

Single Brain Physical Slice Library Proposal

Creating a complete human neural connectivity database via sparse, automatically-directed ultramicroscopic imaging using an automated retrieval brain slice storage system

Ken J. Hayworth
ken@brain.jpl.nasa.gov
BrainCAD@hotmail.com

Abstract

The complete and accurate determination of all the regions, axonal pathways, and neuronal circuits of the human brain would form a touchstone, a foundation of certainty, with which to constrain all other theories of how the brain functions. In this paper it is argued that current neuroscience methods and instruments, as well as the new neuroinformatic computerized databases of the Human Brain Project, are inadequate for this task. Instead, we propose the development of a new microscopic imaging system that would make possible the imaging of a single human brain across all size scales from the macroscopic anatomy to the synaptic ultrastructure.

This proposed system, called a single brain physical slice library, would make use of an automated retrieval brain slice storage system, which would provide a suite of teleoperated ultramicroscope instruments random access to any requested $1\mu\text{m}$ thick slice of brain tissue. This ability to randomly access thin slices of brain tissue is the key to bridging (in a single imaging experiment) the tremendous range of size scales inherent in neuronal circuits. Random access to slices allows for new experimental techniques (discussed in this paper) that can map all the regions, pathways, and neuronal circuits in a brain while imaging only a tiny fraction of that brain's total volume.

The proposed system would make extensive use of image processing algorithms to automatically direct the imaging sequence in real-time. This sparse, automatically-directed imaging would allow neural connectivity to be determined by literally following microscopic neuronal processes as they course macroscopic distances across the brain. For example, this directed imaging could be applied to 100,000 widely-distributed neurons carefully chosen as being representative of all the brain's regions and pathways. From this directed imaging (tracing axonal projections, dendritic arbors, and local synaptic contacts), a detailed statistical map of all of the information flows and regional synaptic circuits in the brain could be determined after having required the imaging of only 0.01% of the brain's total volume. It is this four-order-of-magnitude reduction in imaging volume that makes the prospect of mapping a human brain's connectivity at the ultrastructure level feasible in time and cost.

Such an automated and teleoperated system would make possible the ultimate level of collaboration amongst neuroscientists by allowing researchers from around the world the ability to perform their own custom-designed, detailed telemicroscopy experiments on the same physical brain. Unlike separate specialized experiments, the results from these brain-mapping experiments would immediately and easily be integrated because they are all performed on the same physical brain. Developing the proposed system represents an ambitious undertaking; however, its concept is firmly based on existing technologies.

Table of Contents

1	Executive Summary.....	4
2	Outline of Paper.....	9
3	Requirements for a Complete Human Neural Connectivity Database.....	11
3.1	Introduction.....	11
3.2	Overview of database structure.....	12
3.3	Sensory surfaces, discrete sensors, and discrete effectors.....	12
3.4	Gross neuroanatomical regions.....	15
3.5	Modules: cortical regions, subcortical nuclei.....	15
3.6	Neural layers.....	16
3.7	Pathways and projective fields: the brain’s block diagram.....	17
3.8	Neurons and neuron components.....	18
3.9	Synapses and microcircuit-multiports.....	19
3.10	Basic circuit statistics: pathways revisited.....	19
3.11	Database as neuroembryology.....	20
3.12	Summary.....	20
4	Creating a Physical Slice Library: Operation, Technical Challenges, and Reasonable Milestones.....	21
4.1	Introduction.....	21
4.2	Overview of a single brain physical slice library.....	21
4.3	Operation.....	23
4.4	Types of experiments.....	23
4.4.1	Single neuron and neural subunit reconstruction.....	24
4.4.2	Regional basic-circuit reconstruction.....	24
4.4.3	Region and region boundary determination.....	25
4.4.4	Interregional pathway following.....	26
4.5	Major technical challenges.....	26
4.6	Human brain physical slice library.....	28
5	BrainCAD: Software Tool for Creating Neural Databases from Slice Library Data.....	29
5.1	Overview.....	29
5.2	Operation.....	30
5.2.1	Single brain physical slice library raw experimental data.....	30
5.2.2	Database entry using BrainCAD.....	30
5.2.3	Visualizing neural connectivity using BrainCAD.....	31
5.2.4	BrainCAD interfaces with neural simulation programs.....	32
5.2.5	The “firewall” between raw image data and interpretations.....	33
5.3	Summary.....	33
6	Conclusions.....	34
7	References.....	35

1 Executive Summary

The cognitive sciences, including neuroscience, cognitive psychology, and artificial intelligence (AI), have made incredible strides recently. Cognitive psychology and classical AI are today providing a unified and coherent model of human intelligence. That model combines high-level, goal-directed reasoning, learning, and complex symbolic representation, with lower-level sensory and motor behaviors [Newell 1990] [Rosenbloom 1993] [Hayworth 1998]. Of equal importance, the connectionist AI community is joining efforts with neuroscience researchers, who are studying the biophysics of computation in real neurons [Koch 1999]. Together they are creating computational models of how vast networks of biologically plausible neurons can perform complex information processing of the type posited for particular brain regions [Marr 1982][Kohonen 1997]. These top-down cognitive theories and bottom-up neuronal theories are at an increasingly high level of maturity. There is great anticipation that when they finally link in the middle, a deep, mechanistic theory of the human mind will be the result.

Many researchers believe that the “middle” at which these theories can be made to meet will be a precise and complete map of all information flows in a human brain, dictated by the macroscopic and microscopic neuroanatomy. **The complete and accurate mapping of this macroscopic and microscopic neuroanatomy will form a touchstone, a foundation of certainty, with which to constrain all other theories of how the brain functions.** Such a map would necessarily include structural (and pharmacological) data at all size scales, even down to the ultrastructure of dendritic spines. Of equal importance, it must reduce and impose order on this raw data; explicitly showing the statistics and information flows in the brain. A block diagram of these information flows is needed upon which functional models can be constrained, compared, and collaborated on. The exact requirements and structure of such a map (or more accurately, a highly structured database) have not been adequately elucidated elsewhere; therefore, a large section of this paper will be devoted to giving a detailed exposition of what information a complete database of brain connectivity must contain.

Unfortunately, such a complete database of brain connectivity does not yet exist, even for lower mammals. Experimental techniques in neuroscience research, which have advanced almost unimaginably in the past decade, are more than adequate to answer specific connectivity questions (at least in animal models) [Windhorst 1999]. However, the wash of thousands of papers coming from different labs, using different protocols, different animals, and different species, are not adding up to the complete database of brain connectivity described here. This problem was fully realized in the early 1990’s by a National Academy of Sciences panel of neuroscience experts evaluating the necessity and viability of a National Neural Circuitry Database. The recommendations of that panel eventually led to the creation of the NIH’s Human Brain Project in 1993. That major NIH project has as its primary goal the creation of a “neuroinformatics” infrastructure that would allow the integration of specialized neuroscience facts into national databases able to disseminate these specialized facts as an integrated whole. During the last decade the Human Brain Project has funded many successful neuroinformatics database programs at a variety of excellent institutions; however, there remains no database approaching the

precision and scope envisioned here. In this paper it is argued that the current programs will **not** be able to achieve this scope and precision in their mapping tasks using their current methodology of simply gathering together research data from disparate laboratories. **It is argued instead that a complete database of brain connectivity requires the development of specific new techniques in automated ultramicroscopy¹. Furthermore, in keeping with the immensity of the task, collaboration between multiple laboratories is critical; however, all laboratories involved must be mapping the same physical human brain.**

Specifically, it is proposed that the many separate techniques now used for mapping the macroscopic and microscopic information flows in animal models (in vivo tracer transport, LM (light microscope) 3D reconstruction via microinjection of dye, and many others) are abandoned for this particular task. Those techniques, in general, require a new animal brain for each new connectivity experiment, which necessitates the two traditional headaches of neuronal mapping. The first headache is correlating results between experiments and the second is correlating results with human brain data on which a vast amount of valuable cognitive studies have been done.

As substitution for these techniques, it is proposed to take advantage of the all-purpose abilities of uniform-stained 3D microscopy. In *non*-uniform techniques such as combined 3D LM and TEM of neurons microinjected with dyes, a complete *single* neuron can be visualized over an enormous range of size scales [Soto 1994]. However, this single, stained neuron can be visualized only at the expense of millions of neurons, surrounding and connected to it, being kept transparent. In contrast, use of uniform-stained 3D microscopy, such as 3D serial TEM using membrane-staining OsO₄, allows entire volumes of convoluted neural matter to be imaged in total [Goodhew 2001][Hayat 2000]. Every neural process in the volume can be recorded simultaneously because physical slicing, tomographic slicing, or both prevent one neuron from occluding another. In theory, uniform-stained 3D ultramicroscopy techniques could single-handedly provide a complete database of brain connectivity, at all relevant size scales, by systematically slicing and imaging the entire brain. Long range axonal projections, regional basic neuronal circuit connectivity, and synaptic ultrastructure would all be contained in the vast volume of voxel data generated.

All the relevant technologies have been demonstrated on small volumes (<1000μm³) of brain tissue in lower mammals [Martone 2000][Fiala 2001] [Perkins 1997]. However, these minute volume-imaging experiments were performed mostly to determine dendritic tree morphology statistics. The reason this technique has not yet been adopted for general use in determining intra- and interregional connectivity involves issues of scale [Fiala 2002]. For instance, even 1mm³ of brain tissue represents an enormous volume compared to the experiments performed so far (a factor of 1x10⁶), and simple extrapolation of times involved shows that mapping an entire brain using this

¹ The term ultramicroscopy is used here to denote any of the class of instruments that can image tissue ultrastructure (beyond the range of LM resolution). This includes the traditional SEM, TEM, and variants, as well as other classes of instruments. One such instrument is the biological 3D SIMS (Secondary Ion Mass Spectrometer), which recently has been applied to imaging biological tissue at tens of nanometers resolutions. SIMS offers the advantage of imaging chemical and isotopic differences as well as structural detail at this resolution.

technique would take virtually forever. (Although, see [Merkle 1989]¹ for a discussion on the feasibility and a more optimistic time estimation of this technique.) In this paper, a full brain mapping using current poorly automated techniques is not proposed. **Instead, the novel concept of a single brain physical slice library is introduced, which would allow the ultimate level of online collaboration in neuroscience research and would result in a complete neural connectivity database after imaging only a tiny fraction of the brain's volume.**

An individual, well-preserved brain (and spinal cord) would be ultramicrotomed into ~1 μ m thick slices, each having appropriate lateral dimensions for examination in the specimen chambers of specially modified ultramicroscopes. These slices would be mounted on special carriers and stored in a specially built slice library storage system. This system would allow automated retrieval and insertion of any slice into any one of a set of ultramicroscope instruments². The goal is to allow a researcher to image any part of this brain (at the ultrastructure level) all via computer control. The automated retrieval system and integrated instruments would automatically load the appropriate slice, pan to the (x,y) coordinates commanded, and begin imaging that tiny volume. Successive pans and slice retrievals would allow the researcher to image as large a volume as is deemed necessary to answer the questions of particular concern to this researcher during this experiment. Afterwards, the research is published and the imaged brain volume is made available online.

