Facts about Brain Preservation

Version: 245 8197

PHOTO CREDITS. Part I: The Bibliotheca Alexandrina's interior, from near the top, by Carsten Whimster. Part IV: Plato's "Allegory of the Cave", drawing by Markus Maurer. Back cover: The Fountain of Eternal Life in downtown Cleveland, Ohio, with 200 Public Square in the background by ricky rhodes.

Abstracts

Abstracts iii

I Can I choose how long I'll live? 1

1 A Possible Cure for Death 3
Charles B. Olson

2 A Door to the Future 4
K. Eric Drexler

3 Broca's Brain 5
Carl Sagan

II How would brain preservation work in practice? 6

4 Aldehyde-Stabilized Cryopreservation 8
Robert L. McIntyre, Gregory M. Fahy

5 High-resolution Whole-brain Staining for Electron Microscopic Circuit Reconstruction 9
Shawn Mikula, Winfried Denk

6 Electron Imaging Technology for Whole Brain Neural Circuit Mapping 10
Kenneth Hayworth

7 Whole Brain Emulation: A Roadmap 11
Anders Sandberg, Nick Bostrom

8 **Vitrifying the Connectomic Self: A case for developing Aldehyde Stabilized Cryopreservation into a medical procedure** 12
Kenneth Hayworth

III How are memories stored in the brain? 13

9 **Memory Systems of the Brain: a brief history and current perspective** 15
Larry R. Squire

10 **Do Thin Spines Learn to be Mushroom Spines that Remember?** 16
Jennifer Bourne, Kristen M. Harris

IV Will a clone of me still be *me*? 17

11 **Personal Identity and Uploading** 19
Mark Walker

12 **The Fallacy of Favoring Gradual Replacement Mind Uploading Over Scan-and-Copy** 20
Keith B. Wiley, Randal A. Koene

V Isn't preserving my brain selfish? 21

13 **The Ethics of Exponential Life Extension through Brain Preservation** 23
Michael A. Cerullo

14 **Overcoming Objections To Brain Preservation** 24
John Smart

Appendices 25

A **Vitrifying the Connectomic Self: A case for developing Aldehyde Stabilized Cryopreservation into a medical procedure** 26
Kenneth Hayworth
A.1 Introduction 26

A.2 A vision of the near (and far) future 28
A.3 What this fictional story is designed to address 40
A.4 Separating facts from personal opinions 41
A.5 The core of the scientific argument: I am my connectome and ASC preserves the connectome 42
A.6 How 'you' are encoded in your connectome 43
A.7 Conclusion 45
A.8 References 46

B Overcoming Objections To Brain Preservation 54
John M. Smart
B.1 The Scientific Advancement Defense 55
B.2 The Exceptional Cases Defense 56
B.3 The Social Benefits Defense 56
B.4 The Religious Objection and Defense 60
B.5 The Natural Aging and Death Objection and Defense 61
B.6 The Patternism Hypothesis 63
B.7 A Number of Non-Obvious Proposals 65
B.8 A Brief Bibliography 66
B.9 Nine More Objections and Defenses—For the Scientists and Philosophers 67
B.10 Challenges for the Future 76

Part I

Can I choose how long I'll live?

Bibliotheca Alexandrina

1 A Possible Cure for Death

Chemical preservation of the brain may prevent death. Life for an individual human being is inextricably linked to the existence of his or her mind. It is widely accepted that the mind is a product of the functioning of the brain, which, according to this view, is nothing more and nothing less than a fantastically complicated machine. Chemical preservation of the brain (promptly after the cessation of vital functions) preserves not only the neuronal configuration but also a great deal of molecular structure. Thus, it is plausible that a chemopreserved brain contains within it the information of the design of the "brain machine". If so, then technology of the distant future may be able to extract that information and construct a new functionally identical brain machine (as well as a body), thereby allowing the corresponding individual to wake up and live again. It is argued that one's identity is defined by what the brain does rather than how it does it or what it does it with, and therefore that replacement of one's brain with a functionally identical machine does not affect one's identity. Some advantages of chemopreservation relative to cryopreservation as a possible means of preventing death are discussed.

OLSON, Charles B. A possible cure for death. *Medical hypotheses*, 1988, vol. 26, no 1, p. 77–84.
`https://www.ncbi.nlm.nih.gov/pubmed/3398793`

2 A Door to the Future

Benjamin Franklin wanted a procedure for stopping and restarting metabolism, but none was then known. Do we live in a century far enough advanced to make biostasis available—to open a future of health to patients who would otherwise lack any choice but dissolution after they have expired?

We can stop metabolism in many ways, but biostasis, to be of use, must be reversible. This leads to a curious situation. Whether we can place patients in biostasis using present techniques depends entirely on whether future techniques will be able to reverse the process. The procedure has two parts, of which we must master only one.

If biostasis can keep a patient unchanged for years, then those future techniques will include sophisticated cell repair systems. We must therefore judge the success of present biostasis procedures in light of the ultimate abilities of future medicine. Before cell repair machines became a clear prospect, those abilities—and thus the requirements for successful biostasis—remained grossly uncertain. Now, the basic requirements seem fairly obvious.

DREXLER, K. Eric. *Engines of creation.* Anchor Books, 1986.
`http://e-drexler.com/p/06/00/EOC_Cover.html`

3 Broca's Brain

It was difficult to hold Broca's brain without wondering whether in some sense Broca was still in there—his wit, his skeptical mien, his abrupt gesticulations when he talked, his quiet and sentimental moments. Might there be preserved in the configuration of neurons before me a recollection of the triumphant moment when he argued before the combined medical faculties (and his father, overflowing with pride) on the origins of aphasia? A dinner with his friend Victor Hugo? A stroll on a moonlit autumn evening, his wife holding a pretty parasol, along the Quai Voltaire and the Pont Royal? Where do we go when we die? Is Paul Broca still there in his formalin-filled bottle? Perhaps the memory traces have decayed, although there is good evidence from modern brain investigations that a given memory is redundantly stored in many different places in the brain. Might it be possible at some future time, when neurophysiology has advanced substantially, to reconstruct the memories or insights of someone long dead? And would that be a good thing? It would be the ultimate breach of privacy. But it would also be a kind of practical immortality, because, especially for a man like Broca, our minds are clearly a major aspect of who we are.

SAGAN, Carl. *Broca's brain: Reflections on the romance of science.* Presidio Press, 1980.
`https://en.wikipedia.org/wiki/Broca%27s_Brain`

Part II

How would brain preservation work in practice?

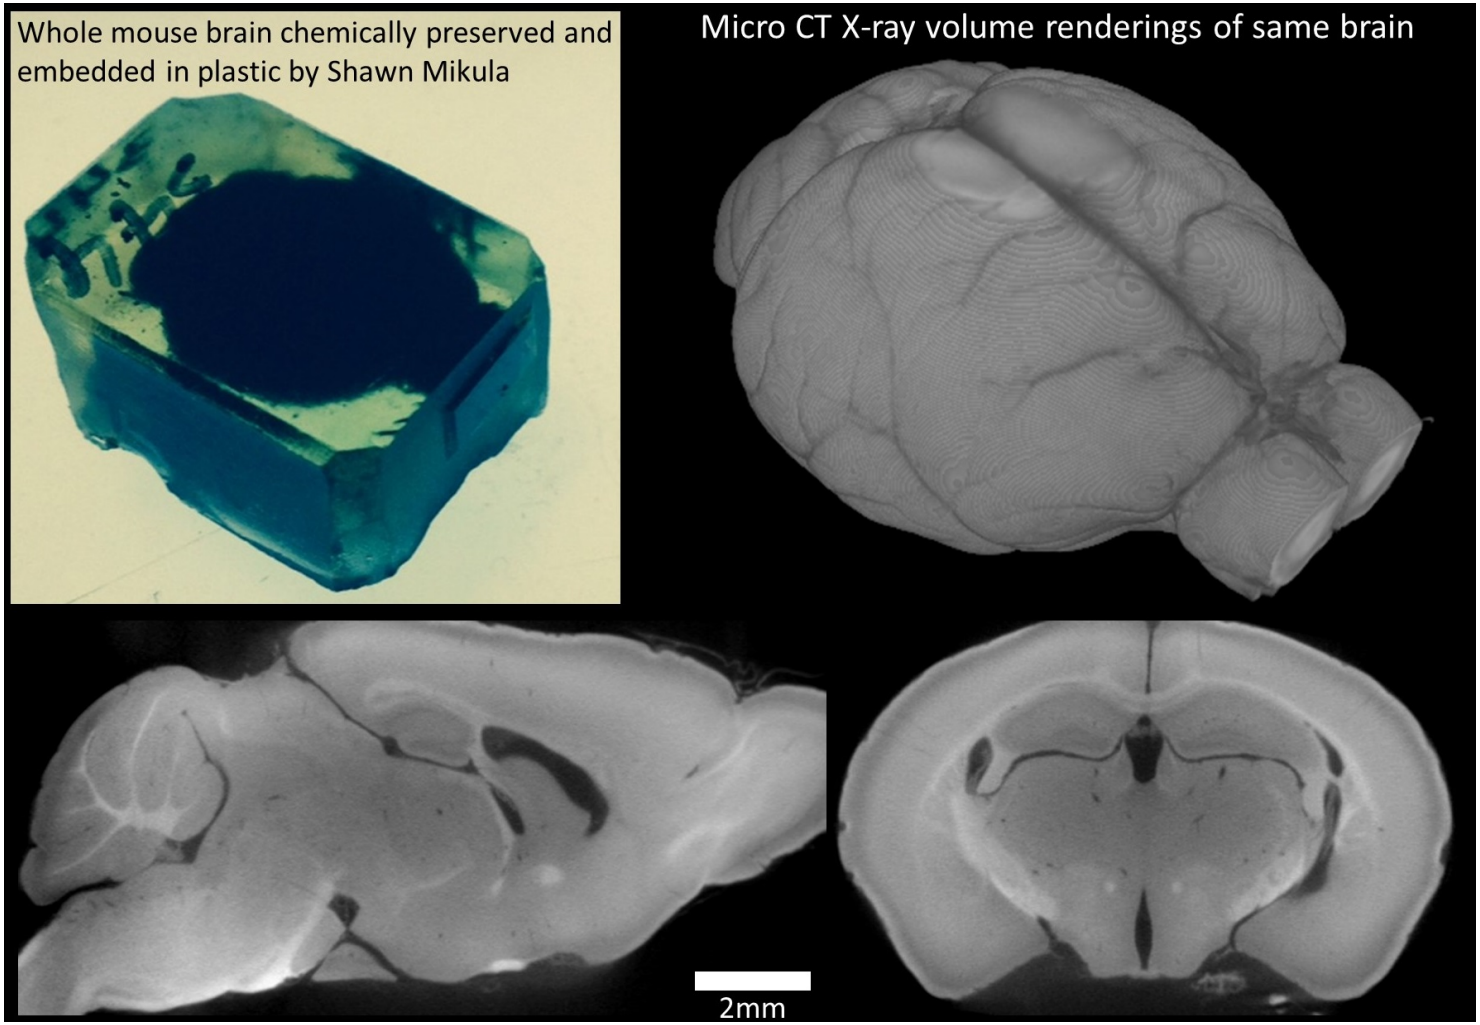
Whole mouse brain chemically preserved and embedded in plastic by Shawn Mikula
Micro CT X-ray volume renderings of same brain
2mm

4 Aldehyde-Stabilized Cryopreservation

We describe here a new cryobiological and neurobiological technique, aldehyde-stabilized cryopreservation (ASC), which demonstrates the relevance and utility of advanced cryopreservation science for the neurobiological research community. ASC is a new brain-banking technique designed to facilitate neuroanatomic research such as connectomics research, and has the unique ability to combine stable long term ice-free sample storage with excellent anatomical resolution. To demonstrate the feasibility of ASC, we perfuse-fixed rabbit and pig brains with a glutaraldehyde-based fixative, then slowly perfused increasing concentrations of ethylene glycol over several hours in a manner similar to techniques used for whole organ cryopreservation. Once 65% w/v ethylene glycol was reached, we vitrified brains at −135 °C for indefinite long-term storage. Vitrified brains were rewarmed and the cryoprotectant removed either by perfusion or gradual diffusion from brain slices. We evaluated ASC-processed brains by electron microscopy of multiple regions across the whole brain and by Focused Ion Beam Milling and Scanning Electron Microscopy (FIB-SEM) imaging of selected brain volumes. Preservation was uniformly excellent: processes were easily traceable and synapses were crisp in both species. Aldehyde-stabilized cryopreservation has many advantages over other brain-banking techniques: chemicals are delivered via perfusion, which enables easy scaling to brains of any size; vitrification ensures that the ultrastructure of the brain will not degrade even over very long storage times; and the cryoprotectant can be removed, yielding a perfusable aldehyde-preserved brain which is suitable for a wide variety of brain assays.

MCINTYRE, Robert L. and FAHY, Gregory M. Aldehyde-stabilized cryopreservation. *Cryobiology*, 2015, vol. 30, p. 1–11. `http://dx.doi.org/10.1016/j.cryobiol.2015.09.003`

5 High-resolution Whole-brain Staining for Electron Microscopic Circuit Reconstruction

Currently only electron microscopy provides the resolution necessary to reconstruct neuronal circuits completely and with single-synapse resolution. Because almost all behaviors rely on neural computations widely distributed throughout the brain, a reconstruction of brain-wide circuits—and, ultimately, the entire brain—is highly desirable. However, these reconstructions require the undivided brain to be prepared for electron microscopic observation. Here we describe a preparation, BROPA (brain-wide reduced-osmium staining with pyrogallol-mediated amplification), that results in the preservation and staining of ultrastructural details throughout the brain at a resolution necessary for tracing neuronal processes and identifying synaptic contacts between them. Using serial block-face electron microscopy (SBEM), we tested human annotator ability to follow neural 'wires' reliably and over long distances as well as the ability to detect synaptic contacts. Our results suggest that the BROPA method can produce a preparation suitable for the reconstruction of neural circuits spanning an entire mouse brain.

MIKULA, Shawn and DENK, Winfried. High-resolution whole-brain staining for electron microscopic circuit reconstruction. *Nature methods*, 2015, vol. 12, no 6, p. 541.
http://dx.doi.org/10.1038/nmeth.3361

6 Electron Imaging Technology for Whole Brain Neural Circuit Mapping

The goal of uploading a human mind into a computer is far beyond today's technology. But exactly how far? Here I review our best cognitive and neuroscience model of the mind and show that it is well suited to provide a framework to answer this question. The model suggests that our unique "software" is mainly digital in nature and is stored redundantly in the brain's synaptic connectivity matrix (i.e., our Connectome) in a way that should allow a copy to be successfully simulated. I review the resolution necessary for extracting this Connectome and conclude that today's FIBSEM technique already meets this requirement. I then sketch out a process capable of reducing a chemically-fixed, plastic-embedded brain into a set of tapes containing 20×20 micron tissue pillars optimally sized for automated FIBSEM imaging, and show how these tapes could be distributed among a large number of imaging machines to accomplish the task of extracting a Connectome. The scale of such an endeavor makes it impractical, but a version of this scheme utilizing a reduced number of imaging machines would allow for the creation of a "Connectome Observatory"—an important tool for neuroscience and a key milestone for mind uploading.

HAYWORTH, Kenneth J. Electron imaging technology for whole brain neural circuit mapping. *International Journal of Machine Consciousness*, 2012, vol. 4, no. 1, p. 87–108.
http://www.brainpreservation.org/wp-content/uploads/2016/01/ElectronImagingTechnologyForWholeBrainNeuralCircuitMapping_Hayworth2012.pdf

7 Whole Brain Emulation: A Roadmap

It appears feasible within the foreseeable future to store the full connectivity or even multistate compartment models of all neurons in the brain within the working memory of a large computing system.

Achieving the performance needed for real-time emulation appears to be a more serious computational problem. However, the uncertainties in this estimate are also larger since it depends on the currently unknown number of required states, the computational complexity of updating them (which may be amenable to drastic improvements if algorithmic shortcuts can be found), the presumed limitation of computer hardware improvements to a Moore's law growth rate, and the interplay between improving processors and improving parallelism. A rough conclusion would nevertheless be that if electrophysiological models are enough, full human brain emulations should be possible before mid-century. Animal models of simple mammals would be possible one to two decades before this.

SANDBERG, Anders, BOSTROM, Nick. *Whole Brain Emulation: A Roadmap*, Technical Report #2008-3, Future of Humanity Institute, Oxford University, 2008.
`https://www.fhi.ox.ac.uk/brain-emulation-roadmap-report.pdf`

8 Vitrifying the Connectomic Self: A case for developing Aldehyde Stabilized Cryopreservation into a medical procedure

. . . [b]ut the main point of this paper is to persuade the scientific and medical community that now is the time to develop this ASC procedure into a reliable medical procedure that can be offered to terminal patients. This is a radical proposal that can easily be misunderstood. This misunderstanding often manifests itself in questions like: "Why on earth would a terminal patient desire such an option in the first place?", "How would such a procedure work on a practical level?", "Are patient safeguards even possible for a procedure whose final success won't be known for decades or centuries?", "Can we even imagine the technologies that would allow future revival?"

