World Scientific
www.worldscientific.com

ELECTRON IMAGING TECHNOLOGY FOR WHOLE BRAIN NEURAL CIRCUIT MAPPING

KENNETH J. HAYWORTH

Center for Brain Sciences, Harvard University,
52 Oxford St. Northwest Bldg. 347.20, Cambridge, MA 02138, USA
hayworth@fas.harvard.edu

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

Keywords: Connectome; mind uploading; FIBSEM.

1. Introduction

On May 25, 1961 John F. Kennedy delivered his now famous speech to a joint session of congress stating "... I believe that this nation should commit itself to achieving the goal, before this decade is out, of landing a man on the moon and returning him safely to the earth". This was an incredibly ambitious goal especially given that the only American space flight prior had been Alan Shepard's sub-orbital flight aboard a Redstone rocket less than a month before the speech. However Kennedy's science advisors were reasonably sure that the difficulty lay mainly in great engineering challenges — much more powerful versions of the liquid fueled Redstone rocket, more sophisticated control and guidance electronics — and did not require radically new science or technologies to achieve. Today we may be in a similar situation relative to the goal of uploading a human mind into a computer. It is unlikely that radically new science and technologies are needed to achieve the goal of "landing a

1   man in cyberspace and returning him safely to consciousness". I will argue that the
2   key technology that needs to be developed to make mind uploading a reality is the
3   ability to create a 3D volume electron microscope image of an entire chemically-fixed
4   and plastic-embedded brain at sufficient resolution to allow the tracing of synaptic
5   connections between all neurons. As in the Apollo program, such a feat represents an
6   enormous engineering challenge but the basic technology for 3D electron imaging of
7   brain tissue at this resolution exists today, and I will lay out one possible path for how
8   this imaging technique could be robustly scaled to the size of an entire human brain.
9   In this article I first review today's most well supported cognitive and neuroscience
10 model of the human cognitive architecture (ACT-R), and show how it implies that a
11 successful mind uploading could be accomplished based only on a static map of all
12 neural connections in a brain (along with sufficient structural information on synapse
13 size, location, cell morphology, etc. to allow correlation with functional experiments
14 performed on other brains). Such a static map of all neural connections has been
15 termed the "Connectome" [Seung, 2012] and a new field called "Connectomics" has
16 sprung up over the last few years [Eisenstein, 2009] developing automated electron
17 microscopy techniques for the 3D imaging of brain tissue at resolutions sufficient to
18 reliably map such connections over long distances.
19   I review one of these techniques in particular, the Focused Ion Beam Scanning
20 Electron Microscopy (FIBSEM) technique [Knott *et al.*, 2008], which has demon-
21 strated the ability to image tissue at resolutions approaching $5 \times 5 \times 10$ nm. The
22 finest axonal and dendritic process in the brain can shrink down to ~ 40 nm in
23 diameter and researchers using traditional electron microscopy techniques based on
24 50 nm thick serial sections have found it difficult to reliably trace such processes over
25 long distances using such thick sections. In contrast, the new FIBSEM images have
26 such high resolution that all neuronal processes should be reliably traceable even
27 using computer automated algorithms. Furthermore, FIBSEM's use of a focused ion
28 beam (as opposed to traditional physical sectioning with a diamond knife) gives it the
29 potential to achieve the reliability and large-scale automation that would be
30 necessary for a human mind uploading attempt. In these regards, FIBSEM represents
31 a truly enabling technology for mind uploading. Continuing with the rocket science
32 analogy, we might say that today's FIBSEM resolution demonstration is equivalent
33 to the first demonstration of a liquid fueled rocket engine achieving a specific impulse
34 theoretically sufficient to achieve orbital velocity.
35   A fundamental physical limitation of the FIBSEM ablation process is that this
36 high resolution can only be obtained for very small samples on the order of 20 microns
37 across. To overcome this limitation I have developed a technique [Hayworth *et al.*,
38 2010] which can section large blocks of plastic-embedded brain tissue into 20 micron
39 thick strips optimally sized for high-resolution FIBSEM imaging. This thick sec-
40 tioning procedure results in such high-quality surfaces that the finest neuronal pro-
41 cesses can be traced from strip to strip. This lossless thick-sectioning technique opens
42 up the possibility of applying the high resolution of FIBSEM to arbitrarily large
    tissue volumes (potentially even whole brains), and allows for the use of massively

1   parallel imaging by spreading the thick strips across multiple FIBSEM imaging
2   machines.

3   Sectioning an entire human brain into 20 micron thick strips with 100% reliability
4   will be a difficult engineering challenge, but one that appears quite achievable. Below
5   I discuss the technical challenges in preserving an entire human brain at the ultra-
6   structure level suitable for electron microscopic imaging. I then sketch out an
7   industrial process which could reduce such a chemically-fixed and plastic-embedded
8   brain into a set of tapes containing 20 micron $\times$ 20 micron tissue pillars optimally
9   sized for automated high-resolution FIBSEM imaging. Finally I discuss a set of
10   technical milestones (which I believe are reachable in the next few years) which would
11   demonstrate all of the crucial imaging steps necessary for uploading a human mind.

2. Evidence that We are Our Connectome

A central thesis of this paper is that an electron microscope volume image of an entire
chemically-fixed and plastic-embedded human brain (with sufficient resolution to
allow the tracing of synaptic connections between all neurons) would contain suffi-
cient information to fully encode the unique memories, skills, and personality of that
particular individual. Further, given this static image (and a thorough understanding
of the functioning of individual neuron types gleaned from side experiments not
performed on this particular brain) it is hypothesized that a computer simulation
could be constructed that when turned on would behave so similarly to the original
that close friends and relatives would be hard pressed to distinguish (from behavior
alone) this simulation from the original. Finally, and crucially, it is hypothesized that
this simulation would exhibit first-person conscious experiences very similar to those
experienced by the original person. This is, after all, what is meant by the phrase
"mind uploading". There are many levels at which one could disagree with this thesis
so I will briefly cover the logic behind it before moving on to the technical discussion
of how to obtain such a high-resolution volume image of an entire brain.

The most fundamental assumption on which this thesis depends is that the human
mind and its conscious experiences are purely a computational phenomenon. We are
broadly similar to an intelligent robot. Our body is analogous to the robot's mech-
anical body. Our brain is analogous to the robot's computer hardware including
the physical hard drives which contain its control software. Our mind is analogous to
the computational processes that result by the running of this control software on the
robot. Our first-person conscious experiences are analogous to the particular stream
of tokened data structures that the robot uses to model its own perceptions, goals,
and actions as part of its self-model (used as a higher-order guide for reasoning about
and planning its future behaviors). What would "mind uploading" mean for such a
robot? It would simply mean that the robot's software is copied to another computer
(possibly of different design) controlling another robot body. This of course would
result in a similar behaving robot with a similar tokening of internal data structures.
Although in principle there may be some deep flaw in this analogy between a human