The true advance is what comes *after* this data is published. Let us suppose that another researcher, perhaps at a quite remote institution, reads the published results and studies the online brain volume. Perhaps she identifies axon collaterals of a particular neuron in the online volume that are not currently included in the basic circuit theory of that particular region of the brain. The particular online volume is not extensive enough to show where these collaterals lead, so she initiates (by remote) another imaging session **on this same brain and contiguous with the original imaged brain volume**. She can use the automated retrieval and imaging system to follow a particular axon collateral to see where it finally arborizes and which neuron types it synapses with. Discovered, this data is published and the online imaged volume of this brain grows significantly.

If the system can be made reliable and relatively user-friendly, then this type of collaboration could substantially change the way neuroscience research is done. Many researchers from around the world would see the advantages of this intensified collaboration (as if all were peering through the same microscope). A bottleneck could

¹ This paper, written by Ralph Merkle in 1989 while at Xerox PARC, offered a prescient vision of what the nascent technique of electron tomography could offer the study of the brain. Its thorough and well-documented exposition of the issues involved in the automated imaging of large-scale neural structures at ultrastructural resolution has contributed greatly to the current proposal.

² Automated retrieval and imaging on this scale is complex but appears tractable. Commercial automated slide loaders and X-Y-Z stages for light microscopes are beginning to be used in the clinical diagnosis industry [Visioninst]. For vacuum instrumentation, the semiconductor industry has solved many technical difficulties, including the transfer of specimens between high vacuum imaging instruments like TEMs and SIMSs. Automation for drug discovery research has produced fully automated microplate systems able to store and randomly access 600,000 samples in a single, low-temperature storage system. Also, the National Center for Microscopy and Imaging Research (NCMIR) has invented a suite of tools enabling online "telemicroscopy", demonstrating the ability for a researcher to remotely control an IVEM at NCMIR while receiving feedback of TEM images online [Fan 1993]. Elsewhere, extensive work has been performed in automating the steps of single axis tilt electron tomography.

soon form, since the automated retrieval and imaging system can only perform so fast. Seeing the need, other institutions could install their own versions of this automated physical slice imaging system and the *original* brain slices could be divided up between these institutions. There would still be only a single brain, but the slices would be distributed in a fashion that would allow parallel imaging using all the resources of this collaborative network.

With this growing level of infrastructure and growing number of researchers imaging particular brain regions and pathways, the online brain volume would begin to look like an extensive underground mining operation, with many narrow tunnels crisscrossing in 3D space connecting larger excavation sites. Individual researchers would have *sparsely* imaged mm³ sized sections of cortical regions and subcortical nuclei. Other researchers would have concentrated on imaging ten-micron-wide filamentous strands of neural tissue in order to follow particular axons as they course from one well-imaged cortical region to another well-imaged cortical region. This sparse directed imaging, which bridges the tremendous range of size scales in the nervous system, is depicted in *Figure 1* on the following page.

If we assume the brain contains on the order of 1000 distinct processing regions, and the equivalent¹ of 100 full neurons needed to be imaged in each of these regions in order to completely determine the statistical connectivity of synaptic circuits and pathways, then only 100,000 of the human brain's 100 billion neurons would need to be imaged in order to generate a complete neural connectivity database. Assume that the directed (process-following) imaging of a neuron is only 1% efficient (i.e. that 99% of the volume imaged while following a neural process does not belong to that neuronal process itself). This would mean that a volume equivalent of 10 million neurons would need to be imaged to complete this task. By this very crude estimate, only 0.01% of the human brain's volume would need to be imaged. This is an approximate volume of only 140mm³ or about the size of a frog's brain. **It is this four-order-of-magnitude reduction in imaging volume that makes the prospect of mapping a human brain's connectivity at the ultrastructure level feasible in time and cost.**

The conclusions drawn about neural connectivity would be much more detailed and much more certain than any current techniques can produce. Even new concepts of interregional basic circuits might be in the offing. The complete and precise conclusions drawn about neuron morphology, intraregional basic circuit connections, and interregional projections would finally allow the creation of a complete database of brain connectivity and successfully meet the mandate set under the Human Brain Project. As stated before, such a complete database of brain connectivity would bring about a new era of collaboration amongst all cognitive science researchers, and this could finally allow the bottom-up theories of brain function (connectionist and biophysics of computation) to meet the top-down cognitive theories. It is from *this* collaboration that a deep, mechanistic theory of the human mind may result.

¹ Imaging complete neurons is not necessary to determine morphology and connectivity statistics. When we say that the equivalent of 100 full neurons will be imaged we mean that many partial pieces of thousands of neurons will be imaged in each region such that the total imaged pieces add up to an equivalent of 100 complete neurons being imaged.

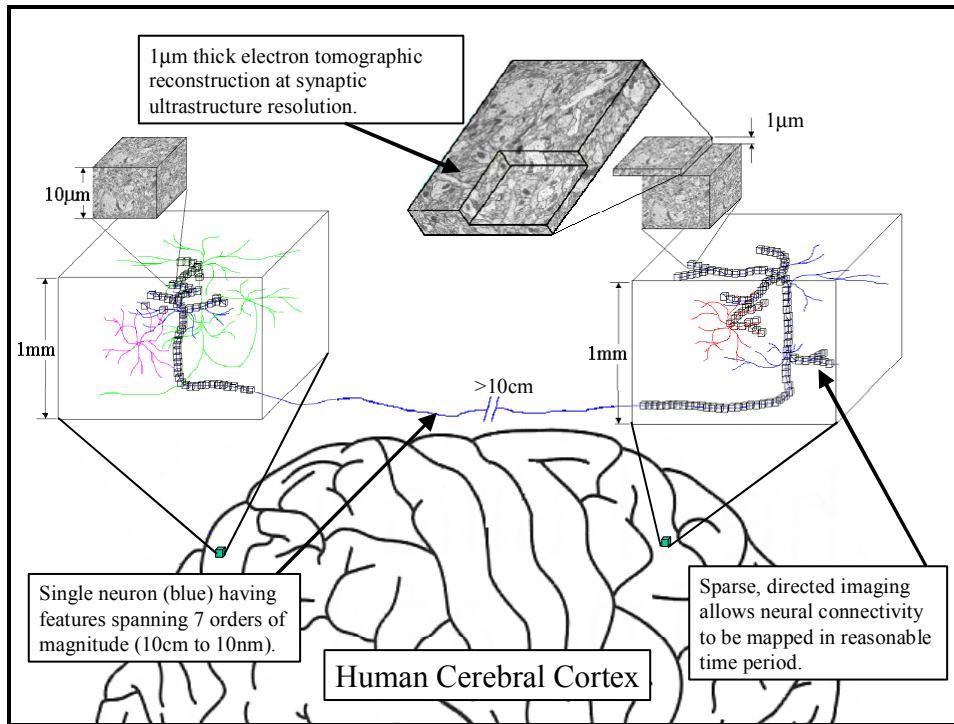


Figure 1: Bridging of size scales in the nervous system using the random-access imaging abilities of a single brain physical slice library. This drawing depicts the sparse, directed imaging of two 1mm³ sized regions of cortex that are quite removed from each other, yet which share a common projection neuron (blue). The projection neuron's cell body is in the right volume, and its axonal arbor is in the left volume. By directing the electron tomographic imaging along specific neural processes, both the local region's basic circuit connectivity and the interregional pathway projections can be imaged. Even though these neurons have processes spanning many cubic millimeters in volume, many orders of magnitude *less* volume need actually be imaged. Section #4 describes this type of experimental procedure, as well as several others, which can be performed using the random-access imaging abilities of a single brain physical slice library.¹

¹ TEM volumes depicted in the drawing are adapted from the online brain volumes on the Synapse Web site www.synapses.bu.edu. The 2D images from that site were voxel displayed using a custom program. Cerebral cortex drawing is adapted from [Gupta 1997].

2 Outline of Paper

The rest of this paper consists of two main exposition sections plus some short follow up sections.

In section #3, the requirements for a complete neural connectivity database are outlined. By drawing only on currently well-accepted theories of how the brain is organized [Kandel 2000][Shepherd 1990], a database structure is outlined which would make explicit all the neuroanatomical details that are relevant to building models of brain function. The database structure described in this section is crucial to the proposal since it explicitly states the *types of questions* that the proposed technology of a single brain physical slice library must be capable of providing answers to.

In section #4, the technical aspects of creating and operating a single brain physical slice library and automated imaging infrastructure are addressed. A key part of that section is a list of the types of experiments that could be performed using a single brain physical slice library. These experiments are:

- Single neuron and neural subunit reconstruction
- Regional basic-circuit reconstruction
- Region and region boundary determination
- Interregional pathway following

The section addresses how each of these experiments could be performed using the physical slice library. It also shows that the results of these experiments would provide the level of neural detail required for a complete neural connectivity database of the type outlined in section #3.

A list of the main technical challenges foreseen and a set of reasonable initial milestones are given. Some of the major technical hurdles do not have to do with automated ultramicroscopy per se, but have to do with developing new fixation, staining, and slicing techniques that can reliably prepare a full brain for imaging at the ultrastructure level. Since current techniques are geared toward imaging tiny neural volumes, the fixation, staining, and slicing steps traditionally used render most of the non-imaged brain volume useless. Current techniques must be modified (or new techniques invented) to prevent this, and this will involve a great deal of basic research and experimentation in itself.

The goal of mapping *human* brain tissue presents its own set of problems, not the least of which has to do with preservation of neural ultrastructure without the ability to perform in vivo (or in block) perfusion of fixatives before death (as is the standard protocol in animal ultrastructure imaging). Nearer-term milestones proposed avoid this issue by developing all techniques first on rat or mouse models on which the majority of current 3D neuronal ultrastructure reconstruction is performed [Synapse][Jensen 1989].

