Vitrifying the Connectomic Self

A case for developing Aldehyde Stabilized Cryopreservation into a medical procedure

Kenneth Hayworth (February 2018)

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 at -130^{o}C, 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.

[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).

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 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

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

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

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 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

[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

[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)

and Go at superhuman levels[10]. When your job is to build applications based on artificial brains it becomes easier to imagine yourself 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.

[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

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 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

[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 (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.

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 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

[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.

[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.

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].

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].

---

[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 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 prepared for electron microscopy but they do not use perfusion to do so. Other

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 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

---

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.

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

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 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

[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).

[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.

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 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

[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).
[35] E.g. Wilson & Wearing 1995

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 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 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

[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.

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]

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.

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.

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)

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.

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 support 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.

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 mm^3 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.