Journal of Evolution & Technology

A peer-reviewed electronic journal published by the Institute for Ethics and Emerging Technologies

ISSN 1541-0099

26(2) – October 2016

The Future of Brain-Computer Interfaces: Blockchaining Your Way into a Cloudmind

Melanie Swan
Philosophy and Economic Theory, New School University, New York NY

m@melanieswan.com

Journal of Evolution and Technology - Vol. 26 Issue 2 – October 2016 - pgs 60-81

Abstract

The aim of this paper is to explore the development of brain-computer interfacing and cloudminds as possible future scenarios. I describe potential applications such as selling unused brain processing cycles and the blockchaining of personality functions. The possibility of ubiquitous brain-computer interfaces (BCIs) that are continuously connected to the Internet suggests interesting options for our future selves. Questions about what it is to be human, the nature of our current existence and interaction with reality, and how things might be different could become more prominent. I examine speculative future scenarios such as digital selves and cloudmind collaborations. Applications could be adopted in tiers of advancing complexity and risk, starting with health tracking, followed by information seeking and entertainment, and finally, self-actualization. By linking brains to the Internet, BCIs could allow individuals to be more highly connectable not just to communications networks but also to other minds, and thus could enable participation in new kinds of collective applications such as a cloudmind. A cloudmind (or crowdmind) is the concept of multiple individual minds (human or machine) joined together to pursue a collaborative goal such as problem solving, idea generation, creative expression, or entertainment. The prospect of cloudminds raises questions about individual versus collective personhood. Some of the necessary conditions for individuals to feel comfortable in joining a cloudmind include privacy, security, reversibility, and retention of personal identity. Blockchain technology might be employed to orchestrate the security, automation, coordination, and credit-assignation requirements of cloudmind collaborations.

1. Brain-computer interfaces

Brain-computer interfaces (BCIs) are any manner of technology that might link the human brain to communications networks such as the Internet. In more detail, a brain-computer interface, brain-machine interface (BMI), neural prosthesis, etc., is typically a computational system implanted in the brain that allows a person to control a computer or other electronic device using electrical signals from the brain (Peters 2014). Individuals use BCIs to generate alphanumerical characters on a computer screen in the following way. The BCI equipment registers the electrical output of the brain when the eyes are focused on a particular location or quadrant of a computer screen – on the "q" in a stretched-out string of letters, for example – and outputs the letter onto the computer monitor (Mayo Clinic 2009).

The primary aim of BCIs at present is repairing human cognitive and sensorimotor function. One of the most widely adopted uses is cochlear implants, where a small computer chip is substituted for damaged control organs in the inner ear. The chip transforms sound waves into electrical signals that are interpretable by the brain (Friehs 2004). Vision restoration is another application: here, implantable systems transmit visual information to the brain. Two-way BCIs are another form of the technology under development, using both output and input channels for communication between the brain and the external world. There would be the usual BCI communication output from the brain in the form of translating neuronal activity into electronic commands to move robot arms, wheelchairs, and computer screen cursors. In addition, feedback from this activity could be input back into the two-way system via electrical brain stimulation that delivers signals into the brain (Bo 2015).

BCIs comprise an active area of research and could start to integrate advances from adjacent fields such as neuroscience, nanomaterials, electronics miniaturization, and machine learning. For example, one neuro-imaging research project is starting to make guesses as to what participants see during brain scans, purporting to be able to distinguish between a cat and a person (Smith 2013). Merging this kind of functionality with BCIs might produce new applications. Other experimental BCI projects have been proposed. One is Neocortical Brain-Cloud Interfaces: autonomous nanorobots that could connect to axons and neuronal synaptic clefts, or embed themselves into the peripheral calvaria and pericranium of the skull (Boehm 2016). Another project, Brainets, envisions linking multiple organic computing units (brains) to silicon computing networks (Pais-Vieira 2015). A third project is Neural Dust, in which thousands of 10-100 micron-sized free-floating sensor nodes would reside in the brain and provide a computing processing network (Seo 2013).

2. Future applications of BCIs

So far BCIs have been conceived primarily as a solution for medical pathologies. However, it is possible to see BCIs more expansively as a platform for cognitive enhancement and human-machine collaboration. The BCI functionality of typing on a keyboard with your mind suggests the possibility of having an always-on brain-Internet connection. Consider what the world might be like if each individual had a live 24/7 brain connection to the Internet. Just as cell phones connected individual people to communications networks, BCIs might similarly connect individual brains to communications networks. I propose a variety of BCI applications and concepts throughout the rest of this paper, all of which are speculative and not in development to my knowledge.

In one sense, ubiquitous BCIs are expected. It is contemplated that communications technology, already mobilized to the body via the cell phone, could be "brought on board" even more pervasively. BCIs are merely a next-generation improvement to the current situation of people constantly staring at their phones. In another sense, though, BCIs are not only a "better horse" technology: they are also a "car" in that it is impossible to foresee the full range of future applications that might be enabled from the present moment. BCIs pose a variety of practical, ethical, and philosophical issues. Life itself and the definition of what it is to be human could be quite different in a world where BCIs are widespread. Some of the immediate practical concerns of BCIs could include invasiveness, utility, reversibility, support, maintenance, upgradability (hardware and software), anti-hacking and anti-virus protection, cost, and accessibility. Beyond practical concerns, there are ethical issues regarding privacy and security. For example, neural data privacy rights are an area where standards need to be defined (Swan 2014a).

There could be at least three classes of BCI applications introduced in graduated phases of risk and complexity: biological cure and enhancement; information and entertainment; and self-actualization (realization of individual cognitive and artistic potential). Each of these merits separate discussion.

2.1 Health and enhancement BCI applications

One first class of BCI applications could relate to cure and enhancement. These applications can be framed as providing an "Apple HealthKit or Google Fit for the brain." The idea is to employ BCIs as a constant health monitor, pathology resolver, and neural optimizer. One of the great promises of BCI technology is that applications such as daily health checks might run automatically in the background to improve our lives. Periodic health checks could be orchestrated seamlessly by ambient quantified-self smart infrastructure (essentially the next generation of unobtrusive sensors worn on the body such as smartwatches). Personal biometric data could be transmitted to longitudinal health profiles in electronic medical records. This could facilitate the development of advanced preventive medicine. Preventive medicine is maintaining a state of health by detecting and resolving potential conditions in the 80 per cent of their life cycle before they become clinically diagnosable (Swan 2012b). Neural data streamed from BCIs to secure health data banks could finally start to allow amassing of "big health data," datasets large enough to study the longitudinal norms of brain patterns and cognitive well-being. A variety of health management and neural performance enhancement applications could ensue.

Personal biometric data collected by cellular telephone applications are an example of how personal data from BCIs might be treated. Norms for collecting and storing personal biometric data are starting to be codified. Ostensibly, neural data is just a special case of personal biometric data, with additional sensitivities. Apple HealthKit, for example, automatically captures 200 health metrics per day via the iPhone and seamlessly uploads them to the Internet cloud for subsequent on-demand analysis (Swan 2015d). Google Fit on the Android platform performs a similar function (Welch 2014). However, despite the potential benefits of automated health data collection, appropriate social and legal contracts with technology providers are not yet completely in place. Individuals may not fully grasp how their personal data is being collected, stored, and used (including being sold to third parties). This is important since personal medical information is a valuable asset. Health data may be worth ten times more to hackers than financial data such as credit card numbers and personal identity (Humer 2014). Even though cell phone users "must explicitly grant each application the permission to read and write data to the HealthKit store" (iOS Developer Library 2015), health-tracking data may be collected without full user awareness (other than by having agreed to the initial phone activation agreement). When installing applications, it can be easy for users to accede quickly to requested permissions without fully understanding what they entail in terms of granting access to personal biometric data.

2.2 Information and entertainment BCI applications

A second class of BCI applications is related to information and entertainment. One application could be brain-based information requesting. Information query could be both pushed and pulled: automatically pushed to users per pre-specified settings, or pulled (requested) by users on demand. Data notifications could be presented in the mind's visual space. This would be the analog to information cards or short data messages being posted to a phone or smartwatch. Here, the information could be presented in the brain, for example as an unobtrusive notification in the lower visual field. BCIs could have Google Now type functionality, making contextual guesses about information that might be relevant in the moment (such as transportation delay information).

