Published in final edited form as:

Brain Imaging Behav. 2014 June ; 8(2): 323–331. doi:10.1007/s11682-013-9255-y.

Human Neuroimaging as a “Big Data” Science

John Darrell Van Horn* and Arthur W. Toga

The Institute for Neuroimaging and Informatics Keck School of Medicine of USC University of Southern California 2001 North Soto Street - Room 102, MC 9232 Los Angeles, CA 90089-9235, toga@usc.edu

Abstract

The maturation of *in vivo* neuroimaging has lead to incredible quantities of digital information about the human brain. While much is made of the data deluge in science, neuroimaging represents the leading edge of this onslaught of “big data”. A range of neuroimaging databasing approaches has streamlined the transmission, storage, and dissemination of data from such brain imaging studies. Yet few, if any, common solutions exist to support the science of neuroimaging. In this article, we discuss how modern neuroimaging research represents a multifactorial and broad ranging data challenge, involving the growing size of the data being acquired; sociological and logistical sharing issues; infrastructural challenges for multi-site, multi-datatype archiving; and the means by which to explore and mine these data. As neuroimaging advances further, e.g. aging, genetics, and age-related disease, new vision is needed to manage and process this information while marshalling of these resources into novel results. Thus, “big data” can become “big” brain science.

Introduction

Images of the human brain, in form and function, seem to be everywhere these days - on television, in glossy magazines, and on internet blogs worldwide. This is due, in many respects, to the incredible amount of information these images present and the sheer number of brain imaging research studies being performed to spy on the brain in action or at rest, to examine how it is built and wired, and what happens when things go wrong. Indeed, neuroimagers routinely collect more study data in a few days than was collected in over an entire year just a decade ago. These data are a rich source of information on detailed brain anatomy, the subtle variations in brain activity in response to cognitive stimuli, and complex patterns of inter-regional communication. Taken individually, these various data types would have once formed the basis for entire research programs. Now, with interests not only in multi-modal neuroimaging but the inclusion of co-occurring biological and clinical variable collection requiring linkage between geographically distributed researchers, neuroscience programs are rapidly becoming the brain-focused versions of projects more akin to those involving particle physics. The methods by which these data are obtained are themselves contributing to this growth, involving finer spatial and temporal resolution as MR physicists push the limits of what is possible and as brain scientists then rush to meet

*Corresponding author Phone: (323) 442-7246 jvanhorn@usc.edu.

those limits. It is safe to say that human neuroimaging is now, officially, a “big data” science.

Such examples of large-data, their promise and challenges, have not gone unnoticed. In the US, The National Science Foundation, the National Institutes of Health, the Defense Department, the Energy Department, Homeland Security Department as well as the U.S. Geological Survey have all made commitments toward “big data” programs. The Obama Administration itself has even gotten in on the act. In response to recommendations from the President’s Council of Advisors on Science and Technology, the White House sponsored a meeting bringing together a cross-agency committee to lay out specific actions agencies should take to coordinate and expand the government’s investment in “big data”, totaling \$200 million in support (see http://www.whitehouse.gov/sites/default/files/microsites/ostp/big_data_fact_sheet_final.pdf). Among the examples of “big data” featured at the meeting was – no surprise - human neuroimaging. Additionally, the recent announcement of the Brain Research through Advancing Innovative Neurotechnologies (BRAIN) Initiative (<http://nih.gov/science/brain/index.htm>) forms part of a new Presidential focus aimed at revolutionizing understanding of the human brain. Initiatives surrounding large-scale brain mapping are also underway in Europe (<http://www.humanbrainproject.eu/>; Frisoni 2010) and examples of large-scale brain data sets have been on full display at recent annual meetings of the *Organization for Human Brain Mapping* (OHBM; <http://www.humanbrainmapping.org>) in Beijing, China in 2012 and Seattle, Washington in June 2013.

However, as the richness of brain data sets continues to grow and the push to place it in accessible repositories mounts, there are many issues to be considered on how to handle the data, move it from place to place, how to store it, analyze it, and share it.

How Big is “Big”?

While size is a relative term when it comes to data, medical imaging applied to the brain comes in a variety of forms which each generating differing types and amounts of information about neural structure and/or function. Moreover, *in vivo* neuroimaging is not unimodal but, rather, remarkably diverse - examining brain form, function, and connectivity, and rapidly improving its ability to resolve finer spatio-temporal scales. Further advancements in magnetic resonance scanner field and gradient strength are an ongoing area of research (Barry, Strother *et al.* 2011). A brief look at the history of some central parameters for specific forms of functional and structural neuroimaging is illustrative.

Since its inception the collection of blood oxygenation level dependent (BOLD) functional time series has often been among the data types having the biggest storage footprint. While in the 1990’s, the earliest BOLD imaging studies used sampling intervals of one volume every 4 seconds for a modest number of slice planes needed to image the full brain. During the mid-2000’s, with advances in multi-channel coil technologies and the emergence of 3T imaging systems, 2 second intervals were the most frequently used for the same in-plane image size but having finer slice resolution. More recently, improved sampling methods

have made it possible to routinely obtain multiple BOLD image volumes of the whole head per second (Feinberg, Moeller *et al.* 2010).

Likewise, diffusion tensor imaging (DTI) has undergone its own steady progress both in the time to acquire data but also the degree of angular resolution by which to measure patterns of water molecule diffusion. Early diffusion imaging studies measured diffusion along 6 directions seeking to utilize diffusion signal to infer neural fiber orientation, quantifying this as a diffusion tensor, and to then link these tensor orientation patterns to estimate neural fiber pathways (Basser, Pajevic *et al.* 2000). Interest in examining finer degrees of angular difference in fiber orientations and to detangle crossing fiber pathways lead to using more diffusion directions through high angular diffusion approaches having 32, 64, or greater MR gradient directions (Zhan, Leow *et al.* 2010). More recent approaches have decomposed the spectrum of diffusion which is capable of then resolving upwards of 512 fiber directions in approximately the same time and volume sizes as the original DTI imaging sequences (Wedeen, Wang *et al.* 2008).