Also included in section #4 is a discussion of the various instruments that may be used to examine slices in the single brain physical slice library. Of course, imaging of 3D neuron morphology is of primary concern; however, pharmacological issues, such as neurotransmitter type and membrane excitability properties, are also necessary parts of any complete neural connectivity database. The technique of serial electron tomography

seems the best suited for 3D reconstruction; however, it has some automation issues which need resolving, and does not provide significant information on chemical properties without special (difficult to automated) staining steps. SIMS (Secondary Ion Mass Spectrometer) imaging offers 3D scanning abilities as well as chemical and isotopic imaging at nanometer resolutions [Stelly 1995] [Cameca]. SIMS might also prove easier to automate for scanning large area brain slices since it avoids the issues of TEM support grid occlusion. Ideally, all these instruments and others would be available to researchers using a physical slice library, and be fully integrated with the automatic slice retrieval system.

Following these main exposition sections, section #5 addresses some extensions of this core project, specifically a proposed software tool called BrainCAD. BrainCAD is envisioned as a software tool for creating object-oriented databases of neural connectivity information of the type described in section #3. BrainCAD is also envisioned as a graphical CAD-like program that would directly interface with the intermediate resolution (LM) atlas of the physical slice library as well as any imaged ultrastructure resolution volumes. In this way, BrainCAD would provide the *link* between the imaged raw data and the creation of descriptive neural connectivity databases.

3 Requirements for a Complete Human Neural Connectivity Database¹

'One thing must be stressed quite firmly: henceforth functional localisation of the cerebral cortex without the lead of anatomy is utterly impossible in man as in animals. In all domains, physiology has its firmest foundations in anatomy. Anyone wishing to undertake physiological localisational studies will thus have to base his research on the results of histological localisation. And today with greater reason than ever, one must recall the words of the past master of brain research, Bernhard Gudden, spoken three decades ago in the face of a dangerous tendency to specialise in extirpation experiments: "Faced with an anatomical fact proven beyond doubt, any physiological result that stands in contradiction to it loses all its meaning ... So, first anatomy and then physiology; but if first physiology, then not without anatomy".'

-Korbinian Brodmann in Localisation in the Cerebral Cortex (1909)
translated by Laurence Garey [Neuroscience]

3.1 Introduction

We wish to create a database structure that explicitly and concisely provides all neuroanatomical details of the brain². This structure must not overstep its mandate by presupposing functional hypotheses; however, it must assume a great deal of uniformity in the structural organization of the brain in order to concisely convey the neuroanatomy. There are well-accepted theories of how the brain is organized at the macroscopic and microscopic size scales [Kandel 2000][Shepherd 1990]. These theories posit cortical regions, subcortical nuclei, axonal pathways, topological projections, layers of neurons, neuron types, basic circuits within regions, and repeating microcircuits involving just a few synapses each³. These entities form the standard language of modern neuroscience and underlie experimental hypothesis of neural function. These are the structures that a neural connectivity database must describe.

¹ The following proposed neural database structure is one possible way for meeting the requirements of a concise, precise, and complete description of the brain's anatomy. There are many subtleties of brain connectivity; therefore, this database structure should not be assumed fully finished nor be assumed to provide a unique taxonomy. This section's purpose is simply to impress upon the reader the level of detail such a database should provide.

² Filling in this database structure with the detailed facts of neuroanatomy is, of course, quite another issue. The single brain physical slice library described in section #4 is an experimental technique designed to provide these facts. In this section we are only concerned with what *types* of neuroanatomical facts are needed not with the facts themselves or the methodologies to arrive at them.

³ Entities like cortical columns will not be included in this proposed database structure. Not because there is any doubt that they exist, but because they are really the purview of functional models, not neuroanatomy. In general, any brain regularity that can be described without knowledge of the response of neurons (but that is presumed to have direct influence on neural responses) is pure neuroanatomy, and thus belongs in the database.

3.2 Overview of database structure

The core of this neural connectivity database is a parcellation of the brain into *neural modules*. A neural module is defined here as contiguous region of neural tissue with uniform (or smoothly varying) basic circuit¹ and interregional connectivity properties. Each module is assigned a 2D or 3D local coordinate system. This local coordinate system provides a way to describe the topological projections of axon bundles between (and within) modules. Beyond that use, this local coordinate system also allows parameterized, statistical functions to be defined for the density of cell bodies within a module, dendritic arbor shapes and sizes, axonal arbor shapes and sizes, etc. All these important neuroanatomical statistics are allowed to vary smoothly across a module according to the cell body's location within that module. **This methodology of using statistical, parameterized functions (defined with respect to a module's local coordinate system) to concisely describe the microanatomy, interregional, and basic circuit connectivity of millions of neurons is what gives this database structure its unique precision and completeness.**

Instead of simply listing the proposed structure of such a database here, its structure is built up step by step in the following exposition starting from the sensory inputs and motor outputs.

3.3 Sensory surfaces, discrete sensors, and discrete effectors

The logical place to start our database is with a full cataloging of the sensory inputs and motor outputs. We recognize several different sensory modalities and submodalities. These can be further subdivided into two categories: discrete and spatially continuous². A partial listing follows:

Discrete sensors:

- Golgi tendon organs (proprioceptive sense)
- Muscle spindle fibers (proprioceptive sense)
- Joint capsule sensors (proprioceptive sense)
- Utricle and Saccule receptors (vestibular sense)
- Semicircular canal receptors (vestibular sense)
- Nociceptive receptors (pain sense, at least four types)
- Thermal receptors (temperature sense, at least four types)

¹ The term "basic circuit" used here is taken from [Shepherd 1990]. It is a diagram showing types of neurons and their specific synaptic connections within a particular region of the brain. Thus, a region's basic circuit describes its intraregional neural connectivity.

² Discrete is taken to mean any sensory input where a single sensory cell's activation contains explicit information upon which to base behavior. For example, a Golgi tendon organ's activation signals a stretched muscle and is used to directly control the stretch reflex. In contrast, in spatially continuous sensory surfaces, each sensory cell activation must be compared to other nearby cell activations in order to provide information to base behavior off of. For example, retina photoreceptors involved in vision (as opposed to those involved in pupillary reflexes) convey information only when compared to relative activation levels of many nearby photoreceptors to create, for example, an on-center ganglion cell response.

- Retinal photoreceptors involved in pupillary reflexes
- Taste receptors (at least 5 types)
- Smell receptors

Spatially continuous sensory surfaces:

- Meissner's corpuscles (tactile sense)
- Merkel disk receptors (tactile sense)
- Pacinian corpuscles (tactile sense)
- Ruffini's endings (tactile sense)
- Rod photoreceptors (visual sense)
- Cone photoreceptors (visual sense, three types)
- Cochlear tonotopic membrane hair cells (auditory sense)

It is not enough to simply list these different types of input receptors. A useful neural connectivity database must also provide precise information about the numbers, locations, and spatial distributions of each type of receptor.¹ As an example, in order to catalog the spatial distribution of each class of tactile sensor, a 3D skin surface model must be included in the database. A coordinate system, (x,y) , is defined upon this skin surface, and a sensory cell density function, $\rho(x,y)$, is defined for each sensor class. This statistical distribution function satisfactorily specifies the distributions of perhaps tens of thousands of surface receptors in a concise and accurate manner. Another function, $r(x,y)$, should also be added which specifies the single sensor receptive field² radius as a function of the spatial location. **These functions, and all others described for this database, can be concisely specified by parameterized multidimensional splines.** Retinas and cochlear surfaces are handled in a like fashion.

The discrete receptors must be handled in a different fashion. The proprioceptive sensors are intrinsically involved with the sensing and control of individual muscles or joints. Therefore, each muscle, bone, and joint must be included in the database in the form of a 3D model as well as a database listing. The distribution of receptors within a muscle has been shown to be of little consequence, so the 3D placement of receptors within each muscle and tendon is not cataloged; however, the total number of receptors in each muscle is.

Muscles have already been cataloged above; however, muscles are composed of many thousands of muscle fibers. From a nervous system viewpoint, neither the muscle nor the individual muscle fibers are discrete units. Groups of 10 to 1000 muscle fibers are controlled together by the same spinal motor neuron's efferent projection. Together, the group of muscle fibers and its associated motor neuron are referred to as a motor unit.

¹ A functional database would go on to include the exact response characteristics of each receptor class to the physical stimuli to which it is specially attuned. However, the database described here is fully anatomical. Its goal is to provide the neuroanatomical framework upon which to hypothesize function.

² The "single sensor receptive field" is defined as the spatial extent (i.e. area on skin for a tactile sensory cell) to which a single sensory cell would respond if there were no other neural connections to that cell. This is to separate it from the functional concept of a sensory receptive field where connections with other sensory cells and neurons can result in such complex behaviors as changing the sensory cell's response radius in real-time to adapt to the current stimuli (e.g. as is the case in retina photoreceptors). For example, the *single cell* receptive field of a photoreceptor is only that part of the visual field that directly illuminates the particular photoreceptor cell itself.

These motor units are the most logical primitives to base muscle effector statistics upon in the database. Each muscle entry should, therefore, also include information about number and type of motor units. Subentries for types of motor units include classes of muscle fibers (e.g. slow twitch vs. fast twitch) and total numbers of fibers. In order to maintain a unity of expression in the syntax of this database, each type of discrete sensor or effector will be assigned a pseudo-coordinate systems to mirror the (x,y) coordinate system of the continuous sensory surfaces.¹ This will round out the sensory and motor entries in our complete human neural connectivity database.²

One might question the fact that we have already included so much information in our “neural” connectivity database and have not yet entered a single neural connection. However, providing a complete description of all sensors and effectors is absolutely crucial to specifying the rest of the information flows in the brain and nervous system. As an example, there is no way to ground the somatosensory maps in Brodmann’s areas 1, 2, 3a, 3b, and 5 without setting up a proper coordinate system on the skin surface and specifying receptor densities. Similarly, there is no way to ground motor maps in the motor and premotor cortices without explicitly including information about muscles, joints, and bones.

To summarize, our database structure now has:

- Sensory surface entries: 3D skin surface, retina surfaces, and cochlear surfaces models with 2D coordinate systems defined.
- Spatially continuous sensor type entries: One entry for each of the dozens of sensor types listed above. Each entry has fields for defining a parameterized density function, $\rho(x,y)$, and single cell receptive field function, $r(x,y)$, defined over the appropriate sensory surface’s coordinate system.
- Skeletal, joint, and muscle entries: Including 3D hierarchical model and relations.
- Discrete sensor entries and subentries: Sensor type entry with subentries for each muscle, joint, or tendon they are associated with. Subentries contain number statistics.
- Muscle effector entries: One entry for each muscle, and one subentry for each type of motor unit.
- Other entries: The bullets above are not meant to be exclusive nor comprehensive, they are meant to show the depth of information necessary for a complete neural connectivity database. They are also meant to display the quantitative, statistical, and parameterized nature of such a database.