BCIs could be the interface for immersive experience, conceptually similar to internalizing virtual reality headsets inside the body. The idea would be to have an onboard Oculus Rift, Meta 2 (Jabczenski 2016), or MindMaze (Lunden 2016). This could allow "HUD-sharing (heads-up display)," as is possible in video games now, and beyond: deeper levels of experience sharing and the "transparent shadowing" (Boehm 2016) of others for purposes ranging from learning to entertainment. A variety of contexts for experience sharing have been suggested, for example apprenticeship, scientific discovery, sports matches, music concerts, political rallies, and sex (Kurzweil 2006). In one example, Greg Bear's science fiction novel *Slant* (1998) explores an updated version of *Brave New World* feelies (movies with sense and touch, not just audio and video). Individual experience feeds could become marketable not just for entertainment purposes, but also for personal and societal record keeping.

Consider that in the future you might grant different levels of access to your personal experience feed. This could include selecting the payment models based on the situation, for example, fee-based or open-source contributions. Live events could provide interesting situations of sharing personal experience into the group memory feed. Computing algorithms could aggregate arbitrarily many contributor threads into a single summary. The crowdfile (e.g., a group experience file) for an event could be a new means of recording human history. After the fact, the event crowdfile could be accessed just as Wikipedia, Twitter, and YouTube are now. During the event, remote participants could join the live crowdfile, similar to live-streams now. To the extent that individual experience files or their "diffs" (salient differences from the aggregate file) could be stored expediently, various accounts of history could be kept simultaneously. Finally, multiple accounts of events could be available from different standpoints (similar to instant replay from different cameras). Any assessment of public opinion such as political polling could undergo a substantial shift as many more individual and collective reactions might be known in detail, and also in real-time. With BCIs, to the extent shared by the owner, experience files could become like any other digital content: a creative object for others to take up, reformulate, repurpose, reinterpret, "mash up," and share back out to the public venue. Just as the Selfie (a self-taken photograph) was the killer app of photos, and moment-showing was the killer app of video blogs (Dedman 2007), some form of "Brain Selfies," "Brainies," "Experiencies," or "myWorld-ies" might be the killer app of BCIs.

2.3 Self-actualization BCI applications

A third class of BCI applications is related to self-actualization. This refers to a full realization of one's potential for self-development. Per Abraham Maslow's theory (1943), self-actualization represents the growth of an individual toward the fulfillment of the highest level of needs, those related to meaning. Carl Rogers (1961) further posited a human drive or tendency for self-actualization. Here, this is understood as becoming one's potentialities, expressing and activating all of the capacities of the human organism. This could include the expression of one's creativity, a quest for spiritual enlightenment, the pursuit of knowledge, and the desire to give to society: anything an individual self-determines as meaningful. Actualization is not merely experiential but generative; it is developing oneself actively and bringing this into concrete expression in reality. There are fascinating possibilities for how BCIs might help with intellectual, creative, and artistic self-actualization. Beyond health tracking and entertainment, one of the strongest aspects of what might be at stake with BCI technologies is the possibility of realizing more of our human potential, and this could be a strong motivation for adoption (Swan 2016a).

3. Cloudmind

The potential future applications of BCIs discussed so far relate primarily to *individuals*; however, BCI technologies might be mobilized similarly into other classes of applications to support *group* activities. We are inherently social creatures and lead interactive lives with others in the context of a social fabric, and new technologies could continue to facilitate these interactions. One of the most potent applications for BCI group applications could be the speculative notion of the *cloudmind* or *crowdmind*. Most broadly, a cloudmind would be, as the term suggests, a cloud-based mind, a mind in the Internet cloud. This would be some sort of processing or thinking capability (hence “a mind”) that is virtual, located in Internet databanks without having a specific body or other physical corporeality. A crowdmind might comprise large numbers of minds operating together.

There could be different kinds of cloudminds. One might be a basic machine mind: algorithms quietly crunching in the background, maybe as the result of the next generations of big data analysis programs. Other types of cloudminds might be different forms of human-machine minds (e.g., a person plus a cloud-based thinking assistant or companion such as Siri or Her (Jonze 2013)). There could be different forms of multiple minds pooled together (*mindpools*), combinations of human minds, human-machine minds, or machine minds. The use of the word “mind” in the expression *cloudmind* could be misleading since the familiar example of mind is the human mind, and machine intelligence is not in possession of the full range of capacities of the human mind such as general purpose problem solving, volitionary action (free will), and consciousness. However, “mind” is meant generally here to denote an entity that has some sort of capacity for processing and “thinking,” perhaps initially in the narrow sense of finding solutions to specified problems, but possibly expanding as processing tasks become more broadly “thinking” oriented. The general definition of a cloudmind is a cloud-based thinker with some sort of analytic processing power.

3.1 Prototypical cloudminds

The notion of a cloudmind is perhaps not so much a new idea as a new label that connotes a greater range of functioning. Prototypical cloudminds already exist in the sense of automated cloud-based systems that coordinate the processing activity of multiple agents. One such prototype is Mechanical Turk, an algorithmic system for organizing individuals to perform online tasks that require human intelligence. In this category of crowdsourced labor marketplaces, there are many other examples such as Topcoder, Elance, and Upwork (formerly Odesk). A second cloudmind prototype is the notion of humans as a community computing network. The idea is that humans, in their everyday use of data, perform a curation, creation, and transfer function with the data. Humans actively transform, mold, steward, and produce data in new forms by interacting with it. Data is active and living, dynamically engaged by humans as a community computer, each person a node operating on data and re-contributing the results back into the network (Swan 2012a). A third kind of cloudmind prototype is “big data,” the extremely large data sets that are analyzed computationally to reveal patterns (such as Amazon and Netflix recommendation engines). This “algorithmic reality” is an increasingly predominant feature of the modern world (Swan 2016b). Big data takes on entity-level status in the notion of the cloudmind, where big data is envisioned as a whole, quietly crunching in the background. The dual nature of technology (having both “good” and “evil” uses) can be seen in big data. On one hand, big data might be seen as contributing to our lives in helpful ways including by reducing the cognitive load required to deal with administrivia. On the other hand, a worry is that big data may not be just guessing our preferences but starting to manufacture them for us (Lanier 2014).

3.2 Cloudmind starter application: Sell unused brain processing cycles to the cloud

In the future, cloudminds involving human brain power might be facilitated by BCIs or other ways of linking human cognitive processing to the Internet. The key feature is the live 24/7 connection, not just generally to the Internet, but specifically to other brains and machine thinkers. One way that individuals might start to explore and adopt BCI cloudmind applications is in a “starter application” idea of selling permissioned braincycles to the cloud. This is a parallel concept to selling self-generated electricity from solar panels back into the power grid. This initial and basic cloudmind application might involve the sharing of unused brain processing cycles. The structure could be timesharing cognitive processing power during sleep cycles or other down time, conceptually similar to participating in community computing projects such as SETI@home or protein Folding@home. The idea would be to securely and unobtrusively share one's own unused resources, downtime braincycles. There could be diverse compensation models for this, including remuneration and donation.

3.3 Cloudmind health app: Virtual patient modeling

More advanced cloudmind applications could correspond to the three classes of individual BCI applications discussed above: health tracking, information and entertainment, and actualization. With pathology resolution and enhancement applications, the daily health check could include longitudinal neural data-logging to Electronic Medical Records (EMRs), which could be integrated into virtual patient modeling systems. The personal health simulation could include different possible scenarios of how patient wellness could evolve from the simulated impact of various drugs or lifestyle choices. The system might model any variety of responses to personal health questions, such as recommending a nootropics stack to maximize cognitive enhancement given a particular individual's genomic profile. Virtual patient simulations could be part of any future EMR (Bangs 2005; Uehling 2004), and instantiated as a cloudmind application with a cloudmind's full range of intelligent processing capabilities. Virtual patient EMR files could be shared more widely (by permission) with family groups and health databank repositories for remunerated research studies and clinical trials. There could be a new concept of "virtual clinical trials" to accompany any physical-world clinical trial. Simulated patient responses could be a supplemental mode of information, particularly to model safety and efficacy. Collecting data initially for medical purposes is already practical and cost effective enough to justify the effort. There is the additional benefit of creating valuable digital health assets that might be mobilized later for many other purposes, for example, to invite participation in user-permissioned cloudmind projects.