K. J. Hayworth

1   and a mechanical robot, it is an analogy that rests at the core of all of our research in
2   cognitive science, neuroscience, and even biology itself.
3   If we accept this general "computational theory of mind" assumption there is still
4   the question of whether the human mind's computational process is even remotely
5   similar to what researchers in the Artificial Intelligence (AI) community would
6   recognize. Decades of research in AI have produced incredibly sophisticated com-
7   putational systems capable of all manner of intelligent behavior. These results pro-
8   vide the bedrock for understanding in general how a physical machine for
9   manipulating symbols can produce a wide variety of intelligent goal-directed beha-
10 viors. However, the algorithmic structure of most of these AI systems is completely
11 incompatible with what is known about the computational limitations of the human
12 brain. An often cited example of this is the 100-step constraint [Feldman and Ballard,
13 1982]. Since computational results are signaled between neurons at timescales no
14 faster than a few milliseconds, human behaviors like visual object recognition (which
15 can be completed within a few hundreds of milliseconds) must not be being
16 implemented by algorithms requiring more than 100 sequential computational steps.
17 It is clear that the vast majority of symbolic AI algorithms for vision, problem
18 solving, speech recognition, etc. violate this and other known nervous system con-
19 straints. These ubiquitous violations are sometimes taken to imply that the human
20 brain should not be viewed as a sequential symbol processor at any level [Fodor and
21 Pylyshyn, 1988], and that therefore the core insights from AI research are irrelevant
22 to how the biological brain functions. If true, this would mean that the results about
23 how to create and manipulate structured symbolic representations and use these as
24 the basis for goal-directed decisions would be immaterial to our understanding of the
25 human mind. And in turn, this would raise serious questions about what we even
26 mean by "copying" the software of the brain and ensuring that the simulated mind
27 produces a "similar tokening of internal data structures". In a digital system like a
28 computer we can be very clear about what constitutes a successful copy of an
29 algorithm, for an arbitrary complex analog system this is less clear.
30   In fact however, the most empirically successful and sophisticated models of the
31 human cognitive architecture stem directly from symbolic AI research and are fun-
32 damentally digital at their uppermost algorithmic level. Starting in the 1970's the
33 cognitive psychologist and AI researcher Allen Newell looked very seriously at the
34 constraints that biological neurons put on the types of algorithms that AI researchers
35 were proposing at the time. His answer to these constraints was to suggest that the
36 brain actually uses a *production system* framework to carry on its symbolic compu-
37 tations [Newell, 1973]. In a production system there is a structured-symbolic working
38 memory representation that holds all information about the current state of the
39 program. At each step of a computation, this working memory representation
40 is pattern-matched against a set of "production rules". Each production rule is an
41 "if-then" statement whose "if"-clause looks for particular (syntactically specific)
42 patterns in the current contents of working memory and if a match is detected its
    "then"-clause "fires" producing a specific change to the contents of working memory.

1   Unlike a program written in a traditional software language like C, a production
2   system program consists of an unordered set of possibly thousands of production rules
3   which are all pattern matched against working memory simultaneously and in par-
4   allel at each step of computation. This was the key insight since even though the
5   brain's circuitry is too slow to support algorithms with tens of thousands of *sequential*
6   steps, it seems perfectly tailored to allow the *parallel* matching of tens of thousands of
7   patterns simultaneously.

8   Newell and colleagues showed that such a neurologically-plausible production
9   system framework could support AI algorithms for general problem solving and that
10   a host of complex human skills and learning behaviors could be modeled within the
11   framework [Newell, 1990]. The results of this research have guided the field of cog-
12   nitive science ever since and today's most sophisticated and empirically supported
13   model of the human cognitive architecture — the ACT-R model — is a direct
14   decedent of this production system research.

15   Simply put, the ACT-R model is the most comprehensive theory available today
16   for how the human mind and brain function at an algorithmic level [Anderson *et al.*,
17   2004; Anderson, 2007]. It has been used in literally hundreds of cognitive science
18   publications from dozens of laboratories to model behaviors ranging from simple
19   stimulus response, visual perception and attention, and accuracy and timing of
20   memory recall, to natural language understanding, problem solving and decision
21   making, and the learning of complex skills [Anderson, 2011]. It would be impossible to
22   fully describe the ACT-R model here but a basic understanding of what it says about
23   the composition of the human cognitive architecture is crucial to any intelligent
24   discussion about the possibility of human mind uploading.

25   The ACT-R model proposes that the human brain is a production system as
26   Newell suggested. Cortical systems like the ventral visual stream process raw sensory
27   information into patterns of neural activity over specific high-level sensory cortical
28   areas (e.g., the lateral occipital complex) which act like buffers in a global working
29   memory representation of the environment. Because of the highly processed nature of
30   these sensory representations they act as symbols (e.g., I see a coffee cup) rather than
31   collections of basic features. Other cortical areas (e.g., certain parietal lobe and
32   frontal areas) also act as buffers holding symbols in the global working memory
33   workspace, but these buffers are filled as a result of internal computations as opposed
34   to sensory input. There are also global working memory buffers which connect to
35   cortical and subcortical motor control systems. The high-level symbols in these
36   buffers represent primitive motor actions like "pick up the object" or "saccade to the
37   red object" whose contents are eventually interpreted by downstream motor centers
38   to produce finely-timed sensory-motor behaviors.

39   This network of cortical buffers is the global working memory workspace for the
40   brain and it is looked at by pattern matching circuits within the basal ganglia. The
41   ACT-R model of the basal ganglia postulates that it encodes the full complement
42   of the "production rules" that we have acquired over our lifetime. The basal ganglia
     act as a straightforward pattern associative memory, and if a pattern is recognized

K. J. Hayworth

(i.e., if a production fires) then it activates an associated set of fibers projecting from the basal ganglia (and other structures) back to the cortex to overwrite or route patterned activity between cortical buffers [Stocco *et al.*, 2010]. Every hundred milliseconds or so the basal ganglia perform such a match-and-route operation, and the sequential result of many of these operations explains all of our actions from the most basic learned reactions (e.g., hit the brakes when you see a red light) to the most sophisticated (e.g., reasoning through a new scientific theory).

Another central part of the ACT-R theory is the presence of a dedicated declarative memory storage and retrieval system. The hippocampus and temporal lobe structures can on-demand produce a snapshot of part of the global working memory buffers (via a form of Hebbian learning) producing a new episodic or declarative memory "chunk". As part of any production firing, the production system (i.e., basal ganglia switching circuits) can request a retrieval of a declarative memory chunk based on presenting a partial pattern (e.g., "$2 + 2 = ?$" would retrieve the declarative memory chunk "$2 + 2 = 4$").

Learning in ACT-R consists of forming new procedural memory productions in the basal ganglia and forming new declarative memories in the hippocampus and temporal lobe. There is also an analog component to this learning which can strengthen or weaken particular procedural and declarative memories based on experience, thus making them more or less likely to match or be retrieved.

That is the 10,000 foot overview of ACT-R — the most detailed and empirically supported model of the human cognitive architecture we have available to us today. The key thing for us to note is that the ACT-R model of the brain suggests that our "software" is mainly digital. Working memory buffers hold discrete symbols (modeled as stable attractor patterns in the cortical neural buffers). Declarative memories are also to be thought of as containing discrete symbols (which in this case would be modeled as stable attractor states in the hippocampus and temporal lobe). And each production rule is also a discrete entity (in this case an associative pattern memory stored in the circuits of the basal ganglia along with specific routing circuits in the cortex).

This means that, according to ACT-R theory, we are justified in using our analogy that mind uploading is a matter of copying the mostly-digital program associated with an individual brain to another computational substrate. We can go further and precisely specify what this program consists of — it consists of the unique collection of production rules, declarative memory chunks, and perceptual and motor memories that the human has accumulated over their lifetime.

Furthermore, the ACT-R theory actually suggests neural substrates for each of these, allowing us to evaluate the difficulty of recovering these unique components of identity from a static brain scan. Neural models of how stable attractor states are stored in the hippocampus and cortex suggest that such attractor states are robustly and redundantly stored in the specific pattern and strength of synaptic connections between neurons [Anastasio, 2010]. Just as one can read off the full complement of

1   stable attractors from a Hopfield network [Hopfield, 1982] if given its weight matrix
2   (i.e., which neurons connect to which others and the strengths and signs of their
3   connections), one should be able to do the same in a real biological network of
4   neurons similarly constructed to store memories as stable attractor states.

5   These models imply that the full complement of discrete stable attractor states of
6   a network of hippocampal or cortical cells should be determinable if it is known
7   precisely which neurons synapse with each other and if their relative strengths can be
8   determined. Can the relative strength of a synaptic connection be determined from
9   an electron micrograph? Yes, electron micrograph studies of hippocampal synapses
10 following the induction of memory via a theta-burst stimulation protocol show clear
11 structural changes in synapses that are correlated with their strengthening [Bourne
12 and Harris, 2011]. These results and others show that a synapse's size and area of post
13 synaptic density (both easily measurable in a $5 \times 5 \times 10$ nm volume electron image)
14 are well correlated with its measured strength. Similarly, associative neural network
15 models (of the type assumed present in the basal ganglia) suggest that production
16 matching rules would again be robustly and redundantly stored in the specific pat-
17 tern and strength of synaptic connections between neurons and should thus be
18 determinable by electron micrographs.