Perhaps the best way to answer all of these questions is to offer a speculative short story meant to summarize and clarify this vision. The following fictional story follows a man diagnosed with Alzheimer's dementia in the year 2030 who chooses to undergo ASC preservation in the hopes of future revival. Extensive footnotes throughout this fictional story briefly explain the science behind key steps and point to references that support the science and technology discussed.

HAYWORTH, Kenneth. Vitrifying the Connectomic Self: A case for developing Aldehyde Stabilized Cryopreservation into a medical procedure. *The Brain Preservation Foundation Website*, 2018.
See Appendix A for the full text.

Part III

How are memories stored in the brain?

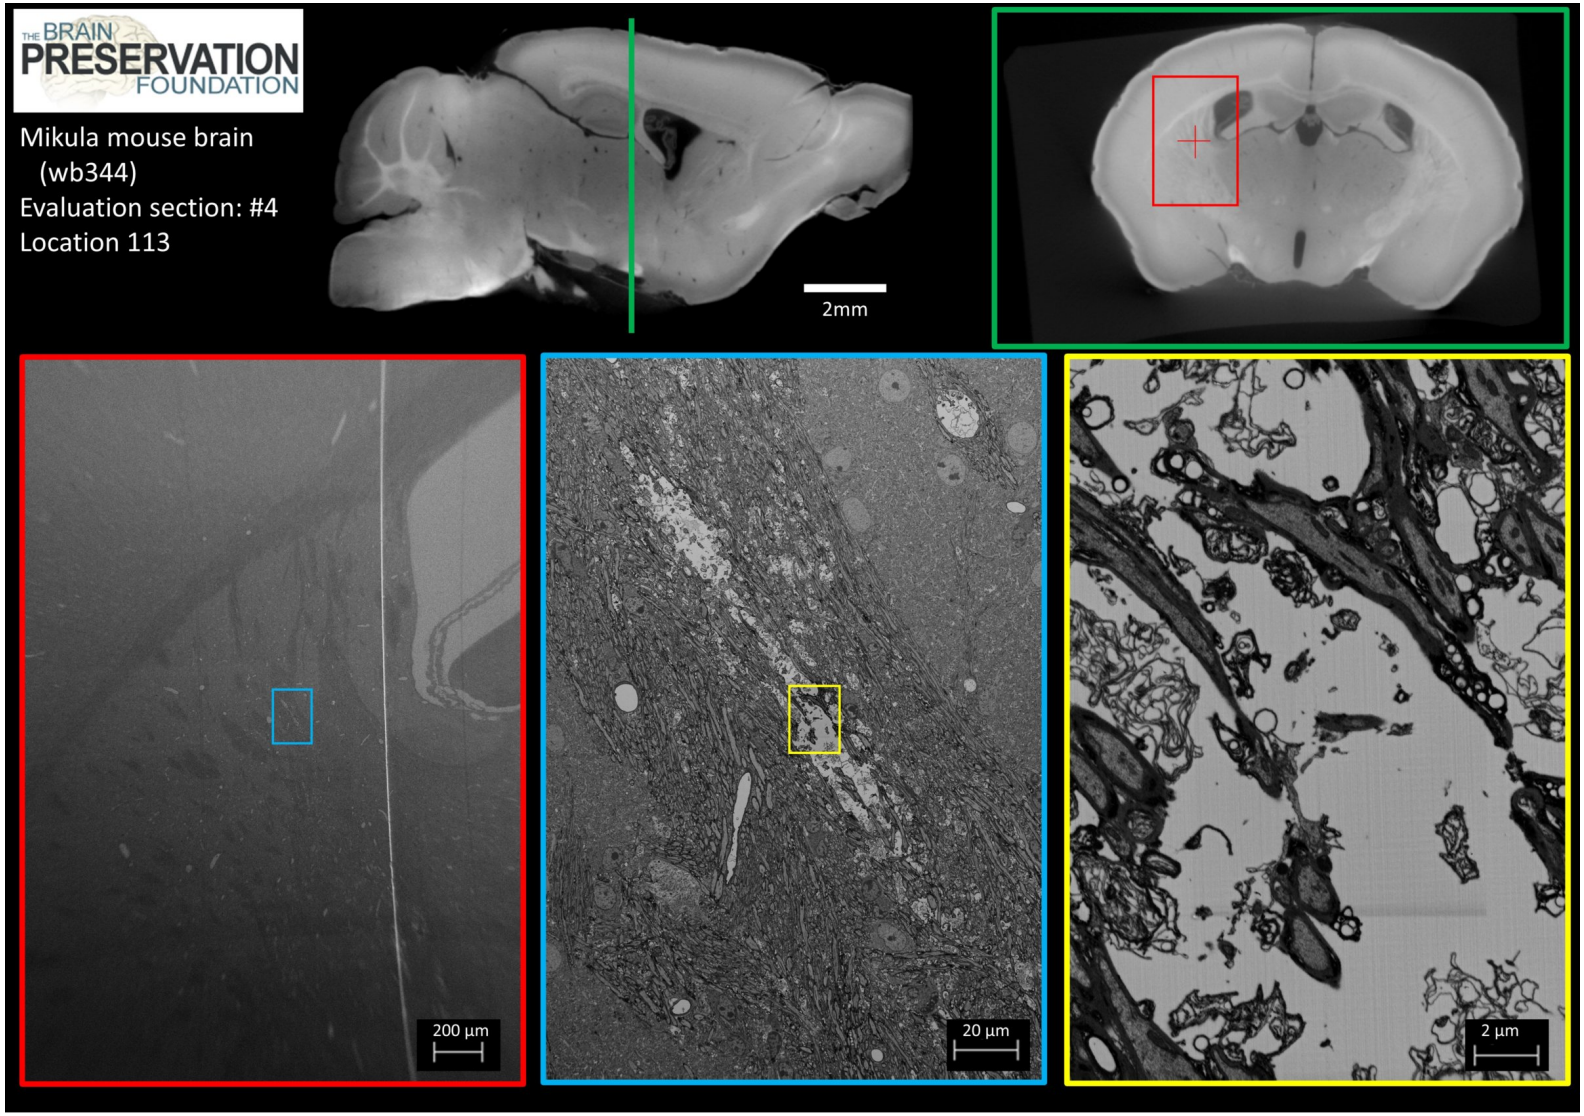
THE BRAIN PRESERVATION FOUNDATION
Mikula mouse brain
(wb344)
Evaluation section: #4
Location 113
2mm
200 µm
20 µm
2 µm

9 Memory Systems of the Brain: a brief history and current perspective

The idea that memory is composed of distinct systems has a long history but became a topic of experimental inquiry only after the middle of the 20th century. Beginning about 1980, evidence from normal subjects, amnesic patients, and experimental animals converged on the view that a fundamental distinction could be drawn between a kind of memory that is accessible to conscious recollection and another kind that is not. Subsequent work shifted thinking beyond dichotomies to a view, grounded in biology, that memory is composed of multiple separate systems supported, for example, by the hippocampus and related structures, the amygdala, the neostriatum, and the cerebellum. This article traces the development of these ideas and provides a current perspective on how these brain systems operate to support behavior.

SQUIRE, Larry R. Memory systems of the brain: a brief history and current perspective. *Neurobiology of learning and memory*, 2004, vol. 82, no 3, p. 171–177.
http://dx.doi.org/10.1016/j.nlm.2004.06.005

10 Do Thin Spines Learn to be Mushroom Spines that Remember?

Dendritic spines are the primary site of excitatory input on most principal neurons. Long-lasting changes in synaptic activity are accompanied by alterations in spine shape, size and number. The responsiveness of thin spines to increases and decreases in synaptic activity has led to the suggestion that they are 'learning spines', whereas the stability of mushroom spines suggests that they are 'memory spines'. Synaptic enhancement leads to an enlargement of thin spines into mushroom spines and the mobilization of subcellular resources to potentiated synapses. Thin spines also concentrate biochemical signals such as Ca^{2+}, providing the synaptic specificity required for learning. Determining the mechanisms that regulate spine morphology is essential for understanding the cellular changes that underlie learning and memory.

BOURNE, Jennifer and HARRIS, Kristen M. Do thin spines learn to be mushroom spines that remember?.*Current opinion in neurobiology*, 2007, vol. 17, no 3, p. 381–386.
https://doi.org/10.1016/j.conb.2007.04.009

Part IV

Will a clone of me still be *me*?

Plato's "Allegory of the Cave."

11 Personal Identity and Uploading

Objections to uploading may be parsed into substrate issues, dealing with the computer platform of upload and personal identity. This paper argues that the personal identity issues of uploading are no more or less challenging than those of bodily transfer often discussed in the philosophical literature. It is argued that what is important in personal identity involves both token and type identity. While uploading does not preserve token identity, it does save type identity; and even qua token, one may have good reason to think that the preservation of the type is worth the cost.

WALKER, Mark. Personal identity and uploading. *Journal of Evolution and Technology*, 2011, vol. 22, no 1, p. 37–51.
`https://jetpress.org/v22/walker.htm`

12 The Fallacy of Favoring Gradual Replacement Mind Uploading Over Scan-and-Copy

Mind uploading speculation and debate often concludes that a procedure described as *gradual inplace replacement* preserves personal identity while a procedure described as *destructive scan-and-copy* produces some other identity in the target substrate such that personal identity is lost along with the biological brain. This paper demonstrates a chain of reasoning that establishes metaphysical equivalence between these two methods in terms of preserving personal identity.

WILEY, Keith B. and KOENE, Randal A. The Fallacy of Favoring Gradual Replacement Mind Uploading Over Scan-and-Copy. *arXiv preprint arXiv:1504.06320*, 2015.
`https://arxiv.org/abs/1504.06320`

Part V

Isn't preserving my brain selfish?

Percentage of the globe with access to brain preservation: **0%**

13 The Ethics of Exponential Life Extension through Brain Preservation

Chemical brain preservation allows the brain to be preserved for millennia. In the coming decades, the information in a chemically preserved brain may be able to be decoded and emulated in a computer. I first examine the history of brain preservation and recent advances that indicate this may soon be a real possibility. I then argue that chemical brain preservation should be viewed as a life-saving medical procedure. Any technology that significantly extends the human life span faces many potential criticisms. However, standard medical ethics entails that individuals should have the autonomy to choose chemical brain preservation. Only if the harm to society caused by brain preservation and future emulation greatly outweighed any potential benefit would it be ethically acceptable to refuse individuals this medical intervention. Since no such harm exists, it is ethical for individuals to choose chemical brain preservation.

CERULLO, Michael A. The ethics of exponential life extension through brain preservation. *Journal of Evolution and Technology*, 2016, vol. 26, no 1, p. 94–105.
https://jetpress.org/v26.1/cerullo.htm

14 Overcoming Objections To Brain Preservation

This page discusses some common objections to and defenses of the value of brain preservation as a social option. We humans are only now coming to understand ourselves as informational entities. In so doing, we are learning the use and value of our information. In considering whether brain preservation is a wise and ethical use of resources, one must ask under what circumstances information itself, both generally and within unique human minds, is worth preserving as an individual choice in free societies.

SMART, John. Overcoming Objections To Brain Preservation. *The Brain Preservation Website*, 2015.
See Appendix B for the full text.

Appendices: Full Text of Select Articles

A Vitrifying the Connectomic Self: A case for developing Aldehyde Stabilized Cryopreservation into a medical procedure

A.1 Introduction

This paper advocates for the development and eventual deployment in hospitals of a medical procedure designed to preserve human brain connectome information for extremely long-term storage. Terminal patients electing to undergo such a procedure would do so with the hope of being 'revived' decades or centuries later via whole brain emulation[1] (i.e. mind uploading). Specifically, this paper advocates that terminal patients be given the option of electing to undergo vascular perfusion with the deadly chemical fixative glutaraldehyde—the same procedure used today to preserve the brains of laboratory animals for the highest-quality electron and immunofluorescent microscopy (e.g. Hayat 2000; Hua et al. 2015; Mikula & Denk 2015; Collman et al. 2015; Murray et al. 2015). Of course perfusion with glutaraldehyde results in near-instant death by any of today's standards; but this near-instantaneous cessation of metabolic activity and crosslinking of biomolecules is precisely what makes glutaraldehyde the optimal choice for preserving the nanoscale structure of whole brains for scientific study (Hayat 1986). A new technique called "Aldehyde Stabilized Cryopreservation" (ASC) (McIntyre & Fahy 2015) has recently demonstrated that such brains can be further perfused with cryoprotectants up to sufficiently high-concentrations to allow ice-free vitrification (Fahy et al. 2004) and storage

[1] I use the term emulation (as opposed to simulation) to make clear that the goal is to produce a fully functioning substitute for the original brain (see Sandberg & Bostrom 2008).

at -130oC, stopping all further decay. Time has essentially stopped for ASC brains stored solid at such a low temperature.

Why is this significant? Because, as I outline in detail below, glutaraldehyde fixation appears to preserve the full range of structural and molecular features that modern neuroscientific theories postulate underlie the encoding of all of the types of long-term memories that make a person unique. A terminal patient electing to undergo an ASC procedure is electing to "hit-pause", halting further disease-related damage to their brain and vascular system, in order to optimally preserve the full informational content of their brain. The similarity to cryonics (Lemler et al. 2004) is obvious, but in this case the dubious possibility of biological revival is dismissed and focus is instead directed toward provably preserving the information content of the brain by the absolute best method known to today's science, i.e. perfusion fixation with glutaraldehyde.

This paper will briefly review what neuroscience knows about how long-term memories are encoded in the brain, and will make the case that ASC is capable of preserving this information. This paper will also discuss the scientific and technological advances that will likely be needed to 'revive' a person by destructively scanning and computationally emulating their preserved brain. But the main point of this paper is to persuade the scientific and medical community that now is the time to develop this ASC procedure into a reliable medical procedure that can be offered to terminal patients. This is a radical proposal that can easily be misunderstood. This misunderstanding often manifests itself in questions like: "Why on earth would a terminal patient desire such an option in the first place?", "How would such a procedure work on a practical level?", "Are patient safeguards even possible for a procedure whose final success won't be known for decades or centuries?", "Can we even imagine the technologies that would allow future revival?"

Perhaps the best way to answer all of these questions is to offer a speculative short story meant to summarize and clarify this vision. The following fictional story follows a man diagnosed with Alzheimer's dementia in the year 2030 who chooses to undergo ASC preservation in the hopes of future revival. Extensive footnotes throughout this fictional story briefly explain the science behind key steps and point to references that support the science and technology discussed.

A.2 A vision of the near (and far) future

The year is 2030 and you go in for a neurological exam after your spouse notices that you are displaying mild memory loss. MRI and blood tests verify that you are experiencing the early stages of Alzheimer's dementia. This is devastating news, especially since you know what is in store. Years before you had been the primary caregiver for your mother during the last five years of her life and watched as the same disease robbed her of her memories to the point where she was unable to recall even her closest loved ones, robbed her of her cognitive abilities to point where the once proud teacher could no longer tie her own shoes, and altered her personality so remarkably that it was unrecognizable[2]. Every year you would take her in for an MRI scan and watch as her doctors showed you the progression of the disease. Looping through the yearly scans, you could literally see the disease shrinking her brain. The doctors would verify this quantitatively: "Her loss of brain volume this year was 3.1%"[3]. At the start of this grueling five year experience you had been comforted by the thought that your mother's immaterial soul would rise to heaven when the time eventually came. But in the end there was no such comfort since you had literally witnessed her soul eaten away a piece at a time in perfect synchrony with the loss of her brain tissue. Now you face that same fate and there is still no cure in sight.

Even a few years ago you would have had only two options: An early exit via euthanasia, or letting the disease take its course. But your doctors now offer you a third option: euthanasia by vascular perfusion with glutaraldehyde followed by long-term cryostorage—a procedure known as Aldehyde Stabilized Cryopreservation (ASC). Glutaraldehyde is a deadly chemical fixative that is used by neuroscientists to preserve the brains of animals prior to processing for electron and immunofluorescence microscopy. Perfusion of glutaraldehyde through the brain's vasculature almost instantly stops metabolic processes by covalently crosslinking cellular proteins into a sturdy mesh. Since life is a set of ongoing biochemical reactions this crosslinking results in immediate death, but it does so in a way that almost perfectly preserves the nano-scale structure of the brain. Fixation by glutaraldehyde is known to preserve the patterns of synaptic connections among neurons[4], preserve the ultrastructural details of synapses[5], and preserve the primary structure and relative locations of most proteins[6]. As a results of this crosslinking, a glutaraldehyde fixed brain is immune to biological

[2] E.g. Lyketsos et al. 2011; Rosenberg et al. 2015

[3] e.g. Chan et al. 2003

[4] E.g. Knott et al. 2008; Briggman, Helmstaedter & Denk 2011; Mikula & Denk 2015; Kasthuri et al. 2015; Lee et al. 2016;

[5] E.g. Hayat 2000; Bell et al. 2014; Bourne & Harris 2011

[6] E.g. Migneault et al. 2004; Murray et al. 2015; Collman et al. 2015

decay processes and will remain 'stable' for months, but eventually diffusion would result in the slow dislocation of biomolecules (e.g. membrane lipids) that were not crosslinked. For extremely long-term storage the glutaralehyde fixed brain is further perfused with a very high concentration of a cryoprotectant agent and brought to a temperature low enough to provide essentially indefinite storage[7].