3.4 Cloudmind information app: Crowdminding an IoT collaboration archipelago

One of the most obvious information-related BCI cloudmind applications could be thought-commanded Google searches and information look-up. Consider how many steps can be required now to obtain simple data elements such as a weather forecast or a movie time. This can involve having to turn on a phone and go through a series of screens, with variable response times as the phone negotiates network connectivity.

Beyond information query, another cloudmind application could be commanding Internet-of-Things (IoT) connected objects in the environment. There are numerous examples of individuals feeling as if they are one with objects and equipment, for example a submarine commander or airline pilot experiencing the ship or aircraft as an extension of their own body (Takayama 2015). A modern example is remote workers piloting telepresence robots in the main office, where again the robot feels like an extension of the individual's physical body (Dreyfus 2015). IoT cloudminds could provide an expanded version of this: using one's mind to control physical objects in a local or remote environment via BCI. A security guard could command a whole building, for example. An IoT home security system could be operated remotely via BCI cloudmind. The science fictional idea of linking humans together as one, in one recent example using Naam's drug Nexus (Naam 2012), could be extended to include linking humans and objects. There could be a joint agent that is a human plus IoT objects, functioning together as one cloudmind entity. On one hand, my being a cloudmind with my IoT objects is merely an extension and formalization of the human-machine fusing phenomenon that already occurs in intensive machine operation ("better horse" technology). On the other hand, the functionality of cloudminding myself into a collaboration archipelago of intelligent action-taking capability with IoT-enabled smart objects is a revolutionary new kind of concept ("car" technology).

3.5 Cloudmind actualization app: Digital self

One implication of a simulated digital patient *self* as a standard part of health records is the possibility of having a *digital self* more generally. There could be a more fully embodied digital self, a version (or versions) of me that exists electronically. Already there are many versions of *digital me's* as digital selves existing online for many purposes. There are digital profiles for different websites: avatars, digital personae, and "fake me" accounts. Any form of my digital profile could be said to reasonably comprise some version of me, including those that explore dimensions of me otherwise not manifested in the physical world. There are digital self projects, such as CyBeRev and Lifenaut, which explicitly aid in the creation of digital selves (Zolfagharifard 2015). Even now, it is possible that algorithms could assemble digital selves of people from existing online footprints such as photos, social media, academic and blog writing, email communication, file storage, and other aspects of digital presence.

Digital selves might be mobilized for many online operations, including eventual participation in cloudminds. The lowest-risk starter applications for digital selves could be related to backup, archival, and storage – a digital self as a biographical record and memory-logging tool. My digital self could become more active as a digital assistant self, a virtual agent version of me deputized to conduct a certain specified slate of online activities. These activities could include purchase transactions, information search and assembly (for example an automated literature search), and more complicated automatable operations such as drafting email, blog entries, and forum posts based on previous content. Digital selves could be an interesting way to extend and monetize one's own self as a computing resource, and provide a possible solution for the transition to the automation economy (one's digital self engages in remunerable online work).

In the scope of their activity, digital selves could participate in computing projects that are increasingly complicated and remunerative, and might eventually lead to cloudminds. Joining a cloudmind project through a limited digital self could be a comfortable and gradual adoption path to cloudminds that builds trust and familiarity. Participating in a cloudmind with a digital copy, including one with an expiration date, could be less risky than participating with one's "real" physical self. Over time, the digital self could incorporate more richness and fidelity from the underlying person, in order to be more active as an agent with volition and decision-making, not just passive storage. Eventually, BCI neural-tracking data could be integrated to produce an even more fidelitous digital self that includes the neural patterns of how an individual actually experiences and reacts to the world. The longer-term conceptualization of the digital self could be an entity that records, stores, simulates, and runs a full "me" node: a digital agent, and eventually a clan of digital agents, operating just like me.

3.6 Cloudmind actualization app: Subjectivation

Regarding personal development and actualization applications, this could be a central motivation for joining a cloudmind, participating with either a traditional "meatspace-mind" or (with less risk) a digital self mind. Of the extensive slate of BCI cloudmind applications, including health tracking, life logging, archiving, sharing, information requesting, fun and entertainment, and IoT archipelago control, one of the real killer apps might be the personal actualization potential that BCI cloudminds could deliver. Cloudminds could be employed in different levels of elective engagement that is productive, generative, and creative: pooled productive activity toward a goal. There could be a wide range of reasons for joining a cloudmind including compensation, fun, new experience, productive use of one's mind, contributing to problem-solving efforts, and self-actualization (growth and development opportunities). Cloudmind collaborations could consist of meaningful and remunerable work, and personal mental engagement and development that is fun, creative, and collaborative. One of the deepest incentives for exploring improved connection and cognition through BCI cloudmind collaborations could be the possibility that they facilitate our individual growth and development as humans, our subjectivation or self-forming (Robinson

2015). BCI cloudminds could allow us to actualize our potential to become "more" of who we are and might be, more quickly and effectively, thereby hastening and accelerating our capacities to be more intelligent, capable, and creative participants in life (Swan 2016a).

3.7 Crowdminds to remedy possibility space myopia

From a practical perspective, one hope or assumption might be that problem solving could be made more effective with technology tools such as BCI cloudminds. Problem solving is a central activity we engage in as humans, and any means of improving our capability to do this might be useful and valuable. Minds (irrespective of type) collaborating together might solve problems more expediently than the "classical" (e.g. current) human methods of individual breakthroughs, competition, and team striving. Cloudmind infrastructure might support improved human problem solving, and enable new kinds of human-human brainstorming, human-machine collaboration, and progress in the development of machine minds. It may be that we have tackled only a certain circumscribed class of problems so far, one limited to human understanding and articulation, whereas the universe of problems and problem-solving techniques could be much larger. We should understand the possibility space (the full universe of possibilities) for many phenomena to be much larger than the part that is human-viewable or human-conceivable. A simple example is the electromagnetic spectrum, where only a small portion is viewable, but where our tools have vastly expanded our reach. This "possibility space myopia" has been documented in many domains, for example in computing algorithms (Wolfram 2002), intelligence (Yudkowsky 2008), perambulatable (i.e. able to walk) body plans (Lee 2013; Marks 2011), mathematics and logic (Husserl 2001), and the size of the universe (Shiga 2008). Collaborative methods between humans and machines, such as BCI cloudminds, might extend our reach into a larger possibility space of problems and their resolution.

4. Cloudmind adoption risks

There are many different kinds of potential risks, limitations, and concerns with cloudminds. Some of the most prominent include privacy and security, how credit is to be marshalled, and how boundaries are to be established so as not to lose one's personal identity. With any radically new technology, especially one that involves a sensitive area of concern such as the human brain, a trustworthy and responsible adoption path should be gradual and begin with specific, limited uses. Some lower-risk, phased adoption strategies already mentioned are sharing unused brain processing cycles (SETI@home for your brain) and participating in a cloudmind with a digital self as opposed to an original self. Another adoption norm might involve giving permission for limited access to certain domains of the brain and human cognitive activity. As one form of boundary control, personal connectome files might be used to demarcate the cortical regions within specified limits . Another boundary might be created by stipulating time blocks during which a brain might be connected to the mindstream (for example, permitting extra braincycles during sleep). The technical details of protected mindcycle sharing, while crucial, are currently impossible to enumerate, given our incomplete knowledge of the brain. The point is that it may be helpful – or even required – to have some sort of structural protection in place to the extent possible.

4.1 Responsible technology design principles

Some of the themes from these adoption strategies might be codified into responsible technology design principles as set forth below.

