

Educating Cyborgs¹

Ben Kotzee

University of Birmingham

INTRODUCTION

In thinking about the nature of humanity, Donna Haraway's image of the cyborg is an arresting one.² Haraway draws attention to the extent to which human intellectual activity has become tied up with technology – to the point that much of human thought is only possible due to technology. To see the point, one only needs to consider how, for most of us, mental calculation has been overtaken by the calculator, how map-reading has become obviated by GPS, and how factual memory is increasingly supplanted by Google. Following Haraway, the idea has taken root that we have witnessed a transformation of human thinking into a hybrid kind of thinking that takes place both in our brains and in our technology, and is “extended” from inside our heads into our technology. In analytic philosophy, the point is made most forcefully by Andy Clark who holds that, with technology influencing and extending our thinking from cradle to grave, we are all, today, “natural born cyborgs.”³

The image of the cyborg has influenced education, too. “Cyborg pedagogy” holds that, because technology is changing the way that children think and learn, we should change our conception of what “learning” is.⁴ George Siemens, for instance, has proposed that:

[l]earning (defined as actionable knowledge) can reside outside of ourselves (within an organization or a database), is focused on connecting specialized information sets, and the connections that enable us to learn more are more important than our current state of knowing.⁵

If Siemens is correct, this sounds an ominous note for the traditional knowledge that we used to expect children to learn in school. Take the case of memorizing facts or mastering mental calculation. If improving technology makes it possible for children to rely more and more on external aids (pocket calculators, Google, etc.) to bolster and extend their own mental resources, is it still necessary to teach traditional school knowledge (such as mental arithmetic, recall of facts, etc.)? If children can calculate using a calculator, or can find facts using Google, why should we insist that they master mental calculation or build up a store of factual knowledge at all? S. Orestis Palermos and Duncan Pritchard ask exactly this question when they stress that education has *always* relied on the use of *some* technological aids to assist thinking.⁶ If we can accept, as a matter of course, that pupils should be able to use pen and paper to work out mathematical problems and – in many settings – calculators, why should we not go further and let them Google during exams?

In this essay, I make a connection between the debate about posthumanism and debates about learning. I hold that, while suggestive, cyborg thinking is tangential to our real concerns in studying education. As an alternative, I endorse an account according to which it is not the *mind* that is extended, but our knowledge gathering *practices*.

FROM CYBORGS TO CYBORG PEDAGOGY

In her “Cyborg Manifesto,” Haraway notes the extent to which technological advances have begun to complicate the nature of human identity, of what it means to be human. What Haraway wrote about in 1991 is even more the case in 2017. Consider how much of what we do depends utterly on technology: I found my way to this venue using the GPS on my phone, on the way I found today’s news on my laptop, in the gym my watch tells me how hard to train (mostly harder). If you wear

a pacemaker or insulin pump, technology helps to keep you healthy or, even, alive. As Andy Clark notes, our technology literally begins to feel part of us. He writes:

As technology becomes portable, pervasive, reliable, flexible, and increasingly personalised, so our tools become more and more a part of who and what we are ...The temporary disability caused by a dead battery is unnerving. It seems we just aren't ourselves today. (The loss of my laptop ... underlined this in a painfully personal way. I was left dazed, confused, and visibly enfeebled – the victim of the cyborg equivalent of a mild stroke.)⁷

Because we have become so dependent on technology, Haraway argues that today, “we are all chimeras, theorised and fabricated hybrids of machine and organism, we are cyborgs.”⁸ Clark and, to a large extent, Haraway stress that this conclusion should not strike us as particularly weird. In *Natural Born Cyborgs* Clark holds that, to a large extent, it was always thus; people have *always* depended on tools to help them think. Take Clark’s “Extended Mind Thesis” (EMT; worked out with David Chalmers). The EMT is the thesis that the human mind is not only contained in the head, but extends outwards into the tools that people use to aid their thinking. The central argument for EMT is found in Clark and Chalmers’s point about the parity of thinking processes that take place in the head and outside it. Clark and Chalmers hold that:

