

Mapping the Human Brain: New Insights from fMRI Data Sharing

John Darrell Van Horn · Alumit Ishai

Published online: 18 August 2007
© Humana Press Inc. 2007

Abstract The sharing of primary data in the field of neuroscience has received considerable scrutiny from scientific societies and from science journals. Many see this as value added for science publishing that can enhance and inform secondary examination of data and results. Still others worry that data sharing is an undue burden for researchers with little long term value to science. But examples of how data sharing can be done successfully do exist. The fMRI Data Center, established at Dartmouth College in 2000 and now based at the University of California Santa Barbara, has worked to facilitate the open sharing of neuroimaging data from peer-reviewed papers to foster progress in cognitive science. The fMRI study on the representation of objects in the human occipital and temporal cortex, published in 2000 in the *Journal of Cognitive Neuroscience* (JOCN), marked the first deposition in the new database. Despite initial concerns about fMRI data sharing, this data set was frequently downloaded. We describe the original results of distributed brain activation patterns elicited by faces and objects in the human visual system, and overview several secondary analyses by independent investigators. A philosopher tested Husserl's temporal components of consciousness, whereas

other brain imagers deployed new analytic tools, from Dynamic Causal Modeling, which estimates the neural interactions between cortical regions, to a novel method for constructing reproducibility maps. These re-analyses revealed new findings not reported in the original study, provided new perspectives on visual perception, generated new predictions, and resulted in new collaborations and publications in high profile journals.

Keywords fMRI · Data sharing · Data base

Introduction

The sharing of neuroscientific data has taken on an increasingly important role in the activities of peer-reviewed journals (Shepherd 2002) and the missions of scientific societies. In an era where more and more experimental data is digitally acquired or represented, the desire to have that information made available so that others may examine and scrutinize it has grown in kind. This is especially true in the context of data that goes into published research articles. The sharing of that primary data in support of the reported results and conclusions in published research articles is now achievable. Online archives and resources that play the role of data brokers accompanied by registries of the emerging and mature tools to analyze that data (Kennedy and Haselgrove 2006) are fast becoming critical elements in everyday neuroscience research.

One such neuroscience database, the fMRI Data Center (fMRIDC; <http://www.fmridc.org>), has been instrumental in promoting data sharing, re-use, and re-analysis. The data from the first study contributed to the archive represent a fine example of how new research and results can be

J. D. Van Horn
Laboratory of Neuro Imaging (LONI),
Department of Neurology, David Geffen School of Medicine,
University of California Los Angeles,
635 Charles E. Young Drive SW, Suite 225,
Los Angeles, CA 90095-7334, USA
e-mail: jvanhorn@loni.ucla.edu

A. Ishai (✉)
Institute of Neuroradiology, University of Zurich,
Winterthurerstrasse 190,
8057 Zurich, Switzerland
e-mail: ishai@hifo.uzh.ch

obtained from contributed data and how these can lead to new insights into human brain function. In this article, we discuss this initial contribution to the fMRIDC archive, how the data from it have been re-used, and the principles of data sharing and dissemination in the large scale archiving of data from published fMRI studies. We provide thoughts on how such efforts can be evaluated and what challenges exist in their implementation. Finally, we comment on the role that leading scientific organizations and science journals have in the open sharing of primary research data.

The fMRI Data Center

In 2000, with the intent of capturing a segment of this literature and the data that underlies it, the fMRIDC was founded at Dartmouth College (Van Horn et al. 2001).¹ This effort sought to facilitate progress in understanding cognitive processes through the collection, archiving and open distribution of neuroimaging data sets in the peer-reviewed literature. fMRIDC project directors reasoned that there could be several positive outcomes to making the complete study data sets available to others. First, the study findings could be independently confirmed, helping to strengthen the findings drawn by the original authors. Second, new statistical methodologies could be applied to the data, providing novel insights into cognitive processes. Different studies could be compared, possibly identifying unanticipated functional homologies between seemingly different cognitive tasks. Moreover, these studies could be used to train the next generation of neuroscientists by using fMRI data that had already undergone interpretation by those who collected it and had published it in leading journals. The focus was on fMRI data from published articles, which allowed collecting and managing the data, as well as constructing an archive that was representative of the field's body of work.

Since the time of the first contributed article, the data from over 120 complete articles have been through this submission process. These data, in turn, have been disseminated to laboratories throughout the USA (Fig. 1) and the world.

We mention the first article contributed to the fMRIDC (Ishai et al. 2000a), as this study, in particular, has enjoyed considerable prominence in the community not simply because of its important findings on the cognitive representation of objects. It has also been subjected to novel re-analyses that have broadened the scope of the original research findings and led to new research collaborations. In the next section,

we describe this study and the several ways in which it has been examined.