1. **Adoption:** Accommodate a wide spectrum of adoption possibilities. This could range from full-fledged adoption to non-adoption, where the exercise of user choice is celebrated and not vilified. Explicit pathways could be defined for gradual adoption in specific phases of increasing levels of comfort and engagement. Adoption should also be freely undertaken (without coercion or incentives) and also reversible, with the ability to leave BCI cloudmind communities easily.

2. **Security and transparency:** Include anti-virus protection, and safeguard against mind hacking as explored in science fiction narratives such as *Slant* (Bear 1998) and *Rainbows End* (Vinge 2007). There should be transparency and opt-out selections regarding data collection, use, privacy, and monetization. Blockchain-based tracking logs could provide independent monitoring of access and activity in BCI cloudmind communities.

3. **Support:** Provide education, community forums, and feedback mechanisms for participants to share their experiences. There could be a social network for the user community to resolve questions and issues. There could be monitoring and support ecosystem provided for users. There could be periodic check-ins and evaluations by external referees as to whether agreed-upon contractual terms are within compliance.

4. **Standards:** Engage with industry standards bodies to establish and steward technology norms as features. At present, for example, BCIs might fall under the IEEE 802.15 working group for Wireless Personal Area Networks (WPANs).

The design principles outlined here could be integrated with other developing precedents for responsible technology design. Some of these parallel efforts include the calm technology movement (principles for non-intrusive design (Case 2016)), and calls for the anticipatory governance of new technologies within existing legal structures (Nordmann 2014).

4.2 Video games, BCI cloudminds, and intelligence amplification

The first class of applications is those related to health. Privacy, security, disclosure, and reversibility could be some of the key adoption concerns in health-related applications, and are fairly straightforward to identify, implement, and track. However, adoption risks might figure more prominently in the case of the second class of applications, those related to immersive experience and entertainment. One fear is that these applications might be so entrancing as to become extremely addictive, possibly to the detriment of otherwise being able to "participate in a meaningful life." Video games are a related example, but here the research results are mixed. There are negative consequences such as susceptibility to addiction (Kuss 2013), player fatigue, and other detriments (Heaven 2015). However, there is also evidence to the contrary: while games interfered with schoolwork, social interaction and other aspects of a "normal life" remained a priority (*New Scientist* staff 2007).

Other encouraging research suggests that humans ultimately prefer novelty to pleasure, and eventually turn away from sustained pleasure-center stimulation out of boredom (Patoine 2009; Yoffe 2009). Keeping a human being engaged is a complex and dynamic process. Certain studies have found that in fact gamers have more grey matter and better brain connectivity than non-gamers (Crew 2015; Bushak 2015; Johnson 2006). Thus perhaps we are doing something right, and new, with video games, not only for fun, entertainment, community, and status garnering through competition, but also for brain development and intelligence amplification. Games might dovetail with actualization in that, as one analyst phrases it, there is something about video games that allows us to "brainjack into our potential" (Heaven 2015). The BCI cloudmind design challenge is to produce applications that safely support our being as humans, while balancing potential risks.

4.3 Human futures: BCI cloudmind fulfillment or virtual reality couch potato?

One worry is that humans might completely tune out of the physical world in favor of the immersive technologies of the virtual world and cease productive engagement in physical reality. However, the definition and meaning of all of these terms might shift. In a situation of remunerative, mentally stimulating, contributive, creativity-expressing possibilities that might be increasingly available in virtual reality, what is it to distinguish situations as "healthy" or "addictive"? The prospects of virtual reality and BCI cloudmind applications raise bigger questions about the meaning and purpose of human life. The parameters by which we assess reality, valorize meaning, and plan our lives might be open to change. The virtual cloudworld might be exactly the venue to provide meaningful engagement opportunities, including for remuneration and sustenance, and fulfillment more generally, especially in a post-scarcity automation economy where labor work is no longer compulsory. It is not clear why "couch-potatoing" into virtually-fulfilling states might be "bad" in a future world of economic abundance, free will, and ample choice.

There are some examples that help to recast the terms of the argument from the position that "too much virtual reality is bad" to "how to constructively transition a greater portion of life to supportive virtual reality environments." One example is the virtual world Second Life, where the individual and cultural reaction has been both celebrated and castigated, but has also helped to clarify that virtual worlds can be a venue for producing meaning in human lives (*Daily Mail* 2008). The science fiction story "The Clinic Seed – Africa" (Henson 2007) provides a positive account of the transition to an upload world, where the operations of human intelligence are increasingly enacted in the cloud. Philosophers such as David Pearce argue for a purpose-driven design approach to technology that focuses on gradients of well-being and aims to eliminate suffering (Pearce 1995). The whole design space of BCI cloudminds is wide open. On one hand, well-designed applications would ideally provide improved human-need fulfillment possibilities that supplement the physical world with the virtual world. However, it would be naïve not to acknowledge that applications are also likely to be designed for the opposite, to cater to the darkest underbelly of human desires to control, dominate, and destroy.

The surprising conclusion might be one of the potential desirability of widespread and persistent use of BCI cloudminds. This could be not just for the recreation, competition, and learning possibilities available via immersive experience, but precisely in order to participate in generative cloudmind collaborations. Crowdminds might provide a crucial vehicle for more humans to experience "work" as a fulfilling and productive activity, as a purpose-directed expression of human energy, and as a means of actualization. Just as "information era work" relies now on the Internet, the next generations of digital era work might use BCI cloudminds as a vehicle. Generative collaborations could be remunerable (via economic resources, reputation, status-garnering, acknowledgement, and other means), and beyond that they might provide a productive, stimulating, mentally engaging venue for fulfilling untapped human potential. While generative collaborations would certainly have their own adoption risks, they might also be an answer to the some of the concerns of entertainment applications regarding immersive experience addiction, boredom, and cognitive fatigue. BCI cloudminds might be fulfilling and actualizing to humans on a sustained basis, and also unleash new tiers of human capability, including safely coordinating collective activity beyond the scope of any one individual.

A bigger topic for consideration is the effective design of group activities, particularly as mediated through technology environments. Articulating and deploying techniques for ensuring that groups are well formed is a future-class concept that has yet to be considered in detail. There is not a clear and readily deployable toolkit or skillset for group formation to enable groups to perform at a higher level. The current understanding of group dynamics still points to narrowband heuristics such as "forming, storming, norming, performing" (Tuckman 1965), and has not evolved to incorporate more robust models. There are some examples that indicate the depth of what is required for effective group operation. One is Convergent Facilitation as a sophisticated, yet simple, way of consistently eliciting and attending to group needs, and keeping individuals engaged in group activity (Swan 2015a). Another is Simondon's notion of successful group formation as *collective individuation*, which arises in moments of the collective structuring of emotion across a group (Swan 2015e).

5. Blockchain cloudmind administration

Adoption risks are one of the biggest concerns with cloudminds, and after that, practical issues, particularly how credit is to be assigned. For credit assignment and other coordination activities, blockchains (cryptographic Internet-based transaction ledgers) might be used (Swan 2015b). Blockchains have the necessary features to administer cloudminds including privacy, security, monitoring, and credit tracking. The properties of blockchains as a universal, secure, remunerative software structure could be ideal for the operation of cloudminds. Getting credit for individual contributions and intellectual property generation in cloudminds is likely to be a concern for some participants, and blockchains could be a solution for unobtrusively tracking line-item transactions at any level of detail. The technology can run undetectably in the background, tracking all contributions, and is always available for full audit trail lookup. Activity can be aggregated easily into consolidated payments for participant contributions. Blockchains could serve as a trustworthy vehicle for assigning credit and assessing future remuneration in all online work, including currently in crowdsourced labor marketplaces.

5.1 Cloudmind line-item tracking and credit-assignation

More specifically, line-item tracking could be implemented in the idea-rich brainstorming environment of cloudminds by using secure deep-learning algorithms to record all participant activity. One current feature of deep-learning systems is the capture and output of activity in interchangeable formats, accommodating audio, video, and text formats. Each participant's most minute thought formulations might be recorded with a neural feed and a time-date stamp that is logged to a blockchain. This would be similar to using a deep-learning algorithm to automatically transcribe Skype audio calls into writing, and posting the result to a blockchain to validate ownership. Important steps in the brainstorming process might thus be tracked, for example whether multiple parties have the same idea simultaneously. There could be additional benefits as well, for example, in the area of novel discovery, elucidating the brainstorming process itself and identifying the precursor factors to actual ideas.