19   In summary, our best model of the human cognitive architecture predicts that a
20 particular human being's unique "software" consists of a discrete set of production
21 rules, declarative memory chunks, and perceptual and motor memories learned
22 throughout their lifetime. These "atomic components of thought" [Anderson, 1998] are
23 robustly and redundantly encoded across the synaptic connections of cortical and
24 subcortical brain areas in a manner that should be retrievable from a static
25 $5 \times 5 \times 10$ nm volume electron image. A successful mind uploading of a particular
26 human being would consist of creating a computer simulation of the brain with just
27 enough fidelity that the same discrete switching events (production firings events,
28 declarative memory retrieval events, perceptual tokening events, etc.) occur that would
29 have occurred in the original biological brain under similar circumstances. This would
30 lead to similar external behavior as well as a similar tokening of internal data structures.

31   Would this simulation be conscious? And would it be the same consciousness as
32 the original biological person? If the computational theory of mind is correct then the
33 answer to both of these questions is yes. In the phenomenal self-model of con-
34 sciousness [Metzinger, 2003] our first-person conscious experiences are a reflection of
35 the particular stream of tokened data structures that our brain uses to model its own
36 perceptions, goals, and actions. Using the ACT-R framework we can identify pre-
37 cisely where this phenomenal self-model resides in the brain. It is the collection of
38 episodic and self-directed declarative memories that encode all of our knowledge
39 about our own perceptions, goals, and actions, along with all of our production rules
40 which create and operate on these self-directed declarative memories. When we are in
41 the midst of a conscious perception it is these production rules which take the current
42 stream of symbols in our global working memory workspace and uses them to create

K. J. Hayworth

additional episodic and self-directed declarative memory chunks incorporating them into our continuously evolving phenomenal self-model. A simulated brain which is based on the same declarative memories and production rules would "reboot" the same phenomenal self-model as the original biological brain, and would exhibit the same tokening of internal data structures (and thus modifications to the phenomenal self-model) that would have occurred in the original brain under similar circumstances. From this perspective continuity of the person's stream of consciousness would have been restored.

3. High-Resolution Reconstruction of Neural Circuits

As just discussed, significant portions of the brain may be well characterized at the algorithmic level as symbol-containing buffers, pattern recognition circuitry, routing circuitry, etc., but at the implementation level the brain consists of hundreds of specialized regions each containing networks of millions of biologically complex neurons. And each brain region is composed of dozens of different neuronal types whose functional properties and connectivity have been finely tuned to ensure robust operation at the algorithmic level. One of the promises of the new field of Connectomics is that mapping the detailed connectivity of neurons at this implementation level will provide a comprehensive understanding of how the algorithmic level functioning is generated.

A typical cortical pyramidal neuron may have a cell body that is only 10 microns in width but will send out an incredibly elaborate dendritic arbor spanning a cubic millimeter in volume. The finest branches of this dendritic arbor are thousands of dendritic spines whose tiny necks can shrink down to just a few tens of nanometers in diameter. This neuron will also send out a long thin axon that in places will shrink down to just a few tens of nanometers in diameter, but will travel many millimeters before branching and arborizing to form synaptic contacts with distant neurons. Any random cubic millimeter of cortical tissue will contain tens of thousands of neurons with overlapping dendritic arbors and will receive tens of thousands of axonal projections whose arbors also overlap. And within this cubic millimeter of densely tangled jungle will be on the order of a few hundred million synaptic connections each with a tiny area of contact on the order of 500 nm in diameter [Shepherd, 1990]. Staining a single neuron (as in the Golgi method) can reveal the gross structure of that single neuron in a light micrograph, but to actually map the connectivity of the entire neural network — determining precisely which neurons synapse on which others — requires high-resolution volume electron microscopy on a scale that has never before been attempted. The field of Connectomics has been formed to do just that by developing new automated methods for obtaining enormous numbers of serial electron micrographs from a volume of neural tissue, and by developing new computer algorithms that can efficiently process the resulting terabytes of image data to automatically reconstruct each neuron's dendritic and axonal arbors and find where they make synaptic contacts with other neurons.

1   As mentioned above, the diameter of neuronal processes routinely shrink to less
2   than 100 nm; for example, dendritic spine necks and fine axons can shrink down to
3   ∼ 40 nm. With standard tissue fixation and embedding protocols the membranes of
4   these tubular structures are made electron dense, hence to resolve the tubular nature
5   of these neuronal processes one requires resolutions on the order of 10 nm or less.
6   Today both the transmission electron microscope (TEM) and the scanning electron
7   microscope (SEM) can easily achieve such resolution while providing a signal to noise
8   ratio sufficient to trace the finest neuropil in osmium fixed, heavy metal stained
9   tissue. However, until relatively recently only the TEM was used for tracing neuronal
10   circuits. This reliance on TEM has been a major roadblock to attempts at large-scale
11   automation since TEM requires that sections be physically cut thin enough
12   (< 100 nm) for electrons to pass through and that they be mounted on gossamer thin
13   plastic films throughout the imaging process. If scanning, as opposed to transmission,
14   electron microscopy could be used it would open up many more possibilities for
15   robust automation since then at least the thin sections could be collected and imaged
16   on a thick sturdy substrate [Hayworth, 2008], but before the widespread introduction
17   of high-brightness field emission electron sources (as opposed to tungsten thermionic
18   sources) one simply could not achieve SEM electron probes of small enough diameter
19   (5 nm) and high enough current to allow quality imaging of neural tissue [Bogner
20   *et al.*, 2007; Joy, 1991].

21   In 2004, Denk and Horstman showed that this reliance on TEM for tracing neural
22   circuits could in fact be overcome. They published a seminal paper demonstrating the
23   SBFSEM (Serial Block Face Scanning Electron Microscopy) method which could
24   robustly automate the process of obtaining a series of electron micrographs from a
25   block of neural tissue [Denk and Horstman, 2004]. Using a high-brightness field
26   emission SEM equipped with a low-energy backscatter electron detector they first
27   showed that one could obtain high-resolution images directly from the face of the
28   tissue block (thus eliminating the need to collect ultrathin sections). Then, following
29   an original design from Leighton [1981], they built an ultramicrotome into the vacuum
30   chamber of the SEM which would repeatedly scrape (using an extremely sharp dia-
31   mond knife) 50 nm layers of material off the surface of the tissue block while the SEM
32   was used to image each freshly revealed block face at high resolution. The result was a
33   fully automated method to volume image a block of neural tissue. They and others
34   have continued to refine this technique so that now it can now achieve 23 nm section
35   thickness and have successfully applied the technique to a reconstruction of the
36   direction selective circuitry of the mammalian retina [Briggman *et al.*, 2011].

37   Both section thickness and lateral resolution are limited in the SBFSEM technique
38   due to the fact that the physical scraping process for removing plastic-embedded
39   block face material becomes unreliable if the material properties of the embedding
40   plastic are degraded due to long exposure to the electron beam during imaging. This
41   means that, in general, a finer lateral resolution will require a greater electron beam
42   dose to the block face and will thus make scrapping a very thin layer of tissue from
     the block face less reliable. This physical scraping process also makes the technique