If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world *is*... part of the cognitive process.⁹

To see the point, consider the following example:

Two friends, Inga and Otto, arrange a meeting for a few days from now. The agreement is that they will meet at the Museum of Modern Art (MOMA) at 11am on a certain day. Inga has a fantastic memory and memorises all her appointments every day. On the day of the appointment, she remembers that she has fixed to meet Otto and goes to meet him at MOMA at 11am. By contrast, Otto suffers from a medical condition that has affected his memory and really struggles to remember appointments. To compensate for this memory problem, Otto keeps a scrupulous diary. The day of the appointment, Otto looks up his activities for the day in the diary. He sees that he has written down that he must meet Inga at MOMA at 11am. He goes there at 11am and the friends meet one another.¹⁰

Clark and Chalmers's argument is this: If you are happy to grant that Inga's *memory* is part of the thinking process that gets her to arrive at MOMA on time, why would you not grant Otto's *diary* that same status? If Inga *knows* that her appointment is at MOMA at 11am (because she remembers it), why should we not say that Otto *knows* this too (based on what he has written down in his diary)? Crucially, note that, if you are happy to say that part of Otto's *knowledge* is contained not inside his head, but in his diary, then you are happy for Otto's very knowledge to extend outside of his head and into his diary. This is Clark and Chalmers's extended mind thesis: the mind is not contained solely in the head, but can be contained in our technology. Today, this means not only our diaries, but our computers, or our mobile phones, or satnavs. If what Clark and Chalmers and Haraway note is true, my mobile phone not only helps me

think, it is part of my thinking. Thinking of humans as cyborgs changes our conception of what it means to be human today; “posthumanism” is a critical and questioning stance to the nature of being human.

I would like to stress two things about the “post-humanist” turn of thought. First, posthumanism has deep roots in phenomenology and in philosophical thinking regarding materiality. Thinkers in the phenomenological tradition (e.g. Merleau-Ponty, De Beauvoir, or, more recently, Dreyfus) have long attacked mentalist pictures of human thinking and stressed that human thinking or what it is to be human is not to be found in the realm of thought or mind, but in the realm of the body. The posthumanist way of thinking takes this to the next level by holding that “materiality” is not only about the biological body, but even includes non-biological extension of the body (e.g. technology). If post-humanist thinking calls into question what we humans are *now*, however, it also draws into question what we ever were. As Clark holds, if technology has changed the very nature of what we are as creatures today, this equally draws into question what human nature is and was all along. Posthumanism is not just a point about minds and technology, but about the involvement of our bodies (and, as it turns out, our tools) in thinking.

A second important matter is that many posthumanist thinkers see this metaphysical claim about the nature of the human – that there is no pure “human” but blurred lines between animal, human, and machine – as a *political claim*. Haraway writes that “a cyborg world might be about lived social and bodily realities in which people are not afraid of the joint kinship with animals and machines, not afraid of permanently partial identities and contradictory standpoints.”¹¹ Also, Charles Garoian and Yvonne Gaudelius, for instance, see this reconception of human identity along “cyborg” lines as an act of *resistance*. They argue that the cyborg

“serves as a complex metaphor that represents the body/technology hybrid while it exposes the cyborg’s dialectical pedagogy of inscription and resistance.”¹²

The political project behind posthumanism is to liberate people from preconceived notions of identity that have, perhaps, been foisted on them. What is called “cyborg pedagogy” (Angus, Cook, and Evans, 2001) is part of this political project. According to Rosemary Klich, cyborg pedagogy:

suggests an approach to learning that considers specific manifestations of the posthuman being within pedagogical process, encourages students to examine their own existence within a “posthuman condition,” and positions the learner as part of a network of material and mediatized components.¹³