Visual Perception, Consciousness and Models of Cortical Connectivity

The functional architecture of the ventral visual pathway is a matter of on-going debate. Brain imaging and electrophysiological recording studies in humans have reported discrete cortical regions in posterior ventral temporal cortex that respond preferentially to faces (Kanwisher et al. 1997), buildings and scenes (Aguirre et al. 1998; Epstein and Kanwisher 1998), letters (Polk and Farah 1998), animals and tools (Chao et al. 1999), and human body parts (Downing et al. 2001). These findings suggest a category-specific, anatomically-segregated modular organization of the object vision pathway. While it may be true that there are dedicated neural mechanisms for certain biologically relevant objects, such as faces, which emerge through evolution, it seems highly unlikely that there are modules for all object categories. An alternative possibility is that the representation of objects in ventral temporal cortex is more widely distributed. To test the modular model, an fMRI study was conducted in which activation elicited by faces and two categories of other objects, namely houses and chairs, were measured using two tasks (passive viewing and delayed matching) and two stimulus formats (grayscale photographs and black and white line drawings). Although each category elicited maximal activation within a specific region (e.g., houses in medial fusiform gyrus, faces in lateral fusiform gyrus and chairs in inferior temporal gyrus), each object category was associated with overlapping and distributed patterns of activation that encompassed a wide expanse of ventral temporal cortex (Ishai et al. 1999). Subsequent analysis revealed that faces, houses and chairs evoked similar patterns of differential activation in ventral occipital cortex (Ishai et al. 2000a). These findings were replicated and extended with additional categories of man-made objects (Haxby et al. 2001). As the representation of an object is not restricted to a cortical "module", the object-form topology hypothesis has been proposed, according to which the functional architecture of the ventral visual pathway is a continuous representation of information about object form that has a highly consistent and orderly topological arrangement. The distributed representation model is not only physiologically plausible, but consistent with single unit recordings in monkeys (Tanaka 1996) and computational models of object recognition (Edelman et al. 1998). The fMRI data from the Ishai et al. 1999; and 2000a publications were deposited in the fMRIDC and to date, the re-analysis of these data resulted in several publications, of

¹ As of January 2007, the fMRI Data Center is based at the University of California Santa Barbara.