1 very sensitive to the embedding properties of the tissue block — something that
2 could make application to very large tissue volumes difficult.
3   In 2008, Graham Knott *et al.* introduced the FIBSEM (Focused Ion Beam
4 Scanning Electron Microscopy) technique for tracing neural circuits [Knott *et al.*,
5 2008]. This technique is also known as Ion Abrasion Scanning Electron Microscopy
6 [Heymann *et al.*, 2009]. FIBSEM works similarly to SBFSEM, but instead of phy-
7 sically scraping a thin layer of material off the block face with a diamond knife, the
8 material is instead ablated away using a focused beam of gallium ions. This change
9 overcomes the lateral resolution limitations of the SBFSEM since the FIB ablation
10 process is relatively insensitive to the material properties of the embedding plastic
11 and therefor a much larger electron dose can be used during imaging. What's more,
12 ion beams can be tightly focused achieving spot sizes in the 10 nm range. Aligning the
13 ion beam parallel to the block face so that it just grazes the surface allows reliable and
14 rapid removal of extremely thin layers less than 10 nm in thickness. What actually
15 limits the resolution of the FIBSEM (for optimally sized blocks) is the depth of
16 penetration of the imaging electrons into the surface of the block [Muller-Reichert
17 *et al.*, 2010] and by using very low voltages ($< 2 \, \text{kV}$), along with sensitive low-energy
18 backscatter electron detectors with energy filtering, recent reports [Knott and
19 Cantoni, 2011] have demonstrated depth resolutions in the range of $\sim 10 \, \text{nm}$ with
20 lateral imaging resolutions of $5 \times 5 \, \text{nm}$. This FIBSEM voxel resolution of
21 $5 \times 5 \times 10 \, \text{nm}$ is more than sufficient to reliably resolve the finest neuronal processes
22 and synapses, and should be sufficient to allow for extremely reliable automated
23 reconstruction algorithms [Lucchi *et al.*, 2010]. Recently Knott has published a
24 Journal of Visualized Experiments online video article covering the entire FIBSEM
25 procedure from tissue processing to automated FIBSEM milling and imaging,
26 demonstrating a final volume image of a piece of cortex imaged with $5 \times 5 \times 5 \, \text{nm}$
27 voxels [Knott *et al.*, 2011]. It is stunning how clearly demarcated the neuronal pro-
28 cesses and synapses are in the resulting FIBSEM volume video.
29   There is, however, a fundamental physical limitation of the FIB ablation
30 process — this slice resolution can only be obtained for very small samples on the
31 order of a few tens of microns across. One cause of this is that the electrostatic lenses
32 of the FIB are able to tightly focus the beam of gallium ions only over a short distance
33 (i.e., its depth of focus range), and outside of this range the beam's shape and its
34 milling efficiency changes. This means that to achieve uniform removal of 10 nm of
35 surface material the block being milled should be only a few tens of microns wide in
36 the direction of the ion beam.
37   A straightforward solution to this sample size limit is to simply subdivide a large
38 tissue block into smaller sub-blocks of optimal ($\sim 20 \, \text{micron}$) size for FIBSEM ima-
39 ging. However, if the goal is to trace neural circuits across the original larger volume,
40 this subdivision process must be such that absolutely no material is lost or damaged
41 during subdivision so that all of the tiny 40 nm diameter neural processes remain
42 traceable across the subdivided faces. This requirement means that such subdivision
    cannot be performed prior to plastic embedding (for example by using a tissue

1 vibrotome) since sectioning wet, soft, unembedded tissue will distort and destroy the
2 fine neuronal processes surrounding the cut.
3   To produce tissue surfaces with nanometer-scale smoothness one typically must
4 use a diamond knife polished to a radius of just a few nanometers, and must cut tissue
5 that has been infiltrated both extracellularly and intracellularly with a plastic resin so
6 that the whole of the tissue block is rigidified to withstand sectioning and to prevent
7 cutting-induced distortion of fine cellular processes at the cut surfaces. Such diamond
8 knife sectioning is used in the SBFSEM to scrape very thin layers off the block face,
9 and in obtaining ultrathin sections for TEM imaging as described above. However,
10 such ultrathin (< 100 nm) sections would be very difficult to reliably obtain in large
11 numbers and to image reliably in the FIBSEM. A more optimal sectioning thickness
12 (for producing sub-blocks of tissue ready for high-resolution FIBSEM imaging) would
13 seem to be closer to the maximum thickness allowed by the FIBSEM milling process
14 itself (i.e., ~ 20 microns). Unfortunately, if one simply attempts to use a diamond
15 knife (which typically must have a large edge angle) to section such a 20 micron thick
16 slab, the differential cutting stresses on the section as it bends are enough to crumple
17 and crack the tissue and severely damage the surfaces.
18   I recently performed a set of experiments to determine if there were cutting par-
19 ameters which would allow for lossless 20 micron thick sectioning of plastic-embedded
20 brain tissue. Adapting a trick from McGee-Russell [McGee-Russell *et al.*, 1990],
21 I used a slightly heated ($\sim 60°$C) diamond knife to section plastic-embedded brain
22 tissue blocks. This heating temporarily and locally softens the plastic during sec-
23 tioning, reducing the local cutting stresses while maintaining overall block rigidity.
24 I also found that I had to use an oil lubricant during the sectioning process to prevent
25 surface damage due to the thick section scraping along the knife face. After a bit of
26 trial and error I determined a set of parameters that could indeed reliably produce
27 20 micron thick brain sections with such high-quality surfaces that the finest neuronal
28 processes could be traced from one face of a thick cut to its matching face. This was
29 verified by taking electron micrographs of both faces and then mirror-reversing one
30 micrograph and aligning it to the other. I also verified that the neural tissue within
31 these thick cut sections was still intact by performing a serial section volume
32 reconstruction of several neuronal processes across two such thick cut sections
33 [Hayworth *et al.*, 2010], Hayworth, in preparation.
34   In summary, the FIBSEM technique represents a truly enabling technology for
35 mapping neural circuits with 100% reliability. It can achieve voxel resolutions of at
36 least $5 \times 5 \times 10$ nm, sufficient to allow straightforward algorithms for computer
37 automated tracing of all neuronal processes and identification of all synaptic con-
38 nections [Merchán-Pérez *et al.*, 2009] along with the morphological parameters which
39 are correlated with their strength (e.g., area of contact and size of post synaptic
40 density). Furthermore, its use of a focused ion beam (as opposed to traditional
41 physical sectioning with a diamond knife) gives it the potential to achieve the very
42 high reliability levels needed for large-scale automation. A serious limitation of
   today's FIBSEM systems is that they can achieve such high-resolution images only

K. J. Hayworth

over tiny volumes — typically on the order of a few tens of microns across. However, this limitation can be relatively easily overcome with the use of a lossless subdivision technique like the one described above. And the availability of such a lossless subdivision technique opens up the possibility of massively parallel FIBSEM imaging, something that would, in any case, be necessary in order to image any relatively large volume of tissue in a reasonable amount of time.

4. Whole Brain Ultrastructure Preservation and Plastic Embedding

Electron imaging of brain tissue at the resolutions necessary to trace neural circuits requires that the tissue's membranes and other key subcellular structures be selectively stained with heavy metals (osmium, uranium, and lead typically). The dense nuclei of these heavy metal atoms deflect the electrons of the SEM's beam back at the electron "backscatter" detector allowing the stained membranes to stand out against the unstained background of the cytoplasm. Electron imaging of brain tissue also requires that this tissue be structurally stabilized by removing all water from within the tissue and infiltrating every nook and cranny of extracellular and intracellular space with a curable plastic resin. This is not only to rigidify the tissue for sectioning (as described above), but also to stabilize the tissue so that it can withstand the intense radiation bombardment of the electron beam.

Today's most typically used tissue preparation protocols for electron microscopy are only able to prepare volumes of less than $1\,mm^3$. An animal's vascular system is perfused through the heart with a mixture of paraformaldehyde and glutaraldehyde. The paraformaldehyde quickly stops cellular degradation and, at a slightly slower rate, the glutaraldehyde provides stronger crosslinks to fix proteins in place. If performed carefully, such perfusion through the vascular system is able to quickly fix the entire nervous system of the animal since every cell is within a few tens of microns of a capillary. However, following this fixative perfusion step the brain of the animal is removed and a very small piece dissected to undergo the remaining tissue preparation steps — which are typically all performed by simple immersion in chemicals. These steps include immersion in osmium tetroxide (to fix membrane lipid molecules in place), immersion in heavy metal staining solutions (e.g., uranyl acetate), immersion in a graded series of alcohols (to remove water from the tissue), and finally immersion in a plastic resin dissolved in an organic solvent (to completely infiltrate the tissue with the heat-curable resin). Because these steps are performed by simple immersion (and thus diffusion) the process fails if attempted on blocks larger than $1\,mm^3$. The result is destruction of tissue ultrastructure and poor staining in the depths of the block.