In each example, as the imaging technology has improved, been extended, or made faster, so too has the amount of brain data which can be obtained. Once these improved methodologies have proven robust and dependable, with analytic methods available with which to utilize them, researchers have been quick to adopt them – doubling or tripling the amount of data they can then gather per subject by doing so. Indeed, a simple examination of fMRI articles from representative issues of the field's touchstone journal, *NeuroImage*, indicates that since 1995 the amount of data collected has doubled approximately every 26 months (Figure 2). At this rate, by 2015 the amount of acquired neuroimaging data alone, discounting header information and before more files are generated during data processing and statistical analysis, may exceed an average of 20GB per published research study. This is likely to be an under-estimate for raw dataset sizes since, as noted above, advances in MRI physics are accelerating the pace at which data can be acquired per unit time.

Individually, the neuroimaging data sets from a single study themselves may not pose major difficulties for processing and analysis using existing algorithms and statistical methods. However, as data sets are amassed into large-scale databases, considerable challenges emerge. How to rectify between study differences in acquisition rates, resolution, scanning parameters, the presence of image artifacts, etc, take on importance – even the same make and model scanner installed at two different sites can generate sufficient differences in image quality to complicate combined analysis. Along with the increasing interest in gathering study data from across the lifespan or focused on specific patient groups, concurrently recorded electrophysiological time courses, and comprehensive phenomic meta-data on each study participant, wrangling, let alone interpreting, these data will not be for the feign of heart.

Big Neuroimaging + Big Genetics = REALLY Big Data

With the ability to obtain genome-wide sets of single nucleotide polymorphism (SNP) information becoming routine and the costs of full genomic sequencing rapidly becoming affordable, interest in linking the influence of genes on the brain is heating up. Soon,

exomic-level genetics will be readily at hand to replace SNP-based approaches. At several gigabytes per genome using Next Generation Sequencing (NGS) methods, for major brain imaging studies such as the Alzheimer's Disease Neuroimaging Initiative (ADNI) (Weiner, Veitch *et al.* 2012), with its initially available sample of 832 subjects, one can expect to store multiple petabytes of genetics data alone. More ADNI data are on their way, too. Coupled with the existing and ongoing collection of multimodal neuroimaging data types, this represents some really big data. As the bond between neuroimaging and genomics grows tighter, with both areas growing at incredible rates, disk storage, unique data compression techniques (like those proposed for genomics (Hsi-Yang Fritz, Leinonen *et al.* 2011); see also <http://www.sequencesqueeze.org/>) and data processing considerations rapidly become a scientific imperative.

Multisite Consortia and Data Sharing

Along with advances in the ability to obtain data there has been an increase in the numbers of multisite consortia like ADNI for examining the healthy and diseased brain. The demand for multiscale data in the investigation of fundamental disease processes has been recognized for several years (Jiang, Liu *et al.* 2008), as well as the need for cooperation across centers and even disciplines to integrate and interpret the data (Van Horn and Toga 2009b). Examples of multisite neuroimaging efforts can be found in the ubiquitous application of neuroimaging in health but also in devastating illnesses such as Parkinson's (Evangelou, Maraganore *et al.* 2009; see also <http://ppmi.loni.ucla.edu/>), psychiatric disorders (Schumann, Loth *et al.* 2010), and also the mapping of human brain connectivity (Toga, Clark *et al.* 2012). In addition to databases of aging and aging-related diseases, large-scale examples from the NIH-based National Database of Autism Research (NDAR; Hall, Huerta *et al.* 2012) and the Federal Interagency Traumatic Brain Injury Research (FITBIR; Bushnik and Gordon 2012) exist to gather neuroimaging, genetic, and phenomic data on autism and brain injury, respectively. The various "grass roots" collections of resting-state fMRI data maintained as part of the "1000 Functional Connectomes" project (http://fcon_1000.projects.nitrc.org/) (see Biswal, Mennes *et al.* 2010) and task-based OpenfMRI (<http://www.openfmri.org>) (Poldrack, Barch *et al.* 2013) are other notable examples.

What is more, pressures to share these data as openly as possible have put the onus on both the data collectors and database curators to store data efficiently and safely while providing data efficiently to anyone who would like to use it. However, inherent in multilaboratory projects are sociologic, legal, and often ethical concerns that must be resolved satisfactorily before they can work effectively or be widely accepted by the scientific community (Beaulieu 2001). While this often involves some expectations on the part of data collectors, etc, the exchange, storage, and computation on brain imaging data has many advantages including 1) maximizing the cost of their collection over the widest possible set of researchers, 2) allowing their use in new methods development, 3) promoting data re-analysis and repurposing, 3) new means for collaboration, 4) hypothesis generation, and 5) clever forms of visualization. In other words, archived data can be subjected to novel approaches which can highlight relationships not envisioned by the original data collectors and shed new light on important mechanisms which are then worthy of additional study (Van Horn and Ishai 2007).

Examples of Big Neuroimaging Databases

Despite the ever present challenges of archiving massive quantities of large brain imaging data, prominent examples for storing, sorting, and mining data exist. Key examples include XNAT Central (<https://central.xnat.org/>) (see Marcus 2007) which contains data from several thousand subjects and is playing a key role in the data management and informatics of the Human Connectome Project (HCP) (Marcus, Harms *et al.* 2013); the SumsDB effort (<http://sumsdb.wustl.edu/sums/index.jsp>) for cortical surface-based atlasing of neuroimaging results (Van Essen 2005); and the NIH MRI Study of Normal Human Brain Development (http://www.bic.mni.mcgill.ca/nihpdp/info/data_access.html) (Evans 2006) containing multisite neuroimaging data of children from ages 6 to 18 years old accompanied by comprehensive neurological assessments and neuropsychological testing results (Almli, Rivkin *et al.* 2007). Such resources represent leading efforts among a growing set of online repositories for the storage and sharing of raw and processed neuroimaging data and results.

Speaking from our own experience, the Laboratory of Neuro Imaging (LONI), formerly based at the University of California Los Angeles (UCLA) and now based at the University of Southern California (USC), has been serving as a repository for single and multi-site neuroimaging research studies for many years (Toga 2002b). The LONI Image and Data Archive (IDA) provides a comprehensive and interactive environment for safely archiving, querying, visualizing and sharing neuroimaging clinical and neurocognitive data. It presents a user-friendly environment to users for archiving, searching, sharing, tracking and disseminating neuro-imaging data. All data is stored on redundant servers with daily and weekly on- and off-site backups. The IDA stores data from ADNI, the Michael J. Fox Foundation, among other prominent multi-site neuroimaging initiatives.

