Journal of EVOLUTION & Technology

A peer-reviewed electronic journal published by the Institute for Ethics and Emerging Technologies
ISSN 1541-0099
26(2) – July 2016

Embodiment in Whole-Brain Emulation and its Implications for Death Anxiety

Charl Linssen
Donders Institute for Brain, Cognition and Behaviour
Radboud University
Nijmegen, The Netherlands

c.linssen@science.ru.nl

Pieter Lemmens
Institute for Science, Innovation and Society
Radboud University
Nijmegen, The Netherlands

p.lemmens@science.ru.nl

Journal of Evolution and Technology - Vol. 26 Issue 2 – July 2016 - pgs 1-15

Abstract

The awareness of death is a central motivating force behind human activity. Their capacities for abstract and symbolic reasoning give human beings a unique foresight of their finite lifetime and forthcoming demise. Because of the overwhelming nature of this realization, we try to cope with the ensuing anxieties by means of various cognitive and existential strategies. One such strategy is to create a meaningful legacy during one's lifetime that will outlive the single individual. Whole-brain emulation (WBE) is another approach, but is unusual because of its literal promise to abolish death. Starting from the premise that WBE is feasible and will advance to such a level that we can speak of uploaded minds, we explore the implications of an allegedly immortal existence in a computational substrate: for our embodiment in the first place, and for death anxiety in the second. We argue that uploading would change the nature of, but could ultimately never abolish, embodiment. Instead, the defining characteristic of all brains are their vital links to the bodies that contain them and their interactions with the environment that are mediated by the body. In this light, we discuss the limits of WBE's potential to mitigate death anxieties: limits related to the (objective) probability of ceasing to exist, but also those that stem from the perception of the body as a proxy for death.

1 Origins of death anxiety

Human beings are a uniquely hybrid species: part biology, part culture (Moravec 1990). We have an unrivalled capacity for symbolic reasoning and thought, expressed in typically human faculties and abilities such as language, logical reasoning, technology, mathematics, and art. This contrasts with the physical or bodily realm, which is an inheritance from our evolutionary ancestors. The process of biological evolution necessarily involves the survival and differential reproduction of organisms, and for human beings bodily survival is among our most deep-seated and elementary drives. However, a new type of self-perpetuation, fundamentally different from its classical counterpart found in biology, is now possible in the realm of symbols.

The body places a person in a more or less standardized species form. Our bodily makeup and basic biology are essentially fixed from the time of conception: an individual cannot choose the particular body envelope they find themselves in. Throughout history, we have tried to run counter to this, to place our bodies in a position subordinate to our rational wills and to exert intellectual control over them. Examples abound: adornment; modification (including plastic surgery); and augmentation through the use of technology. In spite of these efforts, the control that we can exert over our own bodies, given the current state of technology, remains weak. We remain as captive as ever to an essentially unchangeable bodily envelope. To make matters worse, it constantly reminds us of its inevitable demise, through pain, illness and aging, as well as through carnal, species-stereotyped actions such as defecation and sex (Becker 1973).

The symbolic realm, on the other hand, does not constrain us to a predetermined species role or even to the physical world per se. Instead, it allows us to express our desire for perpetuation and continued life through the use of symbol or metaphor in ways that are highly specific, the outcome of each individual's creative personality and personal wishes. The symbolic realm satisfies the ancient inner urge for perpetuation and continued life, not through passing on the genes of oneself or one's tribe, but through passing on technological artifacts and via cultural impact, such as knowing that our various religions, countries, or personal achievements (e.g. scientific discoveries) will outlive us as single individuals (Greenberg et al. 1997).

It is only in the confluence of these two domains – the tangible and the symbolic – that we can know the terror of death: "to have emerged from nothing, to have a name, consciousness of self, deep inner feelings, an excruciating inner yearning for life and self-expression – and with all this, yet to die" (Becker 1973, 87). Only human beings have a real foresight of their impending demise, made possible by their symbolic conception of themselves as beings placed in time. This foresight brings with it feelings of dread and anxiety, because we are programmed (by evolution) against death, yet powerless to stop it. To prevent these overwhelming feelings from interfering too much with our daily functioning, we develop strategies for repression or "denial." These repressions are always imperfect and leave us vulnerable to psychological dysfunction, and they may in fact disrupt everyday life (e.g. in compulsions: see Becker 1973). An alleviation of death anxiety mediated by technology could thus have a major positive impact on psychological well-being.

Note that drawing a distinction between the tangible, physical realm and that of symbolic meaning does not rely upon notions of mind-body dualism (Cartesianism) or non-physicalism (non-materialism). All thought and experience is assumed to be causally dependent on, and fully explained by, the structure and function of our nervous systems. Nonetheless, awareness of and coping with the concept of death (a symbolic construct) are central aspects of human psychology. Their effects can be studied and quantified by, for example, experimentally inducing death salience and recording changes in behavior and physiology. Studies such as these have demonstrated the pervasive nature of death anxiety and its coping mechanisms, even in everyday life (for an overview, see Greenberg et al. 2014). The importance of such defensive mechanisms is highlighted by the finding that eliminating death awareness from working memory is an active cognitive process that competes for resources with other processes (Arndt et al. 1997).

Because of its roots in higher cognitive processes, death anxiety is complex and multifactorial. Some cognitive dimensions of death anxiety have been defined by Neimeyer (1994). These include each of the following: beliefs or ideas about the state of being dead (non-existence); understandings of the dying process or being destroyed; the thought of a radical transformation or separation; the idea of death as a threat to the meaningfulness of life (or as a threat to coming to a full realization of life's basic goals and propensities); images of significant others; fear of helplessness and loss of control; fear of uncertainty and the unknown; fear of missing out on things; and ideas about the body after death. Attitudes toward death and responses to death salience are personal and the result of cumulative life experience (Cicirelli 2006), which makes them susceptible to change in personal and cultural circumstance. Anxiety about death is an ancient motivator of human action (Rendu et al. 2014) and has found expression from cave drawings to the more recent "mindfiles" that promise a type of continued survival in cyberspace after death (Rothblatt 2013).