As Klich explains it, what is so particularly liberating about cyborg pedagogy is that it encourages students to question the essence of being human and, with it, the essence of human knowledge. Indeed, Tim Angus, Ian Cook, and James Evans, who did much to popularize the term “cyborg pedagogy,” advocate:

an approach to teaching and learning which takes as its nemesis the “banking system of education” where, to summarise it crudely, students are encouraged to learn dominant understandings of the world and to repeat those dominant understandings back in assessments which determine their academic progress (Giroux & McLaren, 1994; hooks, 1994). These understandings are, not surprisingly, structured through the binary oppositions that Haraway sees as fundamental to such exclusions ...¹⁴

Instead, they call for a pedagogy that encourages “students to (a) identify and be critical of these binary logics in action ..., and (b) undercut and/or find spaces between them in order to undermine the forms of domination which result from their taken-for-granted use.”¹⁵

Advocates of cyborg pedagogy hold that *the very realization that we may be cyborgs* is educationally important. Really, cyborg pedagogy is not about advocating a particular view regarding technology, or cultivating in students a particular relationship with technology, but about getting students to question human identity and the nature of human thinking. In all of this, one may say, the cyborg is a metaphor or a stalking-horse for what is really at issue: human identity. Because the thrust of cyborg pedagogy – that we should encourage students to question what it means to be human – relies on the metaphor of the cyborg, I have decided to tackle this metaphor head on and question it. Is the central claim behind posthumanism – that we people have literally become cyborgs – *right*? Because my own field is analytic epistemology, I focus in what follows more on Clark’s formulation of the idea than Haraway’s; however, I take their ideas to be parallel.

CLARK’S EXTENDED MIND THESIS

Haraway and Clark hold that, despite the fact that most of us are not literally part-machine (like the Six Million Dollar Man or Bionic Woman), the pervasiveness of technology in our intellectual lives makes us, for practical purposes, cyborgs. As Clark puts it, our minds “extend” into the technology we use.

The most important objection to the Extended Mind Thesis (EMT) is the “objection from cognitive bloat.” The objection goes as follows. Say that one grants that a person’s mind is not located only in

their head, but extends to encompass all of the cognitive aids that they use (like notebooks, calculators, or even the whole internet). Does this mean that we must now imagine that everyone already knows everything that is contained in those cognitive aids? Take the example of looking something up in a telephone book. I do not know the telephone numbers of everyone who lives in my city; however, I can easily look up anyone's phone number (provided it is listed). Would we say that, because I can so easily look it up, I *know* the telephone number of everyone in my city? No – at best we would say that “I know how to look it up.” The objection against Clark's position is this: if we extend our conception of what the mind is too easily, it will lead to our granting people too much knowledge and too easily.

Clark's response to this objection was to propose a restriction on cognitive extension to the effect that a cognitive aid would only extend one's mind if it were properly *integrated* in one's thinking. Clark holds that it is not good enough if one uses the cognitive resource only occasionally. But, he holds, a cognitive aid does become part of one's mind if (i) it is reliably available and typically invoked in one's thinking, (ii) the information in it is automatically endorsed (and not usually subject to critical scrutiny), and (iii) the information is easily accessible when required.¹⁶ These conditions clearly apply for someone like Otto in Clark and Chalmers's diary example, and have come to be known as the “trust and glue” conditions for cognitive integration. Still, it is doubtful whether the “trust and glue” conditions solve the problem of cognitive bloat. To see this, consider that many university students have Wikipedia readily available on their tablets or smartphones, typically cite Wikipedia articles in what they write, and almost unfailingly endorse what Wikipedia says about matters from archery to zoology. Despite the fact that they trust Wikipedia and are cognitively “glued” to it, it is clear that they do not know everything in Wikipedia.