Obviously such volume limitations must be overcome before human mind uploading can be attempted. Although it has never been demonstrated that a whole mammalian brain can be preserved at the ultrastructure level for electron microscopic imaging, there are many results that suggest that a protocol could be developed to do just that. Perfusion fixation with osmium tetroxide has been demonstrated on whole

1 brains [Palay *et al.*, 1962], and serial vascular perfusion first with glutaraldehyde,
2 followed by osmium tetroxide and uranyl acetate, and finally by an alcohol dehy-
3 dration series has been demonstrated on whole organs [Bachofen *et al.*, 1982; Oldmixon
4 *et al.*, 1985] demonstrating sufficient preservation to allow ultrastructure studies.
5 Plastic infiltration of whole mouse brains has also been demonstrated [Mayerich *et al.*,
6 2008], and there are recent reports of the development of a full ultrastructure fixation,
7 staining, and embedding protocol for the mouse brain for use in the mapping of long
8 distance axon trajectories via serial block face SEM [Mikula *et al.*, 2011]. There is even a
9 challenge prize being offered for the first demonstration of such ultrastructure pres-
10 ervation across an entire large mammalian brain [Hayworth, 2011]. Given these results
11 and the fact that the bourgeoning field of Connectomics will require larger and larger
12 volumes of high-quality prepared tissue, it is likely that a protocol to preserve and
13 plastic-embed an entire human brain at the ultrastructure level will be perfected
14 relatively soon.
15
16
17 ## 5. Industrial-Scale Thick Sectioning for Random-Access,
     Massively Parallel FIBSEM Imaging
18
19 To review, the FIBSEM technique can already achieve voxel imaging resolutions
20 sufficient to allow the computer automated tracing of neural circuits and the
21 identification and strength estimation of synaptic connections. In order to apply
22 this technique to a whole human brain the brain must first be chemically fixed and
23 plastic embedded. It must then be subdivided into sub-blocks of optimal size for
24 FIBSEM imaging — each approximately 20 microns in width. Above, I described a
25 method (using a heated, oil-lubricated diamond knife) which can losslessly sub-
26 divide plastic-embedded brain tissue into such 20 micron thick slabs, but it is an
27 open question whether such a technique could be successfully applied to anything
28 as large as an entire human brain. Here I will briefly sketch out an industrial-scale
29 process that might accomplish such a feat. The general plan is to section a human
30 brain into over 30 million "tissue pillars", each 20 microns wide ×20 microns
31 deep ×180 mm long. At each step in the sectioning process the tissue to be sec-
32 tioned is first securely adhered to a collection tape prior to being sliced off by the
33 heated diamond knife — in this way the tissue slices are always robustly con-
34 trolled. As a result, the final tissue pillars are securely adhered to the edge of a
35 sturdy tape ready for high-resolution FIBSEM imaging. This tape-collection
36 methodology is loosely based on the automatic tape-collection technique I devel-
37 oped for collecting ultrathin sections for electron microscopic imaging [Hayworth
38 and Hayworth, 2004; Hayworth, 2008].
39 A plastic-embedded human brain is first mounted on a gantry-style Thick-Strip
40 Sectioning and Tape Collection Machine (Fig. 1(a)). This device has two diamond
41 knives (both heated and oil-lubricated) which are used in tandem to section the brain
42 into a tape containing 3 mm wide 20 micron thick tissue strips. First a heated
   "furrowing" knife is used to create a set of losslessly-cut furrows in the block face at
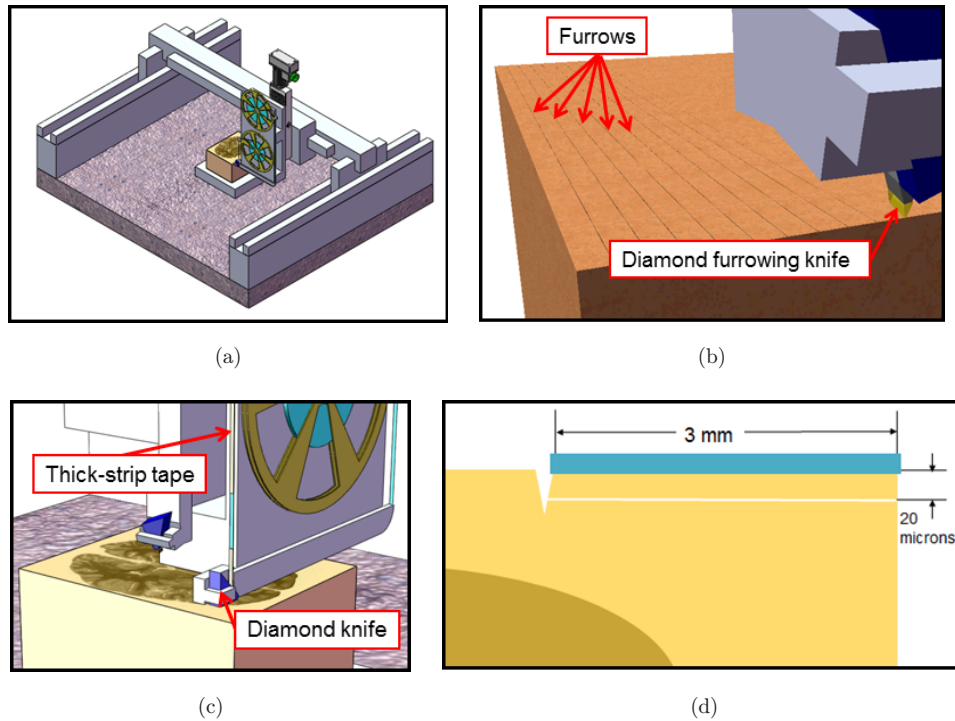
K. J. Hayworth

(a)

(b)

(c)

(d)

Fig. 1. CAD drawings depicting whole brain thick sectioning process. (a) Human brain block mounted in gantry-style Thick-Strip Sectioning and Tape Collection Machine. To collect a 20 micron thick layer of the brain first a set of furrows are made in the block face at 3 mm intervals (b). Then a heated platen is used to adhere a collection tape to the block face between two furrows. And finally a diamond knife removes a 3 mm wide 20 micron deep strip of brain tissue immediately under the tape (c) while simultaneously reeling the tissue strip up with the attached tape. This is repeated 50 times to remove the full 20 micron thick surface layer across the brain block; and the entire process is repeated again and again to remove a total of 5000 such 20 micron thick layers covering the entire brain's depth. (d) Relation between the furrow, tissue strip, and block face-adhered collection tape.

3 mm intervals (Fig. 1(b)). Each furrow goes a little deeper than 20 micron into the face of the plastic block. Following this, a heated platen is used to adhere (using a thin layer of heat-activated adhesive) a 3 mm wide collection tape to the block face, aligning this tape precisely between two furrows. Once adhered, the entire platen and tape reel assembly move forward while a second heated diamond knife slices off the 3 mm wide 20 micron deep strip of brain tissue immediately under the tape (Fig. 1(c)). This results in a brain strip (180 mm in length) being reeled up securely adhered to the collection tape. Figure 1(d) shows the relation between the furrow, tissue strip, and collection tape. This strip-collection process is repeated to collect each of the strips comprising the top 20 microns of the brain block (a total of 50 strips for a 150 mm wide brain). Then the entire furrowing and striping process is repeated again and again to collect each sequential 20 micron thick layer of the brain block. At the end of this process (which would take several years at a 1 mm/s cutting speed)

1   the entire 100 mm deep brain block is reduced to set of "thick-strip tapes" containing
2   a total of 250,000 tissue strips.

3   These 20 micron thick tissue strips are still not in a form that will allow efficient
4   high-resolution FIBSEM imaging. To allow this, they must undergo additional sec-
5   tioning into 20 micron × 20 micron pillars. The thick-strip tape is first adhered to a
6   semi-rigid backing tape using a Semi-rigid Backing Applicator machine (Fig. 2(a)).
7   This is done to make it sturdy enough to orthogonally section in a reel-to-reel fashion.
8   Then this tape is put into a Pillar Sectioning and Tape Collection Machine (Fig. 2(b)).
9   It is within this machine that the original 3 mm wide strips are repeatedly sectioned in
10   the orthogonal direction creating the final 20 micron × 20 micron × 180 mm long
11   tissue pillars, each securely adhered to the edge of a new collection tape.