You are not surprised that your doctor offers you this ASC option. The controversial new procedure has been all over the news for the last few years and, after a heated legal battle, ASC had recently been declared an acceptable method of euthanasia in the state you live in. On the face of it, it is an outlandish idea: fix your brain with a deadly chemical and store it in a static state for decades in the hope that future technology might be able to scan in your brain's information and revive you as a computer-emulated brain controlling a robotic body. Since childhood you had been fascinated by the idea of cryonics, intrigued by the idea of waking up in the far future to experience its wonders firsthand, and you vividly remember how disappointed you were when you learned how difficult real cryonics was—how much damage it caused to the brain's ultrastructure. But this new ASC technique was designed to overcome these limitations by chemically fixing the brain prior to the cryonics procedure, allowing the perfusion of cryoprotectants to be performed at room temperature over an extend length of time, thereby ensuring complete and uniform cryoprotectant concentration in every cell[8].

And the idea that you might wake up in the future as an emulated brain controlling a robotic body? When you initially heard of this idea, while watching the debates over ASC's legal adoption, it seemed patently absurd. "If such an emulated brain was even possible wouldn't it be 'just a copy' of me?", "I would still be dead wouldn't I?"[9] But the idea caught fire among the early-adopter 'Silicon Valley' crowd—the crowd you happen to work with. At work you are immersed in the world of artificial deep neural networks, networks that learn to drive cars, translate languages, recognize faces and objects, and that learn to play Chess and Go at superhuman levels[10]. When your job is to build applications based on artificial brains it becomes easier to imagine yourself

[7] McIntyre & Fahy 2015

[8] McIntyre & Fahy 2015

[9] I and others have addressed the philosophical questions regarding the preservation of personal identity with respect to brain preservation and mind uploading in papers like 'Personal Identity and Uploading' (Walker 2011); 'Killed by Bad Philosophy' (Hayworth 2010), 'Uploading and Branching Identity' (Cerullo 2015), 'The Fallacy of favoring gradual replacement mind uploading over scan-and-copy' (Wiley & Koene 2015)

[10] E.g. Hassabis et al. 2017; LeCun et al. 2016; He et al. 2015; Taigman et al. 2014; David et al. 2016; Silver et al. 2016

upgrading to an artificial substrate.

You decide to discuss your options with your coworkers. Unsurprisingly, for them the idea of waking up as a fully computer-emulated brain controlling a robotic body is literally the most attractive part of the ASC idea, and they proclaim, in all seriousness, that if the technology for mind uploading was available they would immediately sign up for the procedure. You ask them how they wrestle with the philosophical implications. Again, unsurprisingly, they embrace the idea that self-copies would be possible. They even discuss how being an emulated brain will allow one to 'fork' one's mind into two copies in the morning, live separate lives with separate conscious points of view during the day, and later in the evening 'merge the deltas' back into a single conscious self. After hours of discussions you admit that their enthusiasm has infected you as well. You decide that you will opt for the procedure, and, in consultation with your doctor, you set a tentative date for your ASC euthanasia. You set it for two years from now, before the most devastating decline will begin.

2032

The year is 2032 and two years have passed since your initial Alzheimer's diagnosis. During that time you and your spouse have kept track of your cognitive decline, and, via regular MRI scans, you have witnessed the gradual shrinkage of your brain. There is no doubt that your decline is accelerating with every passing month. You have, of course, kept track of the latest research toward finding a cure. Unfortunately achieving such a breakthrough in time to help you seems increasingly remote. But over this two years you have put your life in order. You completed the book you had been working on, trained your successor at work, and spent copious amounts of time with your friends and loved ones including a 'going away' party where they celebrated your life and discussed the possibility that all present would reunite in the future. Now the day has finally arrived for your ASC euthanasia procedure.

In the hospital you are given a general anesthetic that induces unconsciousness, you will remain unconscious for the rest of the procedure. Just before the anesthetic is injected you think: "If I experience anything ever again it will be waking up in the future." Unconscious, you are wheeled into an operating suite specially designed to handle ASC procedures. A vascular surgeon opens your chest and cannulates key blood vessels hooking them up to a perfusion apparatus. A rotating fluoroscope is mounted near your head making possible real time evaluation of the brain's perfusion.

The word is given and the apparatus begins perfusing an oxygenated buffer solution through your body's vascular system, displacing the blood in order to prevent clotting. Within a minute the buffer solution is replaced with a fixative

solution containing glutaraldehyde. As the fixative diffuses through your brain's capillaries it almost instantly halts all metabolic activity and starts to glue each cell's proteins into a sturdy meshwork. You are now dead according to all previous standards.

Over the course of the next half hour, fixative continues to flow through your vasculature. This flow is monitored via the fluoroscope by periodically injecting boluses of x-ray opaque contrast agents. If there are any parts of the brain that are deemed to not be receiving adequate perfusion then the pressure, flowrate, and duration can be adjusted. In cases of vascular blockage surgical intervention may be required to achieve adequate fixation of an area. But in your case the monitoring instruments show that the perfusion has gone flawlessly. The surgeon now gives the word to start gradually introducing cryoprotectant into the fixative solution. Over the course of several hours this gradual increase continues until a 65% ratio of cryoprotectant is reached—sufficient to prevent ice crystal formation at all temperatures[11].

Your body is now transferred out of the hospital and released into the hands of a third-party evaluation organization. They transfer your body into a cold storage unit used to test whether the procedure was successful. Its temperature is lowered to -130 degrees C, the temperature it will eventually be long-term stored at. Over the course of several days it is cycled several times between this temperature and room temperature, mimicking what may occur during a many-decades long storage[12]. Following this, your body is returned to room temperature and an x-ray CT scan is performed to check for any telltale damage to the brain or spinal cord. Using the results of this scan, the evaluation organization drills several small holes in your skull and uses these to take a set of tiny needle biopsy samples from any brain region that they suspect might not have been adequately perfused[13]. These needle biopsies are processed for chemical analysis and electron microscopy. Again it looks like your procedure went flawlessly. 3D electron microscopy of biopsy samples taken from a range of cortical regions and from the hippocampus, striatum, thalamus, and brainstem clearly show that your brain's pattern of neuronal connectivity and the ultrastructural details of its synapses have been well preserved—fixation was good and no ice crystal damage was seen[14]. A subset of biopsy samples

[11] What has been described is the sequence of steps which were outlined by McIntyre & Fahy (2015) and that have been demonstrated to preserve the structural connectome of whole rabbit and pig brains.

[12] Such a temperature cycling test was reported by McIntyre & Fahy (2015).

[13] E.g. Aghayev et al. 2007

[14] Such electron microscopic evaluations were performed by McIntyre & Fahy (2015) for their rabbit and pig brains following cold storage, and I performed independent electron microscopic evaluations as part of the Brain Preservation Prize challenge (

are processed for immunofluorescence microscopy to verify that the expected distribution of receptor proteins and ion channels has been preserved as well[15].

The third-party evaluation organization notifies your hospital and your surgeons of these results (giving them the feedback they need to ensure high quality in future cases), notifies your spouse (providing comfort that the information contained within your brain has been preserved), notifies the designated government regulatory office in charge of licensing the hospital and surgeons, and notifies your health insurance company to certify that full payment for the procedure is warranted[16]. Your body is then transferred to a dedicated long-term storage organization which stores thousands of similarly prepared bodies in large, refrigerated underground caverns.

2098

The year is 2098 and the technology to upload minds has not only been perfected, it is now so routine that healthy biologically-born humans often opt to undergo the procedure, a procedure that still starts with the vascular perfusion of glutaraldehyde. Over the intervening years neuroscience has learned precisely how the brain works, and how to decode the preserved brain's structure to create a faithful computer emulation containing the same memories and personality. As expected, it was verified that most memories are stored as physical changes to synaptic connections. And it was discovered that these memories could be reliably decoded by mapping the pattern of connections among the brain's neurons and by estimating the strengths of synapses based on their size[17]. This level of information is termed the structural connectome and it was found that a reasonably accurate brain emulation could be made based on the information

www.BrainPreservation.org).

[15] E.g. Murray et al. 2015

[16] This passage is meant to address the question: "Are patient safeguards even possible for a procedure whose final success won't be known for decades or centuries?" The answer is that plenty of regulatory safeguards can be put in place based on independent verification of the quality of preservation of the brain's ultrastructure.

[17] The neuroscience literature is filled with research and review articles supporting this conclusion. Here is a selection of articles that I think are particularly relevant to the proposal at hand. Review articles: Kasai et al. 2003; Hoshiba et al. 2017; Bailey et al. 2015; Josselyn et al. 2015; Poo et al. 2016; Lisman 2015; Bourne & Harris 2007; Yuste 2010; Segal 2016; Maren 2005; Lamprecht & LeDoux 2004; Tonegawa et al. 2015; Primary research articles: Matsuzaki et al. 2001; Matsuzaki et al. 2004; Noguchi et al. 2011; Bourne & Harris 2011; Trachtenberg et al. 2002; Liu et al. 2012; Liu et al. 2014; Ryan et al. 2015; Carrillo-Reid et al. 2016; Kitamura et al. 2017; Hayashi-Takagi et al. 2015.

in this structural connectome alone[18].

But neuroscientists discovered that some types of memories are difficult to reliably decode based solely on information available in the structural connectome. They found that if one could annotate this structural connectome with the membrane densities of a small number of key ion channel types then a range of physiological parameters could be estimated much more precisely, leading to a more faithful emulation requiring less post-revival parameter tuning[19].

[18] As the plethora of references cited in the previous footnote show, there is a growing consensus that learned knowledge is encoded via modifications to the strengths of synapses—strengths that should in principle be able to be estimated based on the electron microscopically-imaged structural connectome alone. The most detailed computational models of neural function today use compartmental models of neurons with estimates of ion channel and receptor densities based on neuronal type. These estimates are based on physiological recordings and morphological reconstructions of neurons in hundreds of 'side' experiments (e.g. Markram et al. 2015). With these facts in mind, the most straightforward path to emulating a brain would be to create a compartmental model-level simulation like the one described by Markram (2015), but one based on the electron microscopically-imaged structural connectome. The morphological type of each neuron can easily be determined based on the structural connectome and its compartmental model's ion channel densities would then be filled in based on its morphological type. The main free parameters of such a model are the synaptic strengths which would be estimated based on the sizes and ultrastructural details of the individual synapses in the structural connecotme.

There is good reason to believe that such a procedure would be able to estimate the receptive field properties of individual neurons. For example, long-standing models of visual cortical cells clearly suggest that it is the pattern and strengths of their synaptic connections that define their receptive field properties, not differences in the ion channel densities (e.g. Huble & Wiesel 1962). Existing models suggest that large-scale network-level phenomena like visual object recognition, sensorimotor control, associative memory recall, etc. should be even less sensitive to variations in ion channel densities between neurons of the same morphological class. This is because more global phenomenon like the inhibitory competition among neurons would tend to cancel out these variances. As a concrete example, consider how robust attractor-based models would be to variations in individual neuronal biases (e.g. Rolls and Kesner 2006). The attractor neural networks presumed to underlie much of cortical processing are known for their robustness to noise, their ability to perform pattern completion, and their robustness to damage, all of which would suggest they would also be robust to small inaccuracies in the estimation of ion channel densities between neurons of the same morphological type.

[19] I am assuming here that it will be found that there are some exceptions to the above footnote. I.e. that there will be found some cases in which information crucial to obtaining an accurate simulation is stored not through morphological changes to synapses but through neuron-specific changes to ion channel distributions—information that may not be adequately inferable from morphological correlations like those described above. One possible example of this has already been found in the case of cerebellar Purkinje cells (Johansson et al. 2014). Again, ASC preserves the locations and primary structures of the proteins (ion channels and receptors) that even the most detailed compartmental

The time has come to revive you from your long slumber. The robotic surgeons that will perform this feat begin a complex multistep process that will eventually result in preparing your brain and spinal cord for 3D electron microscopic mapping. Your body is warmed to room temperature and again your vasculature is cannulated and hooked up to a perfusion apparatus. This time the perfusion is used to slowly wash out the cryoprotectant agent over the course of a few hours, bringing you back to the 'freshly' glutaraldehyde fixed state[20].

Then a set of specially designed heavy metal stains are introduced into the perfusate. These stains differentially tag the different ion channel types mentioned previously so that it will be possible to estimate their membrane densities in later electron micrographs[21]. Following this, a different set of stains is introduced to fix and stain membrane lipids and to differentially stain the proteins present at synaptic junctions[22]. At this point all of the key structures and molecules necessary for decoding your memories have been differentially stained in a manner that will make them clearly distinguishable during later electron microscopic imaging.

Now ethanol is introduced into the perfusate and it is ramped over the course of hours to 100% concentration. This is done to extract all of the water from your brain and spinal cord tissue. Then an organic solvent is introduced and slowly ramped to 100% concentration. This is in preparation for the final perfusion which slowly infiltrates every nook and cranny of the brain and spinal cord with a plastic resin. The resin-infiltrated brain and spinal cord is allowed to cure into a solid plastic block over the course of a few days[23].

---

models suggest may be important. So if it is found that some additional information is indeed needed beyond what can be inferred by the structural connectome alone then this does not pose a fundamental objection to ASC preservation, it just means that the revival procedure may require imaging more than the structural connectome.

[20] The washout of cryoprotectant by perfusion was demonstrated by McIntyre & Fahy (2015).

[21] This is the most speculative part of this proposal so far. Currently the tagging of proteins like ion channels is done by immunostaining which is often more difficult to perform in glutaraldehyde fixed tissue (but see Collman et al. 2015 and Murray et al. 2015) and is typically not performed by vascular perfusion. I am speculating that future neuroscientists would be able to develop some sort of tag that would specifically bind to those select proteins that are deemed necessary to annotate the structural connectome and that would differentiate them in subsequent electron micrographs.

[22] The staining described typically involves osmium tetroxide (OsO4), uranyl acetate (UA), and lead citrate (LC). These are typically not perfused but there is precedence in the literature (Palay et al.1962; Bachofen et al 1982).

[23] This is the standard procedure used to prepare brain tissue for electron microscopy (Hayat 2000) except that the vasculature is being perfused with these chemicals. Mikula & Denk (2015) provide evidence that volumes the size of a whole mouse brain can be

Robotic surgeons carefully remove your plastic-embedded central nervous system from its bony, dura mater-wrapped enclosure and mount the plastic block in a special apparatus that will section your brain and spinal cord into 20 micron thick slabs. This apparatus uses a large, ultra-sharp synthetic diamond blade which is heated and lubricated so that it can smoothly cut through the plastic block[24]. Your brain is reduced to a few thousand 20 micron thick sections, each of which is mounted on its own silicon wafer. Your spinal cord is similarly sectioned and mounted.

These wafers are then shipped to a massive imaging facility that resembles a semiconductor fabrication plant. Within this imaging facility the wafers containing your brain's slices are simultaneously imaged across thousands of scanning electron microscopes (SEMs) each utilizing hundreds of electron beams[25]. Looking in on the imaging of one of your brain slices we would see hundreds of electron beams scanning across its surface each creating a 10nm resolution image of the heavy metal stained tissue beneath its beam. Once the surface has been imaged, the entire wafer is robotically transferred to a broad ion milling machine that gently removes the top 10 nm of the tissue surface. This cycle is repeated (image top 10 nm, remove top 10 nm, image top 10 nm, remove top 10 nm, ...) until the entire 20 micron depth has been imaged[26].

After a considerable length of time[27], the robotic imaging facility finally

prepared for electron microscopy but they do not use perfusion to do so. Other papers have explored perfusion dehydration (Oldmixon et al. 1985) and perfusion infiltration with plastic resins (Krucker et al. 2006).

[24] This thick sectioning procedure is based on the one described in (Hayworth et al. 2015).

[25] Two groups have recently demonstrated such multibeam scanning electron microscopes (Eberle et al. 2015; Zuidema et al. 2017). The Zeiss multibeam SEM is already commercially available and has been integrated into the work flow of connectomics imaging (Schalek et al. 2016). It uses 91 electron beams scanning in parallel to dramatically increase overall imaging speed (Kemen et al. 2015). Expansion to at least an order of magnitude more beams appears possible (e.g. Slot et al. 2009) as does mass production.

[26] What is described resembles the well-established focused ion beam-scanning electron microscopy (FIB-SEM) process that is currently used to image structural connectomes (e.g. Xu et al. 2017). But in the above description broad ion beam milling replaces the focused ion beam. FIB-SEM routinely gives less than 10nm isotropic resolution allowing for the automatic tracing of neuronal processes (Plaza et al. 2014) and the automatic identification of synapses (Merchán-Pérez et al. 2009). The author (Hayworth) has demonstrated that FIB-SEM-like datasets can be acquired using broad ion beam milling in small scale laboratory tests (not yet published).

