior cannot be explained by genetic factors [9]. The examples given of memes by Dawkins are broad-ranging: "tunes, ideas, catch-phrases, clothing fashions, ways of making pots, of building arches" [10] (p. 206). As given by pots and arches, technology is part and parcel of culture that can also be both imitated and selected, and so thought of as mimetic. Although further argumentation needs to be offered, we can consider that one type of meme is a protocol, which specifies how a particular computational artifact can be implemented and so reproduced.

Internet programs face selection pressure, as technical artifacts that are more successful in fulfilling whatever tasks are necessary to overcome the problems facing them in the environment will be replicated in proportion to their success. To return to the example of the map, many online maps have emerged since the advent of the web, and many have been reasonably successful at delivering directions. It was not until Google managed to collate vast amounts of mapping data and take advantage of new tweaks in the underlining method of how web pages delivered content to browsers that the mapping data became sophisticated enough, and able to respond quickly enough to users, to allow the program to be integrated into daily life. This caused the usage of Google Maps to be massively replicated, and other new programs such as Apple Maps adopted the same principles as closely as possible, while programs that failed to do so, such as Mapquest, fell into disuse. In this way, the technical programs were varied, but one was massively replicated. The underlying protocol started as just code run by Google, using previous standardized protocols such as HTML, and this may lead to more standardization in the future. Note that the evolution of Google Maps is not driven only by the code, but also its embodiment in mobile phones, and the utility of these mobile phones to humans.

As human users adapt technology and change its underlying protocols, leading to variations in and the selection of the program itself, it can be considered that the human and the technical artifact are engaged in co-evolution. It is precisely this co-evolutionary dynamic that leads to our own tendency to become technically dependent via cognitive extension.

4. Tentative Conclusions

So far, we have argued that merely using a smartphone with a mobile phone map fulfills the conditions for the Extended Mind, and thus achieves cognition integration with technology. While the whole argument has been made elsewhere [11], it can be seen that achieving cognitive integration is both more common and more difficult than simply sticking a wire or chip into a skull. First, cognitive integration in terms of the Extended Mind only holds under certain limited conditions, and these conditions are actually problem-specific. While a notebook might be good enough for Otto in navigating to a museum, it would not be robust enough to use in real time to escape from a sabre-tooth tiger or a drone strike. Yet the open question for the future is whether or not the co-evolution of technology that led to smartphones being robust, portable, adaptive, and dove-tailed will continue, such that a latency reduction—and thus a physical merger—between us and our technical machines will become somehow predestined, just like the singularity [3], the law of neo-dialectical materialism related to AI.

What is under-appreciated is that much internet technology has been guided more or less explicitly by some variation of Licklider's vision of Man–Computer Symbiosis, and that this process has been selected by both intentional (in the case of ARPA funding the internet) and unintentional (the commercial success of the web) factors. Licklider's ARPA funding made computing via the internet more robust and portable by eliminating distances in time and space, while Engelbart and Berners-Lee evolved the internet from teletype terminals connected to time-sharing machines into augmented and dove-tailed cognitive technology, such as smartphones. For the argument outlined here to be complete, the most important task in future will be to outline the historical argument for co-evolution in detail. The mind is no longer "in the head" of intelligent individuals, but is ultimately better conceived as a co-evolving, and historically contingent, collective intelligence.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Conflicts of Interest: The author declares no conflict of interest.

References

1. Schmidt, E.; Cohen, J. *The New Digital Age: Transforming Nations, Business, and Our Lives*; Penguin: New York, NY, USA, 2014.
2. Clark, A.; Chalmers, D. The Extended Mind. *Analysis* **1998**, *1*, 7–19. [CrossRef]
3. Kurzweil, R. *The Singularity Is Near*; Viking Press: New York, NY, USA, 2005.
4. Bostrom, N. *Superintelligence*; Oxford University Press: Oxford, UK, 2014.
5. Clark, A. *Supersizing the Mind*; Oxford University Press: Oxford, UK, 2010.
6. Waldrop, M. *The Dream Machine: J.C.R. Licklider and the Revolution That Made Computing Personal*; Penguin: New York, NY, USA, 2001.
7. Dennett, D. *Darwin's Dangerous Idea: Evolution and the Meanings of Life*; Simon and Schuster: New York, NY, USA, 1995.
8. Basalla, G. *The Evolution of Technology*; University of Delaware Press: Providence, RI, USA, 1985.
9. Dawkins, R. *The Extended Phenotype*; Oxford University Press: Oxford, UK, 1982.
10. Dawkins, R. *The Selfish Gene*; Oxford University Press: Oxford, UK, 1976.
11. Halpin, H. Does the Web Extend the Mind? In Proceedings of the ACM Web Science Conference, Paris, France, 2–4 May 2013; ACM Press: New York, NY, USA; pp. 139–147.