The remainder of this paper explores how a potential future technology, whole-brain emulation (WBE), might interact particularly strongly with death anxiety and associated coping strategies. Clearly, these interactions are complex, and involve more than the simple continued existence of the body. After introducing WBE in Section 2 we shall therefore take a closer look at the role of embodiment in Section 3, before discussing the implications for issues related to death in Section 4.

2 Whole-brain emulation

There is a converging interest in cognitive function from many fields of science. In the last two decades, we have witnessed spectacular methodological and technological developments in microbiology and associated competencies such as lab automation, information storage and retrieval, data analysis, and computational modeling. Although still in their infancy, these developments are allowing us to measure, analyze, and understand many attributes of the nervous system that are known to be relevant for its function. These attributes, such as ultrastructure (Briggman and Bock 2012) and in-situ genetic expression (Lee et al. 2014), can in principle be acquired across the entire volume of the brain, thereby capturing all the variability that is unique to an individual. The recording process is then followed by a mapping or reconstruction where the data is analyzed and converted into a format that is more convenient for computer manipulation.

In addition to a static analysis of the brain data, the computer can implement a phenomenological (descriptive) model of physical processes that take place in the real, biological brain. By numerical evaluation of the model, the simulation would evolve in time, replicating the functional dynamics of normal physiological brain function (Koene 2012). It is still a matter of debate what level of detail is required in the simulation. It stands to reason, however, by extrapolating past scientific progress, that it will become feasible to build computational models that simulate the biological brain to such a degree of accuracy that the simulation can be considered functionally indistinguishable from the original. Given a non-dualistic view of the mind-body problem, this implies that the simulation captures the complete identity and individuality of the person whose brain was scanned. This reasoning has given rise to the term “mind uploading” (Sandberg and Bostrom 2008).

To perform the upload of an individual, a destructive scanning and sectioning process, applied to their biological brain, is currently the most feasible and best investigated method. Alternative approaches could include replacing a single cell at a time by engineered equivalents (Moravec 1990) or hypothetical non-destructive scanning methods analogous to magnetic resonance imaging. Other than uploading, WBE could also be achieved by means of ontogenetic software that emulates the growth and development of a few cells or a single cell (e.g. a zygote) into a human central nervous system. The resulting individuals could be said to be “born” as emulants.

Once the process of uploading is complete, the uploaded person has acquired *substrate independence* (Koene 2013): the computational model is universal and does not rely on the specific type of computer that is used to implement it. Everything that is computable (such as the model) can be computed on any general-purpose processing unit. A computer that is supporting the emulation will have a particular architecture, but it can just as well be replaced by one with an entirely different architecture: perhaps one that does not use silicon, but carbon nanotubes, for its computation elements, or a neuromorphic chip that performs vast amounts of local computations in parallel, instead of a traditional general-purpose processor with a sequential load-store architecture. The temporal evolution of the emulation is invariant to the substrate that is used: although the supporting hardware can be very different, we are not adjusting the neural network that is being simulated, but only the simulator (Sandberg and Bostrom 2008).

By making use of the wealth of information implicit in biology, WBE can in principle be achieved without much knowledge of how the activity of networked cells leads to behavior. However, unless the emulation is going to be held in a state of complete isolation, some information about the *neural code*, that is, how information is represented by firing patterns of neurons, is going to be necessary. Knowing the code allows interfaces between the spiking signals of neurons and the digital electrical signals in which information is represented in conventional computer formats. An example of such an interface could be that between a peripheral nerve and an artificial muscle. Peripheral codes are relatively simple and well-understood compared to codes in the central nervous system; contemporary applications of neuronal coding and decoding already include prostheses for motor control, exteroception, and proprioception (Wander and Rao 2014).

In the rest of this paper, for the purpose of discussion, we shall axiomatically assume the feasibility of WBE. We assume that all the functionally relevant information can be acquired and composited into a simulation; when run, that simulation then emulates the original person, with a subjective experience and external identity indistinguishable from the original (Blackford 2014).

3 Embodiment in whole-brain emulation

In its most basic sense, the term "embodied" is synonymous with "incarnate," referring to the fact that we are each invested with a body. The body is both host to the brain and positioned to mediate all of the individual's experience with the world. These two elements will be taken up separately in Sub-section 3.1. In Sub-section 3.2, we then argue that having a body is inevitable, even in WBE. In Sub-section 3.3, we review the implications of WBE for what it means to have a body.

3.1 Distinguishing embodiment and substrate

In the process of uploading, a change of substrate occurs. Initially, the physical processes that are necessary for the operation of the brain are carried out by biology. In biology, there is an inextricable link between substrate and (human) body: the body is host to the brain, offering metabolic support and structural protection. When the brain is uploaded, these functions of the body become superfluous to brain function: an emulation no longer requires oxygen and nutrients for sustenance. Thus, the literal "incarnation" has changed from a biological substrate to the substrate of a computer.

However, this transition does not necessarily alter our perceptions of our bodies and the environment around us. The brain can perceive only indirectly, by receiving (and sending) sequences of neuronal action potentials ("spikes") that encode properties of the world. The brain does not "see"; it only receives spiking signals from the retinas and other senses. But this is not what we experience: we experience light and objects as existing objectively and veridically, independent of perception, and without any reference to action potentials or neural coding. The term *embeddedness* refers to this environment, the phenomenal world "out there" that we can feel and touch and interact with.

The physical body forms an essential part of this interaction, because the distribution of its senses and musculature dictates the ways in which we can even begin to interact with our surroundings. A body is on the one hand clearly a demarcated entity in the world, but on the other hand, being a part of the same world, is subject to the same physical laws and constraints, being thus "embedded" in the world. The term "embodiment" refers to this body, which constitutes a necessary bridge between the objects and phenomena that exist in the environment and the spike signals of the nervous system.