[27] Based on today's multibeam SEM technology alone (Kemen et al. 2015) imaging a single human brain at 16nm isotropic resolution would require several thousand machines operating in parallel for several years. If electron imaging is the eventual technology used then one would expect many more beams per SEM would be used (e.g. Slot et al. 2009) and each of these machines would be assembled using large-scale robotic mass production.

completes imaging the thousands of slices that used to make up your brain and spinal cord. The broad ion milling part of the imaging process has literally vaporized them a layer at a time, and the atoms that once made up your brain have now been carried away by the milling machines' vacuum pumps. But the information that your brain contained still exists, stored on the hard drives of the imaging facility. The images of all of the sections of your brain and spinal cord are now computationally stitched together into a single volume with 10x10x10nm voxel resolution[28].

Now massive computers go to work interpreting this electron microscopic volume. They first map out your entire structural connectome—computationally reconstructing the morphology of every neuron, every axonal and dendritic process, and every synapse in your brain and spinal cord[29]. Then the computers estimate the functional type and strength of every synaptic connection based on measurements of its ultrastructural features[30]. Then this structural connectome is annotated with estimates of the membrane densities of the specially-labeled ion channel types[31]. This molecularly-annotated structural connectome will form the blueprint for emulating your brain in a computer.

A computer emulation is created based on your molecularly-annotated structural connectome. This emulation will not be modeling the brain at the ion channel level, instead it will model only what is needed to capture the computational features of your mind. After decades of neuroscience research, and after extensive experiments on the first humans that volunteered to be uploaded, it is known precisely how to interpret this molecularly-annotated structural connectome, and it is known precisely what level of abstraction is necessary for emulation. In general the emulation will model each neuron in your original brain as an electrical compartmental model, with the detailed properties of individual ion channels, receptors, protein transcription etc. all subsumed by a simplifying set of approximating equations[32].

Prior to 'starting up' the emulation, a specially designed set of algorithms

[28] Stitching of such separately-imaged thick sections is described in (Hayworth et al. 2015).

[29] Deep neural network-based tracing algorithms like the one described by Januszewski et al. (2017) are demonstrating that fully automated reconstruction should eventually become achievable.

[30] The correlation between the size of a synapse and its functional strength (e.g. the number of expressed AMPA receptors) has been extensively researched, for example: Matsuzaki et al. 2001; Kasai et al. 2003; Bourne & Harris 2007; Bartol et al. 2015; Hayashi-Takagi et al. 2015.

[31] For an example of what such annotation of the EM connectome might look like see Collman et al. (2014)

[32] For an example of what this level of simulation might look like see Markram et al. (2015).

is used to decode your brain's function at a computational level. This involves estimating the receptive field properties of every neuron in spinal and subcortical sensorimotor circuits and in all cortical sensory hierarchies. This mapping is needed to allow the severed nerves going into and out of your brain and spinal cord to be fitted properly to your robotic body's sensory inputs and motor outputs. This will reduce the time you need to spend in rehabilitation learning to control your new robotic body.

Similar algorithms are used to decode your brain's higher mental functions at the symbolic computational level. Each possible attractor state in your cortex's many specialized regions is mapped and assigned a symbolic label[33]. This creates an approximate map of the mental vocabulary you use to distinguish colors, shapes, faces, places, patterns of motion, sounds, words, emotions, individual persons, etc. This crude symbolic mapping can be very useful for 'debugging' your emulation once it is started up—it will allow the specialists overseeing your synthetic revival to literally 'read your mind' in real time, allowing them to quickly adjust key simulation parameters if needed.

This procedure also allows them to coarsely decode your life's memories without the need to bring you back to consciousness. This can be done by mapping out all attractor states in the long-term memory circuits of your medial temporal lobe. Each of these temporal lobe attractor states is a long-term memory that is associated with a particular state of cortical activation. By decoding the cortical state associated with each of these long-term memory attractors you can roughly decode, at the symbolic level, all of the episodic memories of a person's life[34]. These symbolically-decoded memories can be used to create non-conscious ancestor simulations—avatars that descendants and historians can interact with but which do not support conscious thought or goal-directed action. Over the intervening decades many people have been preserved by ASC specifically with this application in mind and have put clauses in their preservation contracts specifically forbidding the creation of a conscious

[33] Many modern neuroscience experiments record the activity of collections of neurons (using optical imaging or multichannel electrode recording) and decode these patterns in order to get a high-level 'symbolic' description of what is being represented. 'Symbolic' in the sense that we external observers can successfully interpret them as representing specific aspects of the external world, or of the animal's internal state, that drive behavior. For example, Pfeiffer & Foster (2013) recorded the activity of hundreds of hippocampal 'place cells' and found that they could decode their sequences of firing as representing spatial trajectories that predicted the immediate future behavior of the animal. In a similar way, Chang & Tsao (2017) were able to decode the activity in face patches of primate visual cortex in order to precisely understand its representational vocabulary.

[34] This type of memory decoding should, in principle, be possible if the current attractor-based models of cortical-medial temporal lobe interactions are roughly correct (e.g. Rolls & Kesner 2006; Lisman 2015).

emulation based on their brain. But you have chosen to go through the entire process including the revival of consciousness.

Revival

Finally the time comes for your computational revival. The emulation of your brain is put into a state approximating waking up after a long sleep. As you return to consciousness the specialists carefully monitor your mental states looking for any problems: "Cortical activity is not increasing rapidly enough, adjusting reticular parameters... An epileptic seizure is developing in the left temporal lobe, adjusting local inhibition parameters... The striatal activity is 20% below expected, adjusting dopamine circuits..." Soon you are awake again looking out at the world through robotic eyes. You immediately try to recall who you are but are unable to retrieve any of your past episodic memories[35]. You start to panic but the specialists quickly see the problem and perform more adjustments to your cortical circuit parameters. After this you begin to remember: "I was diagnosed with Alzheimer's and opted for a crazy-sounding euthanasia procedure in 2032... Am I an upload?!?...I can recall my name, my spouse and children, my childhood... I can see... and hear... and feel... and I can move my... my... robotic arms!"

The specialist overseeing your revival brings you up to speed. She tells you that the year is 2098 and that you have indeed been uploaded. Your brain is being emulated on a million-node cluster computer tucked inside your robotic skull. Your temporary robotic body is a basic class-3000 model that you will be able to customize in form and function later. She tells you that for the next few weeks you will be undergoing mental and physical rehabilitation designed to get you as close as possible back to your 'pre-upload baseline'. And she tells you that she will not be able to answer many of your questions regarding what life is like in this future world until after your initial rehabilitation—it would be too confusing and might potentially disturb the rehabilitation process.

It will take several weeks of painstaking rehabilitation until you are able to master the control of your robotic body, and until they have tweaked your brain's emulation parameters sufficiently that you 'feel like your old self'. Prior to this tweaking some experiences just seem wrong. The color of the roses in your room don't quite match your memories of roses, so, with a little tweaking, now they do. The taste of vanilla ice cream is a bit sour, but with a little tweaking it now tastes just like you remember. The feel of silk running over your robotic finger seems a bit off, but after a bit of parameter tweaking it now

[35] E.g. Wilson & Wearing 1995

feels 'smooth as silk'. "Yes that pin prick hurts just like I remember it used to, but the cold of that ice cube in my hand feels a bit off... that's better now."[36]

A series of virtual simulations allow you to safely regain your previous sensorimotor skills. It takes a few awkward hours to get comfortable walking and balancing again, several days to regain sufficient hand-eye coordination to play a game of basketball. But with each parameter tweak you feel more comfortable with your virtual body as its control signals are adjusted to match your brain and spinal cord's sensorimotor memories. And all of these adjustments learned in virtual reality work just as well when you are again transferred back into your robotic body.

They test your emotional responses by first showing you short movies and then by putting you in some mock situations in virtual reality. Nervousness, sexual attraction, fear, joy, boredom, love, humor, everything is put through its paces and crosschecked with your long-term memories of like experiences. They even monitor your dreams while you sleep. The specialists overseeing your rehabilitation are meticulous. They can even tell when you are lying about something feeling different, after all they have direct access to your mental state at the symbolic level. They keep explaining to you that the goal of this rehabilitation is to get you back to your 'baseline self', after that you can decide on your own what you want modified.

Three weeks of rehabilitation have passed and you now feel just like your old self. You have even upgraded to a robotic body specially designed to resemble your original biological body, but you have opted for one approximating you at 25 years of age and with a considerably better physique than you had even at that time in your life. You are now ready to be reintroduced to friends and family members that are still alive or that have been uploaded like you.

The reunion is held in virtual reality. You are pleased to see many of the same faces that were at your going away party, unfortunately your spouse is not among them. It was previously explained to you that she is still in ASC storage and is due to be revived next year. You try to pry information out your friends regarding what life is like in 2098, but they have also been told to avoid that discussion. They all just say that it is awesome beyond your imagination and that you should get ready for a really wild ride once you 'graduate' from rehabilitation. It also becomes clear that this reunion is part of that graduation as your former friends and loved ones get a chance to evaluate how faithful this emulation is with respect to their own memories of you.

After your successful 'graduation' from rehabilitation you are enrolled in a

[36] This passage is designed to point out how one's own internal memories form a highly-redundant set of checks that can be used to interactively fine-tune emulation parameters.

set of courses designed to bring you, and a dozen fellow ASC uploads from the 2030's, fully up to speed on life in 2098. Science, history, technology, psychology, philosophy, economics, culture—it is like being a wide-eyed child again learning about the world for the first time! Thankfully they have begun to allow you to upgrade your intelligence, slowly, which makes everything much easier. After a few months you have mastered subjects that seemed impossible to your old biological self: "General relativity and quantum mechanics are a breeze! Really so simple now that I think of it."

But this boost in I.Q. and flood of new knowledge is just the start. Everyone in your cohort is excited to finally be enrolling in the most interesting class in this 'reintegration' school—Varieties of mental experience available to the uploaded mind. Your class starts and the teacher begins: "Class, I will now unlock your emulation's base-level safety protocols so that you can adjust your own state of mind. You will see a list of built-in settings. Let's start with the first one labeled 'Nirvana'..."

[End of story]

A.3 What this fictional story is designed to address

One question the above fictional story was designed to answer is: "Why would a terminally ill patient desire the option to choose ASC as their method of euthanasia?" Hopefully the answer is now crystal clear.

First, why would someone choose euthanasia? Many of us have watched loved ones suffer through the excruciating final stages of a fatal disease, or through the decrepitude of extremely old age, or through soul-eating dementia, and have decided that if we ourselves are one day facing a similar fate that we would like the option to choose an early exit, dying with dignity on our own terms.

Second, why would someone choose ASC as their method of doctor assisted euthanasia? Most people that choose euthanasia do not do so because they are sick of living, they do so because they are sick of suffering and today's medical science simply offers them no long-term hope. If ASC is developed into a reliable medical procedure and offered as an option in hospitals then it will represent hope to these suffering patients. As depicted in the story above, ASC offers not only the possibility that their suffering

will be halted, but that their health and youth will be restored as well, and that they will wake up in a future significantly more advanced than today. To experience the far future firsthand is perhaps the greatest adventure one can imagine, and there is no question that many adventurous people will line up for even a slim chance to do so, especially terminally ill patients whose only alternative is oblivion. So it is likely that many terminally ill patients would choose euthanasia by ASC if it was available in hospitals. The real question to be debated is whether such an option should be made available.

A.4 Separating facts from personal opinions

Of course there are many reasons why an individual might reject ASC for themselves even in the face of a terminal illness. Perhaps the most common reason would be that the individual's religious beliefs provide them a different sort of hope for revival. Another reason might be philosophical in nature, e.g. they believe that revival through a synthetic copy is not 'real' survival even if that copy retains their memories and personality. Or they might dismiss the technological possibility of revival by mind uploading or by any other means. Or they might think the chances of future revival are so vanishingly small as to not be worth the trouble or expense. They might reject ASC for themselves because they feel it will require too much of an adjustment to get used to living in so different a world. They might reject ASC because they are afraid that life in that future world might be unsatisfying or perhaps even unbearable. And they might reject ASC because of their sociological views, or because they feel they have lived long enough and it is proper to let nature run its regular course.

All of these are perfectly proper reasons for an individual to reject ASC preservation for themselves, but none represents a good reason to withhold the option of ASC from someone else who truly desires it. These are personal opinions not facts. You personally might not consider it survival to have an emulation based on your brain's connectome awaken in the future, but many other people would consider it a form of survival, and a highly desirable one at that. Who is right? There is likely no definitive answer to this question because there is no agreed upon definition of self-identity. You may have strong opinions on this matter, but please do not misinterpret your personal opinions as facts that give you the right to withhold ASC from a terminal patient who considers it their only chance at survival.

In my assessment, the option of choosing ASC should be withheld only if the available science does not support the possibility of future revival.

Serious debate should be focused squarely on this question and not be clouded by personal opinions.

A.5 The core of the scientific argument: I am my connectome and ASC preserves the connectome

Discussions on the future possibility of mind uploading are often prematurely terminated when one party proclaims: "We understand almost nothing about how the brain works, therefore it is impossible to speculate on what it would take to upload a mind, or even whether it is possible in principle." This is the key objection I have received to this proposal from colleagues in the scientific and medical fields and the burden is on me to clearly address it here. In short my answer is: "We understand enough to know what needs to be preserved."

Fifty years ago dismissing medical brain preservation on the grounds of insufficient knowledge might have been prudent, but the cognitive and neurosciences have progressed enormously in recent decades. Real progress has been made at all levels and general principles have come into focus. Most importantly, it is now a bedrock assumption in the field, supported by a wide range of experimental evidence, that the connectome is the brain's fundamental computational and memorial substrate. The following quotes from experts across the neuroscience field testify to this:

> "One of the chief ideas we shall develop in this book is that the specificity of the synaptic connections established during development underlie perception, action, emotion, and learning."
> — Principles of Neural Science Textbook (Kandel et al. 2000)

> "[E]verything you know is encoded in the patterns of your synaptic weights..."
> — Computational Cognitive Neuroscience Textbook (O'Reilly et al. 2012)

> Memories are thought to be encoded as enduring physical changes in the brain, or engrams. Most neuroscientists agree that the formation of an engram involves strengthening of synaptic connections between populations of neurons
> — Finding the Engram Review Article (Josselyn et al. 2015)

> "There is now general consensus that persistent modification of the synaptic strength via LTP and LTD of pre-existing connections represents a primary mechanism for the formation of memory engrams." — What is memory? The present state of the engram Review Article (Poo et al. 2016)

> "[The] predicate... of all modern neuroscience is that cognitively important functions can be explained as an emergent property of neurons and their network connections... Perhaps 20 years ago, one could have argued that the emergence of cognitive function from interconnected neurons was deeply mysterious. That does not seem true today. What has changed is that we now have a feel for how networks can produce cognitively relevant computations" — The Challenge of Understanding the Brain: Where We Stand in 2015 Review Article (Lisman 2015)

> "I am my connectome." — Connectome: how the brain's wiring makes us who we are (Seung 2012)

A.6 How 'you' are encoded in your connectome

Explaining how the neuroscience community arrived at this tentative consensus would require a much, much longer paper, one that would need to review a large fraction of modern neuro- and cognitive science. But let's try to at least outline such a paper and provide some references:

First we would need to review the different memory systems of the brain. Squire (2004) offers a taxonomic review of these memory systems most of which I very briefly summarize here:

Hippocampus Neural circuits within the hippocampus and other medial temporal lobe structures support the initial learning of what is colloquially referred to as memory: specifically declarative or episodic memories (Squire, Stark & Clark 2004). This memory system specializes in rapid 'one-shot' learning with little generalization in order to provide maximal discriminability among distinct episodes (Atallah, Frank & O'Reilly 2004).

Striatum Circuits within the striatum support the initial phase of procedural learning (Ashby, Ennis & Spiering 2007)—the learning of sequences of motor or cognitive actions (Aldridge et al. 1993). Learning

in the striatum is modulated by dopaminergic inputs from the brain's reward system, which in turn is modulated by the striatum itself. This arrangement creates a joint system optimized for reinforcement and temporal difference learning (O'Reilly et al. 2007).

Cortex The knowledge initially learned within both of the above systems is, over time, consolidated in the cortex which is specialized for generalization (Pasupathy & Miller 2005; Kitamura et al. 2017). For example, learning within cortical sensory hierarchies (visual, auditory, etc.) can be thought of as creating 'perceptual memories'. Repeated exposures to sensory stimuli train these cortical hierarchies to categorize the raw sensory signals along a myriad of different perceptual dimensions (e.g. shape, size, orientation, movement, color, texture, etc.) (Kanwisher 2010; DiCarlo, Zoccolan & Rust 2012).

Amygdala Circuits within the amygdala support emotional memories–learning that associates high-level cortical states with more primary motivational inputs (Janak & Tye 2015).

Next we would need to review cognitive architecture models (e.g. Anderson et al. 2008; O'Reilly, Hazy & Herd 2012; Eliasmith et al. 2012) that show how these different memory systems can interact to create the amazing flexibility of human cognition. Perhaps the best example of such an overall cognitive architecture today is the ACT-R model which is summarized excellently in the book 'How can the human mind occur in the physical universe' (Anderson 2009).

At this point we would have a pretty good top-level overview of how the cognitive science and systems neuroscience communities view the mind-brain relationship. But it would be appropriate to also discuss how cognitive models of consciousness and self-identity are mappable onto such cognitive architectures. Some appropriate references for that might be: Dehaene & Naccache 2001; Anderson 2009; Metzinger 2004; Dennett 1991.

