

Note: This is an author's draft. The published article can be found at <https://www.skeptic.com/magazine/archives/21.2/>)

An argument for the scientific and technical plausibility of mind uploading

By Kenneth Hayworth (December 2015)

Peter Kassan's article in this issue of Skeptic argues that the idea of mind uploading is "...science fantasy, based on a misunderstanding both of the overwhelming complexity (and our near-total ignorance) of the brain, and of what computer models are...". Do any real neuroscientists believe that mind uploading might be possible? Kassan's article mentions one that does –me. So I have been given the honor to write this rebuttal.

Any discussion regarding mind uploading must be about what can reasonably be assumed possible in the distant future, not what is achievable today. I am certainly not arguing uploading is easy, or that it will occur within the next few decades; but I will argue it is a technically achievable, potentially desirable, long-term goal. I will present evidence that current neuroscience models support the possibility. I will cover recent developments in electron microscopy that hint at the technology needed. I will touch on cognitive models which directly support the mind-as-computation hypothesis, and I will delve deep into the consciousness debate. Finally, I will discuss a recently developed method for long-term brain preservation that seems sufficient to support future mind uploading, a fact that makes this discussion not just 'academic'.

Possible vs. Impossible

Hopefully we can all agree that it is physically possible to one day colonize the planet Mars with vibrant, self-sustaining encapsulated cities. And we can also all agree that such colonization would be incredibly difficult, requiring enormous resources and advancements. If colonization were ever to occur it is reasonable to assume it would take centuries. Reasonable people can disagree on whether Mars colonization is even a desirable long-term goal. They can also disagree on whether the first small self-sustaining colony will be achieved by the year 2050, 2150, or 3050. And they can certainly disagree on how best to prioritize today's resources with respect to that goal. But it would be unreasonable to level (unsupported) statements that such colonization is physically impossible, especially in the light of our successful 'baby steps' toward that goal (e.g. landing men briefly on our moon). And it would certainly be unreasonable to ridicule the scientists and engineers (e.g. at NASA and SpaceX) who, motivated by such lofty long-term goals, have decided to devote their lives to tackling some of its technical obstacles.

Kassan's argument is analogously structured, asserting that mind uploading is theoretically impossible but backing up his assertion only by pointing out its great technical challenge. His argument ignores the significant 'baby steps' that have already been made (e.g. automated, reliable methods that scan neural tissue at the nanometer scale). And ignores the fact that all of today's neuroscience models are fundamentally computational in nature, supporting the theoretical possibility of mind uploading. Kassan declares "our near-total ignorance of the brain", a quite

inflammatory statement which I find perplexing given neuroscience's enormous advances. A small counter example: a flurry of papers have recently tested decades old computational models of how memories are formed by genetically tagging (in mouse) only those neurons active during a fear memory formation. Optogenetic reactivation of those same 'engram' cells was sufficient to recall the memory¹, and to 'incept' a false memory², and have confirmed long-suspected aspects of the synaptic connectivity underlying memories³.

Is it possible to image synaptic connectivity?

Kassan's article points out the obvious: our current technology is not sufficient to scan an entire brain at the required resolution. This is equivalent to arguing that Mars colonization is forever impossible because we don't currently have the rockets. Kassan actually goes so far as to claim something as "impossible" that I and other scientists have been doing at a small scale for years. He states: "A proposed (but never implemented) method of gathering the information involves preserving your brain in some kind of plastic and then sectioning it... This is unlikely to work. However thin each slice is, it would be impossible to section your brain without destroying a countless number of synapses."

"Impossible" is a strong word. As far back as 1958 electron microscopists have been embedding brain tissue, sectioning and imaging it to create 3D models of every synapse⁴. Recently 3D electron microscopy (3DEM) has undergone a revolution⁵. It is no longer necessary to collect fragile sections, instead a technique called Serial Block Face Scanning EM is much more reliable and automated⁶. The surface of a plastic-embedded sample is imaged with a scanning electron microscope (SEM), then a 25nm layer is automatically scraped off with a diamond knife, and the fresh surface is again imaged. Repeating this cycle thousands of times creates 3D images of neural circuits^{7,8}.