Blockchains could thus be a useful technology for aggregating multi-threaded experience into a cohesive whole. Blockchain cloudmind tracking would be conceptually similar to applying a software version control system (such as Github, SVN, or CVS); this is literally a line-item tracking system for a brainstorming session. This kind of feature could permit even more precise assignment of credit in cloudmind collaborations. Further, BCI cloudminds might facilitate brainstorming with new functionality such as recognizing and consolidating generative threads and posing them back to participants. Questions could arise about upward size limits on cloudminds and how very large cloudminds are to be organized so that they do not descend into confusion and chaos. This is precisely one reason that blockchains might be a well matched administration tool, as they can allow for arbitrarily many participants, simultaneously tracking and merging all contribution threads. Optimal cloudmind-sizing and parameterizing could be a focal point in the future, as various optima could emerge and "the more minds the better" might ultimately not be a relevant maxim.

5.2 Private brainstorming facilitates mediation and diplomacy

Blockchains might offer more advanced functionality, too. For example, there could be private problem-solving sessions in sensitive situations such as conflict resolution and international diplomacy. There could be software-mediated private brainstorming where algorithms synthesize ideas and try to find common ground, while individual participants do not access or know the main thread of progression, but only their own experience and what is mediated in response. All positions could be encoded and logged to blockchains with time-date stamping, where contents are confirmed but kept private by running a hashing algorithm run over them (proof-of-existence functionality).

5.3 Enumerating the idea-generation process

As mentioned, one implication of the digital tracking of idea formulation in cloudminds could be a more detailed understanding of the process and structure of idea genesis more generally. This knowledge could in turn be fed back into BCI cloudmind systems to stimulate, catalyze, and facilitate idea generation. Existing research studies have tried to isolate the process of idea generation in affective phenomena in humans, for example those that precede the "Aha!" moment of creative inspiration. Some examples include gamma wave bursts, feeling a sense of challenge and persistence, and engaging in collaborative improvisation (Swan 2015f). Participant thought formulations might be recorded with a time-date stamp, including both linguistic and pre-lingual traces based on biophysical response and neural activity detected by BCIs and other quantified-self devices. This could lead to a manipulable understanding of how idea formulation and problem solving function at the individual and group level.

5.4 Gradual opt-in BCI cloudmind adoption phases

I propose the following implementation blueprint for cloudminds using blockchain technology as the secure automation and tracking mechanism. The first step could involve facilitating selling braincycles to the cloud, akin to selling generated electricity back to the grid (SETI@home for your brain). Blockchains could securely track and remunerate these contributed cycles. Once the idea of crowdsharing one's mind into cloud computation is understood and securely implemented, cloudmind participations could be expanded. The next level of applications could include data analysis, problem solving, creative expression, and idea generation. Blockchains could track line-item contributions in cloudmind collaborations, acknowledging and rewarding new ideas in a trustable annuity stream, in a ledger that is open for scrutiny, but inconspicuous, not detracting from the idea generation process itself. Assured that that the administrative details of credit assignment and remuneration are being handled competently and fairly in the background, participants could have more ease in directing energy and focus toward the projects at hand.

Blockchains might be helpful not just in overall administration and line-item credit tracking, but also in the ongoing safety and security monitoring of cloudminds. Blockchains could implement, monitor, and enforce relevant safety measures such as anti-viral provisions (to prevent mind hacking) and anti-crowdmind provisions (to prevent not being able to leave a crowdmind voluntarily). There could be limits on the amount of time spent in cloudminds. Feature norms for responsible cloudmind technology could include roll-back to any previous version (standard Wiki functionality), and always being able to reinstantiate an original non-networked digital copy of "you." Perhaps one of the biggest fears is being mind-controlled (such as in the *Matrix* movies). Therefore the design task for setting up cloudmind computing networks could be framed as contributing user-controlled resources into the computational infrastructure. Each human participant might want to maintain his or her own agency at every step, contributing cognitive compute-time as an asset to cloudminds. Human societies have given much thought to the issues of owning and accessing assets, but partially collaborating them into a bigger entity involving our own cognitive resources is a frontier.

Beyond sharing SETI@home-type mindcycle processing, blockchains could function as a trustable checks-and-balances system for administering BCI cloudminds more generally. In the application of group IP-generation, "ideachains" could track individual contributions. Further, blockchain-based smart contracts could be employed as independent advocates to monitor cloudmind activity. A digital safeguard norm could be launching a smart-contract DAC (distributed autonomous corporation), essentially a network security agent, automatically with the launch of any cloudmind. The DAC could serve as a third-party advocate to monitor cloudmind activity, for example running exploitation checks for security breaches. More sophisticated "identity sanctity anti-viral checks" could canvas brain patterns for warning signals of cult-like or prisoner-like behavior, neurodegeneration, and other situations of compromised cognitive liberty. As with any new technology, we could expect that the three areas of functionality, risk, and response would evolve in lockstep (Swan 2015c).

6. Strongest BCI concern: Retention of personal identity and fear of groupmind

If concerns related to adoption risks and the practical administration of cloudminds might be allayed, perhaps one of the biggest remaining issues a potential participant might have is the fear of being irreversibly incorporated into a groupmind. This could be the sense of joining a cult, entranced into losing the ability or desire to leave, unable to make clear and rational decisions or to advocate for oneself. Responsible technology design could help acknowledge and address this fear. Step-by-step ingress-egress processes could feature prominently in cloudmind documentation and training orientations so that humans could become familiar with the process and feel comfortable about joining and leaving cloudminds. The core design requirement could be ensuring freedom and extricability.

6.1 Personal identity

One way to understand the fear of being irreversibly incorporated into a cloudmind is through the lens of personal identity. Personal identity is a concept that might be rethought in the context of human-technology collaboration. Personal identity is often taken as a given, and while there are many arguments in favor of the human as a unit and personal identity as a property, it is not clear that personal identity should be assumed to be persistent and foundational in the future (Swan 2014b). Personal identity is meant (per the Merriam-Webster dictionary definition) as "the persistent and continuous unity of the individual person."

6.2 Biology: Multiple levels of organization

In the biological sciences, there are many "unit levels" of organization of which individual organisms are just one. An ongoing topic of scientific debate is precisely the question of the correct unit-levels at which to understand nature's systemic phenomena. There are multiple tiers that appear to be natural units of organization including genes, genomes, organs, tissues, operating mechanisms, organisms, phenotypes, family groups, haplotypes, gene pools, and ecologies (Hull 1980). Organisms have a privileged role in some cases, but not always, and the concept of the individual is only one model for studying biology – one that is perhaps unjustifiably anthropocentric. This is partly due to the fact that, as investigators, we ourselves are individual units, and thus might take it as a given that human organisms are a privileged or at least default unit. Another possible reason for preferencing the individual as a unit is that individuals are tangible, clearly distinct at our visible level of macroscale, and have obvious demarcation points that are easy to classify such as birth and death. On the other hand, one factor of human life, sexual reproduction, does not support the privileging of just one organism, as gene mixing and matching in recombination across the population confers evolutionary fitness.

In principle, any level of organization can be studied in biology, and there is a contemporary trend away from focusing exclusively at the organism level. New methods such as complexity science allow for a broader and more comprehensive focus beyond the traditional reductionist and linear cause-and-effect scientific method, and may correspond more fully to the underlying phenomena that exist simultaneously at multiple organizational levels. Biology is complex and seems to incorporate myriad levels of processes within systems, including invisible, yet potent, rules to produce macroscale behavior such as protein expression and bird flocking (Mazzocchi 2008). One example of complexity science informing neuroscience conceptualizations is a new understanding of the brain as fields of resonance, as opposed to point-to-point neuronal firings (Boyden 2015, 2016). Biology is a reminder of the dynamism of evolution, so that even if the human organism might seem to be a "final" body plan now, with human-technology integration there could be subsequent "evolutions" of the human form.