Next we would need to review the existing neuroscience models of each of the memory system above, in order to understand how memory is physically encoded in each:

Hippocampus Relies on Long Term Potentiation/Depression (LTP/LTD) in glutamatergic synapses onto the dendritic spines of hippocampal dentate, CA1, and CA3 cells (Lisman 2015; Rolls & Kesner 2006).

Striatum Relies on dopamine modulated LTP/LTD in glutamatergic synapses onto the dendritic spines of striatal medium spiny neurons (Kreitzer & Malenka 2008; Yagishita et al. 2014).

Cortex Relies on LTP/LTD in glutamatergic synapses onto the dendritic spines of cortical pyramidal cells (Holtmaat & Svoboda 2009; Matsuzaki et al. 2004).

Amygdala Relies on LTP/LTD in glutamatergic synapses onto the dendritic spines of lateral amygdala pyramidal cells (Maren 2005; Johansen et al. 2010).

There is a clear pattern here. The best neuroscience models of all of these different memory systems propose that their disparate types of memories are encoded through the same process of LTP/LTD at a particular class of synapses, specifically glutamatergic synapses onto dendritic spines. Of course there are hundreds of important details that are being glossed over in this brief synopsis, but those details do not alter the general consensus that is being conveyed—there seems to be a common mechanism for long-term memory storage in the brain that involves structural changes to synapses. This common mechanism is what the above quotes are referring to, and there is now a considerable literature reviewing the experimental evidence underlying this (e.g. Kasai et al. 2003; Hoshiba et al. 2017; Bailey et al. 2015; Josselyn et al. 2015; Poo et al. 2016; Lisman 2015; Bourne & Harris 2007; Yuste 2010; Segal 2016; Maren 2005; Lamprecht & LeDoux 2004; Tonegawa et al. 2015).

This is the fundamental body of evidence that supports the conclusion that 'I am my Connectome': Hundreds of painstaking neuroscience experiments that have uncovered the synaptic basis of memory across all of the different memory system that make up our brain's cognitive architecture.

A.7 Conclusion

ASC demonstrably (McIntyre & Fahy 2015) preserves the patterns of synaptic connections that these quotes and references suggest store the majority of the brain's learned knowledge. But we must not forget that ASC, because it is based on glutaraldehyde fixation, preserves far more than simply the structural connectome. Glutaraldehyde fixation preserves the locations and identities of a wide range of biomolecules important to neuronal function. Given all the correlations and redundancies present in the brain, it seems clear, at least to me, that ASC is almost certainly preserving the vast majority of the information content in the brain that makes each person unique.

As stated earlier, the option of choosing ASC should be withheld from terminal patients who desire it only if the available science does not sup-

port the possibility of future revival. The above references seem, to me, to strongly suggest that ASC does preserve the information content of the brain, and therefore it should support at least the possibility of future revival.

As I see it, the next steps are clear: The neuroscience and medical communities should begin an open debate regarding ASC's ability to preserve the information content of the brain. If an argument can be made that ASC does not preserve crucial information stored in the brain, information that cannot be inferred from the many ultrastructural and molecular details that ASC does preserve, then that argument should be brought forward now. If such an argument is not forthcoming, then the scientific and medical communities should immediately start developing ASC into a reliable, regulated medical procedure that can be offered to terminal patients.

A.8 References

Aghayev, E., Thali, M. J., Sonnenschein, M., Jackowski, C., Dirnhofer, R., & Vock, P. (2007). Post-mortem tissue sampling using computed tomography guidance. Forensic science international, 166(2), 199-203.

Aldridge, J. W., Berridge, K. C., Herman, M., & Zimmer, L. (1993). Neuronal coding of serial order: syntax of grooming in the neostriatum. Psychological Science, 4(6), 391-395.

Anderson, J. R., Fincham, J. M., Qin, Y., & Stocco, A. (2008). A central circuit of the mind. Trends in cognitive sciences, 12(4), 136-143.

Anderson, J. R. (2009). How can the human mind occur in the physical universe?. Oxford University Press.

Ashby, F. G., Ennis, J. M., & Spiering, B. J. (2007). A neurobiological theory of automaticity in perceptual categorization. Psychological review, 114(3), 632.

Atallah, H. E., Frank, M. J., & O'reilly, R. C. (2004). Hippocampus, cortex, and basal ganglia: Insights from computational models of complementary learning systems. Neurobiology of learning and memory, 82(3), 253-267.

Bachofen, H., Ammann, A., Wangensteen, D., & Weibel, E. R. (1982). Perfusion fixation of lungs for structure-function analysis: credits and limitations. Journal of Applied Physiology, 53(2), 528-533.

Bailey, C. H., Kandel, E. R., & Harris, K. M. (2015). Structural components of synaptic plasticity and memory consolidation. Cold Spring Harbor Perspectives in Biology, 7(7), a021758.

Bartol Jr, T. M., Bromer, C., Kinney, J., Chirillo, M. A., Bourne, J. N., Harris, K. M., & Sejnowski, T. J. (2015). Nanoconnectomic upper bound on the variability of synaptic plasticity. Elife, 4, e10778.

Bell, M. E., Bourne, J. N., Chirillo, M. A., Mendenhall, J. M., Kuwajima, M., & Harris, K. M. (2014). Dynamics of nascent and active zone ultrastructure as synapses enlarge during long-term potentiation in mature hippocampus. Journal of Comparative Neurology, 522(17), 3861-3884.

Bourne, J., & Harris, K. M. (2007). Do thin spines learn to be mushroom spines that remember?. Current opinion in neurobiology, 17(3), 381-386.

Bourne, J. N., & Harris, K. M. (2011). Coordination of size and number of excitatory and inhibitory synapses results in a balanced structural plasticity along mature hippocampal CA1 dendrites during LTP. Hippocampus, 21(4), 354-373.

Briggman, K. L., Helmstaedter, M., & Denk, W. (2011). Wiring specificity in the direction-selectivity circuit of the retina. Nature, 471(7337), 183.

Carrillo-Reid, L., Yang, W., Bando, Y., Peterka, D. S., & Yuste, R. (2016). Imprinting and recalling cortical ensembles. Science, 353(6300), 691-694.

Chan, D., Janssen, J. C., Whitwell, J. L., Watt, H. C., Jenkins, R., Frost, C., ... & Fox, N. C. (2003). Change in rates of cerebral atrophy over time in early-onset Alzheimer's disease: longitudinal MRI study. The Lancet, 362(9390), 1121-1122.

Chang, L., & Tsao, D. Y. (2017). The Code for Facial Identity in the Primate Brain. Cell, 169(6), 1013-1028.

Collman, F., Buchanan, J., Phend, K. D., Micheva, K. D., Weinberg, R. J., & Smith, S. J. (2015). Mapping synapses by conjugate light-electron array tomography. Journal of Neuroscience, 35(14), 5792-5807.

Cerullo, M. A. (2015). Uploading and branching identity. Minds and Machines, 25(1), 17-36.

David, O. E., Netanyahu, N. S., & Wolf, L. (2016). DeepChess: End-to-End Deep Neural Network for Automatic Learning in Chess. In International Conference on Artificial Neural Networks (pp. 88-96). Springer International Publishing.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. Cognition, 79(1), 1-37.

Dennett, D. C. (1991). Consciousness Explained. Boston (Little, Brown and Co) 1991.

DiCarlo, J. J., Zoccolan, D., & Rust, N. C. (2012). How does the brain solve visual object recognition?. Neuron, 73(3), 415-434.

Eberle, A. L., Mikula, S., Schalek, R., Lichtman, J., Tate, M. K., & Zeidler, D. (2015). High-resolution, high-throughput imaging with a multibeam scanning electron microscope. Journal of microscopy, 259(2), 114-120.

Eliasmith, C., Stewart, T. C., Choo, X., Bekolay, T., DeWolf, T., Tang, Y., & Rasmussen, D. (2012). A large-scale model of the functioning brain. Science, 338(6111), 1202-1205.

Fahy, G. M., Wowk, B., Wu, J., Phan, J., Rasch, C., Chang, A., & Zendejas, E. (2004). Cryopreservation of organs by vitrification: perspectives and recent advances. Cryobiology, 48(2), 157-178.

Hassabis, D., Kumaran, D., Summerfield, C., & Botvinick, M. (2017). Neuroscience-inspired artificial intelligence. Neuron, 95(2), 245-258.

Hayashi-Takagi, A., Yagishita, S., Nakamura, M., Shirai, F., Wu, Y. I., Loshbaugh, A. L., ... & Kasai, H. (2015). Labelling and optical erasure of synaptic memory traces in the motor cortex. Nature, 525(7569), 333.

Hayat, M. A. (1986). Glutaraldehyde: role in electron microscopy. Micron and Microscopica Acta, 17(2), 115-135.

Hayat, M. A. (2000). Principles and techniques of electron microscopy. Biological applications.

Hayworth, K. (2010). Killed by bad philosophy. www.brainpreservation.org/wp-content/uploads/2015/08/killed_by_bad_philosophy.pdf

Hayworth, K. J., Xu, C. S., Lu, Z., Knott, G. W., Fetter, R. D., Tapia, J. C., ... & Hess, H. F. (2015). Ultrastructurally smooth thick partitioning and volume stitching for large-scale connectomics. Nature methods, 12(4), 319-322.

He, K., Zhang, X., Ren, S., & Sun, J. (2015). Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In Proceedings of the IEEE international conference on computer vision (pp. 1026-1034).

Holtmaat, A., & Svoboda, K. (2009). Experience-dependent structural synaptic plasticity in the mammalian brain. Nature reviews. Neuroscience, 10(9), 647.

Hoshiba, Y., Wada, T., & Hayashi-Takagi, A. (2017). Synaptic Ensemble Underlying the Selection and Consolidation of Neuronal Circuits during Learning. Frontiers in Neural Circuits, 11.

Hua, Y., Laserstein, P., & Helmstaedter, M. (2015). Large-volume en-bloc staining for electron microscopy-based connectomics. Nature communications, 6.

Janak, P. H., & Tye, K. M. (2015). From circuits to behaviour in the amygdala. Nature, 517(7534), 284.

Januszewski, M., Kornfeld, J., Li, P. H., Pope, A., Blakely, T., Lindsey, L., ... & Jain, V. (2017). High-Precision Automated Reconstruction of Neurons with Flood-filling Networks. bioRxiv, 200675.

Johansen, J. P., Hamanaka, H., Monfils, M. H., Behnia, R., Deisseroth, K., Blair, H. T., & LeDoux, J. E. (2010). Optical activation of lateral amygdala pyramidal cells instructs associative fear learning. Proceedings of the National Academy of Sciences, 107(28), 12692-12697.

Johansson, F., Jirenhed, D. A., Rasmussen, A., Zucca, R., & Hesslow, G. (2014). Memory trace and timing mechanism localized to cerebellar Purkinje cells. Proceedings of the National Academy of Sciences, 111(41), 14930-14934.

Josselyn, S. A., Köhler, S., & Frankland, P. W. (2015). Finding the engram. Nature Reviews Neuroscience, 16(9), 521-534.

Kandel, E., Schwartz, J., & Jessell, T. (2000). Principles of Neural Science.

Kanwisher, N. (2010). Functional specificity in the human brain: a window into the functional architecture of the mind. Proceedings of the National Academy of Sciences, 107(25), 11163-11170.

Kasai, H., Matsuzaki, M., Noguchi, J., Yasumatsu, N., & Nakahara, H. (2003). Structure–stability–function relationships of dendritic spines. Trends in neurosciences, 26(7), 360-368.

Kasthuri, N., Hayworth, K. J., Berger, D. R., Schalek, R. L., Conchello, J. A., Knowles-Barley, S., ... & Roberts, M. (2015). Saturated reconstruction of a volume of neocortex. Cell, 162(3), 648-661.

Kemen, T., Garbowski, T., & Zeidler, D. (2015, July). Multi-beam SEM technology for ultra-high throughput. In Photomask Japan 2015 (pp. 965807-965807). International Society for Optics and Photonics.

Kitamura, T., Ogawa, S. K., Roy, D. S., Okuyama, T., Morrissey, M. D., Smith, L. M., ... & Tonegawa, S. (2017). Engrams and circuits crucial for systems consolidation of a memory. Science, 356(6333), 73-78.

Knott, G., Marchman, H., Wall, D., & Lich, B. (2008). Serial section scanning electron microscopy of adult brain tissue using focused ion beam milling. Journal of Neuroscience, 28(12), 2959-2964.

Kreitzer, A. C., & Malenka, R. C. (2008). Striatal plasticity and basal ganglia circuit function. Neuron, 60(4), 543-554.

Krucker, T., Lang, A., & Meyer, E. P. (2006). New polyurethane-based material for vascular corrosion casting with improved physical and imaging characteristics. Microscopy research and technique, 69(2), 138-147.

Lamprecht, R., & LeDoux, J. (2004). Structural plasticity and memory. Nature reviews. Neuroscience, 5(1), 45.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436-444.

Lee, W. C. A., Bonin, V., Reed, M., Graham, B. J., Hood, G., Glattfelder, K., & Reid, R. C. (2016). Anatomy and function of an excitatory network in the visual cortex. Nature, 532(7599), 370

Lemler, J., Harris, S. B., Platt, C., & Huffman, T. M. (2004). The arrest of biological time as a bridge to engineered negligible senescence. Annals of the New York Academy of Sciences, 1019(1), 559-563.

Lisman, J. (2015). The challenge of understanding the brain: where we stand in 2015. Neuron, 86(4), 864-882.

Liu, X., Ramirez, S., Pang, P. T., Puryear, C. B., Govindarajan, A., Deisseroth, K., & Tonegawa, S. (2012). Optogenetic stimulation of a hippocampal engram activates fear memory recall. Nature, 484(7394), 381-385.

Liu, X., Ramirez, S., & Tonegawa, S. (2014). Inception of a false memory by optogenetic manipulation of a hippocampal memory engram. Phil. Trans. R. Soc. B, 369(1633), 20130142.

Lyketsos, C. G., Carrillo, M. C., Ryan, J. M., Khachaturian, A. S., Trzepacz, P., Amatniek, J., ... & Miller, D. S. (2011). Neuropsychiatric symptoms in Alzheimer's disease. Alzheimer's & Dementia, 7(5), 532-539

Maren, S. (2005). Synaptic mechanisms of associative memory in the amygdala. Neuron, 47(6), 783-786

Markram, H., Muller, E., Ramaswamy, S., Reimann, M. W., Abdellah, M., Sanchez, C. A., ... & Kahou, G. A. A. (2015). Reconstruction and simulation of neocortical microcircuitry. Cell, 163(2), 456-492.

Matsuzaki, M., Ellis-Davies, G. C., Nemoto, T., Miyashita, Y., Iino, M., & Kasai, H. (2001). Dendritic spine geometry is critical for AMPA receptor expression in hippocampal CA1 pyramidal neurons. Nature neuroscience, 4(11), 1086.

Matsuzaki, M., Honkura, N., Ellis-Davies, G. C., & Kasai, H. (2004). Structural basis of long-term potentiation in single dendritic spines. Nature, 429(6993), 761.

McIntyre, R. L., & Fahy, G. M. (2015). Aldehyde-stabilized cryopreservation. Cryobiology, 71(3), 448-458.

Merchan-Perez, A., Rodriguez, J. R., Alonso-Nanclares, L., Schertel, A., & DeFelipe, J. (2009). Counting synapses using FIB/SEM microscopy: a true revolution for ultrastructural volume reconstruction. Frontiers in neuroanatomy, 3.

Metzinger, T. (2004). Being no one: The self-model theory of subjectivity. MIT Press.

Migneault, I., Dartiguenave, C., Bertrand, M. J., & Waldron, K. C. (2004). Glutaraldehyde: behavior in aqueous solution, reaction with proteins, and application to enzyme crosslinking. Biotechniques, 37(5), 790-806.

Mikula, S., & Denk, W. (2015). High-resolution whole-brain staining for electron microscopic circuit reconstruction. Nature methods, 12(6), 541-546.

Murray, E., Cho, J. H., Goodwin, D., Ku, T., Swaney, J., Kim, S. Y., ... & McCue, M. (2015). Simple, scalable proteomic imaging for high-dimensional profiling of intact systems. Cell, 163(6), 1500-1514.

Noguchi, J., Nagaoka, A., Watanabe, S., Ellis-Davies, G. C., Kitamura, K., Kano, M., ... & Kasai, H. (2011). In vivo two-photon uncaging of glutamate revealing the structure–function relationships of dendritic spines in the neocortex of adult mice. The Journal of physiology, 589(10), 2447-2457.

Oldmixon, E. H., Suzuki, S., Butler, J. P., & Hoppin Jr, F. G. (1985). Perfusion dehydration fixes elastin and preserves lung air-space dimensions. Journal of Applied Physiology, 58(1), 105-113.

O'Reilly, R. C., Frank, M. J., Hazy, T. E., & Watz, B. (2007). PVLV: the primary value and learned value Pavlovian learning algorithm. Behavioral neuroscience, 121(1), 31.