Another technique, Focused Ion Beam SEM (FIB-SEM), dispenses with physical cutting altogether, replacing it with ion milling⁹, which works somewhat like sand blasting but with individual atoms. Even a relatively broadly focused beam can remove layers just 5nm thick, and is easily kept 'sharp' and correctly positioned with voltage adjustments. In our lab, we have developed custom FIB-SEMs that scan brain tissue at 10x10x10nm¹⁰, running automatically 24 hours a day for months.

A single FIB-SEM could never, by itself, image an entire brain. Is it possible to divide the task across many FIB-SEMs? In a recent paper¹¹ I showed that such a strategy is plausible, demonstrating reliable, near-lossless sectioning of plastic-embedded brain tissue into 20 micron thick slabs which were imaged separately and computationally stitched back together for tracing.

Is it possible to stain and embed an entire brain for 3DEM? A recent paper¹² by Shawn Mikula and Winfried Denk shows that it is, at least for a mouse brain. They developed their protocol in support of their larger goal of imaging an entire mouse brain at the synapse level. Their plan involves modifying the fastest SEM in the world which uses 91 beams scanning simultaneously¹³. Even so, it is estimated to require years to image. A human brain would have to be divided and

imaged in parallel. I estimate it would require roughly 3,000 of these ultrafast SEMs 10 years to image a single human brain.

Yes, such a project would be incredibly expensive using today's technology. Yes, it is laughable to propose starting such a project today (no one is). But how can Kassan justify saying it is "impossible", especially since there is every reason to expect brain imaging technologies will continue to improve?

Are connections enough? How does the brain work anyway?

So far I have argued that it will eventually be possible to scan a whole brain at 10nm resolution. Why 10nm? Because at that resolution one can unambiguously trace the brain's connections¹⁴. But are connections enough?

Obviously we do not yet have a complete theory of the brain, but decades of research have narrowed in on its outlines. I will attempt to cover its basics in order to show how memory and cognition are thought to be encoded by connectivity.

Perhaps the most complete cognitive-level architecture of the human mind today is ACT-R which has served as the basis of hundreds of publications across dozens of labs¹⁵. There are ACT-R models of problem solving, attention, language, etc. These models perform on the same experimental tasks that human subject do, and make testable predictions on the fine scale of reaction times, learning rates, fMRI activation patterns, etc. ACT-R's features are motivated and constrained by neuroscience, thus ACT-R functions as a bridge between high-level cognitive models and neural models.

ACT-R models are written in a high-level symbolic format. It assumes that the brain represents the world by tokening symbols in a global workspace of cortical memory buffers. Cognitive processes are modeled as a set of if-then rules called 'productions' whose if-clauses are all simultaneously pattern-matched against the memory buffers every 50ms. When a match occurs, that production's then-clause 'fires' resulting in particular manipulations of the buffers. In addition, ACT-R posits a set of perceptual and sensorimotor modules (allowing it to control a virtual body of sorts) and a declarative memory module (allowing storage and retrieval of structured memories called 'chunks'). ACT-R models are centered on particular sets of 'productions' and 'chunks' which are so crucial that they have aptly been called "the atomic components of thought". There is a detailed theory of learning in ACT-R which describes how such 'chunks' and 'productions' are created (and weighted) during cognition.

There, in a nutshell, is our best theory of the cognitive-level architecture of the human mind. Like the "standard model" in physics, it requires a lot of experience and practice to understand its implications, and it is certainly not 'complete' or universally accepted in its particular details. But it stands as a summary model of what we think we know. I have often been confronted by people, like Kassan, who casually state we have a "near-total ignorance" of how the brain gives rise to intelligence. A reasonable question to ask is: "Have you studied cognitive models like ACT-R

(and its competitors) and reviewed their successes and failures carefully?" If not, may I suggest the excellent book "How can the human mind occur in the physical universe?" which summarizes decades of cognitive science research through the lens of ACT-R models¹⁶.