In fact, there can be map-territory conflicts in trying to conform the human-constructed social apparatus to underlying biological phenomena. Nearly all of human law and governance is based on the human individual. Rights, responsibilities, and property are typically vested in the individual. However, biology is not so clean. For example, attempting to ground "personalized genomics" in the unit of a person does not always make sense. The supposedly personal genome is vastly shared, beyond the scope of an individual organism, with parents, siblings, extended families, haplotype communities, and ethnic groups. Ethics practices might be most appropriately deployed in group models, in addition to individual models (Swan 2010). The conclusion from the biological standpoint is that there are multiple tiers of organization and there is not necessarily any support for the privileging of the organism as a unit. Individuals have been a primary ordering mechanism in the world so far, perhaps because there have not been other alternatives. No collective, plural, or cloudmind forms have existed yet, but this does not imply that they cannot.

6.3 Sociality: Humankind is inherently social

Despite existing as individuals, we have always been social. As Aristotle notes, "man is by nature a social animal" (Aristotle 2013) even if antagonistically so, as in the Kantian "unsocial sociability of men" (Kant 1963) Recent positions in the philosophy of mind postulate that language, meaning, and the use of reason are social institutions and that the full realization of human beings requires social enactment. Sociality is "a condition for the exercise of intelligent activity" (Descombes 2014), one that is required for "the intentional act of reasoning" (Pettit 1996). In short, sociality is an undeniable feature of life. There are numerous familiar examples of social groups such as corporations, churches, and families. These are plural, non-individual unit entities that centrally comprise society. Sociality is a fundamental aspect of being human, and cloudminds might help to extend what is possible in human sociality.

For example, psychologically and sociologically, there is the notion of the existence of a collective consciousness (Bobrow 2011). There is one's internal experience of consciousness, and often a sense of shared reality and thought perspectives with others. Different modes of spirituality and religion discuss collective consciousness under the rubric of "oneness" or "non-separation" and also assume that the human is not an exclusive standalone unit. In Buddhism, there is the notion of non-self, the idea that the self is a temporary instantiation in physical reality that exists in order to attend to immediate issues such as biological needs, whereas the true nature of reality is an uncoalesced flow (with no subject-object distinction) where something might be considered a self only if it were to be permanent (Albahari 2011). Cloudminds might help to capture, define, and engage the notion of the collective consciousness, useful, at minimum, for assessing levels of shared perception across a group.

6.4 Philosophy: Parfit's claim of relational experience over identity

In philosophy, one of the strongest claims regarding personal identity is articulated by Derek Parfit when he says that continuing personal identity is not required for the survival of the person. Instead, what is relevant is *relational experience* between past and future selves and events (Parfit 1986). What is salient is not a persistence of who "I" am over time, but the timeline, how events are connected in my past, present, and future. More specifically, the being that is "I" has some sense of continuity and connectedness of psychological features such as memories of events and personality traits, but there is no requirement for this "I" to have a personal identity. There is empirical evidence from fMRI studies to support Parfit's claim. In studies, individuals were unable to distinguish between their own future self and any other person (Opar 2014; van Gelder 2013; Pronin 2008). There was no identifiable brain activity corresponding to personal identity in the sense of a person having a different response when recognizing their own future self versus others. The hypothesis was that if personal identity were an important psychological construct, then fMRI activity would be more pronounced when recognizing the self. Another finding that supports the non-criticality and non-persistence of the human self as a unit is that humans can forget that they are a subject, a unit, a self, when immersed in flow state and nondual consciousness experiences. In the flow state, a person becomes engrossed in a task to the extent of losing track of time and a sense of identity (Csikszentmihalyi 2008; Abuhamdeh 2012). In nonduality, the brain registers a lack of subject-object distinction, conceiving of no difference between itself and world objects (Josipovic 2014).

6.5 Custom and familiarity: Knowing no other reality

One understandable reason for preferring personal identity and humans as an organizational unit is custom and familiarity – so far there has been no other alternative. Having known only one mode of existence, we may have developed a natural attachment to it. Thus we might assume that elements such as memory, emotion, and "meaning" can be instantiated only in humans, but it is possible that, in the future, any pattern associated with human brain activity might be elucidated and stored as information in digital networks. This could include intellect, creativity, memory, emotion, consciousness, and the notion of self. Further, the different versions of the idea of digital selves and BCI cloudminds do not necessarily preclude or curtail personal identity, and some rather accentuate and extend it. In the future, the physical "meatspace" body might not be the only hardware substrate on which to run the self; it might be possible to run one's personal identity on other platforms, including Internet cloudminds or other physical-world hardware such as robots or IoT archipelagos.

6.6 One answer: Identity multiplicity

Historically, society has been based on human individuals for several reasons: 1) politically, for the grounding of responsibility and representation; 2) economically, for the vesting of labor efforts and property ownership; 3) reproductively, for species survival through sexual reproduction; 4) internally, for consciousness, emotion, creativity, ideas, and interior life; and 5) liberty-wise, for self-determination and the volitionary pursuit of interests and ideals. However, it is perhaps just an element of efficient construction that the societal apparatus has been organized with human individuals as units, and this might be a model that could be supplemented with cloudmind or other collective personhood formats.

For example, a common principle of new technology is multiplicity, in the sense that more choice is produced. The new technology creates more options, not an either-or situation. Precisely the point of a new technology, then, is that there can be a proliferation of choice and new possibilities. A classic example is the phonograph and the radio, where one did not render the other obsolete but both worked together to expand the overall market for listening to music. Products like eReaders have similarly expanded the book-reading market. Particularly with a concept, such as BCIs and cloudminds, that is tightly integrated with the human form, any potential adoption might be best positioned if it follows this principle of multiplicity. BCI cloudmind products would need to demonstrate how they make more options available to the consumer. In addition, many safeguards might be required initially, including reversibility.

In the case of BCI cloudminds, identity multiplicity might involve many different forms of participation. There could be "classic meatspace brains," one or more digital selves, and different configurations of selves (for example, a team of selves, what Hanson calls a "self clan" or an "em[ulation] clan" (Hanson 2015)). Human beings are currently constrained to an embodied form; however, this may not be the situation in the future. Digital identities might become so distributed, portable, easily copied, open-sourceable, sharable, and malleable that it no longer makes sense to think in terms of distinct entities but rather in some other parameter such as instances. Reputation could still matter, however, and serve as the central coordination element. For example, in the case of digital societies blockchain consensus trust mechanisms could confirm and validate smart network transactions based on agent reputation.

7. Conclusion: BCI cloudminds

In this paper, I have addressed the potential advent of brain-computer interfaces (BCIs) that are ubiquitous and widely adopted, enabling humans to become continuously connected not just to the Internet but also to other minds. I've discussed individual and collective applications in the areas of health tracking and enhancement, information seeking and entertainment, and self-actualization, in particular that of humans and machines digitally linked in cloudminds. Again, a cloudmind will consist of multiple individual minds (human or machine) joined together in pursuit of a collaborative goal such as problem solving, idea generation, creative expression, or entertainment.

I have outlined risks and proposed safeguards that might be needed for individuals to feel comfortable joining a cloudmind. A key fear could be irreversible incorporation into the cloudmind. This fear might be assuaged though the enactment of responsible technology design principles and a plan for how gradual adoption of BCI cloudminds might proceed. Some of the core responsible technology design principles are privacy, security, reversibility, credit assignation, and personal identity retention. Particularly regarding the sensitivity of personal identity, the adoption path could be one of identity multiplicity – initially participating in cloudminds with limited versions or digital copies of oneself. A slow adoption path could involve initially backing up part of one's mindfile digitally for storage, recovery, and archival, and then slowly testing additional functionality by selling unused braincycles to the processing grid, running a digital self with limited operations, and eventually proceeding more fully to cloudmind collaborations.