O'Reilly, R. C., Munakata, Y., Frank, M. J., & Hazy, T. E. (2012). Computational cognitive neuroscience. PediaPress.

O'Reilly, R. C., Hazy, T. E., & Herd, S. A. (2012). The Leabra Cognitive Architecture: How to Play 20 Principles with Nature. The Oxford Handbook of Cognitive Science, 91.

Palay, S. L., McGee-Russell, S. M., Gordon, S., & Grillo, M. A. (1962). Fixation of neural tissues for electron microscopy by perfusion with solutions of osmium tetroxide. The Journal of cell biology, 12(2), 385-410.

Pasupathy, A., & Miller, E. K. (2005). Different time courses of learning-related activity in the prefrontal cortex and striatum. Nature, 433(7028), 873.

Pfeiffer, B. E., & Foster, D. J. (2015). Autoassociative dynamics in the generation of sequences of hippocampal place cells. Science, 349(6244), 180-183.

Plaza, S. M., Scheffer, L. K., & Chklovskii, D. B. (2014). Toward large-scale connectome reconstructions. Current opinion in neurobiology, 25, 201-210.

Poo, M. M., Pignatelli, M., Ryan, T. J., Tonegawa, S., Bonhoeffer, T., Martin, K. C., ... & Mullins, C. (2016). What is memory? The present state of the engram. BMC biology, 14(1), 40.

Rolls, E. T., & Kesner, R. P. (2006). A computational theory of hippocampal function, and empirical tests of the theory. Progress in neurobiology, 79(1), 1-48.

Rosenberg, P. B., Nowrangi, M. A., & Lyketsos, C. G. (2015). Neuropsychiatric symptoms in Alzheimer's disease: what might be associated brain circuits?. Molecular aspects of medicine, 43, 25-37.

Ryan, T. J., Roy, D. S., Pignatelli, M., Arons, A., & Tonegawa, S. (2015). Engram cells retain memory under retrograde amnesia. Science, 348(6238), 1007-1013.

Sandberg, A., & Bostrom, N. (2008). Whole brain emulation.

Schalek, R., Lee, D., Kasthuri, N., Peleg, A., Jones, T., Kaynig, V., ... & Lichtman, J. W. (2016). Imaging a 1 mm3 volume of rat cortex using a MultiBeam SEM. Microscopy and Microanalysis, 22(S3), 582-583.

Segal, M. (2017). Dendritic spines: morphological building blocks of memory. Neurobiology of learning and memory, 138, 3-9.

Seung, S. (2012). Connectome: how the brain's wiring makes us who we are. Houghton Mifflin Harcourt.

Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., ... & Dieleman, S. (2016). Mastering the game of Go with deep neural networks and tree search. Nature, 529(7587), 484-489.

Slot, E., Wieland, M. J., De Boer, G., Kruit, P., Ten Berge, G. F., Houkes, A. M. C., ... & Teepen, T. F. (2008, March). MAPPER: high throughput maskless lithography. In Proc. SPIE (Vol. 6921, No. 1, p. 69211P).

Squire, L. R. (2004). Memory systems of the brain: a brief history and current perspective. Neurobiology of learning and memory, 82(3), 171-177.

Squire, L. R., Stark, C. E., & Clark, R. E. (2004). The medial temporal lobe. Annu. Rev. Neurosci., 27, 279-306.

Taigman, Y., Yang, M., Ranzato, M. A., & Wolf, L. (2014). Deepface: Closing the gap to human-level performance in face verification. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 1701-1708).

Tonegawa, S., Liu, X., Ramirez, S., & Redondo, R. (2015). Memory engram cells have come of age. Neuron, 87(5), 918-931.

Trachtenberg, J. T., Chen, B. E., Knott, G. W., & Feng, G. (2002). Long-term in vivo imaging of experience-dependent synaptic plasticity in adult cortex. Nature, 420(6917), 788.

Walker, M. (2011). Personal identity and uploading. Journal of Evolution and Technology, 22(1), 37-52.

Wiley, K. B., & Koene, R. A. (2015). The fallacy of favoring gradual replacement mind uploading over scan-and-copy. arXiv preprint arXiv:1504.06320.

Wilson, B. A., & Wearing, D. (1995). Prisoner of consciousness: A state of just awakening following herpes simplex encephalitis.

Xu, C. S., Hayworth, K. J., Lu, Z., Grob, P., Hassan, A. M., García-Cerdán, J. G., ... & Hess, H. F. (2017). Enhanced FIB-SEM systems for large-volume 3D imaging. Elife, 6.

Yagishita, S., Hayashi-Takagi, A., Ellis-Davies, G. C., Urakubo, H., Ishii, S., & Kasai, H. (2014). A critical time window for dopamine actions on the structural plasticity of dendritic spines. Science, 345(6204), 1616-1620.

Yuste, R. (2010). Dendritic spines. MIT press.

Zuidema, W., Rahangdale, S., Keizer, P., Hoogenboom, J. P., Kruit, P., Wolters, A., & Giepmans, B. N. G. (2017). 10kfps Transmission Imaging in a 196 Beam SEM. Microscopy and Microanalysis, 23(S1), 586-587.

B Overcoming Objections To Brain Preservation

This essay was first published on the Brain Preservation Foundation's Website[1].

In just the last ten years, neuroscientists have been using powerful new tools in genetics, molecular biology, biotechnology, microscopy, systems biology, data science, and other fields to finally uncover the set of epigenetic, extracellular, and synaptic changes that form the molecular basis of memory. The 2014 Kavli Prize in Neuroscience was won by three neuroscientists for the discovery of specialized brain networks for memory and cognition. The 2016 Brain Prize was won by three neuroscientists for elucidating the molecular mechanisms of long-term potentiation, one of the keys to storing and sustaining lifelong memories in mammalian brains. With all this progress at hand, human brain preservation at the end of our lives, as a personal choice, is a more viable and compelling idea than ever before.

This page discusses some common objections to and defenses of the value of brain preservation as a social option. We humans are only now coming to understand ourselves as informational entities. In so doing, we are learning the use and value of our information. In considering whether brain preservation is a wise and ethical use of resources, one must ask under what circumstances *information itself*, both generally and within unique human minds, is worth preserving as an individual choice in free societies.

I firmly believe that giving each of us more options to preserve any information we individually may want available for the future, including our own memories and identities, is a wise and ethical act, as long as a minority of neuroscientists expect that the act of preservation may have future informational value, and preservation can done sustainably with respect to the environment.

If we had a validated and affordable brain preservation option available to us at the end of our lives, some of us would presently choose to preserve our mental information. Many others would presently not, and many of

[1] http://www.brainpreservation.org/content-2/overcoming-objections/

us have strong views on the nature of death, both for ourselves and others. Nevertheless, if preservation can be validated to preserve even simple episodic memories in well-studied synaptic circuits in any higher model organism, a validation that has already occurred for lower organisms[2] and which may arrive long before we have a complete theory of all the ways memories and personality are stored in the brain, a compelling case can be made that making the preservation option available for all who might want it, and championing the ability of each of us to make our own free choices on this matter, will be a great advance in building the kind of world we want.

In coming to understand ourselves as informational entities, one important insight is recognizing the **patternist nature of self**. What we call our memories, our personality, and our "self" is not our matter but is in fact a complex and special informational pattern held in our biology. We know this because when *sufficiently complex patterns in our brain are replicated in technology*, as happens when a cochlear or retinal implant is integrated into a deaf or blind persons nervous system, this small part of their self now operates as technology. It is thus the pattern that matters, not the "substrate".

As we will argue, our science and technology are presently engaged in an accelerating process of replicating, preserving, and "uploading" our biological patterns into significantly faster, more durable, and capable technological ones. This process of "pattern uploading" from biology to technology may be as natural, useful, and universe-driven as it is human-chosen. If brain preservation works, a question the BPF was founded to investigate, it will only be the latest in a long series of technological advances on Earth that increasingly capture and improve our all our unique and valuable biological patterns.

B.1 The Scientific Advancement Defense

Perhaps the most common argument for developing better brain preservation techniques is to advance our sciences and technologies, including neuroscience, medicine, microscopy, cognitive science, and computer science, and to gain more of the social benefits they provide. This is a good place to begin a discussion of brain preservation, as most will see this value. But while the scientific advancement argument recognizes the value of preservation of samples of medically unique brains (from individuals with mental

[2]http://www.alcor.org/Library/pdfs/Persistence.of.Long.Term.Memory.in.Vitrified.and.Revived.C.elegans.pdf

disorders, or other functional differences or abilities), and a sample of individuals with "typical" brains (to understand the range of healthy function), note that it does not consider the value of preserving memories or identity in these or other individual brains.

B.2 The Exceptional Cases Defense

Some individuals who have no interest in brain preservation for themselves will nevertheless grant the potential value of brain preservation for others who might wish it, in exceptional cases. They may grant its value for a child or young adult who has been struck down early in life by disease or accident (or perhaps for the benefit of the child's parents, to aid their grieving process). They may grant its value for those individuals who feel they have unique and unpreserved culture, history or knowledge they wish to pass on to the future. They may grant its value for someone who believes they have unfinished creative goals that they feel uniquely capable of pursuing, relative to other minds, for many years or decades to come. Albert Einstein's brain, which has been chemically preserved, has been cited by some in this regard. Helping people to understand and support such "exceptional cases" can begin to move society toward acceptance of this technology, and is a reasonable next step. David Ewing Duncan, in **When I'm 164**, 2012, has found that **only 1%** of individuals in developed societies are presently interested in living beyond their biological lifespan. We can call this 1% the "exceptional cases" who might presently consider the brain preservation option, if it were validated, which today it is not. However, this minority could easily grow if neuroscience advances, validation emerges, cost comes down, and social behaviors change. Emergence of the brain preservation option could also have positive effects on the larger society, as we argue next.

B.3 The Social Benefits Defense

If any brain preservation technology can be proven to preserve the key morphological features, and any molecular features, that neuroscientists presently believe contain our memories or identity, and if neuroscience and computer science can show that those features alone are able to preserve and create memories in both animal and computer models, then the availability of affordable and environmentally sustainable brain preservation services, the option and freedom to use them by anyone in society, and their use by

a socially significant minority may begin to *change those societies for the better today*, regardless of how much or how soon anyone's personal neural information is retrieved at a later date.

Specifically, social values in such societies may move measurably toward what we can call a **Preservation Value Set**. Imagine any country where a socially significant minority, let's say 100,000 individuals, have personally made the brain preservation choice at death. Given the conversations that must have occurred in the larger society during the creation and access of this freedom, and the predictable lowering of cost and improvement of access that comes as any technology becomes more widely adopted, we can expect some measurable changes in those individuals social values and perceptions, and to some degree, within the wider society as well. As a result of this level of use and access of the brain preservation choice, a politically-significant fraction of individuals in such societies may become, as measured on social surveys, noticeably more **science-oriented** (more willing to advocate and fund rapid and responsible scientific advances in their society, given the increased personal benefit they may receive), more **progress-oriented** (more willing to see and support signs of social progress, as they desire to be revived in a measurably better world), more **future-oriented** (more comfortable making long-term plans in more facets of their life), more **sustainability-oriented** (less willing to harm their environment today, as they realize they may return in the future), more **preservation-oriented** (more motivated to preserve the unique species in our natural environment and the unique information in human culture and minds), more **truth- and justice-oriented** (better behaved today, as those who have experienced injustice may donate their memories so that present crimes may be righted via future forensics, and so future laws may better match true human behavior), more *diversity-oriented* (more motivated to live in a "usefully unique" way themselves, to increase the value of their memories and mind to future generations) and ultimately, more **community-oriented** (more desirous of living in a way that makes them valuable not only to themselves, but also to loved ones and society). For many, achieving a significant shift toward preservation values in our societies today, regardless of how much neural information is eventually recovered in the future, is the most important reason to support the brain preservation effort.

Why are we suggesting 100,000 adopters of brain preservation before measurable positive social changes may occur? Humanity is naturally set in our ways in our thinking about death, and confronting many aspects of death. This can be observed whether we are talking about the death of an idea, a social theory, a spouse, or of ourselves. Changing our thinking regarding death and the proper ways we should deal with it often roughly

follows the DABDA stages (Denial, Anger, Bargaining, Depression, and eventual Acceptance) identified by the psychologist Elizabeth K ubler-Ross. In the beginning of any proposed new social freedom involving death (in this case, the death of the idea that biological death is and should be final for every individual), some degree of social denial and resistance are to be expected. Thus brain preservation technologies will very likely require a politically significant minority of Early Adopters, willing to lobby for it in political and legal ways, in order for access and affordability to grow. The majority of the public will very likely continue to ignore and discount this personal option for some long time to come, as a natural and expected response. We suggest that 100,000 adopters and their friends is a politically significant minority. Consider the Netherlands, which has a population of 17 million. In that country, the Levensiende ("Life End") group, and advocacy and support organization for the right to physician-assisted suicide and personal choice in confronting death, has slowly grown to its status of 130,000 members in 2012. Physician-assisted suicide, usually with less than six months of life expectancy left, and earlier in special cases, has been legal in the Netherlands since 2002. Roughly 2% of all of those dying in the Netherlands now annually choose this option. But it wasn't until May 2012, after much lobbying from Levensiende and other groups, that it became legal for Netherlanders to choose physician-assisted suicide in their homes, in addition to hospitals and hospices. It took the lobbying and stories of a politically significant minority to change the laws and increase access, and push the general public into acceptance of this expanded social choice regarding death. The same denial and resistance dynamic is likely to be expected in early efforts to increase social access to the brain preservation choice. But at some point, social acceptance occurs, and we have a society considerably more aware of and willing to advocate for "preservation values", whatever those may be.

As another way to speculate on what those values changes will be, some research indicates that changing our perception of the finality and unfairness of death may even make us measurably **less dogmatic** in our beliefs, **more tolerant of social change**, and more willing to champion **cognitive diversity**. As Sam Harris notes in The **Moral Landscape**, psychologists have discovered that merely reminding judges and juries of the fact of death increases their inclination to automatically punish those who have violated the law, and to reward those who uphold cultural norms. Others have replicated this association between death awareness and cognitive dogmatism and intolerance.

Awareness of our eventual death can be a great motivator, as Steve Jobs eloquently reminded us in his Stanford Commencement Speech in 2005,

before his own death from the cancer he had recently acquired. But there are many great motivators to live purpose-filled lives: awareness of our limits, curiosity, passion, honor, duty, ethics, hope, vision, and intelligence, for example. Those who currently choose to preserve their brains at death for the possiblity of future revival are no less motivated to live full and valuable lives.

I do believe that death is ultimately inevitable. Nothing ever lives "forever." All living things have eventually grown into something else, die, or are themselves eventually outmoded and outgrown.

But what is different now is that for the first time, due to our advancing science and technology, we can hope that in the near future, our outdated ideas and behaviors will die appropriately, inside our own minds, to be forgotten and outgrown by the individual who carries them, only when they have *outlived their social usefulness*, as judged by the individual who carries them and by the society in which they were created. An increasing number of us don't want our minds and ideas to *die inappropriately*, often long before their usefulness has ended, due to the current limitations of biological nature.

Since the dawn of civilization, millions have lamented the loss of personal history and unexpressed insights that occurs with their own death. Our lifespan is surprisingly short by contrast to the appropriate lifespan of the unique experiences and ideas we gain and create during our lives, much of which we are not able to express in our behaviors or works prior to death. It seems to some that just as we are reaching an age where experience leads to wisdom, we must end our lives. Much of this unique internal information is presently lost at death, and only some of it is eventually reinvented by others.

Even if our children were to wear a camera their entire lives, as may one day occur, much of their unique subjective personality, thinking style, experience and insights may never be reinvented by anyone in the future. Each human mind has an astronomical number of connections that are unique only to that individual, and the present and future value of that diversity is far beyond our current ability to estimate. If future society continues to have finite computing capacity and limited ability to recreate its diverse history, as it does today, valuable information will always be lost with involuntary death. Soon brain preservation may offer a powerful new way to reduce this loss. For an excellent overview of how the advance of civilization is directly tied to the effectiveness and efficiency of the preservation and exchange of our unique history and ideas, see **The Guardian of All Things: The Epic Story of Human Memory**, by Michael S. Malone, 2012.

Evolution loves diversity, and life's diversity has constantly increased over the entire history of biology, human culture, and technology. Even the great extinctions further increased our genetic and species diversity. When one considers how much unique information presently dies with an individual without being sufficiently shared through that individual's behavior or works, the increase in useful diversity that is promised by brain preservation may be a social advance on par with writing, moveable type, mechanical recording, and other major historical advances in our cultural memory.

For all of the reasons above, brain preservation, if undertaken by a socially significant minority in any society, may become a major social good.

B.4 The Religious Objection and Defense

While there has been very little guidance on this issue from religious leaders so far, adherents to most religions today might think that the preservation of their brains at biological death would go against their beliefs. We must be respectful of and sensitive to such statements, while at the same time recognizing that behavior and ethics here will never be uniform. Within every religion there will always be individuals and communities who do not believe that the preservation choice conflicts with their faith. There are already patients from several religious faiths in cryonic storage. These individuals expect or hope to be revived in the future, if their God or the Universe permits.