But how does ACT-R map onto the actual neural circuits of the brain? ACT-R's perceptual and sensorimotor modules are simply high-level abstractions of existing neural models^{17,18}. ACT-R's declarative memory module is a high-level abstraction of existing models of hippocampal and temporal lobe circuits¹⁹. And ACT-R's production module is based on the basal ganglia²⁰, including how 'procedural rules' that are first learned in the basal ganglia are later stabilized as direct cortical circuits²¹. There are fantastic books^{22,23} covering the full range of such biologically-realistic 'connectionist' models. Of course many questions remain at all levels of this story. One I am particularly interested in is how the basal ganglia's pattern recognition circuits might be sensitive to and manipulate syntax in the way proposed by ACT-R^{24,25}. But the remaining questions should not distract us from the fact that a coherent, satisfying theory of how cognition works has begun to arise as these references attest.

Getting back to the question "Are connections enough?", each of the above neural models argues that computation and learning are encoded mainly in the pattern of glutamatergic neural projections onto the 'spiny' principal neurons of each brain region. There is an enormous body of literature on such 'spiny' synapses, specifically how they are the main site of the long term potentiation / depression (LTP / LTD) underlying learning. A recent review²⁶ summarizes decades of evidence for this. Other papers^{27,28,29} show that a synapse's size and other SEM-visible features are strongly correlated with its functional strength.

Such 'connectionist' models have long been the standard in neuroscience, but today's tools allow us to test these assumptions more directly than ever. One great example is a recent paper titled "Labelling and optical erasure of synaptic memory traces in the motor cortex"³⁰. The authors developed a photoactivatable form of the Rac1 protein and inserted its DNA (into mice) in such a way that it would express only in the dendritic spines of recently enlarged synapses, allowing them to literally see those synapses that encoded a new memory. And they were able to shrink only those same synapses using a flash of light. The result was as predicted: an erasure of the learning seen by decreased performance in the experimental task relative to a control task. This experiment and many others are rigorously testing 'connectionist' theories, and so far they seem to be holding up.

Could a simulation be conscious?

Kassan states: "The notion that a computer model of your brain might be an adequate substitute or replacement for your real brain is a profound misunderstanding ... The computer model can never substitute for the entity being modeled." A simple thought experiment can address this concern: "Could a computer simulation of the human retina be a substitute for the real thing?"

The retina consists of layers of neurons which eventually feed into a final layer of ganglion cells, the axons of which form the optic nerve. Decades of research have revealed the

transformation that is being computed –each ganglion cell’s axon transmits a center-surround spatial (and temporal) filtering of part of the visual image. Given today’s miniaturized cameras and microprocessors it would be trivial to create an artificial retina “simulation”, unfortunately it is much more difficult to tie the outputs of such a simulation into the surviving optic nerve of a patient blinded by retinal disease.

Technical difficulties aside, assume a simulation is made based on recordings of a person’s functioning retina. Following that, most of the person’s retina is removed and replaced with a camera and this simulation, its spiking output tied precisely into the surviving ganglion cells. Would the person be able to “see” with this simulated retina? Every neuroscience experiment and model that I know of says they would, and several research groups have spent decades working through difficult interface challenges to make such retinal prostheses a reality.

This is clearly a counter example to Kassan’s claim that a “computer model can never substitute for the entity being modeled”, and it is clear why. If the function of the thing being modeled is simply to process information then an accurate simulation is just as good as the original. A simulation of the weather will not get one wet, but a simulation of a pocket calculator is just as helpful balancing one’s expenses.

If my retina can be replaced by a simulation then couldn’t also my visual cortex? The simulation would be more complex and would have to deal with bidirectional connections, but the principle is the same. If it was wired to precisely the same input and output axons as the original, and it was fed a simulated view of a sunset, then how would, for example, the language centers of my brain know the difference? Would they not respond that “I” am seeing a sunset?

This is the slippery slope of materialism. If the brain’s functioning is governed by the causal laws of physics then any subset of the brain’s neurons should in principle be replaceable by a computer simulation of those same neurons hooked up to the rest with electrodes. As long as the causal relationships are maintained then the outward behavior of the person must remain the same, even to the extent of verbally claiming to have the same conscious experiences.