¹ As will be seen later, these coordinate systems are crucial to defining pathway axonal projections as coordinate transformations. For example, the pseudo-coordinate system for the fast-twitch muscle fibers of all the skeletal muscles in the body are smoothly mapped by an axonal projection to motor units in the spinal cord (and brain stem) and from there onto the primary motor cortex.

² A truly complete database would also include sensors and effectors associated with the viscera, such as neural control of secretory glands, etc. However, we will neglect these in the current discussion.

3.4 Gross neuroanatomical regions

A useful neural connectivity database must include a 3D model of the gross neuroanatomy in order to allow comparison to functional imaging, recorded experimental data, medical images of pathological conditions, etc. This 3D anatomical perspective is also the perspective neuroscientists are most comfortable with. Several of these 3D gross neuroanatomy databases now exist [Mai 1997]. Our complete neural connectivity database will require a 3D gross anatomy model of the brain in voxel form at a resolution sufficient to define the boundaries of all neural modules (to be defined later). Also note that the peripheral ganglia, spinal cord, and brain stem must be fully modeled in the database along with the brain proper.

The database will define a hierarchical set of “gross neuroanatomical region” entries. Each will consist of a voxel volume delineation of a piece of gross neuroanatomy, such as the thalamus. Regions that are not obviously delineated by the gross neural anatomy (such as the lateral geniculate nucleus of the thalamus) are not included here. They might appear later in the database as neural modules if they meet the strict definition of such. These gross neuroanatomical region database entries are primarily to help orient neural module placement in the 3D volume of the brain.

There has been much research into defining a canonical brain atlas coordinate system so as to make the correlation of imaged regions from different brains possible. This is called brain warping [Toga 1999]. This is an extremely difficult problem in general since different people do not even share the same number of cortical gyri and sulci. **For this project, it is much simpler to just adopt the intermediate resolution voxel atlas from the single brain slice library as the canonical atlas.**

3.5 Modules: cortical regions, subcortical nuclei

The brain and spinal cord can in principle be divided up into well-delineated cortical regions and subcortical nuclei, the exact definition of which will be provided shortly. Cortical regions (e.g. Brodmann’s areas) have a distinct 2D layered structure. Examples are regions of the cerebral cortex and cerebellar cortex. In contrast, nuclei have a less distinct 2D structure and are sometimes best described as 3D blobs. These are important distinctions since we will need to define coordinate systems for each type in the database.¹ This 2D vs. 3D aspect is the only difference recognized in the database structure that follows, so they are lumped together under the term “neuronal module”.

Each neuronal module consists of two types of neurons: intrinsic neurons and projection neurons. Projection neurons are cells in the module that project axons outside of the module (i.e. to other modules). A module’s intrinsic neurons, projection neurons, and incoming axonal input fibers are termed the triad of neuronal elements [Shepherd 1990] and form the basis for all discussions of connectivity within a module.

The definition of a neuronal module used here is as follows: **A neuronal module is a contiguous region of neural tissue with uniform (or smoothly varying) basic circuit**

¹ All coordinate systems in the database are actually 3D in nature; however, they are defined such that two of the coordinates parallel the natural layer organization of the brain region.

and interregional connectivity properties. For example, if one part of a neuronal module receives axonal projections from sensory surface A, then all other parts must receive projections from module A as well. If one part of a neuronal module projects axonal connections to module B, then all part must project to module B as well. If particular interneurons of type C reside in one part of a neuronal module, then type C neurons must reside in all parts of the module. As a corollary then, **module boundaries can be defined simply as places on a cortical surface or within the brain volume across which basic circuit and interregional properties change abruptly.**

A module entry in the database must define the coordinates and boundaries of the module within the gross anatomy atlas described above. It must also define a module coordinate system. This coordinate system is simply a positional (x,y) grid measured in millimeters for 2D cortical region modules, or a positional (x,y,z) grid measured in millimeters for a 3D subcortical nuclei modules¹. Like the coordinate systems setup for sensory surfaces, these coordinate systems will allow the precise description of thousands of axon projections and millions of neural connections using a concise language of parameterized spline functions.

These neuronal modules are the central components of the neural connectivity database. Sensory surfaces can also be considered a type of neural module having a location and extent in the gross anatomy database and having an associated coordinate system. In order to maintain a unity of expression in the syntax of this database, discrete sensors and effectors will be assigned pseudo-coordinate systems (think of the motor maps in the primary motor cortex) and will thus be considered as neural modules.

Having defined sensory and effector modules, all macroscopic information flows can now be defined by module-to-module projections. These projections are termed “pathways” and are defined in a later section.

3.6 Neural layers

Within each module there may be millions of neurons; however, there will be at most several dozen *classes* of neurons. In a cortical region module, each class of neuron will be distributed roughly in a layer a certain distance from the cortical surface. We will define a set of neural layer entries under each module entry in the database, **one neural layer for each distinct class of neurons within the module**². Neural layers adopt the same 2D coordinate system as the parent module. The database entry for a neural layer would use this coordinate system to describe a density function, $\rho(x,y)$, for the neuron cell bodies in that layer. There would also be a distribution function in the z direction defining the statistical range of cell body position from the cortical surface.

¹ Notice the problem of finding and defining cross-individual or cross-species landmarks, in order to define these coordinate systems, is eliminated by populating this database with reference to a single brain physical slice library.

² This definition of neural layer is not the same as is traditionally used in neuroscience. Traditional layers describe gross histological appearance of cortical regions under the light microscope. Thus, it is said that the cerebral cortex has a six-layered structure. This definition is of little use in a complete neural connectivity database; instead, the spirit of the term is kept by defining neural layers as sets of neurons with distinct properties and presumably distinct connectivity statistics. It also preserves the idea that each neural layer forms a 2D sheet with all the computational properties inherent in that form.

Neural layers are the central entities that form pathway connections. Every axonal pathway extends from one neural layer (consisting of some class of projection neurons) of one module to one neural layer of another module.

3.7 Pathways and projective fields: the brain's block diagram

A pathway defines a projection of perhaps millions of axons from one layer in a module to a different layer in a different (or in the same) module¹. The database entry for this pathway concisely summarizes these million axons by defining two parameterized functions. These functions are defined with respect to the two modules' coordinate systems, (x,y) and (x',y') , defined over the projecting and the target layer respectively. First, a topological mapping function, $(x',y') = F(x,y)$, is defined which maps each point on the projecting layer to a point on the receiving layer. Second, an axonal arbor radius function, $r(x',y')$, is defined. That function defines the "projective fields" of each projection neuron onto the target neural layer. We do not need to define the total number of axons or their density function, since these are already defined for the neurons in the projection layer.

It is extremely important to remember that the topological projection function above does not constitute a "cortical map", since cortical maps are functional entities not belonging in this anatomical database. For example, it is well known that the somatosensory maps of the fingers in SI can be modified during adulthood in both monkeys and humans; however, this does not mean that the basic topology of the axonal projections of the dorsal column-medial lemniscus system have been modified. More likely, changes in synaptic strength in the target SI neurons cause the functional response of those cells to change. This database is concerned only with the neuroanatomy projections that *constrain* these functional changes.

Having defined a set of modules, a set of layers within each module, and a set of pathways connecting all layers to all other layers, we have finally defined the brain's complete top-level connectivity block diagram. Notice that this block diagram is complete² in the sense that it not only defines interregional pathways, but also defines all layer-to-layer connections within the same module. These intramodule pathways are the foundation of all regional basic circuits. Even though this block diagram does not yet contain the complete information of the database, it already provides a skeletal framework of information flows in the brain.

Even this level of detail, without the specifics of the neurons and synapses within the modules (discussed below), would prove invaluable for generating and constraining system level models of brain function. It must be stressed that even the most advanced

¹ This does not imply that a set of projection neurons send axonal projections only to a single class of target neurons. There are many examples in the neuroscience literature that contradict that oversimplification. The definition of pathway here requires defining multiple database pathway structures for each target cell population.

² Recall that a neural layer object is defined for every cell type. Pathway objects are defined for every connection between layers (i.e. neuron types), so the database has the statistics of numbers and connectivity for all pairs of connecting neuron types. Put another way, every synapse in the brain would belong to one of these pathway objects.

maps of brain region interconnections available today do not contain this level of detail and fall far short of even providing a comprehensive and certain list of brain modules.

3.8 Neurons and neuron components

We have left out some crucial details of anatomical connectivity in the above discussion. For example, when we enter into the database a pathway connection from projection neurons in one module to neurons in another module, just how do they synapse on these target neurons? Are there hundreds of synapses on the target or just one? Do they synapse on the distal parts of the neuron's apical dendrite or synapse on the basal dendrites? So far we have succeeded in describing billions of neurons and axon projections with perhaps only a few tens of thousands of database entries and parameterized spline functions. Will we now have to resort to describing billions of neurons separately in order to capture neuron morphology and synapse level connectivity? The answer must be no, because that approach would yield a worthless database of too many unorganized facts. Instead, we will continue the methodology of using statistical, parameterized functions to describe, concisely and precisely, the statistics of millions of neurons and connections at once.

Each neural layer object is defined for only a single class of neurons in a single neural module. Assuming the brain has 1000 modules, and each of these modules has on average 10 cell types having distinct connectivity and morphological properties, then there would be a total of 10,000 distinct neuron objects stored in the database (one neuron model associated with each layer object).¹

Each neuron object would first divide the neuron into a set of *neuron components*. Neuron components are functionally distinct regions of a neuron, such as the soma, apical dendrites, and basal dendrites of a cortical pyramidal cell. Depending on the statistics of neural connections seen in a particular brain region, a more specific subdividing of the neuron may be in order. For instance, the basal dendrite component may need to be further subdivided into distal and proximal, or even into several main branch components. The purpose of this subdivision is to allow the database to capture statistics on where synapses occur on neurons.

These particular components themselves would be described using statistical, parameterized functions. For instance, the apical dendrite component for a particular neuron type would have specified its depth with respect to the cortical surface, its radial extent within the layer, its bushiness within this extent, etc. These elements of neuron morphology cannot be specified as single numbers since they may vary across a module. Instead, each morphological parameter would itself be described using a parameterized function defined over the coordinate system of the parent module. Also, the specific branching pattern and membrane characteristics of the subcomponents could be concisely stored as lists of parameters used in generative branching algorithms (i.e. neural growth algorithms).