INTEGRATION AND ASSESSMENT

The debate in this area turns on the nature of the integration of a cognitive aid with one's thinking processes. Of particular importance to education are the implications for assessment. In thinking about any educational process, being able to specify what someone knows and does not know is crucially important. In teaching, teachers must always tailor their teaching to what learners know and do not know. First, they must avoid teaching what learners already know – being told something you already know is to be “reminded” not to be taught. Second, teachers must make sure that learners genuinely come to know the (new) things that they are being taught – that what the teacher teaches is actually absorbed, cognitively speaking. If the teacher realizes that what she is communicating is not being absorbed any longer, it is best for her to stop her teaching or to go back, try again, take a break, or refresh. Simply put: teachers need to be able to assess the state of learners' knowledge in order to know where to start their teaching and where to stop. The problem that EMT creates in the context of education now becomes clear. . If the teacher must count the content of all the cognitive aids (such as the contents of Google) that a student uses as *part* of what their students know already, then it seems that students will have to be counted as “already knowing everything in Google,” making it pointless to try and teach them anything, “having a capacity as great as the internet” for learning, meaning that the teacher will never have to stop her teaching. Accepting Clark's EMT clouds the picture of who exactly learns what from whom and when.

In this vein, Kenneth Aizawa provides an example to probe Clark's thinking.¹⁷ Two students, Otis and Opie, are both due to sit the same chemistry exam at the end of the term. Opie attends every class and spends some time mastering the new material after every class. Before the day of the exam, he diligently studies for the exam and, on the day

of the exam, scores an A. Otis is a much less diligent student. He does not attend class and spends no time mastering the material or studying for the exam. The night before the exam, he copies all the important formulae onto little notecards and hides these away in his clothes. During the exam, he cribs from his cards. He also scores an A. The professor can recall seeing Opie in class, but not Otis; smelling a rat, she calls him in and asks how he managed to score an A on his exam. Aizawa completes the example as follows:

the professor asks Otis how he was able to get an “A” on his first exam. Otis explains that he had basically copied out notecards for the exam, which he read during the exam. Outraged, the professor called Otis a cheater and threatens to change his grade to a “0”. Otis, however, is resourceful and draws upon a philosophy class he did not sleep through one day. He replies that the information in his notecards function just like the information constituting an ordinary non-occurrent belief that his friend Opie had. It just so happens that the information lies beyond his skin.¹⁸

Aizawa points out that, as a matter of educational practice, we find Opie’s epistemic behaviour much better than Otis’s. This is demonstrated in the fact that we award Opie a pass mark on the exam and Otis a fail. The difference between Opie and Otis is easy to explain. While Opie has internalized knowledge of chemistry in such a way that he can answer the question on a chemistry exam independent of any cognitive aids, Otis has not and can answer the exam only with the help of an external aid (his notes).

Unfortunately, Aizawa’s “exam” example will only go so far. As Pritchard and Palermos point out, *some* technology (e.g. pen and paper)

is entirely allowable in all exams and other technology (e.g. a calculator in mathematics and a dictionary in foreign languages) is allowable in most exams beyond the basic level.¹⁹ The question is exactly what aids we should allow and in which exams. Indeed, Pritchard and Palermos go on to suggest that we should allow students more aids today than we used to in the past (including allowing students to Google during exams). Besides the descriptive observation that educators regularly distinguish between intellectual performance aided by and unaided by technology, we need a principled answer to the question what is *better* about unaided intellectual performance (if anything)?

INTEGRATION AS METAPHYSICS AND EPISTEMOLOGY

Clark is a philosopher of mind, and the EMT is a theory in the *metaphysics of mind*. For Clark, our minds consist in a hybrid of our brains and our technology. In this theory, the concept of integration fulfils a metaphysical role; it stipulates when such a hybrid mind – a “cyborg” – has truly come into being. Notice, however, that the question whether a cyborg has come into being is a different question from what the cyborg knows. Take the ordinary way that we speak about persons knowing things. Persons know some things and do not know other things; we establish whether a person knows some truth *p* by asking whether they believe that *p*, whether *p* is true, and whether they have justification for believing *p*. If we follow Clark’s reasoning we should include, in the category of persons, also cyborgs. This is well and good, but it follows that cyborgs are subject to the same epistemic demands as persons are: it is only if a cyborg believes some truth *p* with justification that they know it. We can say that, what Clark’s integration condition establishes is *whether* a being is such that we should count it as a “person” for epistemic purposes. It is a *further* question, however, whether that hybrid person or “cyborg”

actually knows something or does not know it.