12   This pillar tape is now in the optimal configuration to allow high-resolution
13   random-access FIBSEM imaging. A particular pillar tape is mounted in a reel-to-reel
14   FIBSEM system as shown in Fig. 3. The tissue pillar to be imaged is first reeled into
15   position under the electron beam and clamped in place. The SEM then takes a high
16   resolution (5 × 5 nm) image of the top face of the tissue pillar over an area of
17   approximately 20 × 20 microns. Following this, the side-directed focused ion beam is
18   used to ablate away the top 10 nm of material that was just SEM imaged. This
19   process of SEM imaging and FIB surface ablation is repeated 2000 times to image an
20   approximately 20 × 20 × 20 micron cube of tissue[1].

21   Perfecting such an industrial-scale brain sectioning scheme would no doubt be
22   expensive and difficult, but its engineering tolerances are actually quite reasonable.
23   This is because the required sectioning thickness is 20 microns which is 2−3 orders of
24   magnitude thicker than the sections used for serial TEM reconstruction or in the tape-
25   collection for SEM process [Harris *et al.*, 2006; Hayworth, 2008]. I would estimate that
26   a system to reduce a whole brain into FIBSEM-imageable pillar tapes could be
27   developed and built with mainly off-the-shelf components for around $10 million.

28   It is the FIBSEM imaging requirements that actually represent the most difficult
29   and expensive investment for any large-scale imaging project along these lines.
30   Today's commercial FIBSEM systems are very complex machines costing around
31   $2 million apiece, and their imaging rates are quite slow — on the order of 1 MHz
32   typically. Also, current systems are not well suited for automation or for extended
33   imaging times due to limited lifetime ion sources. To make practical large-scale
34   imaging a custom redesign of the FIBSEM will be necessary. Optimizing the tissue
35   staining, detector, and electron beam parameters should allow imaging rates of up to
36   10 MHz (a rate that has already been achieved in the tape-collection for SEM pro-
37   cess). It should also be possible to redesign the system for quick setup of a tape pillar
38   by directing the FIB beam at a scintillator detector during operation — thus pro-
39   viding closed loop control of the FIB milling parameters and milling rate during
40

41   [1] The actual volume imaged must be trapezoidal in shape, with a wider area milled off the top of the volume
42   than the bottom. This is to prevent the milled sidewalls of the tissue from shadowing the areas to be imaged
     in lower slices. This is what is depicted in the figure.

K. J. Hayworth
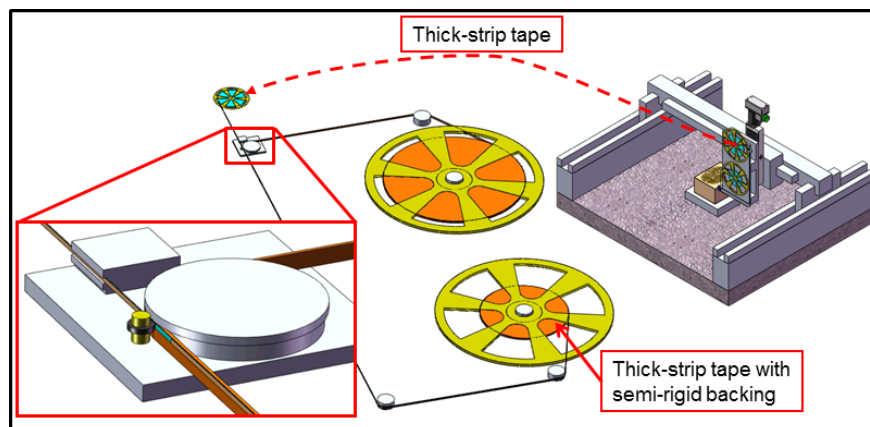
(a)
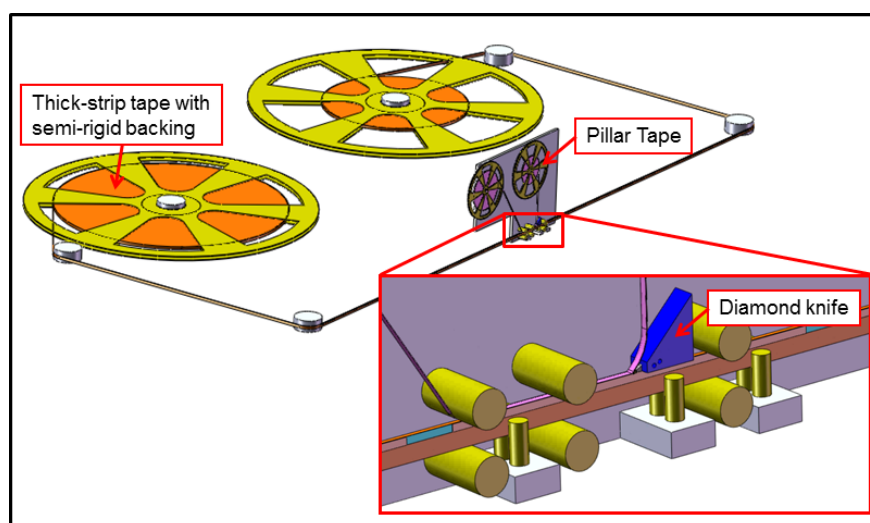
(b)

Fig. 2. CAD drawings depicting processing needed to produce the final tissue pillar tapes ready for FIB-SEM imaging. The strip tape is first adhered to a semi-rigid backing tape using a Semi-rigid Backing Applicator machine (a). Then this tape is put into a Pillar Sectioning and Tape Collection Machine (b) in which the original 3 mm wide strips are repeatedly sectioned in the orthogonal direction creating the final 20 micron × 20 micron × 180 mm long tissue pillars, each securely adhered to the edge of a new collection tape. This process must be repeated multiple times until the original 3 mm wide strip is reduced to a set of 150 pillar tapes.

imaging and setup. Such closed-loop control of milling should also greatly increase the reliability of the entire process.

Given such an optimized FIBSEM with an imaging rate of ∼ 10 MHz (and accounting for setup and FIB milling time), a single 20 × 20 × 20 micron cube could
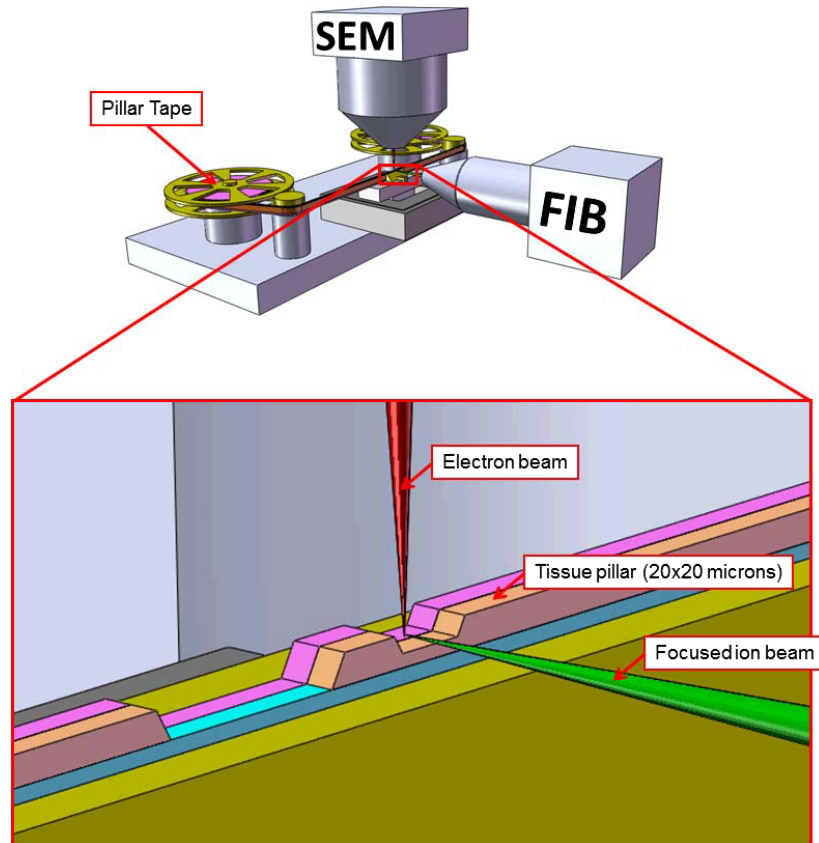
Fig. 3. CAD drawing depicting how a pillar tape is loaded into a reel-to-reel FIBSEM machine. Any $\sim 20 \times 20 \times 20$ micron (trapezoidal) region of the pillar tape can be imaged in the FIBSEM by reeling the tape so that the desired tissue pillar and region is under the electron beam, and then performing ~ 2000 FIB milling and SEM imaging steps on that pillar region until the entire volume in that part of the pillar has been volume imaged at $5 \times 5 \times 10$ nm voxel resolution.