At present, the biological substrate (the brain) is concomitant with the body, given that the brain is situated in the same body that mediates our experience. In WBE, this need no longer be the case (Figure 1). Although the computer chips that are supporting the emulation are still sitting (tangibly) in a room somewhere, they are far removed from the person's direct experience: a person's experience is not confined to that computer room or chip. Thus, in an uploaded existence, the perceived embodiment can be vastly different from the substrate. Embodiment in WBE should thus refer to the body and environment that we perceive to be physically real, even though it may not actually exist in the real world (see Sub-section 3.3).

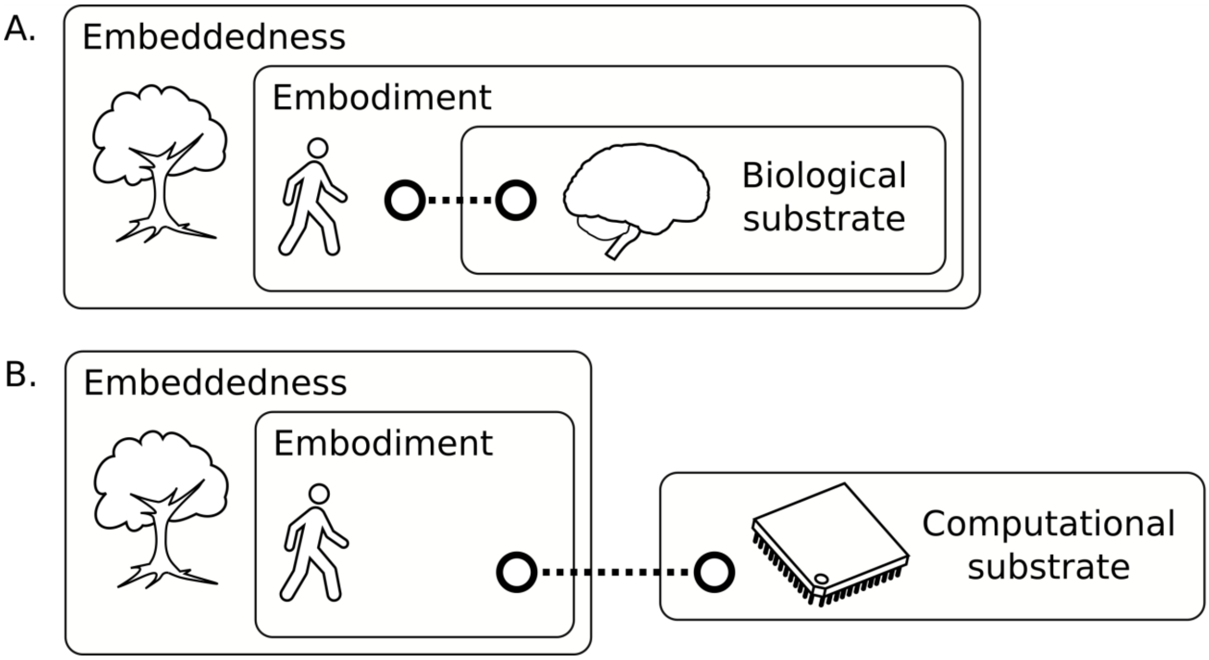

Figure 1: Diagram showing the dependencies between embeddedness, embodiment, and substrate. (A) Present situation: the brain is contained in the body which is contained in the environment. Action potentials are routed between brain and body via nerve fibers (dashed line). (B) Situation for WBE: the substrate is no longer fixed to the body. Action potentials are electronically routed between the emulation and body (dashed line).

3.2 Primacy of the body in the mind

A considerable part of the human brain is dedicated to interacting with the world. Pioneering work in the nineteenth century demonstrated a close correspondence between areas of the neocortex and the sensory and motor groups of the body. By directly stimulating the cortex with electrical current, movement of the muscles can be evoked, as well as percepts (Silverstein 2012). In many areas, a topographical continuity was found to map nearby locations on the cortical surface to nearby points in the periphery (e.g. in primary motor, sensory, or visual cortex; see Lesser et al. 1998). Topographic maps are found extensively in the brain, not just in lower sensory areas, but even in the hippocampus, which sits at the top of the cortical hierarchy, and in frontal areas (Silver and Kastner 2009) that are well known to be involved in higher, quintessentially human thinking. The widespread distribution of topographic correspondences intimates a strong integration of body-related processing and processing related to higher cognition.

Convergent evidence indeed supports this viewpoint. From early electrophysiological results, it would appear that different regions of the cerebral and cerebellar cortex (or brain at large) can be separated by clear boundaries, and that a specific function can be assigned to each region. Such divisions can, however, be made on the basis of countless variations between cells: morphology, synaptic distribution, epigenetic profile, plasticity, connectivity, and so on. Early brain atlases were based only on coarse cytoarchitecture; taking into account the much more detailed data that has only recently become available, dense networks emerge that do not follow a strict and regular pattern, but can be described by only the most general laws, such as power-law scaling and small-worldness (characterized by many "shortcuts"). For example, a prominent association fiber bundle directly connects frontal and primary visual areas (Forkel et al. 2014), possibly involved in attention or prediction. On the whole, anatomical data negates the initial modular view of cognition, which now instead appears to rely on widespread integration between all parts of the neural network.

Similarly, a functional view of the neocortex has emerged that emphasizes its hierarchical organization (Moratti et al. 2014). For example, areas related to vision are organized so that a relatively close correspondence exists between lower areas and the outside world, and between higher areas and symbolic thought. In the visual hierarchy, lower levels are indeed organized in topographic correspondence with the visual field, and extract basic image features related to geometry, motion, and color. Higher regions operate at a more integrative and semantic level: for example, these identify objects in a scene and conceptualize their location (Perry and Fallah 2014). However, the processing that occurs along this hierarchy is complex and interdigitated: any organizational scheme that acknowledges the actual complexity found in nature is forced to intersperse many other areas between the lower and upper, which correspond to varying degrees with low-level or high-level features but can never clearly be associated with either. In such a scheme, it is impossible to maintain a distinction between "pure" sensory processing and "pure" higher cognition. In fact, hierarchically lower areas are known to be actively involved in prediction and analysis of upcoming stimuli (Chaumon et al. 2014), mental imagery (Albers et al. 2013), and the direction of attention (Zumer et al. 2014). Attentional mechanisms are crucial to overcome information bottlenecks in the brain, and are largely orchestrated by frontal (higher cognitive) regions (Katsuki and Constantinidis 2014). Also, the roles of the cerebellum, a prominent part of the human brain that has long been known to be involved in the acquisition of precise motor skills and control, are now known to include higher cognition, language, and affect processing (De Smet et al. 2013).