Historically, there has been much attachment to personal identity and the human as a unit. In part, this has been because our only known mode of lived existence has been as "classic humans." However, there is support from biology and other fields regarding the non-exclusivity of the human. Over time, there could be subsequent progressions of the human form, and possibly a loosening or expanding of the notion of personal identity. The key to transitioning to futures of greater human-technology integration in an empowering manner could be maintaining "the right relation with technology" – one that is enabling not enslaving (Heidegger 1977). Blockchains might be an important singularity-class technology (one that is globally robust and transformative) for maintaining an empowering relation with technology through the safe adoption of BCI cloudminds. One inspiration for these kinds of futures is in Charles Stross's novel *Accelerando*, where blockchain-type trust networks of humans and technology are joined in partnership, and digital copies "watch over their originals from the consensus cyberspace of the city" (Stross 2006, 355).

References

Abuhamdeh, Sami, and Mihaly Csikszentmihalyi. 2012. The importance of challenge for the enjoyment of intrinsically motivated, goal-directed activities. *Personality and Social Psychology Bulletin* 38(3): 317–30.

Albahari, Miri. 2011. Nirvana and ownerless consciousness. In *Self, no self?: Perspectives from analytical, phenomenological, and Indian traditions*, ed. Mark Siderits, Evan Thompson, and Dan Zahavi, 79–113. Oxford: Oxford University Press.

Aristotle. 2013. *An introduction to Aristotle's Ethics Books I-IV*. Ed. Edward Moore. London: HardPress Publishing. (Orig. pub. c. 328 BCE.)

Bangs, Alex. 2005. Predictive biosimulation and virtual patients in pharmaceutical R&D. *Studies in Health Technology and Informatics* 111: 37–42.

Bear, Greg. 1998. *Slant*. New York: Tor.

Bo, Hong. 2015. Brain-computer interfaces. Johns Hopkins School of Medicine. http://pages.jh.edu/SCI/projects/bci.shtml (accessed October 8, 2016).

Bobrow, Robert S. 2011. Evidence for a communal consciousness. *Explore* 7(4): 246–48.

Boehm, Frank J. (unpublished). Facilitating a (neocortical) brain/cloud interface (B/CI) conceptual nanomedical strategies.

Boyden, Edward S. 2015. Optogenetics and the future of neuroscience. *Nature Neuroscience* 18: 1200–1201.

Boyden, Edward S. 2016. How the brain is computing the mind. *Edge.* February 12. https://www.edge.org/conversation/ed_boyden-how-the-brain-is-computing-the-mind (accessed October 8, 2016).

Bushak, Lecia. 2015. Video games may improve brain's connections; Action gamers have more gray matter. *Medical Daily.* April 28. http://www.medicaldaily.com/video-games-may-improve-brains-connections-action-gamers-have-more-gray-matter-331026 (accessed October 8, 2016).

Crew, Bec. 2015. Gamers have more grey matter and better brain connectivity, study suggests. *Science Alert*. April 28. http://www.sciencealert.com/gamers-have-more-grey-matter-and-better-brain-connectivity-study-suggests (accessed October 8, 2016).

Case, Amber. 2016. *Calm technology: Principles and patterns for non-intrusive design*. Sebastopol, CA: O'Reilly Media.

Csikszentmihalyi, Mihaly. 1973. *Flow: The psychology of optimal experience.* New York: Harper Perennial Modern Classics.

Daily Mail. 2008. Revealed. The "other woman" in Second Life divorce... who's now engaged to the web cheat she's never met. November 14. http://www.dailymail.co.uk/news/article-1085412/Revealed-The-woman-Second-Life-divorce--whos-engaged-web-cheat-shes-met.html (accessed October 8, 2016).

Dedman, Jay. 2007. Being grateful: What is momentshowing? *United Vloggers*. April 30. http://www.unitedvloggers.com/2007/04/30/being-grateful/ (accessed October 8, 2016).

Descombes, Vincent. 2014. *The institutions of meaning: A defense of anthropological holism*. Cambridge, MA: Harvard University Press.

Dreyfus, Emily. 2015. My life as a robot. *WIRED*. September 8. https://www.wired.com/2015/09/my-life-as-a-robot-double-robotics-telecommuting-longread/ (accessed October 8, 2016).

Friehs, Gerhard M., Vasilios A. Zerris, Catherine L. Ojakangas, Mathew R. Fellows, and John P. Donoghue. 2004. Brain-machine and brain-computer interfaces. *Stroke* 35: 2702–2705.

Hanson, Robin. 2015. Em Redistribution. *Overcoming Bias* blog, entry posted July 17. http://www.overcomingbias.com/2015/07/em-redistribution.html (accessed October 8, 2016).

Heaven, Douglas. 2015. Death by video game: A power like no other. *New Scientist*, August 19. https://www.newscientist.com/article/mg22730350-700-death-by-video-game-a-power-like-no-other/ (accessed October 8, 2016).

Heidegger, Martin. 1977. *Question concerning technology, and other essays*. New York: Harper Torchbooks. (Orig. pub. 1954.)

Henson, Keith. 2007. The clinic seed – Africa. *Journal of Geoethical Nanotechnology* 2(2): 1–6. Available at http://www.terasemjournals.org/GNJournal/GN0202/henson1.html (accessed October 18, 2016)

Hull, David. 1980. Individuality and selection. *Annual Review of Ecology and Systematics* 11: 311–32.

Humer, Caroline, and Jim Finkle. 2014. Your medical record is worth more to hackers than your credit card. Reuters. September 24. http://www.reuters.com/article/us-cybersecurity-hospitals-idUSKCN0HJ21I20140924 (accessed October 8, 2016).

Husserl, Edmund. 2001. *Logical investigations*. Trans. J.N. Findlay. London: Routledge. (Orig. pub. 1900.)

iOS Developer Library. 2015. The HealthKit Framework. *Apple*. Available at https://developer.apple.com/reference/healthkit (accessed October 8, 2016).

Jabczenski, Marissa. 2016. Meta debuts groundbreaking augmented reality technology at TED 2016. Company Press Release via *BusinessWire.* February 17. http://www.businesswire.com/news/home/20160217005172/en/Meta-Debuts-Groundbreaking-Augmented-Reality-Technology-TED (accessed October 8, 2016).

Johnson, Steven. 2006. *Everything bad is good for you: How today's popular culture is actually making us smarter*. New York: Riverhead Books.

Jonze, Spike (director). 2013. *Her*. Burbank, CA: Warner Bros. Pictures.

Josipovic, Zoran. 2014. Neural correlates of nondual awareness in meditation. *Annals of the New York Academy of Sciences* 1307: 9–18.

Kant, Immanuel. 1963. Idea for a universal history. *On history*. New York: Bobbs-Merrill. (Orig. pub. 1784.)

Kurzweil, Ray. 2006. *The Singularity is near: When humans transcend biology*. New York: Penguin.

Kuss, Daria J. 2013. Internet gaming addiction: current perspectives. *Journal of Psychology Research and Behavior Management* 6: 125–37.

Lanier, Jaron. 2014. The Myth of AI. *Edge.* November 14. https://www.edge.org/conversation/jaron_lanier-the-myth-of-ai (accessed October 8, 2016).

Lee, Suchan, Jason Yosinski, Kyrre Glette, Hod Lipson, and Jeff Clune. 2013. Evolving gaits for physical robots with the HyperNEAT Generative Encoding: The benefits of simulation. *Applications of Evolutionary Computation*. Berlin: Springer-Verlag. Lecture Notes in Computer Science 7835: 540–49.

Lunden, I. 2016. MindMaze, maker of a "neural virtual reality platform," raises $100M at a $1B valuation. *Tech Crunch*. February 17. https://techcrunch.com/2016/02/17/mindmaze-maker-of-a-neural-virtual-reality-platform-raises-100m-at-a-1b-valuation/ (accessed October 8, 2016).

Marks, Paul. 2011. Metamorphosis key to creating stable walking robots. *New Scientist.* January 10. https://www.newscientist.com/article/dn19935-metamorphosis-key-to-creating-stable-walking-robots/ (accessed October 8, 2016).

Maslow, Abraham. 1943. A theory of human motivation. *Psychological Review* 50(4): 370–96.

Mazzocchi, Fulvio. 2008. Complexity in biology. Exceeding the limits of reductionism and determinism using complexity theory. *EMBO Reports* 9(1): 10–14.