be imaged every two hours. At this rate it would take almost 30 years for a full cubic millimeter volume to be imaged using a single machine. Clearly a single FIBSEM does not come close to meeting the requirements for imaging an entire 1×10^6 mm^3 human brain. Even a fleet of 100 such machines operating over 10 years could only succeed in imaging 0.003% of the brain. However, such a fleet of 100 FIBSEM machines would make possible the creation of a "Human Connectome Observatory", something of critical importance to the advancement of neuroscience, and something that I believe is the most crucial milestone on the road to true mind uploading.

6. A Human Connectome Observatory

The astronomical research community excels at the efficient pooling of research funds. No individual astronomer could afford to build the Keck Observatory

K. J. Hayworth

1   telescopes on the summit of Mauna Kea ($\sim$ \$200 million), or the Large Binocular
2   Telescope in Arizona ($>$ \$100 million), or ESO's Very Large Telescope Array in Chile
3   ($>$ \$300 million). And they certainly could not afford to build the Spitzer Space
4   Telescope ($>$ \$800 million), or the Chandra X-ray Space Telescope ($>$ \$1 billion), or
5   the Hubble Space Telescope ($>$ \$3 billion). But any individual astronomical
6   researcher can apply for observation time on any of these grand instruments at
7   relatively reasonable rates. The fantastic gains in our understanding of the cosmos
8   have in large part been a result of this strategy of upfront pooling of resources to build
9   top-of-the-line instruments which are then shared by all.
10   The ability to section an entire human brain into a set of randomly accessible
11   FIBSEM-imageable tapes would make possible a similar pooling of resources within
12   the neuroscience community. The key is random-access. Because the brain has been
13   reduced to a set of tapes, any $20 \times 20 \times 20$ micron cube within the brain can be called
14   up for imaging by: 1) Loading the proper tape into a FIBSEM machine, 2) Reeling
15   the tape so that the desired tissue pillar and region is under the electron beam, 3)
16   Raising the tissue pillar until the FIB scintillator detector indicates that FIB surface
17   milling has begun, 4) Performing $\sim$ 2000 FIB milling and SEM imaging steps on that
18   pillar region until the entire $20 \times 20 \times 20$ micron volume in that part of the pillar has
19   been volume imaged at $5 \times 5 \times 10$ nm voxel resolution. For roughly the cost of a top-
20   of-the-line earth-based optical telescope observatory ($\sim$ \$200 million) a facility could
21   be built containing a fleet of 100 FIBSEM machines coupled to robotic tape loaders
22   which would allow any pillar tape to be loaded into any of the 100 FIBSEM
23   machines. The result would be a "Human Connectome Observatory" in which the
24   highest resolution imaging could be directed anywhere that is desired within a single
25   human brain.
26   What could the neuroscience community do with such a Connectome Observa-
27   tory? From the above, we assume that each FIBSEM can automatically load and
28   image any desired $20 \times 20 \times 20$ micron cube of tissue every two hours. This means
29   that the entire 100 FIBSEM Observatory would have a total imaging throughput of
30   50 cubes per hour, or equivalently it could image a volume of 1 mm^3 every 100 days.
31   Such an observatory could image 4.4 million cubes over a 10 year operational life,
32   representing a total volume of 35 mm^3. Again, the key is that these 4.4 million cubes
33   can be distributed across the brain in any fashion desired. For example, a single tiny
34   axon could be efficiently followed from its cell body in one region of the brain to its
35   terminal synapses several centimeters away in another brain region by sequentially
36   targeting imaging cubes only along its path. The Connectome Observatory could
37   trace such fine axons at an average rate of about 1 mm/hour (i.e., a line of 50 cubes
38   stretches 1 mm in length), and thus over its 10 year lifetime the Connectome
39   Observatory could trace a total of $\sim$ 90 meters of long distance axonal connections. It
40   is more likely that the neuroscience community would collectively decide (in an open,
41   proposal-driven process similar to the one used to govern time on astronomical
42   observatories) to perform large dense reconstructions in key brain regions along with
    directed axon following to trace specific fibers between these dense reconstructions.

1   For example, key circuits of the human visual system could be mapped out by
2   imaging a $0.5\,mm^3$ volume of the lateral geniculate nucleus (LGN), a $1\,mm^3$ columnar
3   volume of the primary visual cortex (V1), and a $1\,mm^3$ columnar volume of extra-
4   striate cortical area (V2), while strategically directing groups of imaging cubes along
5   the axonal pathways between these regions. Such a strategy would produce fantas-
6   tically detailed circuits maps of the individual regions (sufficient to determine all
7   neuronal types and the details of their local circuit connections) while at the same
8   time putting each of these dense reconstructions in context of each other by deter-
9   mining the origins of particular input axons and the destinations of particular
10 outgoing axons.
11   Over a 10 year period the Connectome Observatory would allow researchers to
12 perform large dense reconstructions (each $0.5\,mm^3$) in a total of 40 different brain
13 regions while at the same time putting each of these dense reconstructions in context
14 of each other and of the rest of the brain by allowing dozens of fine long-distance
15 axons (totaling 40 meters worth) to be traced into and out of each of the dense
16 reconstructions — all within the same brain. And, of course, the same infrastructure
17 could be used to map neural circuits across the brains of other species (e.g., fly,
18 mouse, cat, monkey).
19   Such a comprehensive and detailed understanding of a single brain's circuitry is
20 crucial to understanding how the brain works. It will allow us to test high-level
21 algorithmic theories of brain function like ACT-R by literally looking for the brain-
22 wide routing circuitry that it predicts. It will allow us to test theories of how mem-
23 ories are encoded by extracting the complete connection matrix of a volume of cortex
24 or hippocampus and determining the stable attractor states of the network. It will
25 allow us to test theories of visual feature extraction and binding by looking for the
26 intra- and interregional circuits that particular visual theories predict. It will allow us
27 to begin to perform test simulations of the whole nervous systems of small organisms
28 (like drosophila) along with functionally well-characterized regions of larger organ-
29 isms (e.g., the mammalian retina).

30
31 ## 7. "We Choose to go to the Moon..."
32
33 I believe that the successful construction of a Human Connectome Observatory is the
34 crucial milestone in the path to human mind uploading. It is also readily achievable.
35 The expected scientific payoff of such a Connectome Observatory should easily jus-
36 tify a construction cost on the order of an earth-based astronomical observatory,
37 especially given that this cost, and its scientific bounty, will be shared by the entire
38 neuroscience community.
39   The Connectome Observatory's scale is far smaller than what is needed for the
40 ultimate goal of uploading a human mind, but it will demonstrate the brain pres-
41 ervation, large-scale lossless subdivision, and imaging technologies necessary for such
42 an attempt. More important, it will enable a comprehensive vetting of our cognitive
    and neuroscience understanding of the brain in a way that will give us the confidence

K. J. Hayworth

to proceed to this ultimate scale. Just as Alan Shepard's sub-orbital flight was a sufficient demonstration to justify proceeding with the full Apollo program, the Connectome Observatory should provide a sufficient demonstration that the goal of uploading a human mind is achievable.