We can use another example by Aizawa to explain the point: we can imagine a person/device combination that is well enough integrated in the metaphysical sense to be called a cyborg, while at the same time not conducting themselves creditably, epistemically speaking. In Aizawa's example, a keen cyclist called Mort buys a new cycling computer for his bicycle. Mort comes to rely on this computer to show him how fast and far he rides, how hard he is training etc. Over time, Mort's cycling computer becomes integrated in his cognitive character; he always has it to hand when cycling, he consults it all the time to know his speed, and he relies on it implicitly (cycling faster when it shows he is going too slow, cycling slower when it shows he is going too fast, stopping when it shows he has covered the right distance ... and so on). By Clark's standards, Mort-and-his-computer has become a cyborg. However, it is a further question whether Mort-the-cyborg *knows* how far and fast he rides. If what the computer tells him is *false*, for instance, he will not *know* how far and fast he rides. In Aizawa's example, this is exactly the case; because Mort has failed to calibrate the computer to his bicycle's wheel size, the cycling computer gives him systematically false information.²⁰

The Mort example is one where a person can be counted as a cyborg according to Clark's definition, but still fails to know because the cognitive technology he uses is not reliable (due to being set up wrong). The implication is this: whether we want to say that our minds have become melded with our machines is a red herring to the educational matters that concern us. While very interesting from the metaphysical point of view (are all our students really little cyborgs?), the real educational questions that will concern us, such as "what are our students learning?" and "have they learned effectively enough for us to be satisfied with their learning?" are not settled on this metaphysical basis alone. What we need to estab-

lish is whether, for our purposes, children have learned what they need to learn *well enough*. What we need is a way to think about what learners can do with and without technology, and how that should influence our thinking about what “good enough” learning really means.

CREDIT FOR PRACTICE

At bottom, what is driving the problem is one of (what virtue epistemologists call) “epistemic credit.” In Aizawa’s example, does Opie deserve epistemic credit for passing his exam? Yes, because he mastered the material and then solved the problems on the exam independently. Does Otis deserve epistemic credit for passing his exam? Well, not as much as Opie. Otis certainly deserves a degree of credit – after all, he understood the material well enough to make notes about it and he understood the notes well enough to pass the exam with their help. But, he still had to cheat to pass the exam. Whether Otis used notes or a mobile phone, computer, or tablet, Otis does not deserve our epistemic admiration as much as Opie does.

The credit problem in this area is simple: can epistemic credit be attributed to a person based on what a machine tells them or not? For our educational purposes, we don’t (yet!) think it is quite good enough to allow computers to be used completely generally in exams – we still value the student’s performance that is not aided by technology more than that of the student who is aided by technology. Outside our exam halls, however, we do give people quite a bit of credit for what they can achieve using technology. The fact that people in our society know and can do all sorts of interesting things are genuine cognitive achievements enabled by technology: Think of the accountant who knows how much tax I have to pay based on a spreadsheet, the delivery driver who finds my house from the other side of town using his satnav, or the doctor

who can see inside my very bones using an x-ray machine. Denying that people like these genuinely know what cognitive technology tells them would be unrealistic, for it would render their success in accomplishing these important tasks mysterious (“The accountant didn’t really know how much tax I had to pay; luckily his computer told him”).

How, though, does one explain that we can attribute epistemic success to *people* when we all well know that they could never have achieved those successes but for their use of technological aids? Two broad routes are open to us here:

1) Either we can say that person and machine together constitute a person/machine hybrid or cyborg that, for epistemic purposes, can be the subject of the attribution of success as much as an ordinary person can. (This is the approach taken by Clark).

Or

2) We can stick to our ordinary metaphysics of persons and tell a different story about how people can take credit for what machines tell them (even in cases where they are seemingly fully reliant on those machines).

I want to argue for a version of option 2.