One might argue that “peripheral” regions could be simulated but core “conscious” regions could not. But neuroscience has found no such distinction. There is no principled reason, for example, why the memory functions of the temporal lobe could not be replaced by a suitable simulation, in fact there is some limited research success toward that goal³¹. And there is no principled reason why such a simulation could not drive emotional responses, like the optogenetic experiments I mentioned earlier did. Based on all existing evidence, there is no ‘magic’ subset of neurons which could not, in principle, be replaced by a suitable simulation leaving both the outward behavior and internal conscious experience of the person unaffected. And if any subset can be replaced then why not the entire brain?

We have slid all the way down the slippery slope and found that the materialist assumptions underlying neuroscience seemingly force us to accept that a simulated brain would be just as conscious as a biological one. I have met many people who refuse to accept this conclusion arguing that there must be something in a biological brain that could not be replaced by simulation.

Whereas this is in principle a possibility, I see no neuroscience evidence that motivates it, only the unfounded belief that consciousness is so special that a mere computer simulation could never replicate it. In this sense it is similar (perhaps identical?) to a belief in an intelligent designer based only on a gut intuition that natural selection could never be sufficient to create man.

So what is consciousness then? Well, by definition, an agent is conscious if it is 'like something to be' that agent. Any explanation of the 'internal' aspects of consciousness should be intimately tied to those physical processes in the brain that drive the outward behaviors we associate with consciousness. This inexorably leads to positing an internal "self-model" –a set of data structures that summarize an agent's perceptions, affective states, and action decisions as happening to, and coming from, a central "I". Who uses this self-model? It is used by the rest of the agent's information processing system to help intelligently guide future behavior. In a nutshell this is cognitive science's default theory of consciousness. It goes by various names but a particularly clear description of it, the "phenomenal self-model (PSM)" has been articulated by Thomas Metzinger³². It, at least in principle, explains all of the outward aspects of conscious behavior. For example, an injured agent containing a PSM would act like it is conscious of pain because this 'fact' would be explicitly represented in its PSM, not just causing withdrawal but also appropriate modifications of future actions and perhaps verbal reports explaining how the pain feels.

Such a PSM model might explain the outward signs of consciousness, but does it explain the internal ones? I think the answer is yes. It is in fact 'like something to be' an information processing system that has a PSM –it is precisely like what the PSM represents it to be like as interpreted by the rest of the system. If a PSM representation of pain is interpreted by the rest of the system as a highest-priority goal to stop the source of the pain, one that distracts from all other goals, then that is one aspect of what it is 'like' to consciously feel pain. And during a painful experience, this aspect of what it is like to experience pain will itself be recorded in the PSM.

Let me briefly address a few questions regarding consciousness that I am sure many readers will be wondering about:

Q: Where in the brain is the PSM?

A: Using ACT-R terminology, the PSM would comprise the declarative memory 'chunks' used to represent our self-model, and those 'productions' used to record, process, and interpret them. According to neural models, these 'chunks' and 'productions' are encoded in circuits spanning large regions of the cortex, hippocampus, and basal ganglia.

Q: Precisely when do we become conscious of a stimulus? Is it when it first appears in the cortex's global workspace? Is it when it is incorporated into our PSM?

A: There is no definite answer! An agent with a PSM behaves consciously, and any stimulus (S1) that is incorporated into its PSM may eventually impact behavior so significantly that we would clearly say that the agent was conscious of S1. But it is also possible that the trace of S1 may be erased before subsequent ramifications. This realization, that it is a meaningless question to ask precisely when a stimulus becomes conscious, is the brilliant insight offered by Daniel Dennett³³. I consider Dennett's willingness to question, and eventually discard, this previously unquestioned

assumption to be as essential to a scientific theory of consciousness as was Einstein's willingness to question, and eventually discard, the assumption of universal simultaneity. Anyone struggling to understand consciousness would do well to (re)read Dennett's expositions on this and ponder the ramifications.

Q: Even if a person was simulated perfectly wouldn't it be 'just a copy'?