Like Apollo, scaling up to a full brain will be an incredibly expensive and difficult engineering challenge. For example, a 10 year project for mapping a whole brain would require the equivalent of 3 million of the above FIBSEM imaging machines. To fit within an Apollo-size budget (over $100 billion in today's dollars) mass manufacturing would have to bring each FIBSEM's cost down to around $40,000 apiece. This is actually not unreasonable given that we are already seeing the introduction of high-resolution "solid state" SEMs whose electron optical columns are manufactured using the same photolithographic techniques used by the semiconductor industry [Rynne, 2009]—these SEMs are already priced in the $100,000 range. More likely such mass-producible electron columns will be combined with new technologies like multibeam SEMs which can achieve imaging rates in the 1000 MHz range by imaging with dozens of electron beams simultaneously [van Himbergen *et al.*, 2007]. Like computer technology, electron imaging technology seems to be headed toward exponentially decreasing the cost of each megabyte of data acquired. Any large-scale investment in brain mapping is likely to greatly accelerate this trend and may, like the Human Genome Project, result in a situation where the second mind to be uploaded is done at a tiny fraction of the time and cost of the first.

References

Anastasio, T. J. [2010] *Tutorial on Neural Systems Modeling* (Sinauer, Sunderland).

Anderson, J. R. [1998] *The Atomic Components of Thought* (Lawrence Erlbaum Associates, Mahwah).

Anderson, J. R., Byrne, M. D., Douglass, S., Lebiere, C. and Qin, Y. [2004] "An integrated theory of the mind," *Psychological Review* **111**(4), 1036−1060.

Anderson, J. R. [2007] *How Can the Human Mind Occur in the Physical Universe?* (Oxford University Press, Oxford).

Anderson, J. R. [2011] "ACT-R Research Group Website," Department of Psychology, Carnegie Mellon University, http://act-r.psy.cmu.edu/

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

Bogner, A., Jouneau, P. H., Thollet, G., Basset, D. and Gauthier, C. [2007] "A history of scanning electron microscopy developments: Towards "wet-STEM" imaging," *Micron* **38**(4), 390−401.

Bourne, J. N. and 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. and Denk, W. [2011] "Wiring specificity in the direction-selectivity circuit of the retina," *Nature* **471**, 183−188.

Denk, W. and Horstman, H. [2004] "Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure," *PLoS Biol* **2(11)**, e329.

Eisenstein, M. [2009] Neural circuits: Putting neurons on the map, *Nature* **461**, 1149−1152.

Feldman, J. A. and Ballard, D. H. [1982] "Connectionist models and their properties," *Cognitive Science* **6**, 205−254.

Fodor, J. A. and Pylyshyn, Z. W. [1988] "Connectionism and cognitive architecture," *Cognition* **28**, 3−71.

Harris, K. M., Perry, E., Bourne, J., Feinberg, M., Ostroff, L. and Hurlburt, J. [2006] "Uniform serial sectioning for transmission electron microscopy," *The Journal of Neuroscience*, **26**(47), 12101−12103.

Hayworth, K. J. and Hayworth, A. A. [2004] "Methods and apparatuses for the automated production, collection, handling, and imaging of large numbers of serial tissue sections," United States Patent 7,677,289, Filed July 8, 2004.

Hayworth, K. J. [2008] "Automated creation and SEM imaging of ultrathin section libraries: Tools for large volume neural circuit reconstruction," *Neuroscience Meeting Planner* Society for Neuroscience [Washington, DC, 2008]. Online, Program No. 504.4.

Hayworth, K. J., Tapia, J. C., Kasthuri, N., Schalek, R. and Lichtman, J. W. [2010] "Divide and conquer — Lossless thick sectioning of plastic-embedded brain tissue to parallelize large volume serial reconstructions," *Neuroscience Meeting Planner,* PPP4. Society for Neuroscience, (San Diego, CA, 2010). Online, Program No. 516.91.

Hayworth, K. J. [2011] "Brain preservation foundation website," `http://www.brain-preservation.org/`

Heymann, J. A. W., Shi, D., Kim, S., Bliss, D., Milne, J. L. S. and Subramaniam, S. [2009] "3D imaging of mammalian cells with ion-abrasion scanning electron microscopy," *Journal of Structural Biology* **166**(1), 1−7.

Hopfield, J. J. [1982] "Neural networks and physical systems with emergent collective computational abilities," *Proceedings of the National Academy of Sciences* **79**, 2554−2558.

Joy, D. C. [1991] "The theory and practice of high-resolution scanning electron microscopy," *Ultramicroscopy* **37**, 216−233.

Knott, G., Marchman, H., Wall, D. and Lich, B. [2008] "Serial section scanning electron microscopy of adult brain tissue using focused ion beam milling," *The Journal of Neuroscience* **28**(12), 2959−2964.

Knott, G., Rosset, S. and Cantoni, M. [2011] "Focussed ion beam milling and scanning electron microscopy of brain tissue," *Journal of Visual Experiments* **53**, e2588.

Leighton, S. B. [1981] "SEM images of block faces, cut by a miniature microtome within the SEM — A technical note," *Scanning Electron Microscopy* **2**, 73−76.

Lucchi, A., Smith, K., Achanta, R., Knott, G. and Fua, P. [2010] "Supervoxel-based segmentation of EM image stacks with learned shape features," *EPFL Technical Report*, December 2010.

Mayerich, D., Abbott, L. and McCormick, B. [2008] "Knife-edge scanning microscopy for imaging and reconstruction of three-dimensional anatomical structures of the mouse brain," *Journal of Microscopy* **231**(1), 134−143.

McGee-Russell, S. M., De Bruijn, W. C. and Gosztonyi, G. [1990] "Hot knife microtomy for large area sectioning and combined light and electron microscopy in neuroanatomy and neuropathology," *Journal of Neurocytology* **19**, 655−661.

Merchán-Pérez, A., Rodriguez, J. R., Alonso-Nanclares, L., Schertel, A. and DeFelipe, J. [2009] "Counting synapses using FIB/SEM microscopy: A true revolution for ultrastructural volume reconstruction," *Fontiers in Neuroanatomy* **3**(18) 1−14.

Metzinger, T. [2003] *Being No One: The Self-Model Theory of Subjectivity* (MIT Press, Cambridge).

Mikula, S., Helmstaedter, M. and Denk, W. [2011] "Whole-brain ultrastructural staining in the mouse," *Neuroscience Meeting Planner*, Society for Neuroscience (Washington, DC, 2011). Online (Program No. 304.14/XX66).

K. J. Hayworth

1   Muller-Reichert, T., Mancuso, J., Lich, B. and McDonald, K. [2010] "Three-dimensional
2   reconstruction methods for caenorhabditis elegans ultrastructure," *Methods in Cell Biology*
3   **96**, 331−361.
4   Newell, A. [1973] "Production systems: Models of control structures," in *Visual Information
       Processing* (ed.) Chase W. D. (Academic, New York).
5   Newell, A. [1990] *Unified Theories of Cognition* (Harvard University Press, Cambridge).
6   Oldmixon, E. H., Suzuki, S., Butler, J. P. and Hoppin, F. G., Jr. [1985] "Perfusion dehydration
7   fixes elastin and preserves lung air-space dimensions," *Journal of Applied Physiology* **58**(1),
8   105−113.
9   Palay, S. L., McGee-Russell, S. M., Gordon, S. and Grillo, M. A. [1962] "Fixation of neural
       tissues for electron microscopy by perfusion with solutions of osmium tetroxide," *Journal of
10   Cell Biology* **12**, 385−410.
11   Rynne, J. [2009] "A compact field emission SEM for low voltage imaging," *Microscopy Today.*
12   34−36.
13   Seung, S. [2012] *Connectome: How the Brain's Wiring Makes Us Who We Are* (Houghton
14   Mifflin Harcourt, Boston).
15   Shepherd, G. M. [1990] *The Synaptic Organization of the Brain*, 3rd edition (Oxford University
       Press, New York).
16   Stocco, A., Lebiere, C. and Anderson, J. R. [2010] "Conditional routing of information to the
17   cortex: A model of the basal Ganglia's role in cognitive coordination," *Psychological Review*
18   **117(2)**, 541−574.
19   van Himbergen, H. M. P., Nijkerk, M. D., de Jager, P. W. H., Hosman, T. C. and Kruit, P.
20   [2007] "High throughput defect detection with multiple parallel electron beams," *Journal of
21   Vacuum Science Technology B* **25**(6), 2521−2525.
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42