In addition to these detailed neurological considerations, the integration of the body and higher cognition can be approached by studying language. The term "embodied cognition" often refers to the idea that the development of symbolic thought, such as semantics in language, rationality, and imagination, is crucially based in the body and its interactions with the environment. According to this theory, we develop primary metaphors during childhood and early physical and social experience. These include such basic phenomena as the level of water in a cup going up as water is poured in. A primary metaphor is employed when we state that stock prices are "going up." Higher-level concepts are metaphors that are less direct than the primary metaphors but nevertheless rely on bodily experience in the world. For example, conceptual metaphors related to time, with the past or future lying behind or ahead of us, are rooted in an understanding of how we can interact with the world through seeing and moving (Lakoff and Johnson 1980; Fauconnier and Turner 2004). Without a grounding in basic sensorimotor experience, these abstract notions would have no meaning to us.

To summarize, it is misleading to refer to areas of the brain as if they were modular. Instead, it is impossible to draw strict boundaries separating, for example, the more primitive sensorimotor functions from higher cognition such as planning, judgment, decision-making, and direction of attention. Furthermore, these areas are highly integrated, functionally and anatomically, in a complex and dense network. Thus, we conclude that the aforementioned "quintessentially human" (which, employing the parlance of the first section, would fall under the symbolic as opposed to the physical) functions cannot exist independently of brain structures that are devoted to world interaction and body control. Even for an uploaded individual, a body (human-like or otherwise) would remain a necessity.

3.3 Fulfilling the need for a body

In the process of uploading (Section 2), the complete structure of the nervous system is faithfully encoded into a digital model, which could easily include the complete networks of the retina or spinal cord, until the very point where the peripheral nerve fibers enter and exit the spine, or could even include peripheral nerve fibers that connect the central nervous system to muscle endplates and sensory transduction neurons. As remarked in Section 2, the neural code at this level is simple and well understood, so it is feasible for the signals to be generated and interpreted by soft- and hardware.

The body and environment can, in general terms, take on one of two forms: they can be supplied by a simulation or somehow be linked to the real, physical universe.

By the time uploading is feasible, we will have constructed advanced robot implementations of the human body that can serve as the ultimate telepresence device for us (Brooks 2002). They will allow humans to live in environments where traditionally they could not, such as underwater or on a planet without an atmosphere. Even though our brain emulation is taking place remotely, somewhere safe, our experience can be instantiated in these robotic avatars, which are free to roam within the range of the wireless communications link. Avatar robots will thus extend the range of human action and allow us to escape the limitations of any biological body (Moravec 1990).

The alternative to instantiation in a real-world avatar is to opt for instantiation in a virtual environment. In this scenario, a computer would emulate not only the brain itself but also the body and environment. A familiar contemporary example of how this might look is Second Life, a popular internet-based virtual world (for more, see

Linden Research n.d.). Unlike robotic avatars, simulated avatars are not bound by the laws of physics. In the world of Second Life, one can, for instance, teleport instantaneously from one location to another.

The human brain has a remarkable capacity to adapt quickly, effortlessly, and transparently to new types of embodiment, learning to control their dynamics and novel capabilities (Pino et al. 2014). As an example, consider telepresence robots. At present, they essentially consist of a camera and microphone mounted on a mobile platform. If the camera image is projected onto our retinas, and at the same time we are in control of the gaze or movement of the camera (by moving the robot), we very quickly move our egocentric body space to that of the robot, and experience a sense of "presence" at the robot's location. This can be demonstrated by threatening the robot (for instance, holding it over a cliff) and measuring the galvanic skin response or brain activation patterns of the controlling subject, which are equivalent to those measured in situations where the person's real body is threatened (Ramachandran and Blakeslee 1998).

Because embodiment of a simulated brain is a matter of routing sensorimotor signals, these may at any point in time be re-routed to an avatar with different characteristics. The user could select any desired body on demand. This is known as *multiple embodiment* (Moravec 1990; Clark 2003). After training, our brain is expected to accommodate almost instantaneously to a new avatar in the same way that a tool, when picked up, can immediately become part of our extended body schema (ibid.).

Even virtual bodies have their limitations, and the particular architecture of the human brain imposes further restrictions on how the avatar should work. An impoverished sensorimotor connection to the avatar is detrimental and ruins the sense of presence (Clemente et al. 2014), with delay or latency being especially important (Held and Durlach 1991). Failure to integrate multisensory information has the same effect, for example in the form of visual or spatial incongruencies with the avatar's posture or proprioception (Costantini and Haggard 2007; Perez-Marcos et al. 2012). Equally unconvincing are spatial offsets between where an object is seen and where it is felt (Takahashi et al. 2009).

Essential to a body are the notions of a spatial delineation, and a focalized first-person perspective in space from which the world is perceived. Both can be manipulated and are doubly dissociated. Lesions, seizures, or stimulation of target cortical regions can alter the experience of self-location, such as during so-called out-of-body experiences (i.e. displaced self-location), also described as taking a third-person perspective on one's body. Bodies can even be invisible; that is to say, the sense of "I" can be in empty space ("illusory body"; Serino et al. 2013), but throughout these manipulations both the focal perspective and spatial delineation of the body remain present.