¹ Since neurons in different modules often share many characteristics (e.g. cortical pyramidal cells), these 10,000 neuron objects could be efficiently built up in a hierarchical fashion from a much smaller number of neuron prototypes.

3.9 Synapses and microcircuit-multiports

An axon projecting into a module will arborize and synapse on particular component parts of target neurons. In many cases, synapses involve only a single projecting neuron and a single target neuron. However, there are many cases where the neuroanatomy gets much more complicated. Synaptic structures may involve three neurons simultaneously connecting in a very stereotypic, and clearly functionally relevant, manner. This level of neural organization between the level of dendritic trees and that of individual synapses has been termed by Shepherd as microcircuits [Shepherd 1990]. For the purposes of this database, we will group the neuroanatomy of synapses and microcircuits together under the term *microcircuit-multiport*.

Each neuron type will have a set of microcircuit-multiport objects associated with it. The database will contain a structural and pharmacological description of these microcircuit-multiports and label each “port” into this synaptic structure so as to allow pathway objects to reference them.

3.10 Basic circuit statistics: pathways revisited

Pathway objects are defined in the database for every bundle of axons or axon collaterals projecting from one neural layer to another. Using the above definition of neuron components and microcircuit-multiports, each pathway can now be augmented with information on just where on the target neuron synapses are formed.

Let’s summarize the information a pathway object contains. A pathway object references a projection layer and a target layer. For the projection layer, the particular axon collateral type is referenced. For the target layer, the particular labeled port on a targeted microcircuit-multiport is referenced. This set of pathway data records the same information on connectivity that is usually supplied in basic circuit descriptions of brain regions. For example, the standard “cerebellar circuit” specifying connections between purkinje cells, granule cells, basket cells, stellate cells, mossy fibers, etc., could be described in this database by several of these pathway objects.

Recall that pathway objects also specify the mapping topology of axonal projections. The database also provides information on the extent of the axonal and dendritic arbors. Referencing these statistics, a pathway object can now further specify the *probability* of one neuron synapsing on another. For any particular point in the projecting layer, an axonal arbor radius is specified in the target layer. For any particular point in the target layer, a dendritic arbor radius is specified. Now, after recording that information, the area of overlap between axonal and dendritic arbors is easily calculated. The probability of synapse formation can be recorded in the database as a function of this area of overlap¹. **This information forms the basis of all statistics on divergence, convergence, and receptive field characteristics in the brain.** These statistics are crucial to the understanding of information processing in the brain.

¹ The book *Corticonics* [Abeles 1991] provides an in depth discussion on how to describe the probabilistic connectivity of the cerebral cortex using parameterized functions for dendritic and axonal arbor radii, etc.

3.11 Database as neuroembryology

To get a feel for the information that the above database does and does not contain, imagine that we were given such a database with all the neuroanatomical facts filled in. A block diagram of all the brain's information flows could easily be generated from such a database, and this could be used to *constrain* functional hypothesis. The information could not, however, be used to *predict* function.

The database's description of the brain is in terms of regions, pathways, and statistics of connectivity. In this way, it provides the same information that the genes provide. The human genome has less than 100,000 genes, and so clearly it does not specify each connection in the human brain. What the genome does do is to direct the *growth* of billions of neurons, thus giving rise to gross brain structures, regions, pathways, and eventually synapses. In this way, a complex information flow is specified in the genes, but specific computation is not. The genes also encode for complex biological mechanisms that endow each region of neurons with the ability to adapt synaptic strength and number based on local firing history. These genetically determined regions and pathways and local adaptation mechanisms combine in an intricate and almost totally unknown way to slowly generate the computational algorithms seen in adult brains.

A complete neuroanatomical database, like the one described here, has the job of describing the regions, pathways, and statistics of synaptic connections. It is then the job of other cognitive scientists to determine the local adaptation mechanisms for each region and the eventual computational function each region supports.

3.12 Summary

The above proposed database structure is still inadequate to describe all the complexities of the brain's neuroanatomy. Research into the optimal description of brain connectivity must evolve alongside our growing knowledge of neural function. The hope, however, is that this description has given the reader an appreciation for the level of detail that would be addressed by a single brain physical slice library. Current research methodologies using tracers and microinjection in thousands of different animal brains may never add up to the precise and complete description of brain connectivity described here. It is argued, however, that an approach using a single brain and ultrastructure reconstruction could be up to the task.

4 Creating a Physical Slice Library: Operation, Technical Challenges, and Reasonable Milestones

4.1 Introduction

The key concept of a single brain physical slice library is this: **A well preserved human brain can serve as its own exabyte (10^{18} bytes) atlas of neuronal macro-, micro-, and ultrastructure as long as automated instruments exist that can provide random access to any piece in a reliable and reasonably timely fashion.**

The ultimate slice library, as discussed in the executive summary, must be of a human brain. It is human intelligence that we are ultimately interested in unraveling the secrets of. However, if the physical slice library technology can be proven viable on an animal model, then the only obstacles in applying it to a human brain will be ones of scale and of preservation protocols. A logical milestone to shoot for is to first produce a mouse brain physical slice library. A mouse brain is approximately 3000x smaller in volume than a human brain, yet it is a good model of general mammalian brain organization. Large-scale projects like the Harvard High-Resolution Mouse Brain Atlas have already demonstrated the ability to slice and LM image a whole mouse brain at $\sim 10\mu\text{m}$ resolution in all directions [High]. Also, much of the TEM neural reconstruction research has been developed using mouse or rat models [Synapse]. For these reasons, this section will first discuss the details of a mouse brain physical slice library, following with a short discussion of how to scale the technique up to a full human brain.

4.2 Overview of a single brain physical slice library

The technique to be ultimately developed is as follows: A mouse brain (and spinal cord) is fixed *in situ* via intravascular perfusion with mixed aldehydes. The brain is then stained for LM and TEM, and embedded in a resin block. The brain is then sliced axially in a specially designed automated ultramicrotome producing $1\mu\text{m}$ thick slices. The brain volume of a mouse is approximately $1\text{cm} \times 1\text{cm} \times 0.5\text{cm}$, so this procedure would produce 5000 $1\mu\text{m}$ thick, 1cm^2 area slices.

During slicing, high-resolution color light micrographs ($1\mu\text{m}$ resolution) are automatically taken by scanning the entire 1cm^2 area of each slice sequentially with a driven-stage digital microscope. Images obtained are stored in a voxel volume atlas. The final volume will contain 5×10^{11} color voxels requiring a storage space of 1.5 terabytes. The thin slices, staining, and LM techniques used to create this $1\mu\text{m}^3$ resolution atlas should render visible the cell bodies and larger neural processes so as to provide landmark correlations with the ultrastructure imaging instruments.¹

¹ If properly stained, this LM atlas should also allow large neuronal process following and some degree of pathway reconstruction without the need for ultramicroscopic instruments. This LM atlas is designed to both orient and complement the higher resolution ultramicroscopic experimental studies.

Each of the 5000 1cm^2 , $1\mu\text{m}$ thick slices are automatically placed on specially designed carriers during the automated slicing process. These carriers are designed to be robotically shuttled into and out of an environmentally controlled automated storage and retrieval system. The carriers are also designed to allow robotic shuttling between this storage system and the vacuum chambers of specially modified ultramicroscopes¹. For use with a TEM, the carriers would also be specially designed to have rotating TEM support grids, which gently slide the tissue sample a few hundred microns in one direction in order to uncover grid occluded sections of the tissue automatically. This allows all the tissue to be TEM-imaged without the need for remounting slices.²

Included among the ultramicroscope instruments would be an intermediate-high voltage transmission electron microscope (IVEM) such as the JEOL-4000EX with 400 KeV accelerating voltage capabilities. This voltage has been demonstrated to provide sufficient signal-to-noise ratio to enable electron tomographic reconstruction in tissue slices $1\mu\text{m}$ thick [Soto 1994]. The IVEM is augmented with motorized tilt, etc. to allow remote control of all microscope functions during a tomographic tilt series. The operating precision should allow for tomographic reconstructions with resolutions significantly better than 100nm, allowing individual dendritic spines to be imaged.

Also included among the instruments would be an imaging secondary ion mass spectrometer (SIMS) instrument like the Cameca NanoSIMS-50, which can simultaneously record five mass-selected images (i.e. imaging ions with different atomic numbers or different isotopes with the same atomic number) by sputtering atoms off the surface layers of the neural tissue slices. The finest resolution for SIMS instruments applied to biological applications so far is 50nm, with the ability to produce 3D reconstructions by successively sputtering off surface layers during sequential imaging. Imaging the osmium used for TEM membrane staining would allow the structure of neurons to be recorded. Imaging other elements should allow a limited degree of pharmacological data to be collected.³

The robotic shuttling system and special tissue carriers should be designed to allow other instruments to be added as necessary. These may include instruments that have lower intrinsic resolution, but which can provide detailed pharmacological data, such as neurotransmitter types, mRNA expression, membrane protein property data, etc.

¹ Some instruments may not be able to accommodate 1cm^2 slices for imaging, in which case either smaller slices or special modifications will be required. This may be best accomplished by increasing the complexity of the slice carriers to enable carousel action, as opposed to increasing the total number of slice carriers to be handled. Clearly, this is one of the areas requiring technological innovation.

² This technique is proposed as one possible solution to the TEM support grid occlusion problem. It has not yet been demonstrated.

³ Both the TEM and SIMS instruments degrade the neural tissue as it is being imaged. This degradation is one of the main factors limiting resolution. Already imaged tissue sections will need to be tagged by the controlling computer as damaged, and additional requests to image those regions will be rejected or simply referred to stored voxel data. Steps must be taken to prevent the imaging of one section of tissue from degrading the future imaging of a contiguous section of tissue.

4.3 Operation

It is envisioned that the entire $1\mu\text{m}^3$ resolution LM-imaged volume be available online along with all the raw TEM voxel data for regions which have already been imaged at the ultrastructure level. Any 3D reconstructed neurons and mapped pathway structures derived from the raw voxel data could be provided online as soon as they are generated. Software tools for manual and automatic reconstruction of neuron morphology are still in a primitive state of development, and this raw voxel data should greatly stimulate software research in this area.

Online programs would allow a remote researcher to scrutinize the available data, and also enable the researcher to plan future imaging sessions by highlighting volume regions in the LM atlas to be imaged. After suitable financial arrangements are put in place (in order to share operating costs among participating research groups), the remote researcher would submit a batch volume-imaging job to the main controlling computer.