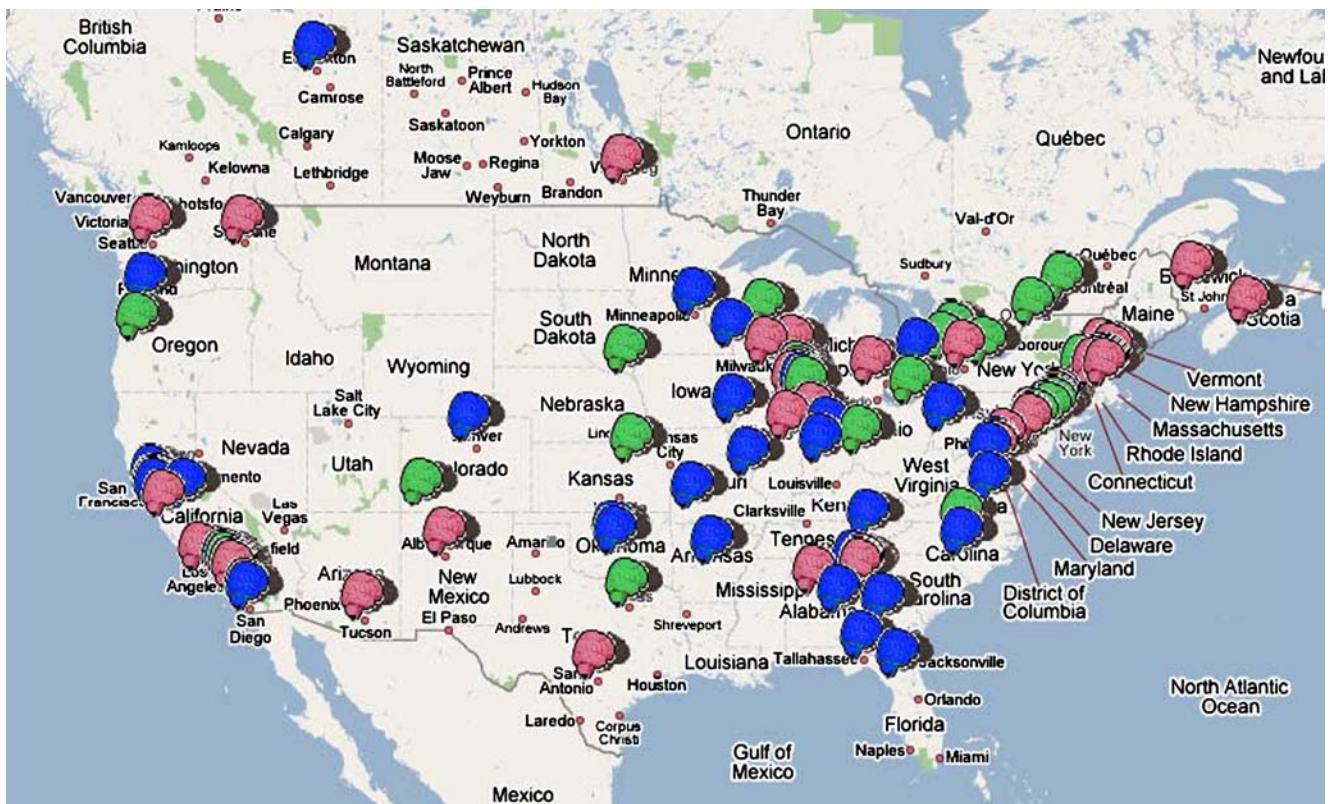


Fig. 1 A Google Maps API (<http://www.google.com/apis/maps/>) plot of locations in the USA and Canada to which one or more complete fMRI study data sets from the fMRIDC archive were delivered. The online maps itself can be viewed at <http://www.fmridc.org/googlemaps/async.html>. Clicking on each *brain icon* will show the fMRIDC study accession number hyperlink for that data set. *Pink icons*

represent study requests from 2000–2004, *blue icons* from 2005, and *green icons* from 2006. This map shows how widely study data sets have been disseminated in North America alone. Requests for one or more data sets and delivered to laboratories in other countries include China, Australia, Germany, UK, Japan, France, Belgium, and Switzerland

which some have received particular recognition for their novel use of shared data (Van Horn 2002).

The first published re-use of the Ishai data were in the context of human consciousness. Daniel Lloyd, a philosopher from Trinity College in Hartford, obtained four data sets from the fMRIDC (Hazeltine et al. 2000; Ishai et al. 2000b; Mechelli et al. 2000; Postle et al. 2000) in order to test specific predictions about human consciousness. Lloyd adopted Husserl's criteria, according to which the phenomenology of consciousness is based on three essential principles: intentionality (the external world as it is experienced and not as it is); superposition (sensory and non-sensory properties are present in perception); and temporality (all objects share perception of present, past, and anticipated future). If indeed these aspects of consciousness are implemented in the brain, the empirical evidence should include temporal flux (with passing time, the multivariate differences between images should increase) and superposition (images sharing task or stimulus conditions should be similar). Using multivariate distance analysis and artificial neural networks, Lloyd showed a time-distance effect (i.e., as the time series progressed, the distance between images increased) and that the past and

future brain states, retention and protention, respectively, are embedded in present brain states. Thus, as time passes, the brain changes “globally, incrementally, and monotonically” (Lloyd 2002). Previous fMRI studies of consciousness compared one state of awareness with another, assumed localization, and ignored the temporal flux (Rees et al. 2002). Lloyd's original approach proposed methodological and conceptual advantages. He was not looking for the loci of consciousness, nor did he identify regions of interest in the human brain. Rather, he analytically defined the characteristic features of consciousness, and tested whether distributed patterns of activation that mediate the phenomenal structures of consciousness could be detected. Because conscious awareness is implicated in all cognitive functions, he utilized data sets from four different studies that included a variety of cognitive tasks (target tracking, passive viewing, delayed matching, reading, and spatial working memory), stimuli (faces, objects, words, pseudo words, 2-D arrays of squares, colored circles), and motor responses (button presses and saccades). Lloyd successfully demonstrated how the neural manifestations of structures of consciousness, concepts that were proposed more than 100 years ago, can be investigated by reanalyzing existing

data sets. Inspired by his fMRI data analysis experience, Lloyd then wrote a literary book (Lloyd 2003) in which he explored paradigms of consciousness while solving the murder of a graduate student at fictional brain imaging laboratory.

The Ishai data set was also used to test novel analytic tools for fMRI research. Carlson and colleagues used these data in one of the first “decoding” analyses of fMRI data. Using linear discriminant analysis, they looked for the voxels that best predicted which object category the subject was presented with. The main results showed that distinguishing one object category from another did not depend on detecting object information from scrambled objects. Moreover, attentional demands, reflected by activation during delayed matching as compared with passive viewing, improved the ability to predict objects from scrambled objects, but did not contribute to object classification (Carlson et al. 2003). Since then, the application of linear discriminant analysis and other classification techniques became very popular in fMRI research, with the ambitious goal of decoding human mental states (Haynes and Rees 2006).