Archiving data in the IDA is straight-forward, secure, and requires no specialized hardware or software. The IDA automatically facilitates the de-identification and pooling of data from multiple institutions, protecting data from unauthorized access while providing the ability to share data among collaborative investigators. Integration of the LONI Debabeler file format translation engine (Neu, Valentino *et al.* 2005) allows users to upload and download image data in a number of common neuroimaging file formats (e.g. DICOM, Analyze, NIFTI, MINC). Once archived, data can be downloaded and/or streamed into automated tools for processing and analysis (Dinov, Van Horn *et al.* 2009; Dinov, Lozev *et al.* 2010).

Presently, the IDA repository contains over 14,000 imaging series for thousands of subjects, with growth averaging ~400 new image series per month, encompassing 130TB of storage. This includes structural, functional, diffusion, and other MRI-based data types, in addition to studies of a range of PET ligands. Since the inception of the IDA, data drawn from it has been utilized in literally hundreds of research articles published in the peer-reviewed literature.

All in all, resources such as these are amassing brain imaging data at already impressive scales. However, will such archives be ready for still more data as investigators acquire more information through scanning as well as other measures of brain form and function?

The next wave for neuroimaging genetic examinations of the brain will certainly test the computational and storage infrastructures of these and other databasing efforts.

The Role of Cyberinfrastructure

Individual desktop computers are now no longer suitable for analyzing potentially petabytes worth of brain and genomics data at a time. What is needed is to combine effort from across multiple, distributed processing elements and leverage the combined power they possess toward massive-scale analyses of neuroscientific resources. The motivating interest in emerging forms of computing for biomedicine is the coordination of resource sharing and problem solving in dynamic, multi-institutional, spatially dispersed, virtual organizations who can gather and exchange data. While the National Science Foundation (NSF) has made major investments in the computer architecture needed for physics, weather, and geological data (e.g. XSEDE, <https://www.xsede.org/>, and the Open Science Grid, <https://www.opensciencegrid.org/>), the NIH has invested warily (e.g. BIRN, <http://www.birncommunity.org/>), despite strongly encouraging data re-use from online data repositories. More than just the sharing of files, this would involve direct access to computer resources, software, data, and services, increasingly required by a range of collaborative problem-solving and resource-brokering strategies emerging from industry, science, and engineering. The Neuroimaging Informatics Tools and Resources Clearinghouse (NITRC; <http://www.nitrc.org>) and the International Neuroinformatics Coordinating Facility (INCF; <http://incf.org>) have begun to deploy local clusters with Amazon EC2 server technology toward this goal but a larger effort will be required involving dedicated processing centers or distributed grids of linked compute centers. Resource availability can be open or highly controlled with providers and consumers defining clearly and carefully just what is shared, who is allowed to use resources, and the conditions under which it occurs (Van Horn and Toga 2009a). It shouldn't come as a surprise that, like so many other forms of data-rich science, big computing is an ideal framework for big neuroimaging.

Models for Data Management and Availability

The organizing, annotating, archiving, and distributing of neuroimaging and biomedical data in useable and structured frameworks have become critical elements across a range of neuroscientific efforts (Toga 2002a). With all this brain data flying around needing someplace to land, several well known efforts to construct and populate large brain anatomy (Mazziotta, Toga *et al.* 2001), function (Van Horn, Wolfe *et al.* 2005), and genetics (Saykin, Shen *et al.* 2010) databases have arisen over the years in order to make such data open to a still broader audience of researchers. In fact, for many recent large-scale neuroimaging projects, such as the Human Connectome Project (HCP; www.humanconnectomeproject.org), their existence is predicated on the expectation that the data obtained will be well organized and available to the community for examination and study.

As one might imagine, there are as many data organizational models as there are laboratories gathering the data. Because of the volume of all the study data itself, it is not uncommon for the imaging data to live apart from the meta-data describing it, e.g. on its own disk partition

or even another computer system altogether. This information, in turn, can itself live apart from other data acquired (e.g. demographics, genetics, phenomics). Organizing and linking all of these distinct and possibly physically separated data types can be a major data management challenge.

What is more, multiple models exist for making data available to others (Van Horn, Grethe *et al.* 2001). In the simplest form, anonymous FTP sites can be used as a place to post datasets where they may or may not be secure and the intentions of those accessing the data uncertain. Databases containing summary results are enormously valuable from a meta-analytic perspective (Laird, Eickhoff *et al.* 2009), however, the original neuroimaging data themselves are likely unavailable. The use of results-only resources may also require care in due to the potential for factors such as publication bias (Jennings and Van Horn 2012). Federated archiving approaches seek to allow data to remain at the sites that collected it but to permit users at a consortia of participating institutions to search and access each other's data (Helmer, Ambite *et al.* 2011). Centralized approaches also exist in which a single center is the focal point for databasing and accompanying informatics for multi-site efforts or multiple-multi-site efforts (Neu, Crawford *et al.* 2012). These and other models have their own unique challenges for data storage, management, and availability which can be expected to become more complex as study data sizes grow. Contributions of raw and processed study data to centralized or federated repositories can provide maximal information and utility to subsequent users (Van Horn and Gazzaniga 2005), where the submission of data can be voluntary, as a condition of publication, or as an agreed upon aspect of multi-site data consortia (Mazziotta, Toga *et al.* 1995; Toga and Crawford 2010; Van Horn and Gazzaniga 2012).

Standards are Frequently Non-Standard

With a diversity of storage models also comes a variety of ways in which study data is managed and meta-data is internally represented. Having common formatting standards for meta-data representation and organization has been an important topic for neuroinformatics over the past decade (Koslow 2000; Helmer, Ambite *et al.* 2011). The creation of data standards enables both intra- and interdisciplinary interaction (Van Horn and Ball 2008), encourages the development of novel software tools for helping to understand relationships within and among data elements, and encourages new investigation of database contents (Neu, Crawford *et al.* 2012). Recent efforts from international data sharing working groups (Poline, Breeze *et al.* 2012) have made considerable headway in helping to develop new frameworks for organizing study meta-data drawing from the best parts of extant meta-data frameworks. Best practices for fMRI results reporting have been recommended (Poldrack, Fletcher *et al.* 2008) though may naturally need to mature further as emerging analytic approaches come into favor. On the other hand, standards designers need to avoid the feature creep that often arises which can result in frameworks which are overly rigid and cannot adapt as new information becomes available and need to be included. Nevertheless, when fully mature, these modern schema will significantly improve the description of neuroimaging data sets, encode the provenance associated with data processing (Mackenzie-Graham, Van Horn *et al.* 2008; McClatchey, Branson *et al.* 2013), and help to populate large-scale archives prospectively, thereby encouraging common analysis frameworks. Such

practical standards will be essential for multi-site trials and major neuroimaging initiatives, where data sharing has been expressly mandated by funding agencies.