The main controlling computer would compile this batch volume-imaging job into an efficient sequence of imaging sessions, attempting to minimize sample loading times, tissue beam damage, etc. The computer would then automatically retrieve slices from storage, load the slices into the vacuum chamber of the desired instrument, scan to the desired X-Y coordinates within the slice, perform any necessary instrument calibration, and begin imaging operations. A batch job may take hours to days, and the researcher would receive periodic updates during this time to allow her to redefine the volume to be viewed.

In the above example, the researcher knew approximately the location and volume extents of brain tissue she desired imaged. This may be the case in some research into neuron morphology and region local connectivity studies. However, another researcher may be interested in following axons or dendrites over relatively long distances. In this situation, the volume to be imaged is not known beforehand. In cases like these, the researcher would be given the option of specifying a particular neural process to follow, and relatively simple axon tracing algorithms would automatically steer the imaging sequence to always image a small volume around this neuronal process. In this way, the computer could automatically and efficiently trace axonal pathways with little human oversight. **This automatic process-tracing mode is envisioned as the main operational mode of the single brain physical slice library.**

4.4 Types of experiments

Four main classes of experiments to be performed using the single brain physical slice library are envisioned. These are:

- Single neuron and neural subunit reconstruction experiments
- Regional basic-circuit reconstruction experiments
- Region and region boundary determination experiments
- Interregional pathway following experiments

4.4.1 Single neuron and neural subunit reconstruction

Neural subunit reconstruction experiments are the type of 3D neuronal reconstruction experiments performed today. For instance, Boston University's Laboratory of Synapse Structure and Function headed by Kristen Harris has forcefully argued for the necessity of 3D ultrastructure reconstruction techniques in order to provide accurate, unbiased statistics of neural structure. Specifically, the vast amount of research into what factors affect hippocampal synapse distribution statistics has been hindered by the field's reliance on 2D electron micrographs to determine differences between experimental and controlled cases, as well as between laboratory data comparisons. The technique of serial electron micrograph 3D reconstruction has proven much superior in determining statistics of spine shape and size, spine density per unit dendrite length, etc. [Harris 1994] [Fiala 1999]

A single brain physical slice library would allow small volume ($<1000\mu\text{m}^3$) 3D reconstructions of this type to be performed anywhere within the brain. In addition, successive small volume reconstructions could focus on different pieces of the same neuron. In this way, complete 3D reconstructions of neurons could be produced. This type of single cell reconstruction experiments could form the basis of a cell-centered database of the type currently being designed at NCMIR [CCDB]. Eventually, every neuron type in the brain would have several prototype reconstructions stored online. These prototypes would form the basis of automatic algorithms that could classify neurons into specific classes.

4.4.2 Regional basic-circuit reconstruction

Imagine obtaining the entire TEM voxel data for a 1mm^3 volume of neural tissue from the primary visual cortex, extending all the way from the cortical surface to the white matter. If this could be imaged at 10nm resolution, it would contain 10^{15} voxels, and would represent an equivalent of 100,000 neurons. This volume would cover the full extent of many neurons' dendritic arbors, their local axon collateral arbors, and the full axonal arbors of many input fibers projecting from distant regions. When properly analyzed, this data would reveal the region's basic circuit connectivity in a level of detail unheard of today.

Using more intelligent volume sampling methods, the same regional connectivity information could be obtained without the need to image a full 1mm^3 volume (which would take a prohibitive amount of time). In a single brain physical slice library, the random access nature of the automated storage and retrieval system would allow the connectivity of a region to be built up using normal search techniques.

Imagine an automated imaging session that begins by tracing an input fiber projecting into the region from the white matter. By imaging successive $10\mu\text{m} \times 10\mu\text{m} \times 1\mu\text{m}$ slice regions around the axon, the controlling computer follows this axon until it arborizes. Following successive branching, the computer eventually comes upon a synapse this axon makes with a neuron local to the region. The computer's process-following algorithm can then *jump across* this synapse and begin following this neuron's dendritic tree inward toward its soma. Having arrived at that neuron's soma, the

computer could then jump to following its axon, etc. The computer would not simply do this depth-first-search through a region's neural connectivity, but would also follow, in parallel, several branches in a breath-first-search manner. **In a sense, the computer is searching the region's neural connectivity by pretending it is following a single action potential along an input fiber and continuing to follow its propagation throughout all the neurons of the region.** Using this method, the full neural connectivity of a region could be determined having only sparsely imaged the 1mm³ volume. Perhaps only one thousandth of that total volume would have been imaged in this way, with a drastic reduction in time required.

4.4.3 Region and region boundary determination

One of the greatest advances in the history of neuroscience was Brodmann's parcellation of the human cerebral cortex into several dozen distinct regions. He did this by comparing only the cytoarchitectonic differences between regions of the cortex as seen in a light microscope. Brodmann simply sampled the vast cortical surface at thousands of points and grouped those nearby points that appeared to have the same layered structure under the LM as being part of the same region. As important, if two contiguous points appeared different, he inferred a region *boundary* existed between them. At the time Brodmann was performing his original research, few neuroscientists would have guessed that the regions discovered in this manner would have so closely paralleled later functional experiments. We know today that this division of the brain into physical bounded regions, each having distinct neural structure and connectivity, forms the basis of computation in the brain.

One of the most exciting experiments that a single brain physical slice library would make possible is a modern day repeat of Brodmann's classical cytoarchitectonic cortical mapping experiment. A random sampling of the cerebral cortex (or of any subcortical structure comprised of many distinct regions like the thalamus) could be undertaken using the techniques described in the above sections. Since these ultrastructure reconstruction techniques provide infinitely more information on a region's structure than the classic LM studies, these random sample points should be easily grouped into distinct regions.

Having grouped these sample points, a pair of points is chosen which are nearby each other but which have been determined to belong to two separate regions. An imaginary line is drawn on the cortical surface connecting these two points. It is clear that somewhere on that line there exists the boundary between these two regions. The random access nature of the physical slice library is used to perform a "binary search" along this line for this region boundary¹. Once a region boundary is found, **the boundary itself can be followed** using the random access abilities of the slice library.

By repeating the above procedure for every region boundary in the cortex, a complete map of all regions in the cerebral cortex could be produced. Boundaries between regions could be defined with micron level precision. Since this entire process

¹ Certainly not all interregional boundaries will be absolutely distinct. Boundaries will, however, be betrayed by a relatively abrupt change in regional structure, which can be discerned using the techniques described above.

would have been performed on a single brain, no between-brain comparisons or coordinate warping procedures would have been required.

4.4.4 Interregional pathway following

Having reconstructed the basic circuit of a region using the techniques discussed above, many input and output fibers of that region would have been identified. Making use of automatic axon following algorithms and the random access abilities of the physical slice library, these projecting fibers could be followed as they course long distances through the white matter or through other brain regions. Eventually they would link up with their source neuron (for fibers projecting into the original region) or be seen to arborize (for fibers projecting from the original region). If this termination point falls within a region whose boundaries have already been identified using the above methodologies, then an interregional pathway between these two regions would have been established.

By performing multiple axon following experiments between these two regions, the topological mapping function of the projection between the regions could be determined. Also, because this technique provides exact knowledge of which neuron in one region connects to which neuron in another region, *interregional circuits* may be found to exist that are currently beyond the techniques of modern day neural mapping to discover.

To fully appreciate the improved pathway mapping abilities that a single brain physical slice library could allow, imagine starting a pathway mapping experiment by following an afferent fiber from a retinal ganglion cell, then automatically following this axon until it synapses on a relay neuron of the LGN, then following the axon of this LGN relay neuron to V1, and further following the axon of this V1 neuron where it leads, etc. In this way, the flow of visual information in the brain could be traced from region to region, eventually completing its journey in primary motor cortex neurons projecting to spinal motor units. **Applying a combination of depth-first-search and breath-first-search at branch points in this pathway tracing experiment would trace out all the parallel information flows and information loops of the brain.**

The above four classes of neuroanatomical mapping experiments, which could be performed using a single brain physical slice library, demonstrate that such a system would indeed be able to provide the detailed neuroanatomical information required for a complete neural connectivity database of the type described in section #3 above.

4.5 Major technical challenges

Today's normal use of ultrastructure imaging instruments can be summed up succinctly: "Just get me one tiny sliver of properly fixed, stained, sliced, and imaged tissue containing an example of the cellular structure I want to study. I don't care if the rest of the brain ends up on the laboratory floor." Most technical challenges to the successful construction of a single brain physical slice library stem from the need to break with this "tiny sliver" mentality. Even though today's neuroscience instruments

have imaged the brain at all size scales from the macroscopic to the molecular, we have not yet truly bridged this incredible range of scales in a unifying way. **The key technical breakthrough that a single brain physical slice library must accomplish is a true union of macro-, micro-, and ultrastructure size scales in a single imaging experiment.**

In order to accomplish this, every step in today's standard neural ultrastructure imaging protocols must be reviewed and possibly modified. The protocols for fixing, staining, embedding, and slicing will have to be modified to produce uniform ultrastructure preservation throughout the entire brain and not just in the region currently under study. Ultramicrotome slicing today is geared towards producing just a few slices having areas of only a few mm² each. Large volumes of a tissue block are sacrificed for support reasons or due to microtome errors. This has presented a major obstacle to larger scale serial electron micrograph reconstruction efforts. The Boston University group, mentioned earlier, has been limited to using only a few dozen 50nm slices for its 3D reconstructions because of this obstacle. The serial electron tomography work of Soto *et al.* uses high voltage electron microscopy and much thicker sections to avoid this obstacle. However, there are resolution limits that the use of thick-slice tomography introduces [Soto 1994]. It seems that the only practical way of making a physical slice library is to use reasonably thick slices, so research into improving current resolution limits may be required.

The need to randomly access any piece of a slice in a TEM also presents problems not addressed in today's methods. The metal grid used to support samples in a TEM occludes large areas of neural tissue in a slice. This is tolerable when imaging a few μm³ of tissue, but completely intolerable when following particular dendritic processes for hundreds of microns as would be done in a single brain physical slice library. There are many possible solutions to this and other obstacles; however, developing these new technologies and protocols will require research.

In order to produce a usable physical slice library, the technology and software must be available to produce imaging results in a timely and user-friendly fashion. Automation of imaging and slice retrieval will go a long way toward providing timely imaging results. There are currently projects going on at NCMIR that aim to allow truly real-time tomographic reconstruction (voxel data from raw tilt-series images) while remotely operating an IVEM. Combining this research with the intelligent, sparse-volume imaging techniques described above should allow researchers to conduct interesting experiments in a reasonable amount of time.