Liou et al. (2003, 2006) introduced a Bayesian method for estimating the reproducibility of activation within voxels. Given an optimal statistical threshold, a voxel was defined as ‘strongly reproducible’ if its active or inactive status was consistent in at least 90% of the sessions. The reproducibility maps were constructed for visually-responsive voxels, namely voxels that responded more to faces, houses, and chairs than to scrambled objects. When the two tasks, passive viewing and delayed matching, were compared, the density of Student’s *t* values as a function of the reproducibility of voxels, revealed a bimodal distribution. When the same comparison was performed on the delayed matching task with different stimulus formats (grayscale photographs and black and white line drawings), a unimodal distribution was found. These results suggest that subjects used different strategies to perform the viewing and matching tasks (Liou et al. 2003, 2006). As the ultimate goal of fMRI research is to understand behavior and to correlate task performance with the underlying neural mechanism, the method suggested by Liou and colleagues could reveal differential strategies used by subjects to perform a variety of cognitive tasks.

The category-related patterns of activation in the ventral stream do not seem to reflect the mere product of a hierarchical, bottom-up, ‘feature’ analysis suggested by early fMRI studies (Malach et al. 1995). Selective attention (Kastner et al. 1999) and visual imagery studies (Ishai et al. 2000b; Ishai et al. 2002) have indicated that face and object perception is also modulated by top-down effects, likely originating in parietal and frontal cortex. Mechelli and colleagues investigated the extent to which category-related responses in the ventral stream are mediated by bottom-up

and top-down effect, using conventional statistical parametric mapping (SPM) with a novel analytic approach, dynamic causal modeling (DCM), which allows, within a Bayesian framework, the assessment of effective connectivity in cortical networks (Friston et al. 2003). The original fMRI data set was re-analyzed in SPM to identify the visual response (i.e., regions that responded more to faces, houses, and chairs than to scrambled pictures), and the category-responsive regions (i.e., regions that responded maximally to one category relative to the other two). Then, for each individual subject, dynamic causal models were constructed for the face-, house- and chair-responsive regions in the ventral stream, as well as two visually responsive, but not category-specific, dorsal regions, namely V3 and parietal cortex. Interestingly, in all subjects, the category-related responses were mediated by input from V3, but not from parietal cortex. For example, when subjects were viewing chairs, the intrinsic effective connectivity from V3 to the chair-responsive region in occipital cortex was stronger than the connectivity to the face- or house-responsive regions (Mechelli et al. 2003). The Ishai data set was used again by this group to compare various dynamic causal models (Penny et al. 2004).

Encouraged by these findings, Mechelli and Ishai established a new collaboration in which they further investigated the bottom-up and top-down effects during perception and imagery of objects. In a study published in *Neuron*, Ishai and colleagues showed that visual imagery from long-term memory of faces, houses and chairs activated small subsets of the category-selective regions in ventral temporal cortex. Moreover, visual imagery activated a network of regions in parietal and frontal cortices, which likely mediate the top-down influence on category-related representations of objects stored in extrastriate cortex (Ishai et al. 2000b). The re-analysis of this data set, using both SPM and DCM, showed that during visual perception, category-selective patterns of activation in extrastriate cortex are mediated by content-sensitive forward connections from early visual areas. Furthermore, during visual imagery, category-selective activation is mediated by content-sensitive backward connections from the prefrontal cortex. Importantly, these novel findings were only revealed in the DCM analysis and were not reported in the original study, which only used conventional SPM analysis. Finally, content-unrelated connectivity between parietal cortex and the category-selective regions were found during both perception and imagery, suggesting that non-selective, top-down processes, originating in superior parietal areas, contribute to the generation of mental images, regardless of their content, and their maintenance in the “mind’s eye” (Mechelli et al. 2004).

A challenge to functional brain imaging studies of the human brain is to identify optimal models that can explain

the information flow between activated regions and move from a semi-modular approach to a more realistic, integrated, cortical networks perspective. Inspired by the novel DCM findings, Ishai and colleagues conducted effective connectivity analyses on the cortical network that mediates face perception. Face stimuli elicit activation within a distributed cortical network in the human brain. The network includes visual (“core”) regions, as well as limbic and prefrontal (“extended”) regions, which process invariant facial features and changeable aspects of faces, respectively (Haxby et al. 2000; Ishai et al. 2005). In the new study, DCM was used to investigate effective connectivity and functional organization between and within the core and the extended systems, and how experience-dependent changes alter the strength of coupling between these regions. The DCM analysis revealed that the “core,” visual system is hierarchically organized in a predominantly feed-forward fashion, and that the fusiform gyrus exerts the dominant influence on the limbic and prefrontal “extended” system. Moreover, emotional faces increased the coupling between the fusiform gyrus and the amygdala, whereas famous faces increased the connectivity between the fusiform gyrus and the orbitofrontal cortex. This study has shown for the first time content-specific dynamic alterations in the functional coupling between visual-limbic and visual-prefrontal face-responsive pathways (Fairhall and Ishai 2007).

Taken collectively, these examples of data sharing demonstrates that re-analyses of published data can reveal new findings that were not reported in the original study, provide new perspectives on visual perception, generate new predictions, and result in new collaborations and high profile publications.