Hayles (1994) recounts two examples of virtual bodies: those of a snake and a crow, animals that are primarily propelled by undulatory locomotion and flight, respectively. These behaviors are normally supported by motor pattern generating circuitry in basal brain areas and the spine – circuitry that is not present in the human nervous system, which is tailored instead for efficient bipedal locomotion. As this circuitry is relatively peripheral, it could be amenable to intervention, where a change of avatar is associated with a change in these peripheral circuits, but this is speculative and more evidence is needed on the efficacy of adapting to avatars increasingly different from human form.

4 Death anxiety in whole-brain emulation

The existence of WBE, or even its conception, can be interpreted as a direct outcome of dealing with death anxiety. Inventing methods and strategies to defeat death is the result of a deep-seated drive toward continued existence. WBE is by its very definition a method by which to transcend biological death. What sets it apart from other, more traditional, strategies for coping with death is that it could be the first to actually fulfill its promise technologically. If WBE is indeed feasible, will it lead, for uploaded minds, to the abolition of death-related anxieties and the associated need for coping strategies? Continuing the distinction between substrate and body that was made in Sub-section 3.1, we take up these aspects in Sub-sections 4.1 and 4.2, respectively.

4.1 Anxiety due to the substrate

Rejecting mind-body dualism implies that, in order to sustain the mind indefinitely, it is necessary and sufficient to sustain the substrate indefinitely.[1] The fact that we remain embodied is no longer an impediment to this, because the substrate and body are now separable (Sub-section 3.1). With WBE, the body becomes an avatar, be it in a simulated environment or in the real world. Avatars free their human users from the lethal repercussions of fate or bad judgment. There need be no risk of dying in a car accident or of old age – or of aging in the first place. Even the destruction of a robot body that exists in the real, physical world does not imply that the person dies, because the avatar can be controlled wirelessly from a safe location.

To perform WBE, some kind of material substrate will always be necessary. The dynamical model describing an individual's brain needs to be instantiated in biological, computational or other hardware that allows its dynamics to unfold over time. Loss of the particular hardware that the emulation is running on need not be tantamount to death, as specific microchips are not essential: underlying hardware is replaceable, and one computer chip is as good as any other (substrate independence). In a suitably designed, redundant system, hardware failure would result in automatic failover, having no effect, or a negligible effect, on the dynamics of the simulation. "Death" for an upload can occur only as the result of an irreversible termination of the dynamical model simulation, for example due to loss of the model data.

The probability of such an irreversible termination could in principle be marginalized by redundancy and careful design, to the point where no such events occur over cosmic timescales. However, this is not the same as achieving immortality as if it were a mathematical (logical) certainty. Reducing the probability of death is not a solution to death anxiety: "The smallest virus or the stupidest accident would deprive a man not of 90 years but of 900 – and would be then 10 times more absurd. […] If something is 10 times more absurd it is 10 times more threatening" (Becker 1973, 267). Furthermore, the model and state of the emulation are themselves vulnerable to extinction:

> [B]rain emulations are extremely vulnerable by default: the software and data constituting them and their mental states can be erased or changed by anybody with access to the system on which they are running. Their bodies are not self-contained and their survival is dependent upon hardware they might not have causal control over. They can also be subjected to undetectable violations such as illicit copying. (Sandberg 2014)

In addition, Rothblatt (2014) mentions the privacy violations that might arise from someone with physical access to the substrate hardware, in particular, the incumbent government. It is an open question who will be endowed with the responsibility of maintaining the physical substrate.

To some degree, these limitations can be mitigated by making off-line, off-site backups of the model data. In case the model data associated with the running emulation is destroyed, it can be re-instantiated from a backup taken at a certain time in the past, so that subjectively the person only experiences a time interval of non-consciousness. Death "forecloses fewer opportunities for uploads" (Sandberg 2014): instead of being binary and irreversible, the concept of death becomes more graded. Having to resort to increasingly older backups (due to an increasingly large magnitude of catastrophe) results in the loss of more memories, with the person retroactively having to forego the period of time between the making of the backup and the time of the accident, in which they were essentially "dead."

Death anxiety, in the most literal sense, is fear brought about by the limitations of a finite lifespan. Let us assume that an uploaded person knows about the substrate that is supporting them, even though it is not normally perceived by them. At present, we cope with knowledge of our biological substrate's numerous vulnerabilities by engaging cognitive strategies that ultimately result in a drive for symbolic immortality. In biology, these fears and strategies are continuously reinforced by such daily reminders as the deaths of others and maladies of the body. In WBE, if the substrate is built to extreme standards of reliability, the frequency of such reminders could be drastically reduced. At the same time, outside threats such as economic crises, political oppression, computer viruses, or war (or even the threat of war) might easily bring the ultimate reliance on a sound substrate back into focus – regardless of their probability of actually eventuating. As an analogy, present-day terrorism is given a great deal of attention and captures the imagination, even though the probability of it happening to a particular individual is infinitesimally small. Given how rapidly humans adapt to the capabilities offered by new technology, it seems likely that any newly achieved levels of redundancy and resilience will simply become the new "normal," at which point we demand the next stage in assuring even higher levels over even longer timescales.

Fundamentally, for an uploaded person, it is clear that elements in the simulation can be destroyed as easily as they can be created. Since an uploaded person is defined solely on the basis of their (model) data, there is, by implication, always the possibility of the impossibility of continued existence. (Without this possibility, uploaded persons would be trapped.) Knowing the ultimate vulnerability of the substrate hardware is, by itself, enough justification for persisting death anxieties, but now there is also the vulnerability that stems from the software (the WBE model data, simulation software, avatar parameters, and so on) and the dependency on hardware that might not be under one's own control. Precise elaborations and emphases in death anxiety will remain subject to circumstance. For example, exposure to events related less to failure of particular physical hardware, but more to the loss of model data or "pattern" – be it from accidental or deliberate action – could induce efforts to invest in cryptographic or other software measures, designed to safeguard model data.