Software for automatic 3D reconstruction and automatic process-following is in its infancy, but some significant strides have been made [Montgomery 1996]. In order to translate the massive amount of raw data generated in volume imaging experiments, fully automatic 3D reconstruction software must be developed. Once again, there are many ways to approach this problem, but a solution will entail much additional research. It is hoped that the eminent arrival of a single brain physical slice library will provide a new momentum for this type of software research.

4.6 *Human brain physical slice library*

The goal of mapping *human* brain tissue presents its own set of problems, not the least of which has to do with preservation of neural ultrastructure without the ability to perform *in vivo* perfusion of fixatives before death. However, the impetus for a human brain physical slice library is definitely there. Consider this quote from a 1997 NIMH call for proposals on new methodologies for determining pathways in aldehyde-fixed human brains [NIMH 1997]:

“There are very few tract-tracing approaches that can be applied to post-mortem tissue, and these approaches have severe limitations such as being capricious, extremely slow, or being able to demonstrate connections only over very small distances. Thus, despite the wealth of knowledge that has accumulated regarding neural connectivity in a variety of species, very little is known about the detailed connections of the human brain.”

Or this more recent (2002) NIMH call for proposals [NIMH 2002]:

“Little is known about the details of the connectivity of the human nervous system, because the best tract tracing techniques are invasive and require the deposit of substances *in vivo*. Methods that would be applicable to post-mortem tissue would allow significant progress in connectional studies of human tissue, as well as non-human tissue, particularly with regard to the development of connections and the connections of structures not easily accessed *in vivo*.”

A single brain physical slice library would provide this ability to map post-mortem brain tissue, and is one of the few options neuroscientists have for in-depth mapping of human brain connectivity. This fact alone should demonstrate the tremendous potential that a single brain physical slice library represents.

Scaling up from a mouse brain to a human brain will increase the number of slices by more than three orders of magnitude, will require redesigning the fixation, staining, and embedding protocols, and will require new automated slicing equipment. However, most other technologies developed on the smaller scale mouse project will be directly transferable to a human brain physical slice library.

5 BrainCAD: Software Tool for Creating Neural Databases from Slice Library Data

In this section we propose a software tool (BrainCAD) that could provide the link between the imaged raw data provided by the single brain physical slice library described in section #4 and the creation of a neural connectivity database like the one described in section #3. The most important point to stress about this software tool for creating brain connectivity databases is that each database represents a *hypothesis* of neural connectivity only. Researchers may image thousands of neurons using the single brain physical slice library, and these images represent *certain* knowledge about the brain. **However, researchers using a tool such as BrainCAD to create brain connectivity databases will be extrapolating from the particular connectivity of these thousands of imaged neurons to the statistical connectivity of billions of neurons.**

5.1 Overview

Section #3 above proposed a structure for a complete neural connectivity database. Many of the entries in that database structure are *geometric* in nature. For instance, modules are defined as geometric volumes with distinct boundaries and are contained within larger anatomically labeled 3D structures. Each of these modules has a local coordinate system (a mathematical manifold) defined upon it, such that any variation of intraregional and interregional neural statistics can be described with respect to the position of a neuron cell body within the module. In the database, pathway objects describe all axonal projections between modules as geometric coordinate mapping functions from one manifold to another. None of these entries can be efficiently entered into a database using text-based entry methods, and none can be efficiently displayed using a text-based display. We envision a CAD-style user interface to this database of neural connectivity, allowing users to graphically define region boundaries, coordinate systems, and pathway mappings by point-and-click mouse operations on a 3D model of the brain.

In contrast to a standard mechanical CAD package, BrainCAD would not start out with an *empty* drawing region. The single brain physical slice library's raw voxel data from the full LM and partial TEM imaged regions would provide a reference backdrop to all user operations in BrainCAD. Gross anatomical regions could be graphically defined by highlighting voxel subsets in the full LM atlas. Using a mouse to pick voxels on the cortical surface, a user could call up a function that automatically interpolates a 2D manifold through these points. This would define a surface object for the cortex. Through further mouse operations, the user could define hypothesized cortical regions, and define hypothesized axonal projections between these regions as coordinate transformations specified by a series of point mappings.

5.2 Operation

We will walk through an example database entry to get a feeling for how neural connectivity databases could be constructed based upon raw voxel data coming from single brain physical slice library experiments.

5.2.1 Single brain physical slice library raw experimental data

Imagine that researchers have performed a set of imaging experiments on two cortical regions. For explicitness, suppose one of the regions (A) is in the primary somatosensory cortex, and the other (B) is a region of the parietal association cortex. Assume that using the techniques described in section 4.4.3, the researchers have identified the exact boundaries of region A; however, assume region B's boundary remains only loosely known. Furthermore, assume researchers choose 10 points uniformly distributed in region A to perform more detailed imaging experiments. At each of these points, initial imaging is performed in order to find a projection neuron at each point of the type known to project to the parietal region B. Assume, in a parallel imaging experiment, that all 10 of these projecting axons are followed all the way to their axonal arbors in region B. Also assume that by following a few branches of each axonal arbor, the computer collects imaging data that provides additional statistics on arborization radius and synaptic termination sites.

5.2.2 Database entry using BrainCAD

From the above experiment, a hypothesis of neural connectivity can be entered into the database in the form of new module, layer, and pathway objects. First, the user would load into BrainCAD the latest version of the single brain physical slice library voxel imaging data. She could then navigate and zoom in on the location of region A, and define a new module object in the database for that region. Let us assume that a gross anatomical object already exists in the database for the cerebral cortex, and a 2D coordinate manifold has already been entered into the database defining a set of coordinates upon the cortex. The user could simply “attach” her mouse pointer to this manifold and draw a contour defining the boundaries of region A so as to align with the boundary exposed during the imaging experiments. The user could also choose to define a new coordinate system for this particular module in order to refine the precision of future entries. In a like manner, a new module is entered for region B; however, for B, since the exact boundary has not been imaged, the user “freehand” draws a hypothesized boundary.

Having defined a new module for region A, the user would now define a set of neuron types that are present in that module. Each of these neuron types also gets associated with a new layer object under the new module. All 10 of the projection neurons imaged are of the same type of neuron, and so all represent the same layer object. A target layer is also defined in module B.

Finally, a new pathway object is defined between these two layers. This is done using a special software interface that mathematically flattens both the projection and target layers and displays them on the computer display side by side. Corresponding pairs of “landmark” points could be specified, via mouse clicks, on this pair of 2D layer displays. These landmark pairs define a coordinate mapping, $(x',y') = F(x,y)$, from module A’s coordinate system to module B’s coordinate system. This mathematical mapping function is simply a smooth extrapolation between the specified corresponding points the user enters. This module-to-module coordinate mapping is the definition of the axon projections from module A’s projection layer to module B’s target layer.

In order to define this particular mapping function, the user can superimpose on this 2D display the 10 neuron cell body locations in module A and the corresponding 10 axonal arbors in module B (both gotten from the raw imaging data). Thus, 10 landmark pairs are displayed side by side. The user can simply click these 10 pairs of points and thus define a crude coordinate mapping between these two layers. The user, again using only mouse clicks on these two 2D layer displays, can also specify an axonal arborization radius function, $r(x',y')$, by using the raw data from the 10 axon arbor extents in module B.

Notice that the above coordinate mapping function defines the projection definitions of perhaps one million axons. However, this mapping function is literally a mathematical extrapolation of only 10 known neural projections. This is why it is important to remember that the entries in this neural connectivity database are provisional hypotheses only. If an eleventh neuron is imaged, that new landmark pair will undoubtedly shift the mapping function. Remember that this method of using statistical parameterized functions to concisely describe the microanatomy, interregional, and basic circuit connectivity of millions of neurons is what gives this database structure its unique precision and completeness.

Finally, note that if data is not available directly from physical slice library experiments, this does not mean that the user cannot enter into the database *hypotheses* on brain regions and connectivity gleaned from elsewhere. A user can use the same graphical entry methods to directly “freehand” draw these hypothetical regions and pathways. **Unlike the centrally stored and maintained single brain physical slice library data, individual researchers will keep and maintain their own neural connectivity databases, containing their current best hypothesis of neural connectivity relating to their own particular research. Researchers will share pieces of their database with other researchers in a collaborative fashion. All of these hypotheses being explored will, of course, lead to another cycle of basic imaging experiments using the single brain physical slice library. By using BrainCAD, the hypotheses have been directly stated with respect to the slice library’s LM atlas. This means that the researcher has in effect already designed the imaging experiment needed to test their hypothesis.**

5.2.3 Visualizing neural connectivity using BrainCAD

Once many modules, layers, and pathways have been entered into the database, a user could use BrainCAD to aid in her understanding of what these information flows

actually mean with respect to brain function. She could do this by using BrainCAD to project sensory surface representations (or pictures) through the various pathways' coordinate transformations in order to see receptive field statistics directly.

Assume that along with the above module entries for the somatosensory cortex and parietal cortex, other modules have been entered into the database, such as the sensory surface for the skin and a layer of Meissner's corpuscle touch receptors defined within that sensory module. Also assume the dorsal column's cuneate nucleus and the VPL thalamic nucleus have been entered as modules and that the pathway from the touch receptors on the skin through these nuclei have been entered into the database.

Given this more complete database, a user testing functional hypotheses could use BrainCAD to pose the following question: "What is the anatomically defined receptive field of a particular parietal neuron projected back on the surface of the skin?" BrainCAD could easily answer this type of query by using all the pathway object's $F(x,y)$ and $r(x,y)$ functions in reverse, and could display the result as a highlighted region on the 3D skin model as the user drags her mouse across the parietal cortex layer display.

One can imagine a whole host of other "visual queries" of the database. For example, imagine that a user clicks a point on the 3D model of the cortical surface, thereby causing the software to instantly highlight (in blue) all of the modules known to project to this selected module. Imagine that this click also causes the software to simultaneously highlight (in red) all of the modules that the selected module projects to. Those modules fitting into both categories are likewise highlighted in purple.

Less exciting graphically, but even more useful, would be the software's ability to generate block diagrams of neural connectivity of particular systems in the brain. A user could click on a subset of the modules in the cortex and print out a block diagram showing all the layer objects in these modules and their information flows.

5.2.4 BrainCAD interfaces with neural simulation programs

BrainCAD is envisioned not only as a tool for neuroanatomists summarizing their findings from experiments on the single brain physical slice library but also as a tool that other cognitive science researchers can use. Specifically, the neural layer objects and pathway objects in the database contain exactly the type of information needed for system level neural simulation programs like USC's NSL (Neural Simulation Language) [Arbib 2001]. It is envisioned that a user could ask BrainCAD to output a NSL description file along with a block diagram query of the type described above.