Practical Neuroimaging Data Exchange

As an experiment in neuroscience databasing, the fMRIDC should be considered a major success. However, by no means is the fMRIDC approach the only framework for the sharing of neuroscience data. There are places for a variety of different data-sharing frameworks within neuroscience, each with their strengths as well as weaknesses. One common expectation from those drawing from these resources should be an adequate description of the data, how it was obtained, details of scanner protocols, from what subjects, under what experimental manipulations were under study, etc. This “meta-data” is the important contextual information needed to reproduce the results of the study with maximal fidelity. Several approaches to standardize meta-data organization and format have been discussed elsewhere (Kotter and Wanke 2005; Dameron and Musen 2007) and an elaboration of their merits is

beyond the scope of this particular article. It suffices to say that carefully constructed frameworks for meta-data description that promote efficient data exchange and re-use need to be developed and put in practice.

However, developers should be mindful that most neuroscientists are not necessarily familiar with the concept of ontologies (i.e. the formal hierarchical frameworks used to describe study meta-data) and the process for providing meta-data should be easy and transparent to them. Not all meta-data is the same and can be loosely classified in terms of micro-meta-data (e.g. the parameters of a particular image file—its dimensions, data type, bits/pixel, etc.—contained in a file header) and macro-meta-data (e.g. the collection of all parameters from a complete study). Distinctions should be made between these types and one should not be used as if it were the other. Moreover, the development of meta-data ontologies for the simple sake of there being a need for ontologies is not sound practice and the proliferation of too many ontologies leads to confusion and needless duplication of effort. Finally, there is a danger that ontologies themselves can become as complicated, if not more so, than the data they are trying to describe and/or run a risk of being so overly inclusive as to lose their ability for practical applications. We do not doubt that such meta-data frameworks are achievable with careful consideration of the needs of the users they meant to serve. With these in hand, researchers will be able to reproduce the published research findings of the original studies as closely as possible and then work to extend them in novel ways.

Evaluating Success

In one sense, the success of any database of shared data is in the degree to which researchers are willing to contribute to as well as draw from that resource. While the exact metrics used to gauge the individual or comparative success of such an effort are open to debate, one need only look to large-scale efforts such as GenBank and related databases. The success of these efforts lies not in how many web-hits they have received, the volume of data they maintain, or even how many download requests they receive. The metric of success that resonates most highly is that these efforts have given rise to a paradigm shift in the way molecular and biological sciences now conduct business. The emerging discipline of bioinformatics was in its infancy a decade ago but now forms the bridge between fields of computer science, statistics, and molecular biology. These researchers troll through volumes of data to craft new knowledge that can lead to novel, testable hypotheses about genes, their function, and dysfunction. It is possible using genomic tools to conduct examinations within and between the genomes of, for instance, several different species of

primates without laying eyes on a single monkey or having to hybridize a single nucleotide. Researchers and their students, skilled in the use of sophisticated software that can process volumes of genetic data, are exploring the vast ways that different organisms are built and do so every day. This is not simply a necessary evil, nor an adjunct to the direct collection of new data—this is their means of research, it is their science. Thus, the impression in which a biologist was someone working in a wet-lab over test tubes and Petri dishes has been augmented to include those informatics-oriented researchers using computer-based tools, looking for unseen or unappreciated relationships between genes that could lead to novel treatments for Society’s most serious diseases. The field of biology and molecular science has been revolutionized due to its willingness to openly share and mine primary data.

This has been the hope for brain imaging—that people would purposefully contribute the data from their peer-reviewed studies such that others can benefit from their efforts and broaden the scope of the original research findings. Overcoming the sociological barriers to sharing data, as demonstrated by the fMRIDC effort, was a first step.

Relation to the Published Literature

The relationship between the neuroscience databases and peer-reviewed journals is an important one. The partnership between the fMRIDC and JOCN is one that is extendable to other journals that publish neuroimaging data for their own contributions to the fMRIDC or other formal data archive. As well, the basic model of this partnership could be