4.2 Anxiety due to embodiment

In WBE, the body and environment that are experienced by a person in day-to-day life are vastly different from the physical hardware that supports their emulation. In spite of a hypothetically indestructible substrate, which would release any a priori constraints on lifespan for the uploaded individual, death anxiety might continue. This is, in part, because much of our current death anxiety derives from the static, predetermined, and confining nature of a body and its embeddedness in the world.

A biological body is experienced as constraining because of its limitations: it is fixed and cannot be molded according to a person's wishes. Millions of years of evolution by natural selection have given us, soberly put, "senses capable of perceiving medium-size objects in a narrow spectrum moving at slow speeds" (More 2014, 223). Needless to say, in WBE this situation could easily be made radically different. Morphological freedom (Sandberg 2013) would be greatly enhanced, since a virtual or robot body could be designed according to individual requirements, as opposed to being one of a given stereotype, with arbitrary and fixed characteristics such as gender and skin color. Even after its initial creation, the form of an avatar would remain plastic and subservient to the demands of the intellectual self, and an uploaded individual is not restricted to a single form (multiple embodiment).

Control over the body also extends to control over the environment. Second Life, for example, contains a separate interface that allows the user to construct objects in the world, to manipulate their size, shape, and appearance, and to endow these with dynamical (scripted) effects. Restrictions are there not because of intrinsic limitations, but because the creators of the software have chosen to enforce them. In principle, even things like gravity and the laws of physics can be altered at will in a virtual environment.

Instantiation in a robot avatar might seem impoverished in comparison, but robot bodies can always be upgraded to make use of the latest technology, and wireless communication can make possible many actions at a distance that would presently appear as magic.

In 1973, Becker said: "There is no strength that can overcome creature anxiety unless one is a god and not a creature" (Becker 1973, 261). But uploaded individuals would find themselves in a position that appears conspicuously close to god-like. By the time uploading can be made available, the boundedness and determinism of the current human body will certainly have been overcome. But does this statement hold for *any* type of body? In spite of the leap in power and control implied by uploading, a bodily instantiation remains unavoidable. Any kind of bodily instantiation is confining because of its very function: sampling states of the world and lending particular degrees of freedom to manipulate it. This problem is exacerbated if the options available for customizing and enhancing the avatar are restricted (for whatever reason, e.g. due to economic or computational cost). The avatar might then be perceived as confining rather than liberating (Slater and Usoh 1994). The same holds for how successful a person is in using the tools available to them (including their body) to manipulate the external environment. If the tools available are primitive or require a high level of skill to operate, the person will experience few degrees of freedom and a type of claustrophobia might ensue. New modes of perception and action require training and practice to master gracefully; due to widespread integration in the brain (Sub-section 3.2), we cannot expect to just "plug in" a skill.

In addition, even with uploading our morphological freedom will remain limited by the need to share common ground with the rest of society and to have meaningful interactions with each other. Bodily appearance is a strong part of one's sense of identity. This holds for whatever avatar we choose; the remarkable plasticity of this mental self-representation has been demonstrated by virtual reality experiments, where a subject infers their expected disposition from their avatar's appearance, and then changes their attitude and behavior to conform to the expectations (Yee and Bailenson 2007). Avatars are also important for a range of nonverbal interactions (Yee et al. 2007), which necessarily rely on shared norms. On a more pragmatic level, it does not make any sense to express oneself using a modality for which the other party has no sensory organ.

For an uploaded mind, any practical restrictions of the avatar and the simulation itself will be encountered on a daily basis, making them much more salient and confronting compared to the now relatively rare events relating to (impending) substrate dysfunction. Embodiment and embeddedness contribute to death anxiety to the extent that the body and the ability to manipulate the physical environment are experienced as limiting the freedom of creativity and self-expression. This makes them contingent on many factors, such as personal ambitions and abilities, societal norms, and technological progress. The question of how specific factors contribute to death-related anxieties can already be addressed by present-day research, for example by measuring the degree to which the anxieties are evoked on the basis of various freedoms or restrictions when exploring a virtual world. Blascovich and Bailenson (2011) describe a virtual reality study in which participants were told that a virtual avatar would be created in their likeness, which would subsequently be stored for generations to come in a university-backed "digital vault." When it was revealed during debriefing that participants' avatars would not actually be saved, they became so irate that the researchers felt it necessary to abort the study.

5. Conclusion

Existence in a silicon substrate would once and for all shift the emphasis from hardware survival to pattern survival: "The pattern is what's important, not the substrate" (Tipler 1994). However, this does not do away with a fight for survival per se. It is fundamentally impossible for WBE to offer a guaranteed, absolute immortality, only a receding probability horizon of death. Even the apparent tautology, that to reduce the probability of death (loss of pattern) will proportionally reduce death anxiety, is by no means vacuously true. Death anxiety is multifactorial and can be interpreted far beyond the straightforward sense of death of the biological body. In this paper we have explored three elements in particular: the substrate, embodiment, and embeddedness in an environment.

Uncertainties related to maintenance of the computational substrate could by themselves form sufficient justification for the persistence of death anxieties. Although the substrate hardware can in principle be made orders of magnitude more reliable and durable than the human body, the essential vulnerability of any hardware not only remains but is compounded by that of the software (WBE model and simulator).

Moreover, in WBE the body and environment will be at our command, but only to the degree of what is made possible by the system's software and hardware, the economy, society, the state, the irreversible passage of time, opportunity costs, and so on. While enhancing freedom with respect to our present situation, WBE ultimately only enforces a new set of constraints and dependencies. There are many further aspects to death anxiety that we have not explored here in detail, but could conceivably allay or exacerbate it. For example, the ability to create other instances of a particular running emulation ("(mind)clones" or "branches") could allay death anxiety analogous to procreation in biology.