Richard Menary holds that cognitive integration should be understood in the context of how the knower manipulates external cognitive resources in order to gain knowledge.²¹ Menary holds that extended-mind-style arguments do not recognize the depth of what is involved in our practices using technology. Being able to use external resources – such as Otto’s notebook or Otis’s notes – presupposes that someone understands how to use these resources and what they communicate. This manipulation of external processes makes available modes of thinking that purely internal representation does not. Manipulating

cognitive devices does not merely extend our cognitive range. Instead, part of what it is to be a specifically human thinker is due not only to our nature as social beings but also to our use of technology. This gives thinking about extended cognition a role to play in explaining how human thought evolved and why it has the specific character that it does at all. As Menary holds: “... we get to be readers and writers, mathematicians and so on by a process of transforming existing cognitive abilities to perform new, cultural, functions.”²²

Menary, then, proposes a different way in which one can view the idea that cognition extends. Study of extended cognition in the style of Clark is the study of the causal relationships between human actors and cognitive technology, which can put cognitive technology on a par (functionally speaking) with “in-the-head” cognitive processes. Menary calls this “artefact extension” (AE). What Menary calls enculturated cognition (EnC) on the other hand, is the study of how our cognitive practice of using cognitive technology works and how this transforms human thinking. On Menary’s account, what a person deserves credit for in using external cognitive resources – and that includes machines – is using it correctly, and well, in accordance with the right *practice* for using it.

Following Menary, we can hold that it is our *practices* of using technology that are extended, not our minds as such. But to say that our practices of using technology are, by their nature, extended, and that we cannot engage in the kind of thinking that we do with technology without it, is already to say much about the nature of human thinking and “the human mind.” It is to say that human thought extended by technology takes on a different character and cannot be realized in any other way than by people working with technology. Because practices are passed on and picked up socially, there is no other way to become part of a practice (or to begin to practice in a certain way) but to *learn* it.

As Meredith Williams holds, saying that a practice is social and that it is learned is two sides of the same coin.²³

If the use of cognitive technology is a form of practice and if any practice must be learned, then *learning* plays a crucial role in understanding the creditable use of technology. To see this, consider that one has to learn to use even the most user-friendly machines. All cognitive technology presupposes at least some background of learning in order to use them. For instance, to use a satnav system when driving presumes some understanding of space, of left and right, of the rules of the road, of the meaning of traffic signs and road markings, etc. To use Google, one needs to understand search terms, search operators, and search results; moreover, one needs to be able to use language and type. These examples show that, even though technology sometimes plays a distinctive enabling role in making possible enhanced human knowledge, the user of the technology brings a crucial background of learning to her interaction with cognitive technology. This is also why I think we should give Otis (in Aizawa's example) *some credit* for what he had learned. He had, after all, learned at least enough to make effective notes and to use these notes effectively to get an A on his exam.

EDUCATIONAL IMPLICATIONS

I started by asking what we should think about the relationship between aided and unaided knowledge in education. Is there any reason to prefer that students be able to perform certain intellectual tasks (that could be performed by cognitive technology) for themselves? Or are aided and unaided knowledge equivalent, and so, we need not prefer unaided intellectual performance in the education system? To answer, it is necessary to step back and consider the motivation of philosophers of mind in proposing the theory of the extended mind at all.

As we saw, Clark notes the functional equivalence in cognition between processes that have their home in the brain and processes outside the body, and concludes that the mind extends outside the head. Fascinating as this is, what does this consideration about the metaphysics of mind tell us about how to conduct education? The observation that the mind extends and that, for many people, cognitive technology is intimately interwoven with their thinking, translates, at best, into a very broad permission to use technology in the classroom. Clark and others (including me) would say that, given how pervasive technology is in modern thinking, it would be absurd and counter-productive to deny school children the chance to learn to use these same tools for thinking. Admitting this still leaves very many pedagogical questions un-answered, however; such as at what age, which tools, in what order they should be introduced, how children should be taught to use them, etc. Saying “the mind extends” is fine. But practically, we want to be able to provide guidance and help to students to extend their minds by using external resources. The instructions “wear a smartwatch” or “use Google” or “get your calculator out” do not tell children in detail how to use any of these pieces of technology or how to use them well.