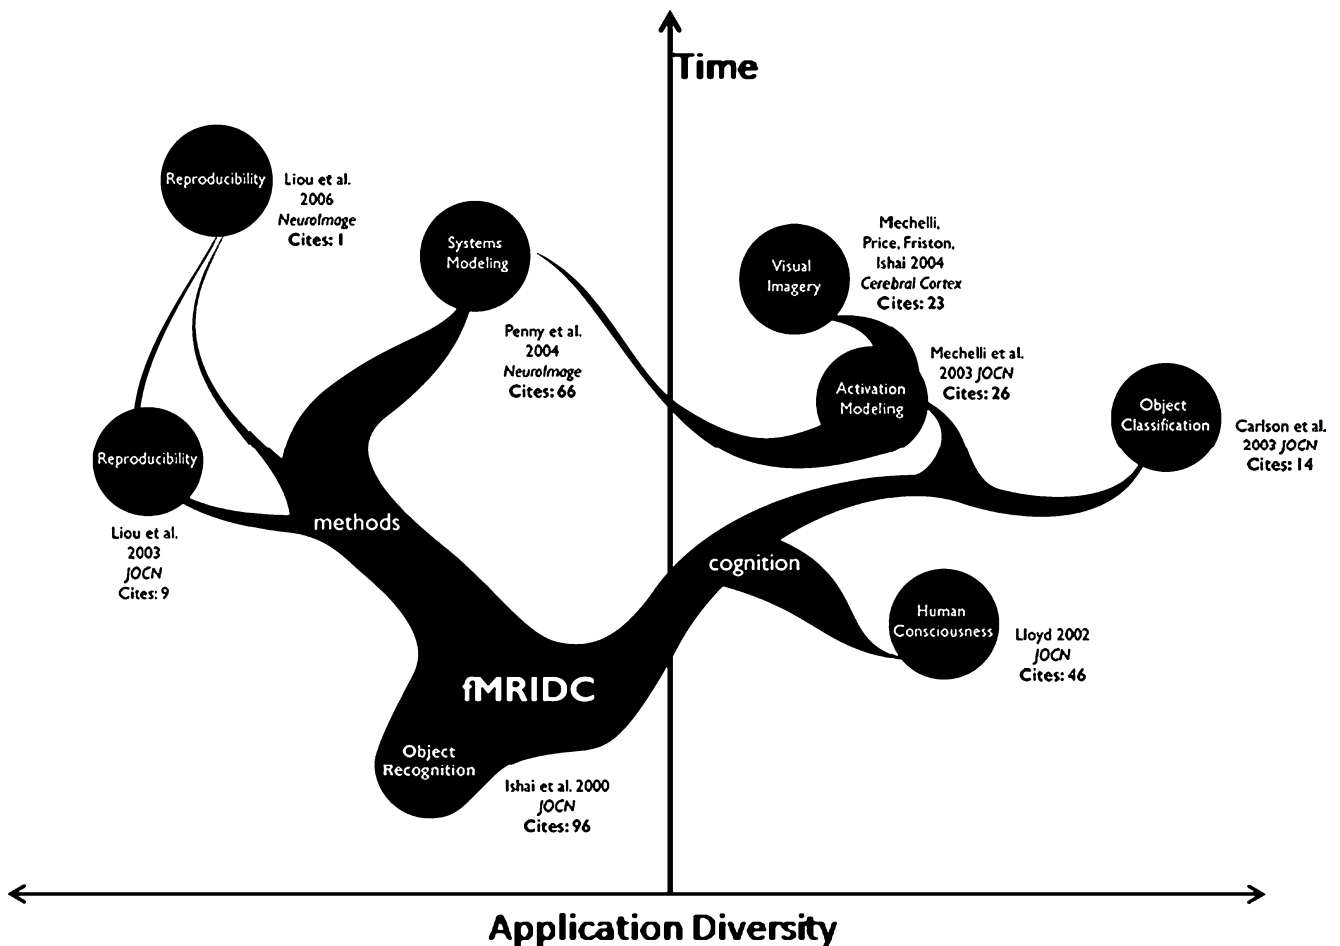


Fig. 2 Family tree of studies using the Ishai data set. The open availability of the data via the fMRI Data Center has allowed that data to inform multiple sub-disciplines within neuroimaging and foster new interactions between researchers. The figure illustrates how the various studies relate to one another, their diversity of application using that data, and the re-use over time. Also noteworthy is the number of citations being accumulated to the studies that have used these data. *Line thickness* is drawn proportional to how frequently that article has been cited. Collectively, as of April 2007, the total number

of citations to articles utilizing these data, including the Ishai et al. JOCN article, is 279 in comparison to the 306 citations for the earlier Ishai et al. PNAS article. Citation values for each article were obtained using Harzing’s Publish or Perish program (<http://www.harzing.com/pop.htm>, Version: 2.0.2673, April 2007). We believe that tracking citation incidences and rates for studies based upon shared data from published studies may be an effective way to measure the utility and long-term impact of neuroscience data sharing

adapted easily to accommodate studies and study data from other neuroscientific domains and modalities. Having software tools that facilitate and make easy the data contribution process are essential. The process and its eventual outcome represents value-added for the journal in terms of enhancing what is being made available with each published article. It also represents the opportunity for researchers in the field to obtain and examine primary data from the published literature itself, confirming results, testing new hypotheses, or exploring emerging analytic approaches. We would conjecture that the Ishai article, for example, has taken on greater importance because of having its data openly available. Figure 2 depicts how, due to the availability of the data from the Ishai JOCN article through the fMRIDC archive, its use has spread across neuroimaging domains, and the collective number of citations across these resulting articles (279 citations), is nearly as great as that for the original data's description in a higher impact journal (e.g. Ishai et al., PNAS 1999; 306 citations). The re-use and re-interpretation of data from published studies helps to inform and energize subsequent published literature. Non-published brain imaging data and archives thereof also have value but in the case of published results, these are the foundation of our field, our collective body of peer-reviewed work, and have the greatest value to scientists.