Kierkegaard claimed that death anxiety ("dread" in his words) can actually enrich life, rather than diminish it: "[…] not that [faith] annihilates dread, but remaining ever young, it is continually developing itself out of the death throe of dread" (cited in Becker 1973, 91). In other words, it could be a good thing for (some) anxiety to remain; it could sustain an inner drive toward (intellectual) self-perpetuation. A guaranteed existence could easily lead to passivity and ennui: if our emulation continues to run indefinitely, we may no longer feel the need to leave behind a symbolic legacy in the form of cultural or technological progress. For some, the absence of a foreseeable death would infringe on the "authenticity" of human existence; they claim that finitude is necessary for genuine individuality and freedom (Bailey 2014). However, as we have seen, there should be no fear of death anxieties, be they literal or symbolic, disappearing in the first place. Instead, central to the human subject appears to lie a powerful drive toward freedom and self-expression, and a move away from the predetermined and fallible mechanics of a biological body, both in its role as a substrate and as an embodiment.

Death anxiety impels us to produce a meaningful symbolic legacy in a given personal context. However, the arrival of WBE will radically change what people value in creating such a legacy. Because of close links to individual life history and cultural standards, the outcomes of these changes are hard to predict, and they are expected to continue to morph over time. Inevitably, then, a great variety in attitudes can be expected. For example, retaining a self-sufficient and self-contained body in the real world could actually be a very safe and comforting thought. It is sometimes tacitly assumed that growing beyond the need for a body is desirable and a key purpose of WBE. Instead, holding on to our bodies, with the additional freedom of designing and modifying them in countless different ways, could be a very fulfilling pursuit in its own right (Vita-More 2004).

Note

1. I.e. with our current neuroscientific knowledge; perhaps the architecture of our brains is somehow fundamentally unsuitable for a lifespan measured in millennia, with respect to, e.g., loss of or interference with encoded memories.

References

Albers, A.M., P. Kok, I. Toni, H.C. Dijkerman, and F.P. De Lange. 2013. Shared representations for working memory and mental imagery in early visual cortex. *Current Biology* 23(15): 1427–1431.

Arndt, J., J. Greenberg, S. Solomon, T. Pyszczynski, and L. Simon. 1997. Suppression, accessibility of death-related thoughts, and cultural worldview defense: Exploring the psychodynamics of terror management. *Journal of Personality and Social Psychology* 73(1): 5–18.

Bailey, J.I. 2014. Enframing the flesh: Heidegger, transhumanism, and the body as "standing reserve." *Journal of Evolution and Technology* 24: 44–62. Available online at http://jetpress.org/v24/bailey.pdf (accessed June 13, 2016).

Becker, E. 1973. *The denial of death*. New York: Free Press.

Blackford, R. 2014. Bring on the machines. In *Intelligence unbound: The future of uploaded and machine minds*, ed. R. Blackford and D. Broderick, 102–118. Malden, MA: Wiley-Blackwell.

Blascovich, J., and J. Bailenson. 2011. *Infinite reality: Avatars, eternal life, new worlds, and the dawn of the virtual revolution*. New York: William Morrow.

Briggman, K.L., and D.D. Bock. 2012. Volume electron microscopy for neuronal circuit reconstruction. *Current Opinion in Neurobiology* 22(1): 154–61.

Brooks, R.A. 2002. *Flesh and machines: How robots will change us.* New York: Pantheon.

Chaumon, M., K. Kveraga, L.F. Barrett, and M. Bar. 2014. Visual predictions in the orbitofrontal cortex rely on associative content. *Cerebral Cortex* 24(11): 2899–2907.

Cicirelli, V.G. 2006. Fear of death in mid-old age. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences* 61(2): 75–81.

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

Clemente, M., A. Rodríguez, B. Rey, and M. Alcaiz. 2014. Assessment of the influence of navigation control and screen size on the sense of presence in virtual reality using EEG. *Expert Systems with Applications* 41(4, part 2): 1584–1592.

Costantini, M., and P. Haggard. 2007. The rubber hand illusion: Sensitivity and reference frame for body ownership. *Consciousness and Cognition* 16(2): 229–40.

De Smet, H.J., P. Paquier, J. Verhoeven, and P. Marien. 2013. The cerebellum: Its role in language and related cognitive and affective functions. *Brain and Language* 127(3): 334–42.

Fauconnier, G., and M. Turner. 2004. *The way we think: Conceptual blending and the mind's hidden complexities.* New York: Basic Books.

Forkel, S.J., M. Thiebaut de Schotten, J.M. Kawadler, F. Dell'Acqua, A. Danek, and M. Catani. 2014. The anatomy of fronto-occipital connections from early blunt dissections to contemporary tractography. *Cortex* 56: 73–84.

Greenberg, J., S. Solomon, and T. Pyszczynski. 1997. Terror management theory of self-esteem and cultural worldviews: Empirical assessments and conceptual refinements. In *Advances in experimental social psychology*, Vol. 29, ed. M.P. Zanna, 61–139. London: Academic Press.

Greenberg, J., K. Vail, and T. Pyszczynski. 2014. Terror management theory and research: How the desire for death transcendence drives our strivings for meaning and significance. In *Advances in motivation science*, Vol. 1, ed. A.J. Elliot, 85–134. Amsterdam: Academic Press .

Hayles, N.K. 1994. Boundary disputes: Homeostasis, reflexivity, and the foundations of cybernetics. *Configurations* 2: 441–67.

Held, R., and N. Durlach. 1991. Telepresence, time delay and adaptation. In *Pictorial communication in virtual and real environments*, ed. S.R. Ellis, 232–46. Bristol, PA: Taylor & Francis.

Katsuki, F., and C. Constantinidis. 2014. Bottom-up and top-down attention: Different processes and overlapping neural systems. *Neuroscientist* 20(5): 509–521.

Koene, R.A. 2012. Fundamentals of whole brain emulation: State, transition and update representations. *International Journal of Machine Consciousness* 4(1): 5–21.

Koene, R.A. 2013. Uploading to substrate-independent minds. In *The transhumanist reader: Classical and contemporary essays on the science, technology, and philosophy of the human future*, ed. M. More and N. Vita-More, 146–56. Malden, MA: Wiley-Blackwell.

Lakoff, G., and M. Johnson. 1980. *Metaphors we live by*. Chicago: University of Chicago Press.