A: If your hard drive crashed erasing a program you had worked many weeks on, it would be a mild tragedy. But if you had a backup copy it would be no tragedy at all. A copy of a program is that program, period. All of our current theories of the human mind are computational and imply that we are like a program in this sense—we can in principle be copied and can have many 'instantiations' running simultaneously. This doesn't even raise serious philosophical issues as Hollywood movies like "The Sixth Day" show, i.e. there are two Arnold Schwarzeneggers, they have the same memories of their life before copying but lead separate lives after -no big deal.

Not an academic discussion

Most people don't spend time worrying about Mars colonization because even if it does occur they will have died long before. A 'proper' skeptical attitude toward mind uploading might be similar: "Yes, as a materialist that has kept up-to-date in cognitive and neuroscience, I accept that mind uploading seems possible in principle, but it is so fantastically complicated in practice that it will not occur in my lifetime." I consider this a very reasonable attitude to take: open minded enough to embrace science in all of its ramifications, but skeptically grounded enough to recognize that anyone claiming that widespread mind uploading will be on offer in the next few decades is either mistaken or lying.

Having said that, I am about to claim something that should ring skeptical alarm bells: An inexpensive cryogenic brain preservation procedure that seems fully compatible with future human mind uploading has recently been demonstrated in an animal model. The implication: widespread medical implementation of this procedure could provide everyone alive today with a bridge to future mind uploading technologies even if those technologies require many centuries to perfect.

Not your father's cryonics

In a recent (open-access) paper³⁴ in the Journal of Cryobiology, researchers Robert McIntyre and Greg Fahy describe a new brain banking procedure called Aldehyde-Stabilized Cryopreservation (ASC). In ASC, an anesthetized animal's carotid arteries are cannulated and its brain perfused with glutaraldehyde fixative. After 45 minutes of perfusion, pumps begin to replace the solution's water with a cryoprotectant agent (CPA). Over the next 4 hours CPA concentration is slowly increased (at room temperature) to a final value of 65%, allowing the brain to then be lowered down to a storage temperature of -135°C, a temperature so low that the brain vitrifies solid without ice crystal formation. Time has essentially stopped for such a brain. They showed that such a brain can be rewarmed, and re-perfused to remove CPA. The result is an intact glutaraldehyde fixed brain with no visible macroscopic defects (e.g. no cracks). Samples taken from across such brains were processed for electron microscopy and showed textbook quality ultrastructure of synapses.

The Brain Preservation Foundation, a non-profit I founded five years ago to advance research and to skeptically challenge the claims of cryonics practitioners, helped fund part of this ASC research and is independently evaluating their claims. I have personally witnessed the entire surgical and storage procedure, and I have performed extensive 2D SEM and 3D FIB-SEM imaging of samples from several brains preserved via their technique. So far I have found that preservation of neuronal circuitry appears uniformly excellent across the entire brain. **Figure 1** shows 3DEM images I acquired from a rabbit brain preserved via ASC and stored at -135°C .




Figure 1. Cryopreserved brains. **(A)** ASC preserved rabbit brain vitrified solid at -135°C . **(B)** ASC preserved pig brain rewarmed after being stored at -135°C . Note the 'goo' covering it is melted CPA. **(C)** Electron micrograph of cortex sample taken from the rabbit brain shown in **(A)**. Synaptic details are clearly visible and well preserved. Scale bar is 1 micron. **(D)** $10 \times 10 \times 8$ micron FIB-SEM volume acquired at 8nm resolution taken from the rabbit brain shown in **(A)**. Neural processes and connections are easily traced in this volume (similar FIB-SEM videos are available online as part of the ASC publication).