Discussion

Large-scale brain imaging data sharing and archival efforts like the fMRIDC have now begun to produce significant scientific rewards for cognitive neuroscience and brain mapping. Other databases, too, hold great promise for linking images of brain activity with other useful biological information (e.g. BrainMap, <http://www.brainmap.org/>; The Laboratory of Brain Imaging (LONI), <http://www.loni.ucla.edu>). The involvement of other researchers as well as multiple scientific communities in examining published brain imaging data must be welcomed and encouraged as this will strengthen and improve the inferences and conclusions that can be made from these data. As a result of these infrastructural and data resources, novel research, hypotheses and education using existing data can reach across scientific disciplines—engaging workers from other fields to apply sophisticated new tools for data analysis and integration.

The human scale of these projects is not insignificant, however, often requiring a dedicated curatorial staff to manage study deposition and to keep computer systems operational. Government mandates to share and archive primary research data cannot be successful unless federal agencies support the infrastructure programs and required personnel that are needed to make such sharing easy and accessible to others. The abrupt demise of the NIMH

Human Brain Project was a black eye for ongoing neuroscience databasing and neuroinformatics efforts and, as yet, has not been sufficiently replaced (Gazzaniga et al. 2006). Further progress will certainly necessitate that funding bodies continue to invest in new means for experimentation, but also in supporting study data repositories, thereby ensuring their survival as essential archival and scientific resources (Bloom 2006). In these times of funding uncertainty, perhaps government investment in these resources is a way to maximize the most science from every research dollar spent. Examples of studies that have been widely re-examined, like the Ishai article, can be used to illustrate the value these efforts have in supporting and extending scientific discovery.

We expect that, over time, because of the enormous scientific and educational benefits, the sharing of neuroimaging and other brain data will simply be an expected part of scientific publishing in high impact factor periodicals. We hope that our experiences with successful data sharing would further encourage the interests of leading scientific societies to move toward greater exchange of data. In so doing, the neurosciences will, like the biological sciences, be able to maximally leverage its collected scientific knowledge into a rich understanding of the human brain with its complex cognition.

Acknowledgements This work was supported in part by the Swiss National Science Foundation grant 3200B0-105278 and by the Swiss National Center for Competence in Research: Neural Plasticity and Repair grant to AI and by the National Institute of Mental Health grant P20 MH072580-03 to JDVH. The authors thank Ms. Amanda Hammond of the Laboratory of Neuro Imaging at UCLA for assistance with preparing Fig. 2.

References

- Aguirre, G. K., Zarahn, E., & D'Esposito, M. (1998). An area within human ventral cortex sensitive to "building" stimuli: Evidence and implications. *Neuron*, *21*, 1–20.
- Bloom, F. (2006). Prying open the black box. *Science*, *314*, 17.
- Carlson, T. A., Schrater, P., & He, S. (2003). Patterns of activity in the categorical representations of objects. *Journal of Cognitive Neuroscience*, *15*, 704–717.
- Chao, L. L., Haxby, J. V., & Martin, A. (1999). Attribute-based neural substrates in posterior temporal cortex for perceiving and knowing about objects. *Nature Neuroscience*, *2*, 913–919.
- Dameron, O., Musen, M. A. (2007). Using semantic dependencies for consistency management of an ontology of brain-cortex anatomy. *Artificial Intelligence in Medicine*, *39*(3), 217–225.
- Downing, P. E., Jiang, Y.H., Shuman, M., & Kanwisher, N. (2001). A cortical area selective for visual processing of the human body. *Science*, *293*, 2470–2473.
- Edelman, S., Grill-Spector, K., Kushnir, T., & Malach, R. (1998). Toward direct visualization of the internal shape representation space by fMRI. *Psychobiology*, *26*, 309–321.
- Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, *392*, 598–601.