At the same time, there are individuals from a variety of faiths who would presently be willing to **donate their memories** to the future, but who would not wish to be **personally revived** in the future, given their particular religious beliefs. Some religious communities may consider brain preservation for memory donation to be acceptable, assuming that such a request is both feasible and would be honored by future society. The mother of one of us (J.S.), a devout Christian, would have gladly preserved her life's memories for her family if affordable (low-cost) brain preservation had been available at the time of her death, but she would not have wished to be revived as an individual in the future. If brain preservation becomes increasingly accessible and affordable in coming years, and if the science and technology continues to improve, we can expect a variety of responses to the brain preservation question, from a variety of faiths.

B.5 The Natural Aging and Death Objection and Defense

Many people today feel that living a long natural biological life is sufficient for them, and they have little to no desire, at the end of life, to extend it beyond what has been given to them by God or nature. Helping us to gracefully accept our biological deaths is the fact that our bodies and minds naturally age and become increasingly frail and feeble after we reach sexual maturity. This makes the sudden cessation of life in our old age much easier to bear. As the American freethinker Robert Ingersoll says in *On the Life Cycle*, 1887:

"There is something tenderly appropriate in the serene death of the old. When eyes are dim and memory fails to keep a record of events; when ears are dull and muscles fail to obey the will; when the pulse is low and the tired heart is weak, and the poor brain has hardly power to think, then comes the dream, the hope of rest, the longing for the peace of dreamless sleep."

But it is also true that the "natural" aging described here is being steadily minimized by advances in science and technology. Consider how sanitation, public health, and medicine have greatly extended our healthspan (the healthy period of our lives) improving average American lifespan from 47 to 77 years over the 20th century. More recently, **longevity research** and **regenerative medicine** are beginning to shorten our frailspan (the physically and mentally frail and enfeebled period of our lives), by slowing the basic processes of aging. For example, a **2011 study**[3] discovered that much of the physiological degeneration that occurs in adulthood, in a mouse population with a premature aging mutation, was due to **a small population of senescent cells** that produce inflammatory proteins. When these cells are removed in middle age or earlier, as in the mice in the study, the body doesn't age "naturally", but retains physical and mental vigor well into old age, with a much more abrupt decline much later in life—a process called "squaring the curve" of aging. The team that accomplished this, Darren Baker and Jan van Deursen at the Mayo Clinic, did it again in 2016[4], this time with genetically ordinary mice, extending their lifespans by 20–30% and making it clear that senescent cells are a key target for healthy life extension. If therapies to remove or block these cells or their inflammatory

[3] http://www.nytimes.com/2011/11/22/science/in-bodys-shield-against-cancer-a-culprit-in-aging-may-lurk.html?pagewanted=all

[4] http://www.nature.com/news/destroying-worn-out-cells-makes-mice-live-longer-1.19287

proteins can be developed for humans, as is now being explored, those who use them will feel like death is a sudden collapse and loss of function at the end of an even longer and more vibrant life than we typically have today. As our social acceptance of natural aging falls, our acceptance of natural death will also be challenged, at least by some.

Some individuals view the brain preservation choice as something that goes against the natural way of (biological) life. They remind us that in life, the old must be removed to make room for the new. Winter clears the way for Spring. This is true from a biological perspective, and yet biology is only part of the story of modern humanity. As our civilization has developed, our minds and our technology have come to play ever-growing roles in the nature of humanity. But increasingly, unique ideas, perspectives, and experiences in modern human minds die *inappropriately*, not archived or retired by conscious choice, as their usefulness fades, but lost because of the limitations of biology.

Anthropologists have observed that the more complex society gets, the greater the social (and economic) value of each individual human life, and the more elaborate our responses of grief and injustice to the loss of life. Society, via our cultural memory, and technology, via writings and recordings and science, are far better at preserving information than our biological bodies, which die on a cyclic basis. When our minds and science were less imaginative and less developed, this cycle of life was more acceptable. Today, the cycle has come under scrutiny, and we can now imagine less informationally destructive ways of life. We may soon have a choice to greatly increase the diversity of mind on Earth.

Information technology in particular is very good at preserving everything that has gone before it, and computers, using less and less physical resources per "bit" of information storage are preserving more and more of our past and present world, and enabling more social creativity, diversity, resiliency, and progress than ever before. As social and technological systems advance, they increasingly learn how to preserve each life's learnings to allow our descendants to do and live better in the next life. In the future, we can imagine ourselves as technological or advanced biological beings, where the only deaths that occur are the "little deaths" that presently happen in our minds every day, when less fit ideas are replaced by better ones, and the old neural connections extinct themselves, making room for new ones—a life of constant growth and change, but no loss of information of value to us or to our communities.

Nature and life continually grow, learn, and change, and we must do the same if we are to understand ourselves as not only biological, but also social and now even technological beings. Today, no one would be considered

fully human without learning the languages and social norms, the cultural technologies, that our society has developed over the last two million years. Our electronic technology, for its part, is not only the fastest new learning system on the planet, it is becoming an increasingly life-like and natural extension of both our environment and ourselves. When we enlarge our definitions of nature and self to include both our culture and our technologies, we can better appreciate and understand the value of brain preservation for all who might desire it. This is natural, but it is a new, more complex nature than the old, just as early life on Earth evolved and developed a new, more complex nature as it grew into our modern forms. In nature, change, growth, learning, and new forms of diversity, adaptibility, resilience and complexity appear to be among the few constants we can depend on.

B.6 The Patternism Hypothesis

Perhaps the greatest challenge to seeing the value of brain preservation today is the need to adopt a "patternist" understanding of the nature of self. The last 150 years of biological science have carefully uncovered the working hypothesis that our individual selves are entirely the result of special complex physical structures and processes, orpatterns in our brains, bodies, and their interaction with the environment. The *patternism hypothesis* proposes that it is a special physical pattern, not the matter, or even the type of matter (computer or biological), that stores the highest level information in living systems. If the special pattern that stores this information can be successfully maintained, and copied as necessary, the information survives.

Remember first that our identities (our selves) are not contained in any *particular* biological matter. All our matter is replaced, or turned over, in our bodies and brains on a moment-by-moment basis. Some **ninety-eight percent of the atoms in our body are replaced every few years (the number varies by study)**, by the food we eat, the air we breathe, the liquids we drink. This is a natural process of pattern copying. We are continually being "uploaded" into new matter with a very similar pattern all the time. Many of our cells (with the exception of the brain) are constantly dividing, replacing old with new. This copying process is never perfect, and certain useful molecular tags (methylation, phosphorylation, ubiquitination) are always lost in this constant and statistical process of molecular renewal and turnover. This copying happens so incrementally, and our patterns are altered so subtly, that we don't notice it, until we study how the process works on the molecular level. Every living thing is constantly having parts of itself "uploaded" into subtly different to very

different physical substrates, as a natural process.

For example, artificial cochleas and retinas replicate and restore sensory aspects of the biological self. Even brain patterns are now being replicated in our technology—see for example Ted Berger's work with the **artificial hippocampus**, and other projects in **neuromorphic engineering**, where chips are designed to replace brain circuitry. This work is simple today, as neuroscientists still do not fully understand all the ways neurons process and store information, but we have every reason to expect continued progress in these efforts. Accepting and understanding the patternist nature of self allows us to realize that one of highest purposes of humanity appears to be a responsibility to continually preserve and improve our best biological, social and technological patterns.

As we have said, what is natural changes as our species changes. As our physical patterns have grown in complexity, humanity's natural abilities and responsibilities have grown in the same measure. Before humanity invented gestural and verbal languages, which were among our earliest "technologies," we had no responsibility to pass on to others, or give extended lifespan to, our individual experiences. But after language arrived, we gained a new responsibility to teach our descendants, and thereby improve our families and culture. Once written language arrived, we gained further responsibilities to physically record and pass on, or give extended lifespan to, our discoveries and experience, and to further improve individual and social wisdom. Today's digital computer and communications technologies are direct extensions of these earlier technologies. We have a new responsibility to improve them as well, to broadly distribute their benefits, to try to minimize their downsides, and to endeavor to use them to increase our ethics, wisdom, awareness, foresight, and resilience.

If inexpensive and validated brain preservation arrives, we will be endowed with new capabilities to pass on, or give extended lifespan to, our memories, learning, and identities to our descendants. In time, we will recognize new social responsibilities to do just this. Whenever we successfully improve the complexity and resilience of our individual and social patterns, and allow them to live for as long as they might be valuable, available to any who might be interested in them, we seem likely to achieve greater individual and social conscience and consciousness, new respect for the value, rarity, and uniqueness of each human life, and new levels of individual and social progress. This is perhaps the greatest potential benefit of the patternist perspective: we can be more effective and aware today, and make better choices in the present moment, choices ideally in greater harmony with the self-improving nature of life and the universe.

In summary, the past century and a half of research in cognitive science

and neuroscience have increasingly established that the entirety of what we call our mind is a complex information processing stream computed by the circuits in our brain, and in the society and technologies in which that brain is embedded. Once we recognize that our critical physical patterns are not only biological, but also social and technological, we can resist the resignation, isolation, and apathy that can accompany biological old age. We can recognize that even as our biological minds begin to fail us, our social and technological ones are growing faster, smarter, and more intimately connected to our biology every year. Furthermore, growing knowledge of brain health and neural plasticity offers us new ways to reduce or reverse "natural" cognitive decline as we age, to restore our mental abilities to more youthful levels and to remain lifelong learners. We learn to see our selves as not just our biology, but also as our social minds and technology, we can become champions of the kinds of scientific and technological developments that will increase innovation, wisdom, resiliency, and social and individual empowerment.

B.7 A Number of Non-Obvious Proposals

If we wish to argue the potential value of brain preservation as a broadly available social option in coming years, it helps to make a number of not-immediately-obvious proposals:

- Brain preservation techniques may soon (perhaps within this decade) be validated to preserve useful neural information, including memories, in model organisms.
- Should brain preservation be validated to preserve neural information at death, this will be a natural process, once we acknowledge that not only our biology, but also our social minds and our technology are natural.
- The preservation of any amount of neural information upon our death could prove valuable to our loved ones and society both today and in the future, if it can be inexpensively preserved today, and if it is reasonable to expect that it could be inexpensively recovered by future technology.
- Low-cost preservation technologies may soon exist, which is relevant to the financial wisdom (expected benefit to cost ratio) of the brain preservation choice, as preservation always involves taking resources from loved ones or society today for an uncertain future return.

- Rapid advances in computing and scanning technologies argues that neural information might be inexpensively and automatically read from preserved brains even a few decades from now, while one's loved ones are still alive.

- Not just memory retrieval, but full revival of the individual, and their indefinite lifespan in the future may also be an outcome of brain preservation, for those who might desire either option.

- Memory retrieval or identity revival will very likely be done in computers in the future, and computer technology is dramatically more miniaturized and resource efficient per computation with each successive generation. If present accelerating, miniaturizing, and efficiency trends continue, technology will support far more living, loving minds in the future than biology ever could, and this ever-increasing diversity of mind, creativity, and intelligence appears to be the long-term trend of nature on Earth.

B.8 A Brief Bibliography

For inspiring evidence of how our biological brains and minds can be continually improved throughout our lifespan, even in advanced age, read Norman Doidge's excellent book, *The Brain That Changes Itself*, 2007. For a general understanding of brains as connectomes, read Olaf Sporns' *Discovering the Human Connectome*, 2012, and Sebastian Seung's *Connectome: How the Brain's Wiring Makes Us Who We Are*, 2012. For more on how our mind and brain are embedded in their social and technological environment, read Andy Clark's excellent general-interest book, *Supersizing the Mind*, 2011. For two good books that discuss our increasingly intimate brain-machine interfaces, and our progress in simulating modular subsystems of the biological brain within our technology, and implanting those systems in living human brains, read Michael Chorost's very accessible *World Wide Mind: The Coming Integration of Humanity*, Machines, and the Internet, 2011, and Miguel Nicolelis's *Beyond Boundaries: The New Neuroscience of Connecting Brains With Machines*, 2011.

For a technical exploration of connectomes, read Sporns' *Networks of the Brain*, 2010, and for a technical understanding of the physical basis of subjective experience and consciousness as emergent and nonmystical processes of neural synchronization, read Gyorgi Buzsaki's excellent *Rhythms of the Brain*, 2006. For an understanding of how organisms are most essentially a type of computer at the genetic, cellular, and physiological levels, you

may enjoy Uri Alon's technical book *An Introduction to Systems Biology: Design Principles of Biological Circuits*, 2006, and Eric Davidson's technical work *The Regulatory Genome: Gene Regulatory Networks (Circuits) in Development and Evolution*, 2008.

B.9 Nine More Objections and Defenses—For the Scientists and Philosophers

We will conclude this page by considering some common scientific and philosophical objections to brain preservation, and suggest some answers that seem reasonable to us. If you have a scientific or philosophical bent and do not presently see the potential value of brain preservation, either for yourself or for others who might choose it, please let us know if you still have questions or critiques after reading this article.

First, hypotheses in science are always conditional, including the patternist hypothesis of self. We may agree to tentatively hold the patternist hypothesis, but to do so also requires us to begin considering its implications with respect to the future of mind and technology. Some of these implications are abstract, unsettling, and not among our normal cultural concepts. Nevertheless, a large body of scientific evidence can be marshalled in favor of the patternist hypothesis, so it makes sense to hold the hypothesis conditionally, and to explore its implications, at least until contrary evidence against it materializes.

Second, many scientifically-literate individuals do not recognize how close our species has come to having the technology to make memory and mind preservation a reality. They are not familiar with the state-of-the-art techniques available for chemically or cryonically preserving neural structure at the synaptic level, and for verifying this preservation and circuit tracing with automated sectioning and volume electron microscopy techniques. Please see the Technology[5] section of this website for references on the current state of the art in these areas. Fortunately, objections based on the lack of our capacity to preserve are easy to define. Overcoming them in a definitive way is one goal of our Brain Preservation Prize.

[5] http://www.brainpreservation.org/tech-prize/

Third, some doubt that we will ever able to decipher the code for long-term memory storage in brains. Fortunately, this doubt seems unreasonable. Neuroscience is rapidly gaining a molecular-level understanding of processes central to long-term memory creation. Our brains store memory in at least three different ways. **Working memory** is stored in conserved electrical patterns, with a persistence of seconds. **Short-term memory** is stored in preexisting hippocampal and cortical synapses and preexisting signaling proteins, with a persistence of a few days. **Long-term memory** is written from our hippocampus to our cortex, primarily during slow-wave sleep every evening. It involves the synthesis of *new* synapses and brain proteins, and modifications to synapse and nuclear proteins, and it has a persistence of a lifetime if it is periodically reinforced. It is long-term memory, encoded in durable synaptic and nuclear changes in neurons, that we particularly care about preserving. If we are revived with the loss of our working memory, as happens after a concussion or anesthesia, this is not of great concern. We are even able to bounce back well if our short-term memory is entirely wiped out, as sometimes occurs in anoxic brain trauma followed by short-term amnesia. There is some tentative evidence, for example, that the hippocampus might be uniquely vulnerable to damage during cryopreservation, unlike the rest of the brain. But as long as our cortical synapses can be well preserved, uploaded, and connected to an artificial hippocampus in the future, we'd likely lose very little useful information and personality, just the last few days of experience prior to preservation. Neuroengineer Ted Berger has been making early versions of implantable artificial hippocampus chips since 2005, for mice. Recall Henry Molaison (HM), the famous memory disorder patient who could not learn new memories after his hippocampi were surgically removed, but who kept all his older long-term memories prior to the surgery. What we care most about in brain preservation is that our long-term memory will survive the preservation process, and can be reinstated from appropriately detailed scans of the preserved brain. One critical proof of this ability will come when neuroscience sufficiently understands part of a model animal nervous system, such as *C. elegans* (the nematode) or *Aplysia* (the sea slug), well enough to train the animal associatively in one of several unique ways while alive, chemically or cryopreserve its brain, scan the relevant bit of brain tissue, and then *correctly predict how it was trained* by reading the scan of the appropriate neural circuits. This will require the ability to model, in a very well-studied behavioral subsystem (neural circuit set), the way both synaptic connections and neuromodulator proteins at these connections bias the

pattern generators in that circuit into a particular set of output patterns[6], a field of research called behavioral plasticity. That demonstration may be ten or more years away, but if and when occurs it will be a major step forward in clarifying how robustly the brain preserves higher information, including memory and experience, in particular synaptic connections and their unique sets of molecular weights.