It is impossible to overemphasize how much better the ultrastructural preservation is in these ASC brains compared with anything previously presented by cryonics researchers. The reason why is clear: ASC starts by perfusing one of the most deadly and aggressive fixatives known – glutaraldehyde. This is anathema to many cryonics advocates who cling to the hope of biological revival. However if one's sights are set on future uploading then structural preservation of synaptic connectivity is the higher priority. Perfusion with glutaraldehyde almost instantly stops metabolic decay and fixes all proteins in place by covalent crosslinks. This stabilizes the tissue and vasculature so that CPA perfusion can be performed at an optimal temperature and rate. The result is an intact brain that can be stored unchanging for millennia if necessary, and whose neural connectivity is preserved as well as fixed-only control brains. In fact, there is every reason to expect that molecular-level details (e.g. receptor proteins and ion channels) are also well preserved by ASC since glutaraldehyde fixation mainly locks such proteins in place preserving their primary structure and, in many cases, their coarse tertiary structure as well³⁵, a fact verified recently by correlated electron and immunofluorescence microscopy of glutaraldehyde-fixed brain tissue³⁶.

Conclusion

I have presented evidence supporting the idea that human mind uploading is a technically achievable goal, albeit one that may take centuries to realize. And I have presented evidence that a potentially inexpensive and reliable preservation technique (ASC) exists already that could, if professionally implemented by the medical establishment, allow everyone alive today to reach that future mind uploading technology by 'hitting pause' for decades in cryogenic storage. I do not expect this short article to have fully convinced anyone, nor should it. These are extraordinary claims that should engender skeptical debate and inquiry into their scientific details and assumptions. The references cited should provide a good starting point for such inquiry.

¹ Liu, X., et al. "Optogenetic stimulation of a hippocampal engram activates fear memory recall." *Nature* 484.7394 (2012): 381-385. [\[LINK\]](#)

² Liu, X., Ramirez, S., and Tonegawa, S. "Inception of a false memory by optogenetic manipulation of a hippocampal memory engram." *Philosophical Transactions of the Royal Society B: Biological Sciences* 369.1633 (2014): 20130142. [\[LINK\]](#)

³ Ryan, T. J., et al. "Engram cells retain memory under retrograde amnesia." *Science* 348.6238 (2015): 1007-1013. [\[LINK\]](#)

⁴ Sjöstrand, F. S. "Ultrastructure of retinal rod synapses of the guinea pig eye as revealed by three-dimensional reconstructions from serial sections." *Journal of ultrastructure research* 2.1 (1958): 122-170. [\[LINK\]](#)

⁵ Briggman, K. L., and D. D. Bock. "Volume electron microscopy for neuronal circuit reconstruction." *Current opinion in neurobiology* 22.1 (2012): 154-161. [\[LINK\]](#)

⁶ Denk, W., and Horstmann H.. "Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure." *PLoS Biol* 2.11 (2004): e329. [\[LINK\]](#)

⁷ Briggman, K. L., Helmstaedter M., and Denk W. "Wiring specificity in the direction-selectivity circuit of the retina." *Nature* 471.7337 (2011): 183-188. [\[LINK\]](#)

⁸ Kim, J. S., et al. "Space-time wiring specificity supports direction selectivity in the retina." *Nature* 509.7500 (2014): 331-336. [\[LINK\]](#)

⁹ Knott, G., et al. "Serial section scanning electron microscopy of adult brain tissue using focused ion beam milling." *The Journal of Neuroscience* 28.12 (2008): 2959-2964. [\[LINK\]](#)