Factors Governing the Utility of “Big Data” Resources

However, simply having “big data” neuroimaging and making it available online is not a means to an end but only the next step in making that data available for others to explore and mine. Several factors contribute to a database’s use and utility, including whether it actually contains viable data accompanied by a detailed description of their acquisition (Van Horn and Toga 2009a); whether the database is well-organized and the user interface is easy to navigate; whether the data are derived versions of raw data or the raw data itself; the manner in which the database addresses the sociological and regulatory issues that can be associated with data sharing; whether the data fully anonymized and open to anyone who wants it, only open to members of a multi-site consortium, or if it is necessary to get IRB approval before access can be granted; whether it has a policy in place to ensure that requesting authors give proper attribution to the original collectors of the data; and the efficiency of secure data transactions. The systems must provide flexible methods for data description and relationships among various meta-data characteristics (Bug, Ascoli *et al.* 2008). Moreover, those that have been specifically designed to serve a large and diverse audience with a variety of needs, represent the types of databases that can have the greatest benefit to scientists looking to study the disease, assess new methods, examine previously published data, or with interests in exploring new ideas (Van Horn and Gazzaniga 2005; Keator, Helmer *et al.* 2013).

Mining Data and Digging for Gold

The successes of molecular biological (Huang, Hu *et al.* 2007), systems biological (Hood, Heath *et al.* 2004), and astrophysics (Gray, Szalay *et al.* 2002) infrastructures for archiving and mining data are well known. With the accumulation of neurobiological data into large databases and the availability of compute cluster enabled means for large-scale data processing, a new form of discovery-oriented neuroscience is on the horizon. Researchers are increasingly mining vast and disparate collections of data and hunting for unseen patterns that might provide clues to underlying biological mechanisms (Phan, Wager *et al.* 2002; Arnone, Cavanagh *et al.* 2009; Frazier and Hardan 2009) and trends in how studies are conducted (Jennings and Van Horn 2012). The terabytes-worth of data provide input for informatics-driven pattern-seeking and other relevant algorithms (Jones and Swindells 2002; Ma, Tromp *et al.* 2002; Schutte, Mitros *et al.* 2002) which can provide further understanding into complex brain processes.

Departing from a purely repository-based approach, a collection of research groups from around the world has adopted what might be considered more of a social networking strategy for data sharing. The Enhancing Neuroimaging Genetics through Meta-Analysis (ENIGMA; enigma.ionu.edu) network brings together interested researchers in imaging genomics, to understand brain structure and function, based on MRI, DTI, fMRI and genome-wide association study (GWAS) data. Among the network’s goals: to ensure promising findings are replicated via member collaborations, in order to satisfy the mandates

of most journals; gain increased statistical power; and to share ideas, algorithms, data, and information on promising findings or methods. They have provided the means for imaging, genetics, as well as computer science, mathematics and statistics experts to make significant collaborative contributions to these fields from which, without the needed infrastructure, most of their expertise would have been excluded (Stein, Medland *et al.* 2012). These successes are not necessarily unique, but building upon them and extending them to a wider set of scientific research arenas is an ever-present theme (Persson 2000; Brookes 2001; Altman 2003).

New Computer Science Designed with Big Data in Mind

While it might be tempting to think that once all the data has been archived, indexed, and is ready to go, that all one would need is to start analyzing it and answers to all our questions about the brain will be revealed. Even examining the contents of an archive to know what data is available to be analyzed requires new, cleverly designed and user friendly software tools and novel approaches for exploratory inspection. Such tools are only now beginning to appear (Bowman, Joshi *et al.* 2012) and their further development will be essential for dealing with existing as well as the expected size of neuroimaging data sets.

Once a selection of data worthy of further analysis has been identified, a new concern is realized – it becomes clear that many software packages for neuroimaging data analysis are ill-suited toward very large data sets involving potentially thousands of subjects. Algorithm optimization is not often considered for when data sets are small or modest in size but as data sets grow memory management is an important factor. New mathematics and informatics approaches will be needed to more completely model multi-modal brain imaging data in the context of cortical anatomy, white matter connectivity, and functional activity. These will need to work fast, be accurate, and be interoperable with other tools so that data processing can be automated as much as possible. Interactive workflow environments for automated data analysis will also be critical for ongoing or retrospective research studies involving complex computations on large multi-dimensional datasets (Dinov, Van Horn *et al.* 2010; Gorgolewski, Burns *et al.* 2011). Yet, few tools, if any, now exist which enable the joint analysis of both genes and brain imaging data which would be capable of efficiently obtaining results while also achieving the requisite degree of statistical power. Moving forward, software engineers will need to create brilliant and innovative ways to tackle the massive amounts of brain and genomic data. These continuous interactions between neuroimagers, geneticists, software creators, and other biomedical scientists will be essential to develop these new, memory-efficient software algorithms and computational tools.

Big Data in the Era of ADNI2, the HCP, the HBP, and BRAIN

As the NIH and other major funding agencies spearhead big picture multi-center brain science efforts, the needs for big data solutions will grow. The HCP consortia are now releasing large datasets which even the best neuroimaging researchers are struggling to analyze. More of that data is on the way. ADNI2, the premier neuroimaging data collection effort for understanding the aging brain, will be providing imaging, phenomic, and exomic information which we have yet to get our heads around how best to analyze. The HBP,

involving data collection sites throughout Europe, will likely eclipse that! And while the basis for infrastructure support is also changing, the simple thought that data and processing will simply be done “in the cloud” is somewhat naïve. The “cloud” must exist somewhere and system failures have been known to happen. Though commercial solutions such as those offered by Google, Amazon, and Microsoft can provide an attractive solution for individual centers, dedicated neuroscience computational databasing and data processing resources remain essential. As the BRAIN initiative takes shape, it is clear that demonstrable access to data, processing workflow methods, running on remote CPU clusters dedicated to this purpose is the best way to ensure that we can keep up in the extraction of real findings as opposed to a collection of results.