Lee, J.H., E.R. Daugharthy, J. Scheiman, R. Kalhor, J.L.Yang, T.C. Ferrante, R. Terry, S.S. Jeanty, C. Li, R. Amamoto, D.T. Peters, B.M. Turczyk, A.H. Marblestone, S.A. Inverso, A. Bernard, P. Mali, X. Rios, J. Aach, and G.M. Church. 2014. Highly multiplexed subcellular RNA sequencing in situ. *Science* 343(6177): 1360–1363.

Lesser, R.P., S. Arroyo, N. Crone, and B. Gordon. 1998. Motor and sensory mapping of the frontal and occipital lobes. *Epilepsia* 39 Suppl 4: 69–80.

Linden Research. n.d. Second Life (official site). http://secondlife.com/ (accessed June 13, 2016).

Moratti, S., C. Mendez-Bertolo, F. Del-Pozo, and B.A. Strange. 2014. Dynamic gamma frequency feedback coupling between higher and lower order visual cortices underlies perceptual completion in humans. *Neuroimage* 86: 470–79.

Moravec, H. 1990. *Mind children: The future of robot and human intelligence.* Cambridge, MA: Harvard University Press.

More, M. 2014. The enhanced carnality of post-biological life. In *Intelligence unbound: The future of uploaded and machine minds*, ed. R. Blackford and D. Broderick, 222–30. Malden, MA: Wiley-Blackwell.

Neimeyer, R.A. 1994. *Death anxiety handbook: Research, instrumentation, and application.* Washington, DC: Taylor and Francis.

Perez-Marcos, D., M.V. Sanchez-Vives, and M. Slater 2012. Is my hand connected to my body? The impact of body continuity and arm alignment on the virtual hand illusion. *Cognitive Neurodynamics* 6(4): 295–305.

Perry, C.J., and M. Fallah. 2014. Feature integration and object representations along the dorsal stream visual hierarchy. *Frontiers in Computational Neuroscience* 8: 84.

Pino, G.D., A. Maravita, L. Zollo, E. Guglielmelli, and V.D. Lazzaro. 2014. Augmentation-related brain plasticity. *Frontiers in Systems Neuroscience* 8: 109.

Ramachandran, V., and S. Blakeslee. 1998. *Phantoms in the brain: Probing the mysteries of the human mind.* New York: William Morrow.

Rendu, W., C. Beauval, I. Crevecoeur, P. Bayle, A. Balzeau, T. Bismuth, L. Bourguignon, G. Delfour, J.-P. Faivre, F. Lacrampe-Cuyaubre, C. Tavormina, D. Todisco, A. Turq, and B. Maureille. 2014. Evidence supporting an intentional Neandertal burial at La Chapelle-aux-Saints. *Proceedings of the National Academy of Sciences* 111(1): 81–86.

Rothblatt, M. 2013. Techno-immortalism 101: What are mindfiles? In *Longevitize! Essays on the science, philosophy & politics of longevity*, ed. F. Cortese, 62–66. N.p.: Center for Transhumanity (Kindle ebook).

Rothblatt, M. 2014. *Virtually human: The promise – and the peril – of digital immortality*. New York: St. Martin's Press.

Sandberg, A. 2013. Morphological freedom – Why we not just want it, but need it. In *The transhumanist reader: Classical and contemporary essays on the science, technology, and philosophy of the human future*, ed. M. More and N. Vita-More, 56–64. Malden, MA: Wiley-Blackwell.

Sandberg, A. 2014. Ethics of brain emulations. *Journal of Experimental and Theoretical Artificial Intelligence* 26(3): 439–57.

Sandberg, A., and N. Bostrom. 2008. *Whole Brain Emulation: A roadmap.* Technical report #2008-3. Oxford: Future of Humanity Institute (University of Oxford).

Serino, A., A. Alsmith, M. Costantini, A. Mandrigin, A. Tajadura-Jimenez, and C. Lopez. 2013. Bodily ownership and self-location: Components of bodily self consciousness. *Consciousness and Cognition* 22: 1239–1252.

Silver, M.A., and S. Kastner. 2009. Topographic maps in human frontal and parietal cortex. *Trends in Cognitive Sciences* 13(11): 488–95.

Silverstein, J. 2012. Mapping the motor and sensory cortices: A historical look and a current case study in sensorimotor localization and direct cortical motor stimulation. *Neurodiagnostic Journal* 52(1): 54–68.

Slater, M., and M. Usoh. 1994. Body centered interaction in immersive virtual environments. In *Artificial life and virtual reality*, ed. N. Thalmann and D. Thalmann, 125–47. New York: John Wiley and Sons.

Takahashi, C., J. Diedrichsen, and S.J. Watt. 2009. Integration of vision and haptics during tool use. *Journal of Vision* 9(6): 1–13.

Tipler, F. 1994. *The physics of immortality: Modern cosmology, God and the resurrection of the dead.* London: Doubleday.

Vita-More, N. 2004. The new [human] genre – primo [first] posthuman. Paper presented at the Ciber@RT Conference, Bilbao, Spain, 23–30 April.
Available at http://www.natasha.cc/paper.htm (accessed June 6, 2016).

Wander, J.D. and R.P. Rao. 2014. Brain-computer interfaces: A powerful tool for scientific inquiry. *Current Opinion in Neurobiology* 25: 70–75.

Yee, N., and J. Bailenson. 2007. The Proteus effect: The effect of transformed self-representation on behavior. *Human Communication Research* 33(3): 271–90.

Yee, N., J.N. Bailenson, M. Urbanek, F. Chang, and D. Merget, D. 2007. The unbearable likeness of being digital: The persistence of nonverbal social norms in online virtual environments. *Cyberpsychology and Behavior* 10(1): 115–21.

Zumer, J.M., R. Scheeringa, J.M. Schoffelen, D.G. Norris, and O. Jensen. 2014. Occipital alpha activity during stimulus processing gates the information flow to object-selective cortex. *PLoS Biology* 12(10): e1001965.