BrainCAD would also have tools for entering neuron type information into the database such as growth rules for prototype neural structural models. This type of database object was discussed above in section 3.8, and the raw data for such neural structural modules would come from experiments of the type discussed in section 4.4.1. We have not discussed this proposed function of the BrainCAD program due to space limitations. However, we note here that BrainCAD could also be directed to output individual neuron structural information to neuron compartmental model simulators such as Caltech's Genesis program [Bower 1998].

5.2.5 The “firewall” between raw image data and interpretations

Also available through BrainCAD menus would be a set of voxel image processing tools for 3D neural reconstruction and process following. These could be applied to the raw data from the single brain physical slice library in order to assist the user in all of the above database entry tasks. **It is envisioned that the raw imaged data be stored and maintained by a central institution, and that all interpretation of that raw data be relegated to BrainCAD-type programs in the hands of individual researchers.** This may seem counter intuitive since the central institution maintaining the physical slice library is required to use voxel image processing algorithms to perform directed (process-following) imaging, and would, therefore, have already reconstructed many neural processes. **However, we believe that the neuroscience research community will be better served by keeping this “firewall” in place between the absolutely certain raw imaging data and any interpretations of this data. Recall that the purpose of a single brain physical slice library is to foster collaboration between neuroscience researchers by providing a common experimental setup on a single shared brain. This collaboration can be fruitful only if each group is assured access to the raw data uncluttered and unbiased by competing hypotheses.**

5.3 Summary

Many programs already exist to assist neuroscientists. If a program like BrainCAD was written for independent use (without its intimate connection to the single brain physical slice library data), then it would be just another one of these programs. Used in this fashion, BrainCAD would allow some collaboration between researchers. However, the old questions of how to map experimental data from several different animal connectivity experiments onto the same database would cripple its intended function. **It is only when a program like BrainCAD is intimately connected with an ongoing connectivity experiment on a single brain that its true power for stimulating collaboration is unleashed.**

6 Conclusions

Today, the mandate of the Human Brain Project remains unfulfilled. The tremendous advancements in neuroscience technologies and methods have produced a flood of excellent results in a variety of specializations. However, the target of producing a complete mapping of the human brain's anatomical connectivity (even without any function-related data) continues to be a remote goal.

This paper has proposed an ambitious research program into the creation of a new microscopic brain-imaging system. We have argued that this single brain physical slice library would be capable of producing the brain imaging data needed to create a complete neural connectivity database. Such a system would require extensive technology development; however, it is firmly grounded in today's existing brain imaging technologies.

The undertaking of mapping the human brain's complete neural connectivity would require a massive collaboration between neuroscience researchers from many specializations. We have argued that a single brain physical slice library would enable this ultimate level of collaboration by allowing researchers from around the world the ability to perform their own, custom designed, detailed telemicroscopy experiments on the same physical brain. Unlike separate specialized experiments, the results from these brain-mapping experiments would immediately and easily be integrated because they are all performed on the same physical brain.

The scale of the proposed project is ambitious; however, the implications of success are staggering. The complete and accurate mapping of the macroscopic and microscopic neuroanatomy of the human brain would form a touchstone with which to constrain all other theories of how the brain functions. Such a complete database of brain connectivity would bring about a new era of collaboration amongst all cognitive science researchers, and this could finally allow the bottom-up theories of brain function to meet the top-down theories of human cognition. It is from *this* collaboration that a deep, mechanistic theory of the human mind may result.

7 References

- [Abeles 1991] Abeles, M. 1991. *Corticonics: Neural circuits of the Cerebral Cortex*. Cambridge University Press, Cambridge.
- [Arbib 2001] Arbib, M.A., Grethe, J.S. 2001. *Computing the Brain: A Guide to Neuroinformatics*. Academic Press, San Diego.
- [Bower 1998] Bower, J.M., Beeman, D. 1998. *The Book of Genesis: Exploring Realistic Neural Models with the General Neural Simulation System*, 2nd ed. Springer-Verlag, New York.
- [Cameca] Cameca NanoSIMS 50 Product Overview, “http://www.cameca.fr/html/product_nanosims.html”
- [CCDB] Cell Centered Database Project, “<http://pamina2.sdsc.edu/CCDB/index.shtml>”.
- [Fan 1993] Fan, G.Y., Mercurio, P.J., Young, S.J., Ellisman, M.H. 1993. Telemicroscopy. *Ultramicroscopy* 52:499-503.
- [Fan 2000] Fan, G.Y., Ellisman, M.H., 2000. Digital imaging in transmission electron microscopy. *Journal of Microscopy*, Vol. 200, Pt 1, pp. 1-13.
- [Fiala 1999] Fiala, J.C., Harris, K.M. 1999 Dendrite Structure. To appear in *Dendrites*, edited by G. Stuart, N. Spruston, M. Hausser.
- [Fiala 2001] Fiala J.C., Harris, K.M. 2001 Extending Unbiased Stereology of Brain Ultrastructure to Three-Dimensional Volumes. *Journal of the American Medical Informatics Association*. 8(1):1-16.
- [Fiala 2002] Fiala, J.C. 2002. Three-dimensional structure of synapses in the brain and on the web. *Proceedings of the 2002 International Joint Conference on Neural Networks*, May 12-17, Honolulu, Hawaii, pp. 1-4.
- [Goodhew 2001] Goodhew, P.J., Humphreys, J., Beanland, R. 2001. *Electron Microscopy and Analysis*, 3rd ed. Taylor & Francis, London and New York.
- [Gupta 1997] Gupta, K. 1997. *Human Brain Coloring Workbook*. Princeton Review Publishing, L.L.C., New York.
- [Harris 1994] Harris, K.M. 1994. Serial Electron Microscopy an Alternative or Complement to Confocal Microscopy for the Study of Synapses and Dendritic Spines in the Central Nervous System. *Three-Dimensional Confocal Microscopy: Volume Investigation of Biological Specimens*, Academic Press.
- [Hayat 2000] Hayat, M.A. 2000. *Principles and Techniques of Electron Microscopy: Biological Applications*, 4th ed. Cambridge University Press, Cambridge.
- [Hayworth 1998] Hayworth, K.J. 1998. *Mechanisms of Human Computation*. JPL Research Document D-16371. Jet Propulsion Laboratory, Pasadena, California.
- [High] High Resolution Mouse Brain Atlas, Havard University, “<http://www.hms.harvard.edu/research/brain/intro.html>”.
- [Jensen 1989] Jensen, F.E., Harris, K.M. 1989. Preservaion of neuronal ultrastructure in hippocampal slices using rapid microwave-enhanced fixation. *Journal of Neuroscience Methods*, 29 217-230

- [Kandel 2000] Kandel, E.R., Schwartz, J.H., and Jessell, T.M. 2000. Principles of Neural Science, 4th ed. McGraw Hill, New York.
- [Koch 1999] Koch, C. 1999. Biophysics of Computation: Information Processing in Single Neurons. Oxford University Press, New York.
- [Kohonen 1997] Kohonen, T. 1997. Self-Organizing Maps, 2nd ed. Springer-Verlag, Berlin Heidelberg.
- [Mai 1997] Mai, J.K., Assheuer, J., Paxinos, G. 1997. Atlas of the Human Brain. Academic Press. San Diego.
- [Marr 1982] Marr, D. 1982. Vision. W.H. Freeman and Company, New York.
- [Martone 2000] Martone, E.M., Hu, B.R., Ellisman, M.H. 2000. Alterations of Hippocampal Postsynaptic Densities Following Transient Ischemia. *Hippocampus* 10:610-616
- [Merkle 1989] Merkle, R.C. 1989. Large Scale Analysis of Neural Structures. Xerox PARC Technical Report CSL-89-10. Xerox Corporation, Palo Alto, California.
“<http://www.merkle.com/merkleDir/brainAnalysis.html>”
- [Montgomery 1996] Montgomery, K. 1996 Automated Reconstruction of Neural Elements from Transmission Electron Microscope Images. Dissertation University of California Santa Cruz.
- [Neuroscience] Neuroscience History on the Web, “www.ibro.org/secondary/sciissues/history_3.htm”.
- [Newell 1990] Newell, A. 1990. Unified Theories of Cognition, Harvard University Press, Cambridge.
- [NIMH 1997] Innovative Approaches For Microscopic Tract-Tracing Mh-97-001, Nih Guide, Volume 26, Number 2, January 17, 1997
- [NIMH 2002] Tract Tracing Methods and Tools. U.S. Department of Health and Human Services Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) Grant Applications PHS 2002-2, “http://grants1.nih.gov/grants/funding/sbirsttr1/section_10b.html”.
- [Perkins 1997] Perkins, G.A., Renken, C.W., Song, J.Y., Frey, T.J., Young, S.J., Lamont, S., Martone, M.E., Lindsey, S., and Ellisman, M.H. 1997 Electron Tomography of Large, Multicomponent Biological Structures. *Journal of Structural Biology* 120, 219-227
- [Rosenbloom 1993] Rosenbloom, P.S., Laird, J.E., Newell, A. 1993. The SOAR Papers (Vols. 1&2). MIT Press, Cambridge.
- [Shepherd 1990] Shepherd, G.M. 1990. The Synaptic Organization of the Brain. Oxford University Press, New York.
- [Soto 1994] Soto, G.E., Young, S.J., Martone, M.E., Deerinck, T.J., Lamont, S., Carragher, B.O., Hama, K., and Ellisman, M.H. 1994. Serial Section Electron Tomography: A Method for Three-Dimensional Reconstruction of Large Structures. *Neuroimage* 1, 230-243
- [Stelly 1995] Stelly, N., Halpern, S., Nicolas, G., Fragu, P., Adoutte, A. 1995. Direct visualization of a vast cortical calcium compartment in Paramecium by secondary ion mass spectrometry (SIMS) microscopy: possible involvement in exocytosis. *Journal of Cell Science* 108, 1895-1909
- [Synapse] Synapse Web, Boston University, “<http://synapses.bu.edu/>”.
- [Toga 1999] Toga, A.W., 1999 Brain Warping. Academic Press.

[Visioninst] SL50 Automated Slide Loader, “http://www.visioninst.com/product_sl50.html”

[Windhorst 1999] Windhorst, U., Johansson, H. 1999. Modern Techniques in Neuroscience Research. Springer-Verlag, Berlin.