Conclusions

Neuroimaging research, by its very nature, is data intensive, multimodal, and collaborative – factors which have been instrumental in its success and growth. Indeed, we contend that neuroimaging is an emerging example of discovery-oriented science, wherein patterns of brain structure and activity present across multiple subjects and dozens of studies can be systematically extracted, examined, and resulting in new knowledge. Yet the infrastructure needed for supporting this advancing form of brain research where data is king is still maturing. The rapid processing of large quantities of data in this way will lead to new scientific outcomes and patterns of results not envisioned during the examination of each study individually. Patterns may suggest fundamental mechanisms. Confirmed mechanisms add to the knowledgebase of neurobiological science and provide the basis for further experimentation and the generation of still more valuable data that can be included in still greater analyses. Greater knowledge about fundamental brain processes then suggests new and testable hypotheses that lead to novel experimentation, the data from which should then be contributed back into a publicly available archive – continuing a healthy and helpful cycle.

The next steps for the development of resources supporting “big data” brain imaging at the Exabyte scale will require the further creation of new tools and services for data discovery, integration, analysis, and visualization. Components for discovering data residing in database architectures must be developed (sort of a PubMed for data discovery). Present examples exist such as the “EB-eye” resource for genomics (<http://www.ebi.ac.uk/ebisearch/>) or the Neuroscience Information Framework (NIF; <http://neuinfo.org>) for neuroscience terminologies. Such meta-resources will need to include contextual information that allows data to be accessed, understood, reused, and the results reproduced. Integrating a broader spectrum of neuroscience data and providing tools for interrogating and visualizing those data will enable investigators to more easily and interactively investigate broader scientific questions.

Beyond just neuroimaging data, architectures for peta-scale biomedical data must be flexible enough to allow integration of additional clinical and biochemical data and analysis results into the database and employ tools for interactively interrogating and graphically visualizing database contents (Bowman, Joshi *et al.* 2012). Frameworks for storing and making available big neuroimaging data, their standards, and the infrastructure for so doing must be

enriched with modern data processing workflow design and execution systems in place to permit exchange of processing methodologies between labs, among consortia members, or by independent researchers. Comprehensive mechanisms to gather, organize, and distribute data, results, and information between and among project participants, but also to the scientific community at large, are worth being examined, developed, and deployed. For the “big data” science of human brain imaging, now is the time to begin.

Acknowledgments

This article was supported by a P41 (RR013642) award from the National Institute for General Medicine (NIGMS), through the National Institutes of Health Roadmap for Medical Research (U54 RR021813, The Center for Computational Biology - CCB), and through an National Institute of Mental Health ARRA award (RC1 MH088194). We express our gratitude the faculty and staff of the Laboratory of Neuro Imaging (LONI) at the University of California Los Angeles.

References

- Almli CR, Rivkin MJ, McKinsty RC. The NIH MRI study of normal brain development (Objective-2): newborns, infants, toddlers, and preschoolers. *Neuroimage*. 2007; 35(1):308–325. [PubMed: 17239623]
- Altman RB. The expanding scope of bioinformatics: sequence analysis and beyond. *Heredity*. 2003; 90(5):345.
- Arnold D, Cavanagh J, Gerber D, Lawrie SM, Ebmeier KP, McIntosh AM. Magnetic resonance imaging studies in bipolar disorder and schizophrenia: meta-analysis. *Br J Psychiatry*. 2009; 195(3): 194–201. [PubMed: 19721106]
- Barry RL, Strother SC, Gatenby JC, Gore JC. Data-driven optimization and evaluation of 2D EPI and 3D PRESTO for BOLD fMRI at 7 Tesla: I. Focal coverage. *Neuroimage*. 2011; 55(3):1034–1043. PMID: 3049844. [PubMed: 21232613]
- Basser PJ, Pajevic S, Pierpaoli C, Duda J, Aldroubi A. In vivo fiber tractography using DT-MRI data. *Magn Reson Med*. 2000; 44(4):625–632. [PubMed: 11025519]
- Beaulieu A. Voxels in the brain: neuroscience, informatics and changing notions of objectivity. *Soc Stud Sci*. 2001; 31(5):635–680. [PubMed: 11811207]
- Biswal BB, Mennes M, Zuo XN, Gohel S, Kelly C, Smith SM, Beckmann CF, Adelstein JS, Buckner RL, Colcombe S, Dogonowski AM, Ernst M, Fair D, Hampson M, Hoptman MJ, Hyde JS, Kiviniemi VJ, Kotter R, Li SJ, Lin CP, Lowe MJ, Mackay C, Madden DJ, Madsen KH, Margulies DS, Mayberg HS, McMahon K, Monk CS, Mostofsky SH, Nagel BJ, Pekar JJ, Peltier SJ, Petersen SE, Riedl V, Rombouts SA, Rypma B, Schlaggar BL, Schmidt S, Seidler RD, G JS, Sorg C, Teng GJ, Veijola J, Villringer A, Walter M, Wang L, Weng XC, Whitfield-Gabrieli S, Williamson P, Windischberger C, Zang YF, Zhang HY, Castellanos FX, Milham MP. Toward discovery science of human brain function. *Proc Natl Acad Sci U S A*. 2010
- Bowman I, Joshi SH, Van Horn J. Visual Systems for Interactive Exploration and Mining of Large-Scale Neuroimaging Data Archives. *Frontiers in Neuroinformatics*. 2012; 6
- Brookes AJ. Rethinking genetic strategies to study complex diseases. *Trends Mol Med*. 2001; 7(11): 512–516. [PubMed: 11689337]
- Bug WJ, Ascoli GA, Grethe JS, Gupta A, Fennema-Notestine C, Laird AR, Larson SD, Rubin D, Shepherd GM, Turner JA, Martone ME. The NIFSTD and BIRN Lex vocabularies: building comprehensive ontologies for neuroscience. *Neuroinformatics*. 2008; 6(3):175–194. [PubMed: 18975148]
- Bushnik T, Gordon W. Updates from the Third Federal Interagency Conference on traumatic brain injury. *J Head Trauma Rehabil*. 2012; 27(3):222–223. [PubMed: 22573040]
- Dinov I, Lozev K, Petrosyan P, Liu Z, Eggert P, Pierce J, Zamanyan A, Chakrapani S, Van Horn J, Parker DS, Magsipoc R, Leung K, Gutman B, Woods R, Toga A. Neuroimaging study designs, computational analyses and data provenance using the LONI pipeline. *PLoS ONE*. 2010; 5(9) PMID: 2946935.