Fourth, even if we understand the code, some doubt we can inexpensively and reliably retrieve memories from a preserved human brain. One doubt arises because of cost. But we are already using automated robotic systems to slice, scan, and upload very small animal brains (including the zebrafish brain, the size of the tip of a pencil) into computers today (these uploads don't reproduce memories because they don't yet have all the critical molecular features, and we still don't understand the code). As technology advances, it is reasonable to expect that the cost of this scanning process will continue to drop exponentially, while capacity continues to grow exponentially. Also, new methods of brain scanning will surely emerge. One promising technology is **molecular-scale MRI**. Recently, MRI machines have been built that can *image individual cell proteins*[7], and there appears to be no theoretical reason these machines could not eventually image whole human brains. Molecular-scale MRI may one day give us the ability to scan plastinated brains inexpensively and nondestructively, and to upload the critical molecular features that encode our memory and identity. Another doubt arises because modern neuroscience suggests that our molecular memories, when remembered, are not simply recalled but are actively "recreated", in a holistic and electrical process, from molecular networks of stored synaptic potentials distributed throughout the brain. But this is not a problem, it is an advantage. We know from artificial neural network models that this holistic way of storing information is robust to damage. Memories are retrieved in a distributed, associational manner from molecular stores. Thus future scans should be able to retrieve memories even from partially damaged brains. Furthermore neuroscientists suspect that humans share a common, or "baseline" brain, in which the vast majority of cellular and molecular structures and processes are highly similar from brain to brain. Simulating this baseline brain is a top goal of current and future neuroscience, in the same way we try to predictively simulate bacteria today, down to the molecular interactions of their metabolome. Such simulations are quite limited today, but they get expo-

[6] https://www.ncbi.nlm.nih.gov/books/NBK20011/

[7] https://doi.org/10.1126/science.1233222

nentially better over over time. On top of our shared baseline brain, we have neural correlates of individuality, or NCI's, molecular stores that comprise our *unique* memories and individuality, and which are persistent, even in the face of chaotic electrical and molecular activity in the brain. Brain preservation is thus about saving the NCI's, which appear to reside almost entirely in our unique synaptic connections, a few associated proteins, and a few nuclear modifications, and later placing these NCI's in a baseline brain emulation in a computer. Fortunately, in addition to being predictable and persistent, molecular NCI's are highly redundant and fault-tolerant. They survive even when the brain temporarily loses all electrical activity in coma, surgery, or cold-water drowning, and through all kinds of trauma and environmental fluctuations. For example, if you forget something because a particular synaptic connection weakens or breaks, you can very often recall and reestablish what you have forgotten simply by thinking of *other aspects* of the memory in question, routing around the damage and reestablishing the memory. All this suggests that memory and identity retrieval from preserved human brains will be a very worthy and exciting scientific and humanitarian endeavor with great chance of future success.

Fifth, while it may be possible to retrieve memories, some doubt that we will be able to retrieve them in a piecemeal, incremental fashion. In the worst case, for example, one might fear it will be necessary to resimulate an entire conscious individual in order to recall even a single memory from that individual's life. Thus, those willing to donate their memories to the future, but who do not wish to be consciously revived in the future, might see brain preservation as undesirable. Fortunately, this fear looks to be unfounded. We can already reconstruct realtime experiences from very small populations of neurons today (e.g., 177 neurons holding visual working memory in a the cat's brain, Stanley et.al. in 1999). Today's early models of consciousness (e.g., Buzsaki's neural synchronization[8] and Tononi and Koch's integrated information theory[9]), though incomplete, are already powerful enough to suggest that this is a number of neurons far too small to be conscious. If long-term neural information is stored in a similar connectionist way to working memory, using small populations of distributed and redundant networks to encode information, we should in the future be able to extract memories and experiences from preserved brains in an incremental, divisible fashion, without restoring higher individual consciousness, if that is what the preserving individual desires. Neural synchrony and feature

[8] http://www.scholarpedia.org/article/Binding_by_synchrony

[9] http://www.scientificamerican.com/article.cfm?id=a-theory-of-consciousness

binding will be required to retrieve memories, but just as you can retrieve memories from local areas of the brain during dreaming, and not be in higher (globally self-aware) consciousness, it seems likely that memories can be retrieved in a similar fashion from a scanned brain, once we understand the long term memory code. Just as an anesthesiologist can prevent consciousness today by administering anesthetics which prevent neural synchronization and allow a neurophysiologist or neurosurgeon to operate without the patient's awareness, future memory donation without individual identity or higher self-consciousness restoration may be a common option in the future, for those who desire this particular choice. Neural synchronization, the current leading candidate for a mechanistic understanding of consciousness, has made great conceptual advances in the last few years. See **Wang's *Physiological Reviews* article**[10] for a recent review of this exciting field. The neural synchronization model of consciousness is consistent with the way disruptions of synchronization with anesthetics remove consciousness, and with the way several patients who have been in a persistent vegetative state for years have been partially **reawakened to consciousness and mental life by administering Zolpidem**, a drug that modulates theta and gamma oscillations in the brain. We are beginning to understand consciousness as an entirely physical process, one we may one day replicate in sufficiently complex technology.

Sixth, some doubt that the full identity and self-consciousness of any particular person could ever be "uploaded," or emulated in a computer or other nonbiological life form. This objection often rests on the **material identity hypothesis**, the belief that the human mind must be indivisibly attached to the particular type of matter, in this case biological matter, that presently generates it. But what comparative psychology and computer science have taught us so far is exactly the opposite. Information processing is independent, to a surprising degree, of the particular physical substrate it is run upon—any substrate of sufficient complexity will do. As biologist Simon Conway Morris states in *Life's Solution*, 2003, both simpler and higher features of the human mind and senses are shared in animals, including insects, with much-simpler and differently-built brains than ours, and a few mental features have already been successfully simulated (replicated) in computer technology. Furthermore these computer technologies, when integrated with biological brains, as in neural, retinal, and cochlear implants in humans, produce replicable components of mind. If we can recreate the relevant patterns of sensation, memory, emotion, experience, consciousness,

[10] https://doi.org/10.1152/physrev.00035.2008

and identity in a computer or robotic body instead of living tissue, we have recreated the mind. If future society scanned your preserved brain at the molecular scale, and could replicate a living brain in a computer that generates sufficiently similar types of patterns, this copy would truly "be" you. Certainly, as several biologists have noted, the ability to replicate all the critical patterns of one material system (wet biology) in another (electronic computers) is not guaranteed. But to date, every new computational substrate that has emerged at the leading edge of universal complexity has not only contained all the capabilities of the previous substrate, it has exceeded them. As universal complexity has journeyed from physics to chemistry to biology to (today's still-primitive and non-autonomous) technology, each new substrate has grown to contain all the physical abilities of the previous, and has introduced powerful new freedoms and abilities as well. Certainly if future science discovered any **pattern insufficiencies (structural or functional)** in our computer simulation of human brains, we could always seek to use advanced nanotechnology to recreate a biological version of the person preserved. Advanced nanotechnology could even repair and reintegrate the same physical matter of the preserved brain into a future repaired biological form. In the very long term future, nanobots created by a society with advanced artificial intelligence might carefully remove the fixative or plastic resin embedding each neuron, repair aging and other damage, and revive the same physical brain that was preserved. Such a course of revival would likely convince even the most skeptical that "they" had been revived as the "same" individual. For examples of such revival scenarios, read Eric Drexlers' excellent *Engines of Creation: The Coming Era of Nanotechnology*, 1987, or this **"realistic" scenario for nanotechnological repair of the frozen human brain**[11], written by an anonymous biologist in 1991. While these scenarios are both plausible and fascinating, material repair and restoration of the preserved brain may turn out to be a very uncommon pathway for the recovery and reanimation of mind. While many of us might desire to be revived as a biological body, patternism suggests that placing such a restriction on our revival would serve only our own vanity, and might be a hindrance to our rapid revival for ourselves, loved ones, and society. Those not willing to let future society create simulations of themselves first may delay their revival and return to future society by many decades, as nondestructive pattern reading and emulation technology for preserved brains may arrive long before advanced nanotechnology. We may think we presently understand the future optimal course of our revival, but the reality is, there are many ways future science might revive us, and many

[11] http://www.alcor.org/Library/html/nanotechrepair.html

useful and "true" copies of ourselves that could come back. If we don't let future science and future minds advise us on the best pathways for memory donation or reanimation, we are very unlikely to pick them today.

Seventh, while some will grant that all valuable biological structure and function may eventually be duplicated by technology, they believe that there is some metaphysical element of self which must exist independently of physical processes, and which could not be transferred in any material duplication. Such individuals would agree with statements like: "If you made an exact molecule-for-molecule copy of me, that copy might *act* just like me, have my memories and my personality, and would even *think* it was me, but it would still not *be* me." The **independent soul hypothesis** is the belief that the mind is not only an emergent property of the brain, but is also independent from (has an existence separate from) the physical patterns and matter that houses it. This is a tradition of many, but not all, religious, philosophical and cultural heritages. At the same time, there are also subgroups of every one of our major religions, philosophies, and cultures which either do not believe, or have never even considered, the idea of metaphysical independence of mind from matter. Most religious scriptures are silent on this question. As philosophers from Descartes to Whitehead have argued, it is certainly useful and appropriate to see our minds as in a different category from physical things. We can observe an apparently fundamental body/mind, material/virtual dualism in all complex matter on Earth. Certainly complex minds are not only emergent, they do seem particularly special in the universe. As human minds grow, over both individual and historical time, they gain astonishingly greater influence over their local material environments, as is reflected in the popular phrase, "mind over matter". We can even see an emergent dualism in the "virtual reality" that complements today's physical computing technology. Several scholars have argued that our computer games and simulations are components of an emerging and still-primitive "technological mind." Yet in all these examples, the material/mental and the physical/virtual are also fundamentally *integrated* (nondual) phenomena. Human minds have emerged on a smooth and divisible continuum from our physically simpler predecessors. While we can observe simple matter without higher mind, science has never observed, and we cannot reasonably imagine, mind without some physical basis to support its complex patterns.

Eighth, some who grant the scientific plausibility of reanimation of their pattern still have little faith that our future ecological or political environ-

ment will be either able to sustain, or will be socially hospitable to, such reanimation. Our existing population of seven billion humans is presently seriously degrading our planet's environmental systems, while demographers are hopefully projecting an end to human population growth in the mid-21st century. Won't adding more humans, even "virtual" humans, just make our precious planet a worse place? To answer this question, we must carefully consider the history and likely future of computing technology, which has seen exponential improvements in its speed, capability, efficiency, and miniaturization for at least 120 years, across at least five different design platforms, since our first complex mechanical computers, such as the 1890 Hollerith Tabulating Machine. This trend is commonly known as Moore's law. What is less commonly appreciated is that our computers have also become astoundingly more energy efficient over the same time period. As Gene Frantz observed in 2000, and named Frantz's law, digital signal processing power used per computation halves every 18 months in our leading computer chips. As computers continue to miniaturize, they also become exponentially more space-efficient and matter-efficient as well. While there are many short-term engineering blocks, physicists presently see no fundamental physical reason that will prenvent us from continuing to make accelerating advances in **nanotechnology**. If we are able to "upload" billions of human minds into future highly miniaturized computers, planetary resource issues will have little relevance. Resource sustainability is an issue for biological humans, which use roughly the same or more level of resources with each doubling. Physical resource accessibility is an increasingly less important issue for computers, which become ever more miniaturized, resilient to damage (as they are able to easily "back up" their complexity), and as their intelligence grows, increasingly *independent* of energy and material resources, per any standardized measure of complexity we choose (per computation, per mind, per society, per species). A world with widespead artificial intelligence will be both radically miniaturized and have abundances, such as fusion energy, that we can scarcely imagine today. Furthermore, the more minds exist, the more diversity, variety, and specialization society contains, and evolution always seems to maximize diversity, however it can. What about dystopian political futures? They certainly are possible, but as Matt Ridley notes in *The Rational Optimist*, 2010, it has been rational so far to expect, on average, social progress in surviving societies over the long term, even as exceptions always exist, and sometimes blind us to the long-term trend. Steven Pinker, in *The Better Angels of Our Nature*, 2011, makes an even more evidence-based claim with respect to the long-term decline in violence frequency and severity in human society, and the increasing subtlety and sophistication of human ethics. One injunction

that seems necessary for all ethical societies will be the *voluntary and reversible nature* of all copying or reanimation when dealing with conscious organisms. If such procedures are voluntary, and if the person undergoing either of them claimed to be essentially the same or improved in some way at the other end, many of us might one day do them as well, to reap their benefits. We would not consider people uploaded from preserved brains to be "zombies" (fake copies, not "real") and we would not view these technologies as violent or immoral, as long as all of those using them claimed to be real, used the technologies *by choice*, and some degree of *reversibility* (even if it was not a perfect or an inexpensive reversibility) was available to those who decide they do not prefer their new state. In practice however, even a less-than-perfect uploading of a preserved human into a virtual world might be desirable, particularly if the original pattern (the preserved biological tissue) was still available for future use. Any previously biological human not appreciating the benefits of their new digital form, and not willing to live with any drawbacks (such as, for example, some memory loss or other deficits), would ideally have the ability to shut down and suspend their life further, and await the arrival of better revival technology. Such reversible, voluntary, and suspendable uploading scenarios may be reasonably expected in future society if humanity's moral development must also improve as a function of our collective intelligence, as several scholars (Norbert Elias, Ron Inglehart, Robert Wright, Matt Ridley, Steven Pinker, etc.) have proposed.

Ninth, the patternist perspective leads us to anticipate some of the unusual mental capabilities that our future selves and societies may one day possess, and these can seem so strange or unsettling that we may reject them intuitively, or decide they belong to a world that has no relation to our own. Consider the following thought experiment. Imagine that you have the ability to reanimate a true copy of yourself using advanced brain scanning and simulation technology. Notice now that this allows you the ability to create *many* true copies. Recall for example the "duplicate" humans that were occasionally created in the transporter in the *Star Trek* science fiction series. If two copies of yourself were uploaded, and you found yourself in a room with your exact copy, there would, at that moment, simply be *two* self-aware versions of you in that room, no matter how counterintuitive to some that this may seem. Just as biology can today make genetically identical twins, technology will one day be able to make mentally identical twins (triplets, quintuplets, etc.) of individual minds, as strange as this seems. Of course, these twinned selves would begin to

diverge from each other the moment they were created, as they would begin to have different subjective experiences. But at the start they would simply be two identical, true copies of "you." If this duplication process wasn't too costly or difficult, we can also imagine that our future selves might engage in such mental "forking" on a regular basis, to generate two or more slightly different personal perspectives on complex and subtle problems. We might also *reintegrate* (merge) these separate selves later, after the problem was solved or no longer relevant. Science fiction authors like Phillip Jennings, *The Bug Life Chronicles*, 1989 and Charles Stross, *Accelerando*, 2006, are among those who have described this strange idea. We can imagine this future ability as a natural extension of the way we presently argue with ourselves, using slightly different yet largely similar neural structures within our own brain, whenever we are "mentally split" over the course of action on a difficult problem. In fact, we must admit that any human being today is already a *Society of Mind*, a collection of somewhat independent and arguing "mindsets," as Marvin Minsky observed in 1987. We might reintegrate these twinned minds/selves eventually, after some period of exploration and experimentation, and such a process, while it might involve the elimination of less adapted mental structures in the process of reintegration, would very likely be seen as growth, not death. We can understand this in the same way that, after long arguments within our own mind today, one set of synaptic structures may end up prevailing, and one or more of the less-fit synaptic structures end up dying. In this process, the less-fit connections end up being reweighted, in a way that involves *effective information destruction* in the network within our own brains, as the less adaptive behaviors, once ignored long enough, attenuate to extinction. To a healthy and mentally integrated self, this kind of information loss feels simply like creativity and growth, not death. So too we can forsee how a future technological self, which has the ability to make multiple copies, backups, and "instances" of itself, would be a system in which "little deaths" were constantly occurring, but in which deep resiliency, continual learning and growth, indefinite lifespan, and substantially less fearfulness and stress over the consequences of conflict would also be achieved. The inevitable competitions and deaths in such a future should feel far less subjectively violent, and involve far less informational destructiveness, than the world we live in today.

B.10 Challenges for the Future

At present, roughly 57 million unique and precious human beings die every year, or 155,000 people every day. It is hard for us to comprehend the

scale of this catastrophic loss of human experience. Thus today we largely avert our minds from this unparallelled daily loss of diversity, wisdom, social history, and individual life, except on those occasions when it touches us personally. Meanwhile, medical science makes slow progress in preventing biological death and extending our health and lifespan. Fortunately, technology is accelerating in its ability to record and augment our lives, and now the preservation and later revival of human memory and identity appear on the verge of scientific reality.

By advancing the appropriate sciences and technologies we can accelerate the arrival of the brain preservation choice for all of us, and end the tyranny of an unchosen death. Given historical rates of accelerating scientific and technological change, it is even reasonable to expect reanimation technologies to be available not centuries from now, but possibly even within this century, while our loved ones are still alive. Furthermore, all of our friends and loved ones who have also chosen preservation will also return to interact with us. For many, this is one of the most important personal motivations for preservation, the likelihood that one's individual pattern may remain useful to those we know today, and remain connected to and supportive of the social community from which it emerged. Once we understand and have internalized the implications of accelerating change on our science, technology, and economy, we can recognize how extraordinary the human future will be, and by direct extension, how extraordinary and opportunity-filled our own lives are here today.

As we consider our extraordinary present and future, each of us has the ability, regardless of our honorable religious, philosophical, or cultural backgrounds, to internalize the implications of accelerating technological change, to consider some version of the patternist hypothesis of self, to champion scientific and technological progress and evidence-based inquiry, and to gently reform our esteemed religious, philosophical, and cultural communities of heritage until they are in better alignment with apparent evidence and scientific truths.