If what I have suggested above is right, the technological education we want to give children will be designed to help them become part of responsible *practices* (in Menary’s sense) of using technology. This will require teaching students about technology, letting them practice with it, and also inculcating attitudes to the use of technology. Moreover, it will involve teaching students about (and letting them discover for themselves) what the limits of technology are, and when technology is and is not reliable. In order to do so, a space will always remain for teaching students what problems they can solve by thinking alone.

- 1 An extended version of this essay was published as: Ben Kotzee, "Cyborgs, Knowledge, and Credit for Learning," in *Extended Epistemology*, eds. J. Adam Carter, Andy Clark, Jesper Kallestrup, S. Orestis Palermos, and Duncan Pritchard (Oxford: Oxford University Press, 2018). Parts are reproduced here by permission of Oxford University Press.
- 2 Donna Haraway, "A Cyborg Manifesto: Science, Technology, and Socialist-Feminism in the Late Twentieth Century" in *Simians, Cyborgs and Women: The Reinvention of Nature* (London: Routledge, 1991).
- 3 Andy Clark, *Natural Born Cyborgs* (Oxford: Oxford University Press, 2003).
- 4 Tim Angus, Ian Cook, and James Evans, "A Manifesto for Cyborg Pedagogy," *International Research in Geographical and Environmental Education* 10, no. 2 (2001): 195-201.
- 5 George Siemens, *Connectivism: a learning theory for the digital age* (2005), <http://www.elearnspace.org/Articles/connectivism.htm>
- 6 Spyridon O. Palermos and Duncan Pritchard, "To Google, or Not to Google, During Exams?," *Educational Technology and University Business* (2015).
- 7 Clark, *Natural Born Cyborgs*, 10.
- 8 Haraway, *A Cyborg Manifesto*, 150.
- 9 Andy Clark, and David Chalmers, "The Extended Mind," *Analysis* 58, no. 1 (1998): 7-19, 9.
- 10 Ibid.
- 11 Haraway, *A Cyborg Manifesto*, 154.
- 12 Charles Garoian and Yvonne Gaudelius, "Cyborg Pedagogy: performing resistance in the digital age," *Studies in Art Education* 42, no. 4 (2001): 333-347.
- 13 Rosemary Klich, "The 'unfinished' subject: Pedagogy and performance in the company of copies, robots, mutants and cyborgs," *International Journal of Performance Arts and Digital Media* 8, no. 2 (2012): 155-170, 155.
- 14 Angus, Cook, and Evans, *Cyborg Pedagogy*, 198-199.
- 15 Ibid.
- 16 Andy Clark, "Memento's Revenge: the extended mind, extended," in *The Extended Mind*, ed. Richard Menary (London: MIT Press, 2010), 46.
- 17 Kenneth Aizawa, "Extended Cognition, Trust and Glue, and Knowledge," in *Extended Epistemology*, eds J. Adam Carter, Andy Clark, Jesper Kallestrup, S. Orestis Palermos, and Duncan Pritchard (Oxford: Oxford University Press, 2018).
- 18 Ibid., 10
- 19 Palermos and Pritchard, "To Google, or Not to Google, During Exams?."

20 Aizawa, "Extended Cognition," 3.

21 See Richard Menary, "Cognitive Integration and the Extended Mind," in *The Extended Mind*, ed. Richard Menary (London: MIT Press, 2010) and Richard Menary, "Cognitive Practices and Cognitive Character," *Philosophical Explorations* 15, no. 2 (2012): 147-164.

22 Menary, "Cognitive Practices and Cognitive Character," 152.

23 Meredith Williams, *Blind Obedience: The Structure and Content of Wittgenstein's Later Philosophy* (London: Routledge, 2009), 198.