- Dinov I, Van Horn J, Lozev K, Magsipoc R, Petrosyan P, Liu Z, MacKenzie-Graham A, Eggert P, Parker DS, Toga AW. Efficient, distributed and interactive neuroimaging data analysis using the LONI pipeline. *Frontiers in Neuroinformatics*. 2010; 3(22):1–10.
- Dinov I, Van Horn J, Lozev K, Magsipoc R, Petrosyan P, Liu Z, MacKenzie-Graham A, Eggert P, Parker D, Toga A. Efficient, Distributed and Interactive Neuroimaging Data Analysis using the LONI Pipeline. *Frontiers of Neuroinformatics*. 2009; 3(22):1–10.
- Evangelou E, Maraganore DM, Annesi G, Brighina L, Brice A, Elbaz A, Ferrarese C, Hadjigeorgiou GM, Krueger R, Lambert JC, Lesage S, Markopoulou K, Mellick GD, Meeus B, Pedersen NL, Quattrone A, Van Broeckhoven C, Sharma M, Silburn PA, Tan EK, Wirdefeldt K, Ioannidis JP. Non-replication of association for six polymorphisms from meta-analysis of genome-wide association studies of Parkinson's disease: Large-scale collaborative study. *Am J Med Genet B Neuropsychiatr Genet*. 2009
- Evans AC. The NIH MRI study of normal brain development. *Neuroimage*. 2006; 30(1):184–202. [PubMed: 16376577]
- Feinberg DA, Moeller S, Smith SM, Auerbach E, Ramanna S, Glasser MF, Miller KL, Ugurbil K, Yacoub E. Multiplexed Echo Planar Imaging for Sub-Second Whole Brain fMRI and Fast Diffusion Imaging. *PLoS ONE*. 2010; 5(12):e15710. [PubMed: 21187930]
- Frazier TW, Hardan AY. A meta-analysis of the corpus callosum in autism. *Biol Psychiatry*. 2009; 66(10):935–941. [PubMed: 19748080]
- Frisoni GB. Alzheimer's disease neuroimaging initiative in Europe. *Alzheimers Dement*. 2010; 6(3): 280–285. [PubMed: 20451877]
- Gorgolewski K, Burns CD, Madison C, Clark D, Halchenko YO, Waskom ML, Ghosh SS. Nipype: a flexible, lightweight and extensible neuroimaging data processing framework in python. *Front Neuroinform*. 2011; 5:13. PMID: 3159964. [PubMed: 21897815]
- Gray J, Szalay S, Thaker AR, Kunszt PZ, Malik T, Raddick J, Stoughton C, vanderBerg J. The SDSS SkyServer - Public access to the Sloan Digital Sky Survey Data. 2002 ACM SIGMOD.
- Hall D, Huerta MF, McAuliffe MJ, Farber GK. Sharing heterogeneous data: the national database for autism research. *Neuroinformatics*. 2012; 10(4):331–339. [PubMed: 22622767]
- Helmer KG, Ambite JL, Ames J, Ananthakrishnan R, Burns G, Chervenak AL, Foster I, Liming L, Keator D, Maciardi F, Madduri R, Navarro JP, Potkin S, Rosen B, Ruffins S, Schuler R, Turner JA, Toga A, Williams C, Kesselman C. Enabling collaborative research using the Biomedical Informatics Research Network (BIRN). *J Am Med Inform Assoc*. 2011
- Hood L, Heath JR, Phelps ME, Lin B. Systems biology and new technologies enable predictive and preventative medicine. *Science*. 2004; 306(5696):640–643. [PubMed: 15499008]
- Hsi-Yang Fritz M, Leinonen R, Cochrane G, Birney E. Efficient storage of high throughput DNA sequencing data using reference-based compression. *Genome Res*. 2011; 21(5):734–740. PMID: 3083090. [PubMed: 21245279]
- Huang H, Hu ZZ, Arighi CN, Wu CH. Integration of bioinformatics resources for functional analysis of gene expression and proteomic data. *Front Biosci*. 2007; 12:5071–5088. [PubMed: 17569631]
- Jennings RG, Van Horn JD. Publication bias in neuroimaging research: implications for meta-analyses. *Neuroinformatics*. 2012; 10(1):67–80. [PubMed: 21643733]
- Jiang T, Liu Y, Shi F, Shu N, Liu B, Jiang J, Zhou Y. Multimodal Magnetic Resonance Imaging for Brain Disorders: Advances and Perspectives. *Brain Imaging and Behavior*. 2008; 2(4):249–257.
- Jones DT, Swindells MB. Getting the most from PSI-BLAST. *Trends Biochem Sci*. 2002; 27(3):161–164. [PubMed: 11893514]
- Keator DB, Helmer K, Steffener J, Turner JA, Van Erp TG, Gadde S, Ashish N, Burns GA, Nichols BN. Towards structured sharing of raw and derived neuroimaging data across existing resources. *Neuroimage*. 2013
- Koslow SH. Should the neuroscience community make a paradigm shift to sharing primary data? *Nature Neuroscience*. 2000; 3(4):863–865.
- Laird AR, Eickhoff SB, Kurth F, Fox PM, Uecker AM, Turner JA, Robinson JL, Lancaster JL, Fox PT. ALE Meta-Analysis Workflows Via the Brainmap Database: Progress Towards A Probabilistic Functional Brain Atlas. *Front Neuroinform*. 2009; 3:23. PMID: 2715269. [PubMed: 19636392]