Naam, Ramez. 2012. *Nexus*. Nottingham, UK: Angry Robot.

New Scientist staff. 2007. Video games interfere with homework but not family. *New Scientist.* July 3. https://www.newscientist.com/article/dn12180-video-games-interfere-with-homework-but-not-family/ (accessed October 8, 2016).

Nordmann, Alfred. 2014. Responsible innovation, the art and craft of anticipation. *Journal of Responsible Innovation* 1(1): 87–98.

Opar, Alisa. 2014. Why we procrastinate. *Nautilus.* January 16. http://nautil.us/issue/9/time/why-we-procrastinate (accessed October 8, 2016).

Pais-Vieira, Miguel, Gabriela Chiuffa, Mikhail Lebedev, Amol Yadav and Miguel A. L. Nicolelis. 2015. Building an organic computing device with multiple interconnected brains. *Nature Scientific Reports.* 5: 11869.

Parfit, Derek. 1986. *Reasons and persons.* Oxford: Oxford Paperbacks.

Patoine, Brenda. 2009. Desperately seeking sensation: Fear, reward, and the human need for novelty. Neuroscience begins to shine light on the neural basis of sensation-seeking. The Dana Foundation. October 13. http://www.dana.org/News/Details.aspx?id=43484 (accessed October 8, 2016).

Peters, Betts, and Melanie Fried-Oken. 2014. Brain-computer interface (BCI). ALS Association. September. http://www.alsa.org/als-care/resources/publications-videos/factsheets/brain-computer-interface.html (accessed October 8, 2016).

Mayo Clinic. 2009. Brain waves can "write" on a computer in early tests, researchers show. *ScienceDaily.* December 7.
http://phys.org/news/2009-12-brain-early.html (accessed October 8, 2016).

Pettit, Philip. 1996. *The common mind: An essay on psychology, society, and politics*. Oxford: Oxford University Press.

Pearce, David. 1995. *The hedonistic imperative*.
http://www.hedweb.com/hedab.htm (accessed October 8, 2016).

Pronin, Emily. 2008. How we see ourselves and how we see others. *Science* 320(5880): 1177–1180.

Robinson, Bob. 2015. Michel Foucault: Ethics. *Internet Encyclopedia of Philosophy*.
http://www.iep.utm.edu/fouc-eth/ (accessed October 8, 2016).

Rogers, Carl. 1961. *On becoming a person*. New York: Mariner Books.

Shiga, David. 2008. Hints of structure beyond the visible universe. *New Scientist.* June 10.
https://www.newscientist.com/article/dn14098-hints-of-structure-beyond-the-visible-universe/
(accessed October 8, 2016).

Seo, Dongjin, Jose M. Carmena, Jan M. Rabaey, Elad Alon, and Michel M. Maharbiz. 2013. Neural dust: An ultrasonic, low power solution for chronic brain-computer interfaces. *arXiv.* 1307.2196 [q-bio.NC].

Smith, Kerri. 2013. Brain decoding: Reading minds. *Nature* 502: 428–30.

Stross, Charles. 2006. *Accelerando.* London: Ace.

Swan, Melanie. 2010. Multigenic condition risk assessment in direct-to-consumer genomic services. *Genetics in Medicine* 12(5): 279-88.

Swan, Melanie. 2012a. DIYgenomics. Citizen Science Health Research Studies: Personal wellness and preventive medicine through collective intelligence. AAAI Symposium on Self-Tracking and Collective Intelligence for Personal Wellness.
https://www.aaai.org/ocs/index.php/SSS/SSS12/paper/viewFile/4293/4681 (accessed October 8, 2016).

Swan, Melanie. 2012b. Health 2050: The realization of personalized medicine through crowdsourcing, the quantified self, and the participatory biocitizen. *Journal of Personalized Medicine* 2(3): 93–118.

Swan, Melanie. 2014a. Neural data privacy rights. In *What should we be worried about? Real scenarios that keep scientists up at night*, ed. John Brockman, 406–409. New York: HarperCollins.
Available at https://www.edge.org/annual-question/what-should-we-be-worried-about (accessed October 8, 2016).

Swan, Melanie. 2014b. The non-cruciality of personal identity: Immortality as possibility. In *The prospect of immortality: Fifty years later*, ed. Charles Tandy, 385-420, Ann Arbor, MI: Ria University Press.

Swan, Melanie. 2015a. Antidote to holacracy: Blockchain smart assets. *Broader Perspective* blog, entry posted July 12.
http://futurememes.blogspot.com/2015/07/antidote-to-holocracy-blockchain-smart.html
(accessed October 8, 2016).

Swan, Melanie. 2015b. *Blockchain: Blueprint for a new economy*. Sebastopol CA: O'Reilly Media.

Swan, Melanie. 2015c. Blockchain thinking: The brain as a DAC (decentralized autonomous corporation). *Technology and Society Magazine* 34(4): 41–52.

Swan, Melanie. 2015d. Connected car: Quantified self becomes quantified car. *Journal of Sensor and Actuator Networks* 4(1): 2–29.

Swan, Melanie. 2015e. Digital Simondon: The collective individuation of man and machine. "Gilbert Simondon: Media and technics." Special issue, *Platform: Journal of Media and Communication* 6(1): 46–58.

Swan, Melanie. 2015f. Nanomedicine and cognitive enhancement. Society for Brain Mapping and Therapeutics. Conference presentation, Los Angeles, CA, March 6–8. http://www.slideshare.net/lablogga/philosophy-of-cognitive-enhancement (accessed October 8, 2016).

Swan, Melanie. 2016a. *Cognitive enhancement as subjectivation: Bergson, Deleuze and Simondon*. Saarbrücken: Lambert Academic Publishing.

Swan, Melanie. 2016b. Rethinking authority with the blockchain crypto enlightenment. *Edge*. (Response to The Edge Question 2016: What do you consider the the most interesting recent [scientific] news? What makes it important?) https://www.edge.org/response-detail/26720 (accessed October 8, 2016).

Takayama, Leila. 2015. Telepresence and apparent agency in human-robot interaction. In *The Handbook of the Psychology of Communication Technology*, ed. S. Shyam Sundar, 160–75. London: Wiley-Blackwell.

Tuckman, Bruce. 1965. Developmental sequence in small groups. *Psychological Bulletin* 63(6): 384–99.

Uehling, Mark D. 2004. Bio-IT World Best Practices Award 2004: Terry Fetterhoff and his team at Roche Diagnostics worked with Entelos to study virtual patients. *Bio-IT World.* August 18. http://www.bio-itworld.com/archive/081804/roche.html (accessed October 8, 2016).

van Gelder, Jean-Louis, Hal E. Hershfield, and Loran F. Nordgren. 2013. Vividness of the future self predicts delinquency. *Psychological Science* 24(6): 974–80.

Vinge, Vernor. 2007. *Rainbows end*. New York: Tor.

Welch, Chris. 2014. Google challenges Apple's HealthKit with release of Google Fit for Android. *The Verge.* October 28. http://www.theverge.com/2014/10/28/7083977/google-fit-now-available-android (accessed October 8, 2016).

Wolfram, Stephen. 2002. *A new kind of science*. Somerville, MA: Wolfram Media.

Yoffe, Emily. 2009. How the brain hard-wires us to love Google, Twitter, and texting. And why that's dangerous. *Slate*. August 12. http://www.slate.com/articles/health_and_science/science/2009/08/seeking.html (accessed October 8, 2016).

Yudkowsky, Eliezer. 2008. The design space of minds-in-general. *LessWrong* blog, entry posted June 25. http://lesswrong.com/lw/rm/the_design_space_of_mindsingeneral/ (accessed October 8, 2016).

Zolfagharifard, Ellie, and Richard Gray. 2015. The social network that lets you live forever: Lifenaut collects enough information to upload your personality to a computer. *DailyMail.* May 18. http://www.dailymail.co.uk/sciencetech/article-3086670/The-social-network-lets-live-forever-Lifenaut-collects-information-upload-personality-computer.html (accessed October 8, 2016).