- Fairhall, S. L., & Ishai, A. (2007). Effective connectivity within the distributed cortical network for face perception. *Cerebral Cortex* (in press). DOI 10.1093/cercor/bhl148.
- Friston, K. J., Harrison, L., & Penny, W. (2003). Dynamic causal modelling. *Neuroimage*, *19*, 1273–1302.
- Gazzaniga, M. S., Van Horn, J. D., Bloom, F., Shepherd, G. M., Marcus, R., Edward, E. (2006). Continuing progress in neuroinformatics. *Science*, *311*, 176.
- Haxby, J. V., Hoffman, E. A., & Gobbini, M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences*, *4*, 223–233.
- Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, *293*, 2425–2430.
- Haynes, J. D., & Rees, G. (2006). Decoding mental states from brain activity in humans. *Nature Reviews. Neuroscience*, *7*, 523–534.
- Hazeltine, E., Poldrack, R., & Gabrieli, J. D. (2000). Neural activation during response competition. *Journal of Cognitive Neuroscience*, *12*, 118–129.
- Ishai, A., Ungerleider, L. G., Martin, A., Schouten, J., & Haxby, J. V. (1999). Distributed representation of objects in the human ventral visual pathway. *Proceedings of the National Academy of Sciences of the United States of America*, *96*, 9379–9384.
- Ishai, A., Ungerleider, L. G., Martin, A., & Haxby, J. V. (2000a). The representation of objects in the human occipital and temporal cortex. *Journal of Cognitive Neuroscience*, *12*(supp 2), 35–51.
- Ishai, A., Ungerleider, L. G., & Haxby, J. V. (2000b). Distributed neural systems for the generation of visual images. *Neuron*, *28*, 979–990.
- Ishai, A., Haxby, J. V., & Ungerleider, L. G. (2002). Visual imagery of famous faces: effects of memory and attention revealed by fMRI. *Neuroimage*, *17*, 1729–1741.
- Ishai, A., Schmidt, C. F., & Boesiger, P. (2005). Face perception is mediated by a distributed cortical network. *Brain Research Bulletin*, *67*, 87–93.
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, *17*, 4302–4311.
- Kastner, S., Pinsk, M. A., De Weerd, P., Desimone, R., & Ungerleider, L. G. (1999). Increased activity in human visual cortex in the absence of visual stimulation. *Neuron*, *22*, 751–761.
- Kennedy, D. N., & Haselgrove, C. (2006). The internet analysis tools registry: A public resource for image analysis. *Neuroinformatics*, *4*, 263–270.
- Kotter, R., & Wanke, E. (2005). Mapping brains without coordinates. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *360*, 751–766.
- Liou, M., Su, H.-R., Lee, J.-D., & Cheng, P. E. (2003). Bridging functional MR images and scientific inference: Reproducibility maps. *Journal of Cognitive Neuroscience*, *15*, 935–945.
- Liou, M., Su, H. R., Lee, J. D., Aston, J. A., Tsai, A. C., & Cheng, P. E. (2006). A method for generating reproducible evidence in fMRI studies. *Neuroimage*, *29*, 383–395.
- Lloyd, D. (2002). Functional fMRI and the study of human consciousness. *Journal of Cognitive Neuroscience*, *14*, 818–831.
- Lloyd, D. (2003). *Radiant cool*. Cambridge, MA: MIT Press.
- Malach, R., Reppas, J. B., Benson, R. R., Kwong, K. K., Jiang, H., Kennedy, W. A. et al. (1995). Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *92*, 8135–8139.
- Mechelli, A., Friston, K. J., & Price, C. J. (2000). The effects of presentation rate during word and pseudoword reading: A comparison of PET and fMRI. *Journal of Cognitive Neuroscience*, *12*, 145–156.
- Mechelli, A., Price, C. J., Noppeney, U., & Friston, K. J. (2003). A dynamic causal modelling study on category effects: Bottom-up or top-down mediation? *Journal of Cognitive Neuroscience*, *15*, 925–934.
- Mechelli, A., Price, C. J., Friston, K. J., & Ishai, A. (2004). Where bottom-up meets top-down: Neuronal interactions during perception and imagery. *Cerebral Cortex*, *14*, 1256–1265.
- Penny, W. D., Stephan, K. E., Mechelli, A., & Friston, K. J. (2004). Comparing dynamic causal models. *Neuroimage*, *22*, 1157–1172.
- Polk, T. A., & Farah, M. J. (1998). The neural development and organization of letter recognition: Evidence from functional neuroimaging, computational modeling, and behavioral studies. *PNAS USA*, *95*, 847–852.
- Postle, B. R., Berger, J. S., Taich, A. M., & D'Esposito, M. (2000). Activity in human frontal cortex associated with spatial working memory and saccadic behavior. *Journal of Cognitive Neuroscience*, *12*, 2–14.
- Rees, G., Kreiman, G., & Koch, C. (2002). Neural correlates of consciousness in humans. *Nature Reviews. Neuroscience*, *3*, 261–270.
- Shepherd, G. M. (2002). Supporting databases for neuroscience research. *Journal of Neuroscience*, *22*, 1497.
- Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annual Review of Neuroscience*, *19*, 109–139.
- Van Horn, J. D., Grethe, J. S., Kostelec, P., Woodward, J. B., Aslam, J. A., Rus, D., et al. (2001). The Functional Magnetic Resonance Imaging Data Center (fMRIDC): The challenges and rewards of large-scale databasing of neuroimaging studies. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *356*, 1323–1339.
- Van Horn, J. D. (2002). Maturing as a science: The New Perspectives in fMRI research award. *Journal of Cognitive Neuroscience*, *14*, 817.