- Ma B, Tromp J, Li M. PatternHunter: faster and more sensitive homology search. *Bioinformatics*. 2002; 18(3):440–445. [PubMed: 11934743]
- Mackenzie-Graham AJ, Van Horn JD, Woods RP, Crawford KL, Toga AW. Provenance in neuroimaging. *Neuroimage*. 2008; 42(1):178–195. PMCID: 2664747. [PubMed: 18519166]
- Marcus D, et al. The extensible neuroimaging archive toolkit (XNAT): an informatics platform for managing, exploring, and sharing neuroimaging data. *Neuroinformatics*. 2007; 5:11–34. [PubMed: 17426351]
- Marcus DS, Harms MP, Snyder AZ, Jenkinson M, Wilson JA, Glasser MF, Barch DM, Archie KA, Burgess GC, Ramaratnam M, Hodge M, Horton W, Herrick R, Olsen T, McKay M, House M, Hileman M, Reid E, Harwell J, Coalson T, Schindler J, Elam JS, Curtiss SW, Van Essen DC. Human Connectome Project informatics: Quality control, database services, and data visualization. *Neuroimage*. 2013
- Mazziotta J, Toga AW, Evans A, Fox P, Lancaster J, Ziles K, Woods RP, Paus T, Simpson G, Pike B, Holmes C, Collins L, Thompson PM, MacDonald D, Iacoboni M, Schormann T, Amunts K, Palomero-Gallagher N, Geyer S, Parsons LM, Narr K, Kabani N, LeGoualher G, Boomsma D, Cannon T, Kawashima R, Mazoyer B. A probabilistic atlas and reference system for the human brain: International Consortium for Brain Mapping (ICBM). *Philos Trans R Soc Lond B Biol Sci*. 2001; 356:1293–1322. [PubMed: 11545704]
- Mazziotta JC, Toga AW, Evans A, Fox P, Lancaster J. A probabilistic atlas of the human brain: theory and rationale for its development. The International Consortium for Brain Mapping (ICBM). *Neuroimage*. 1995; 2(2):89–101. [PubMed: 9343592]
- McClatchey R, Branson A, Anjum A, Bloodsworth P, Habib I, Munir K, Shamdasani J, Soomro K. Providing traceability for neuroimaging analyses. *Int J Med Inform*. 2013
- Neu SC, Crawford KL, Toga AW. Practical management of heterogeneous neuroimaging metadata by global neuroimaging data repositories. *Front Neuroinform*. 2012; 6:8. PMCID: 3311229. [PubMed: 22470336]
- Neu SC, Valentino DJ, Toga AW. The LONI Debabeler: a mediator for neuroimaging software. *Neuroimage*. 2005; 24(4):1170–1179. [PubMed: 15670695]
- Persson B. *Bioinformatics in protein analysis*. Exs. 2000; 88:215–231. [PubMed: 10803381]
- Phan KL, Wager T, Taylor SF, Liberzon I. Functional neuroanatomy of emotion: a meta-analysis of emotion activation studies in PET and fMRI. *Neuroimage*. 2002; 16(2):331–348. [PubMed: 12030820]
- Poldrack RA, Barch DM, Mitchell JP, Wager TD, Wagner AD, Devlin JT, Cumba C, Koyejo O, Milham MP. Toward open sharing of task-based fMRI data: the OpenfMRI project. *Front Neuroinform*. 2013; 7:12. PMCID: 3703526. [PubMed: 23847528]
- Poldrack RA, Fletcher PC, Henson RN, Worsley KJ, Brett M, Nichols TE. Guidelines for reporting an fMRI study. *Neuroimage*. 2008; 40(2):409–414. PMCID: 2287206. [PubMed: 18191585]
- Poline JB, Breeze JL, Ghosh S, Gorgolewski K, Halchenko YO, Hanke M, Haselgrove C, Helmer KG, Keator DB, Marcus DS, Poldrack RA, Schwartz Y, Ashburner J, Kennedy DN. Data sharing in neuroimaging research. *Front Neuroinform*. 2012; 6:9. PMCID: 3319918. [PubMed: 22493576]
- Saykin AJ, Shen L, Foroud TM, Potkin SG, Swaminathan S, Kim S, Risacher SL, Nho K, Huentelman MJ, Craig DW, Thompson PM, Stein JL, Moore JH, Farrer LA, Green RC, Bertram L, Jack CR, Weiner MW. ADNI biomarkers as quantitative phenotypes: Genetics core aims, progress, and plans, Alzheimer's and Dementia. *Journal of the Alzheimer's Association*. 2010; 6(3):265–273.
- Schumann G, Loth E, Banaschewski T, Barbot A, Barker G, Buchel C, Conrod PJ, Dalley JW, Flor H, Gallinat J, Garavan H, Heinz A, Irtnerman B, Lathrop M, Mallik C, Mann K, Martinot JL, Paus T, Poline JB, Robbins TW, Rietschel M, Reed L, Smolka M, Spanagel R, Speiser C, Stephens DN, Strohle A, Struve M. The IMAGEN study: reinforcement-related behaviour in normal brain function and psychopathology. *Mol Psychiatry*. 2010; 15(12):1128–1139. [PubMed: 21102431]
- Schutte BC, Mitros JP, Bartlett JA, Walters JD, Jia HP, Welsh MJ, Casavant TL, McCray PB Jr. Discovery of five conserved beta -defensin gene clusters using a computational search strategy. *Proc Natl Acad Sci U S A*. 2002; 99(4):2129–2133. [PubMed: 11854508]
- Stein JL, Medland SE, Vasquez AA, Hibar DP, Senstad RE, Winkler AM, Toro R, Appel K, Bartecek R, Bergmann O, Bernard M, Brown AA, Cannon DM, Chakravarty MM, Christoforou A, Domin