-
- ¹⁰ Takemura, S., et al. "Synaptic circuits and their variations within different columns in the visual system of *Drosophila*." *Proceedings of the National Academy of Sciences* 112.44 (2015): 13711-13716. [\[LINK\]](#)
- ¹¹ Hayworth, K. J., et al. "Ultrastructurally smooth thick partitioning and volume stitching for large-scale connectomics." *Nature methods* 12.4 (2015): 319-322. [\[LINK\]](#)
- ¹² Mikula, S., and Denk W. "High-resolution whole-brain staining for electron microscopic circuit reconstruction." *Nature methods* (2015). [\[LINK\]](#)
- ¹³ Eberle, A. L., et al. "High-resolution, high-throughput imaging with a multibeam scanning electron microscope." *Journal of microscopy* (2015). [\[LINK\]](#)
- ¹⁴ Helmstaedter, M. "Cellular-resolution connectomics: challenges of dense neural circuit reconstruction." *Nature methods* 10.6 (2013): 501-507. [\[LINK\]](#)
- ¹⁵ <http://act-r.psy.cmu.edu/publication/>
- ¹⁶ Anderson, J. R. *How can the human mind occur in the physical universe?* Oxford University Press, 2007. [\[LINK\]](#)
- ¹⁷ O'Reilly, R. C., et al. "Recurrent processing during object recognition." *Frontiers in psychology* 4 (2013). [\[LINK\]](#)
- ¹⁸ Pouget, A., and Snyder L. H. "Computational approaches to sensorimotor transformations." *Nature neuroscience* 3 (2000): 1192-1198. [\[LINK\]](#)
- ¹⁹ Rolls, E. T., and Kesner R. P. "A computational theory of hippocampal function, and empirical tests of the theory." *Progress in neurobiology* 79.1 (2006): 1-48. [\[LINK\]](#)
- ²⁰ Redgrave, P., Prescott, T J., and Gurney, K. "The basal ganglia: a vertebrate solution to the selection problem?" *Neuroscience* 89.4 (1999): 1009-1023. [\[LINK\]](#)
- ²¹ Ashby, F. G., Turner, B. O., and Horvitz, J. C. "Cortical and basal ganglia contributions to habit learning and automaticity." *Trends in cognitive sciences* 14.5 (2010): 208-215. [\[LINK\]](#)
- ²² Rolls, E. T., Treves, A. *Neural networks and brain function*. Oxford: Oxford university press, 1998. [\[LINK\]](#)
- ²³ O'Reilly, et al. *Computational Cognitive Neuroscience*. 2012. URL: <http://ccnbook.colorado.edu>
- ²⁴ Hayworth, K. J. "Dynamically partitionable autoassociative networks as a solution to the neural binding problem." *Frontiers in computational neuroscience* 6 (2012). [\[LINK\]](#)
- ²⁵ Stocco, A., Lebiere, C., and Anderson, J. R. "Conditional routing of information to the cortex: A model of the basal ganglia's role in cognitive coordination." *Psychological review* 117.2 (2010): 541. [\[LINK\]](#)
- ²⁶ Bailey, C. H., Kandel, E. R., and Harris, K. M. "Structural components of synaptic plasticity and memory consolidation." *Cold Spring Harbor perspectives in biology* 7.7 (2015): a021758. [\[LINK\]](#)
- ²⁷ Bourne, J., and Harris, K. M. "Do thin spines learn to be mushroom spines that remember?" *Current opinion in neurobiology* 17.3 (2007): 381-386. [\[LINK\]](#)
- ²⁸ Kasai, H., et al. "Structure–stability–function relationships of dendritic spines." *Trends in neurosciences* 26.7 (2003): 360-368. [\[LINK\]](#)
- ²⁹ Matsuzaki, M., et al. "Structural basis of long-term potentiation in single dendritic spines." *Nature* 429.6993 (2004): 761-766. [\[LINK\]](#)
- ³⁰ Hayashi-Takagi, A., et al. "Labelling and optical erasure of synaptic memory traces in the motor cortex." *Nature* (2015). [\[LINK\]](#)
- ³¹ Berger, T. W., et al. "A cortical neural prosthesis for restoring and enhancing memory." *Journal of Neural Engineering* 8.4 (2011): 046017. [\[LINK\]](#)
- ³² Metzinger, T. *Being no one: The self-model theory of subjectivity*. MIT Press, 2004. [\[LINK\]](#)
- ³³ Dennett, D. "Are we explaining consciousness yet?." *Cognition* 79.1 (2001): 221-237. [\[LINK\]](#)
- ³⁴ McIntyre, R. L., and Fahy, G. M. "Aldehyde-stabilized cryopreservation." *Cryobiology* (2015). [\[LINK\]](#)
- ³⁵ Hayat, M. A. *Principles and techniques of scanning electron microscopy. Biological applications*. Fourth Ed. Cambridge University Press, 2000. [\[LINK\]](#)
- ³⁶ Collman, F., et al. "Mapping Synapses by Conjugate Light-Electron Array Tomography." *The Journal of Neuroscience* 35.14 (2015): 5792-5807. [\[LINK\]](#)