M, Grimm O, Hollinshead M, Holmes AJ, Homuth G, Hottenga JJ, Langan C, Lopez LM, Hansell NK, Hwang KS, Kim S, Laje G, Lee PH, Liu X, Loth E, Lourdasamy A, Mattingdal M, Mohnke S, Maniega SM, Nho K, Nugent AC, O'Brien C, Papmeyer M, Putz B, Ramasamy A, Rasmussen J, Rijpkema M, Risacher SL, Roddey JC, Rose EJ, Ryten M, Shen L, Sprooten E, Strengman E, Teumer A, Trabzuni D, Turner J, van Eijk K, van Erp TG, van Tol MJ, Wittfeld K, Wolf C, Woudstra S, Aleman A, Alhusaini S, Almasy L, Binder EB, Brohawn DG, Cantor RM, Carless MA, Corvin A, Czisch M, Curran JE, Davies G, de Almeida MA, Delanty N, Depondt C, Duggirala R, Dyer TD, Erk S, Fagerness J, Fox PT, Freimer NB, Gill M, Goring HH, Hagler DJ, Hoehn D, Holsboer F, Hoogman M, Hosten N, Jahanshad N, Johnson MP, Kasperaviciute D, Kent JW Jr, Kochunov P, Lancaster JL, Lawrie SM, Liewald DC, Mandl R, Matarin M, Mattheisen M, Meisenzahl E, Melle I, Moses EK, Muhleisen TW, Nauck M, Nothen MM, Olvera RL, Pandolfo M, Pike GB, Puls R, Reinvang I, Renteria ME, Rietschel M, Roffman JL, Royle NA, Rujescu D, Savitz J, Schnack HG, Schnell K, Seiferth N, Smith C, Steen VM, Valdes Hernandez MC, Van den Heuvel M, van der Wee NJ, Van Haren NE, Veltman JA, Volzke H, Walker R, Westlye LT, Whelan CD, Agartz I, Boomsma DI, Cavalleri GL, Dale AM, Djurovic S, Drevets WC, Hagoort P, Hall J, Heinz A, Jack CR Jr, Foroud TM, Le Hellard S, Macciardi F, Montgomery GW, Poline JB, Porteous DJ, Sisodiya SM, Starr JM, Sussmann J, Toga AW, Veltman DJ, Walter H, Weiner MW, Bis JC, Ikram MA, Smith AV, Gudnason V, Tzourio C, Vernooij MW, Launer LJ, Decarli C, Seshadri S, Andreassen OA, Apostolova LG, Bastin ME, Blangero J, Brunner HG, Buckner RL, Cichon S, Coppola G, de Zubicaray GI, Deary IJ, Donohoe G, de Geus EJ, Espeseth T, Fernandez G, Glahn DC, Grabe HJ, Hardy J, Hulshoff Pol HE, Jenkinson M, Kahn RS, McDonald C, McIntosh AM, McMahon FJ, McMahon KL, Meyer-Lindenberg A, Morris DW, Muller-Myhsok B, Nichols TE, Ophoff RA, Paus T, Pausova Z, Penninx BW, Potkin SG, Samann PG, Saykin AJ, Schumann G, Smoller JW, Wardlaw JM, Weale ME, Martin NG, Franke B, Wright MJ, Thompson PM. Identification of common variants associated with human hippocampal and intracranial volumes. *Nat Genet.* 2012

- Toga AW. Imaging databases and neuroscience. *Neuroscientist.* 2002a; 8(5):423–436. [PubMed: 12374427]
- Toga AW. The Laboratory of Neuro Imaging: what it is, why it is, and how it came to be. *IEEE Trans Med Imaging.* 2002b; 21(11):1333–1343. [PubMed: 12575870]
- Toga AW, Clark KA, Thompson PM, Shattuck DW, Van Horn JD. Mapping the human connectome. *Neurosurgery.* 2012; 71(1):1–5. [PubMed: 22705717]
- Toga AW, Crawford KL. The informatics core of the Alzheimer's Disease Neuroimaging Initiative. *Alzheimer's and Dementia.* 2010; 6(3):247–256.
- Van Essen DC. A Population-Average, Landmark- and Surface-based (PALS) atlas of human cerebral cortex. *Neuroimage.* 2005; 28(3):635–662. [PubMed: 16172003]
- Van Horn JD, Ball CA. Domain-specific data sharing in neuroscience: what do we have to learn from each other? *Neuroinformatics.* 2008; 6(2):117–121. [PubMed: 18473189]
- Van Horn, JD.; Gazzaniga, MS.; Koslow, SH.; Subramanian, A. Databasing the Brain: From Data to Knowledge. John Wiley and Sons; New York: 2005. Maximizing Information Content in shared Neuroimaging Studies of Cognitive Function.
- Van Horn JD, Gazzaniga MS. Why share data? Lessons learned from the fMRIDC. *Neuroimage.* 2012
- Van Horn JD, Grethe JS, Kostelec P, Woodward JB, Aslam JA, Rus D, Rockmore D, Gazzaniga MS. The Functional Magnetic Resonance Imaging Data Center (fMRIDC): the challenges and rewards of large-scale databasing of neuroimaging studies. *Philos Trans R Soc Lond B Biol Sci.* 2001; 356(1412):1323–1339. PMID: 1088517. [PubMed: 11545705]
- Van Horn JD, Ishai A. Mapping the human brain: new insights from FMRI data sharing. *Neuroinformatics.* 2007; 5(3):146–153. [PubMed: 17917125]
- Van Horn JD, Toga AW. Is it time to re-prioritize neuroimaging databases and digital repositories? *Neuroimage.* 2009a; 47(4):1720–1734. [PubMed: 19371790]
- Van Horn JD, Toga AW. Multisite neuroimaging trials. *Curr Opin Neurol.* 2009b; 22(4):370–378. PMID: 2777976. [PubMed: 19506479]
- Van Horn JD, Wolfe J, Agnoli A, Woodward J, Schmitt M, Dobson J, Schumacher S, Vance B. Neuroimaging databases as a resource for scientific discovery. *Int Rev Neurobiol.* 2005; 66:55–87. [PubMed: 16387200]

- Wedeen VJ, Wang RP, Schmahmann JD, Benner T, Tseng WY, Dai G, Pandya DN, Hagmann P, D'Arceuil H, de Crespigny AJ. Diffusion spectrum magnetic resonance imaging (DSI) tractography of crossing fibers. *Neuroimage*. 2008; 41(4):1267–1277. [PubMed: 18495497]
- Weiner MW, Veitch DP, Aisen PS, Beckett LA, Cairns NJ, Green RC, Harvey D, Jack CR, Jagust W, Liu E, Morris JC, Petersen RC, Saykin AJ, Schmidt ME, Shaw L, Siuciak JA, Soares H, Toga AW, Trojanowski JQ. The Alzheimer's Disease Neuroimaging Initiative: a review of papers published since its inception. *Alzheimers Dement*. 2012; 8(1 Suppl):S1–68. PMCID: 3329969. [PubMed: 22047634]
- Zhan L, Leow AD, Jahanshad N, Chiang MC, Barysheva M, Lee AD, Toga AW, McMahon KL, de Zubicaray GI, Wright MJ, Thompson PM. How does angular resolution affect diffusion imaging measures? *Neuroimage*. 2010; 49(2):1357–1371. [PubMed: 19819339]

Figure 1.

The amount of acquired neuroimaging data reported from published articles in representative issues of the journal *NeuroImage* has doubled every 26 months and can expect to top 20GB of purely raw data on average per study in only a few years. Amassing, curating, storing, and sharing of such data from neuroimaging archives presents a growing big data challenge.